

**THE LARGE HADRON COLLIDER AND THE EXPERIMENTS FOR RUN 3 —
ACCELERATOR AND EXPERIMENTS FOR LHC RUN3****Development of the CMS detector for the CERN LHC
Run 3****The CMS collaboration**

E-mail: cms-publication-committee-chair@cern.ch

ABSTRACT: Since the initial data taking of the CERN LHC, the CMS experiment has undergone substantial upgrades and improvements. This paper discusses the CMS detector as it is configured for the third data-taking period of the CERN LHC, Run 3, which started in 2022. The entire silicon pixel tracking detector was replaced. A new powering system for the superconducting solenoid was installed. The electronics of the hadron calorimeter was upgraded. All the muon electronic systems were upgraded, and new muon detector stations were added, including a gas electron multiplier detector. The precision proton spectrometer was upgraded. The dedicated luminosity detectors and the beam loss monitor were refurbished. Substantial improvements to the trigger, data acquisition, software, and computing systems were also implemented, including a new hybrid CPU/GPU farm for the high-level trigger.

KEYWORDS: Calorimeters; Large detector systems for particle and astroparticle physics; Particle tracking detectors

ARXIV EPRINT: [2309.05466](https://arxiv.org/abs/2309.05466)



Contents

1	Introduction	1
2	Solenoid magnet	6
3	Inner tracking system	8
3.1	Pixel detector	8
3.1.1	Detector design	8
3.1.2	Silicon modules	10
3.1.3	Mechanics	14
3.1.4	Services	16
3.1.5	Detector operation	18
3.1.6	Performance of the pixel tracker	18
3.2	Strip detector	19
3.2.1	Detector description	19
3.2.2	Performance of the strip tracker	21
4	Electromagnetic calorimeter	24
4.1	Experimental challenges	24
4.2	Response monitoring	25
4.3	Noise evolution	25
4.4	Signal reconstruction	26
4.5	Trigger	28
4.6	Channel calibration and synchronization	29
4.6.1	Intercalibration precision	30
4.7	Run 2 operations summary	31
4.8	Run 2 performance	32
4.9	Preparation for Run 3	33
4.9.1	Safety and control system	33
4.9.2	Trigger	33
4.10	Calibration	34
5	Hadron calorimeter	34
5.1	The hadron calorimeter in Run 1 and Run 2	34
5.2	Upgrades	36
5.2.1	Motivation and overview of the upgrade	36
5.2.2	HB/HE/HO photodetector upgrade	39
5.2.3	Readout box	41
5.2.4	Photodetector control	42
5.2.5	Optical decoder unit	43
5.2.6	Frontend readout card	45

5.2.7	Slow and fast-control systems	45
5.2.8	Backend electronics: readout	47
5.2.9	Voltage source upgrades	47
5.2.10	Photodetector and system calibration instrumentation	49
5.2.11	HF upgrade	49
5.3	Trigger	51
5.4	System and beam tests	52
5.5	Performance	53
5.5.1	Endcap photodetector performance in Run 2	53
5.5.2	HCAL performance in Run 2	55
5.5.3	HF performance in Run 2	56
6	Muon system	57
6.1	Drift tubes	59
6.1.1	General description	59
6.1.2	Phase 1 upgrades of the DT electronics	60
6.1.3	Detector longevity for Run 3 and beyond	68
6.2	Cathode strip chambers	73
6.2.1	General description	73
6.2.2	Upgrade of the CSC system since Run 1	77
6.2.3	Longevity studies	80
6.3	Resistive plate chambers	83
6.3.1	General description	83
6.3.2	RPC system upgrades since Run 1	84
6.3.3	RPC system longevity	85
6.3.4	Changes to the RPC system in LS2	95
6.4	Gas electron multiplier chambers	97
6.4.1	Motivation and general description	97
6.4.2	Technical design	99
6.4.3	Gas system	102
6.4.4	Electronics	103
6.4.5	CSC/GEM trigger for Run 3	105
6.4.6	Detector control system	107
6.4.7	Chamber assembly and installation	107
6.4.8	Preliminary commissioning results	108
7	Precision proton spectrometer	109
7.1	Forward protons and the roman pot system	110
7.2	Tracking detectors	112
7.2.1	Detector units	112
7.2.2	Readout electronics	114
7.2.3	Support structure and internal motion system	116
7.3	Timing detectors	119

7.3.1	Detector modules	119
7.3.2	Readout electronics	121
7.3.3	Reference clock distribution	122
7.4	Data acquisition and detector control	123
7.4.1	Pixel detectors	123
7.4.2	Strip and timing detectors	125
7.4.3	Detector control system	126
7.5	Roman pot insertion and running scenarios	126
7.5.1	TCT collimator and roman pot insertion scheme	126
7.5.2	TCL collimator insertion scheme	129
8	Luminosity and beam conditions	131
8.1	Real-time bunch-by-bunch luminometers	132
8.1.1	Pixel luminosity telescope (PLT)	132
8.1.2	Fast beam conditions monitor (BCM1F)	133
8.1.3	The forward hadron calorimeter (HF)	135
8.2	Additional luminometers	135
8.2.1	Tracker	135
8.2.2	Muon system	136
8.2.3	Z boson counting	137
8.3	Beam monitoring instrumentation	138
8.3.1	Beam-halo monitor (BHM)	138
8.3.2	Beam-condition monitor for beam losses (BCML)	139
8.4	Radiation instrumentation and simulation	141
8.4.1	Radiation monitoring	141
8.4.2	Radiation simulation	142
8.5	The BRIL online data acquisition and monitoring	143
9	Data acquisition system	145
9.1	Scope	145
9.2	Evolution	147
9.3	Subdetector readout interface	147
9.4	Event builder	149
9.4.1	FED builder	150
9.4.2	Core event builder	151
9.4.3	Performance	152
9.4.4	Event builder load balancing	154
9.5	Event filter	154
9.5.1	File-based filter farm	154
9.5.2	The HLT software infrastructure	155
9.5.3	The HLT menu and output data streams	155
9.5.4	Monitoring	156
9.5.5	Evolution of the HLT farm	157

9.6	Storage and transfer system	158
9.7	Trigger throttling system	159
9.8	Trigger control and distribution system	160
9.9	Networking and computing infrastructure	162
9.9.1	Virtualization	162
9.9.2	Online cloud	162
9.10	Software, control, and monitoring of the DAQ	163
9.11	MiniDAQ	164
10	Level-1 trigger	165
10.1	Calorimeter trigger	165
10.1.1	Calorimeter layer 1 trigger	165
10.1.2	Calorimeter layer 2 trigger	166
10.2	Muon trigger	168
10.2.1	Barrel muon track finder (BMTF)	169
10.2.2	Overlap muon track finder (OMTF)	169
10.2.3	Endcap muon track finder (EMTF)	170
10.3	Global trigger	171
10.4	Trigger menu	173
10.4.1	Trigger seeds for displaced muons	174
10.4.2	Trigger seeds using new kinematic variables	174
10.4.3	Run 3 trigger rates	175
10.5	Online software and monitoring	175
10.6	The L1 scouting system	177
10.6.1	Architecture of the L1 scouting system	177
10.6.2	Applications of the L1 scouting system	178
11	High-level trigger	181
11.1	Overview	181
11.2	HLT reconstruction	182
11.2.1	Tracking	182
11.2.2	Muons	183
11.2.3	Electrons and photons	184
11.2.4	Tau leptons	185
11.2.5	Jets and global energy sums	185
11.2.6	b jet tagging	186
11.2.7	New HLT paths for long-lived particles	186
11.3	Run 3 HLT menu composition, rates, and timing	188
11.4	Data scouting at the HLT	190
11.5	Data parking	192
11.6	Heavy ion physics	193

12 Offline software and computing	195
12.1 Overview	195
12.2 Detector simulation	195
12.3 Event reconstruction	196
12.4 Computer architectures and platforms	197
12.5 Application framework	197
12.5.1 Multithreading	197
12.5.2 Offloading to accelerators	198
12.5.3 Geometry	198
12.6 Data formats and processing	199
12.6.1 Premixing	200
12.7 Computing centers	201
12.7.1 High performance computing (HPC)	202
12.7.2 Data archive at CERN	203
12.8 Computing services	203
12.8.1 Data management	203
12.8.2 Data transfer protocols	204
12.8.3 Central processing and production	204
12.8.4 Workload management system	205
12.8.5 Distributed analysis	206
12.8.6 DBS database	206
12.8.7 Web services and security	206
12.8.8 Monitoring and analytics	207
13 Summary	207
The CMS collaboration	231

1 Introduction

The CMS detector [1] is a large, multipurpose apparatus located at the CERN LHC. The detector was designed for the study of a variety of physics phenomena, including the search for the Higgs boson, which was discovered in 2012 [2–4], and the measurement of its properties, the exploration of the electroweak sector and vector boson scattering, precision measurements of standard model (SM) particles and interactions, flavor physics, heavy-ion physics, and searches for new physics beyond the SM.

The LHC Run 1 started in 2009 and, until the end of 2012, proton-proton (pp) collision data corresponding to a total integrated luminosity of about 30 fb^{-1} were delivered at center-of-mass energies of 7 and 8 TeV. In addition, CMS successfully recorded data from high-energy lead-lead collisions. After a first long shutdown, referred to as long shutdown 1 (LS1), the second data-taking period, Run 2, followed in 2015–2018 at an energy of 13 TeV, during which an integrated luminosity of about 165 fb^{-1} was delivered with peak instantaneous luminosities up to $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, twice

the original LHC design value. During LS1 and Run 2, the first generation of detector upgrades, referred to as the Phase 1 upgrade program, was implemented.

After the second long shutdown (LS2, 2019–2021), the LHC Run 3 was started in 2022 and is expected to deliver about 250 fb^{-1} of integrated luminosity. In Run 3, the center-of-mass energy for pp collisions is 13.6 TeV. During the third long shutdown (LS3), scheduled to start in 2026, CMS will undergo a major upgrade program, referred to as the Phase 2 upgrade, in preparation for the data taking at the High-Luminosity LHC (HL-LHC), designed to deliver instantaneous luminosities up to $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at a pp center-of-mass energy of 14 TeV. At the end of the HL-LHC, the total integrated luminosity is expected to be 3000 fb^{-1} . The Phase 2 upgrade includes a new inner tracking system and a new endcap calorimeter, as well as substantial improvements for most other subsystems of CMS. Upgrades are also in preparation in almost all other areas of CMS. In this paper, we present the various upgrades of the CMS detector since Run 1 that are designed to optimize the detector for sustained or improved performance at increased luminosity and energy.

The CMS detector has an overall length of 22 m, a diameter of 15 m, and weighs 14 000 tons. A schematic view is shown in figure 1. The detector is nearly hermetic, designed to trigger on [5, 6] and identify electrons, muons, photons, and (charged and neutral) hadrons [7–10]. The central feature of the CMS experiment is a superconducting solenoid of 6 m internal diameter and 12.5 m length that provides a magnetic field of 3.8 T with a stored energy of 2.2 GJ. Within the magnetic volume are a silicon pixel and strip tracker, a lead tungstate crystal electromagnetic calorimeter (ECAL), and a brass and scintillator hadron calorimeter (HCAL), each composed of a barrel and two endcap sections. Forward calorimeters extend the pseudorapidity (η) coverage provided by the barrel and endcap detectors. Muons are measured in gas-ionization detectors embedded in the steel flux-return yoke outside the solenoid.

Events of interest are selected using a two-tier trigger system. The first-level (L1) trigger is composed of custom hardware processors and uses information from the calorimeters and muon detectors to select events at a rate of about 100 kHz within a latency of about $4 \mu\text{s}$ [5]. The second level, known as the high-level trigger (HLT), consists of a cluster of commercial processors running a version of the full event reconstruction software optimized for fast processing. It was originally designed to reduce the event rate to around 1 kHz before data storage [6]. During Run 3, the L1 trigger and HLT operate at typical output rates of 110 kHz and 5 kHz, respectively.

A full description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, is reported in ref. [1]. In the remaining part of this section, the CMS detector components are briefly introduced.

In section 2, the CMS solenoid magnet is described.

The inner tracking system (section 3) is used to measure the trajectories of charged particles produced in the collisions at the LHC. It is located in the innermost part of the CMS detector, closest to the interaction point. Prior to the Phase 1 upgrade, the pixel detector had three barrel layers and two disks in each endcap. In its current form, the pixel detector is composed of four barrel layers and three disks of silicon sensors on each side of the interaction point, with a total of 124 million readout channels. During LS2, the innermost barrel layer was replaced to ensure optimal performance until the end of Run 3. The strip tracker comprises ten layers of silicon strip sensors in the barrel, arranged in a cylindrical shape, and nine disks in the endcaps on each side of the detector. The strip sensors are segmented into long, thin strips, which are used to measure the trajectories of the particles and




Figure 1. Schematic drawing of the CMS detector. Reproduced from [11]. CC BY 4.0.

provide a hit resolution of $20 \mu\text{m}$ for charged particles that cross the sensor perpendicularly. The tracker is designed to have excellent momentum resolution and tracking efficiency. It can detect and track particles with transverse momentum p_T as low as 50 MeV within the range $|\eta| < 2.5$. Tracks with a momentum around 100 GeV in the central region of the detector have an impact parameter resolution of about $10 \mu\text{m}$, and a transverse momentum resolution near 1%.

The electromagnetic calorimeter (section 4) is made of 75 848 lead tungstate (PbWO_4) crystals: 61 200 crystals are located in the barrel (EB) and 7324 in each of the endcaps (EE) that provide a pseudorapidity coverage of $|\eta| < 3$. The lead tungstate crystals have a depth of about 23 cm , corresponding to about 25 radiation lengths X_0 . Preshower detectors, consisting of two planes of silicon sensors interleaved with a total of $3X_0$ of lead, are located in front of each EE detector. The ECAL is designed to identify electrons and photons, and measure their positions and energies, which are reconstructed from energy deposits using algorithms that constrain the clusters to the size and shape expected for electrons or photons. The ECAL also contributes significantly to the reconstruction of jets and missing transverse momentum. The electron momentum is estimated by combining the energy measurement in the ECAL with the momentum measurement in the tracker. The momentum resolution for electrons with $p_T \approx 45 \text{ GeV}$ from $Z \rightarrow e^+e^-$ decays ranges from 1.6% to 5%. It is generally better in the barrel region than in the endcaps, and also depends on the bremsstrahlung energy emitted by the electron as it traverses the material in front of the ECAL. The ECAL also provides information on the arrival time of the electrons and photons which can be used in physics analyses, such as searches for long-lived particles.

The hadron calorimeter (section 5) is designed to measure the energy of charged and neutral hadrons. It contributes to the identification of hadrons and the measurement of their properties. It also aids in the reconstruction of jets and missing transverse momentum, and the identification of electrons and photons. The HCAL comprises four subdetectors: the hadron barrel (HB), hadron endcap (HE), hadron outer (HO), and hadron forward (HF) calorimeters. The HB and HE are located inside the solenoid magnet of the CMS detector and surround the ECAL. They cover the pseudorapidity ranges $|\eta| < 1.392$ and $1.305 < |\eta| < 3.0$, respectively, and are made of layers of brass plates interleaved with layers of scintillating tiles. The HF, constructed from steel and quartz fibers, is located outside the solenoid at ± 11.5 m from the collision point, and covers the $3.0 < |\eta| < 5.2$ range. Finally, the HO, made of plastic scintillator and the first layer of the barrel flux return covers the $|\eta| < 1.26$ range. The HCAL is designed to have a good hermeticity, with the ability to detect hadrons in nearly the full 4π solid angle. The Phase 1 upgrade of HCAL was installed in stages from 2016–2019. In HB and HE, the hybrid photodiode detectors were replaced with silicon photomultipliers, which reduced anomalous signals, improved radiation tolerance, and allowed for finer longitudinal readout segmentation. The HF photomultiplier tubes were also upgraded to reduce anomalous signals. The readout electronics were upgraded to support the increased channel count, improve the precision, and add signal timing information. When combining information from the entire CMS detector, the jet energy resolution typically amounts to 15–20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV.

The muon detectors (section 6) are used in the identification of muons and the measurement of their momenta. The muon system comprises four subsystems: the drift tubes (DTs), the cathode strip chambers (CSCs), the resistive-plate chambers (RPCs), and the recently added gas electron multiplier (GEM) detector. Altogether, the CMS muon detectors comprise almost one million electronic channels.

The DTs consist of chambers formed by multiple layers of long rectangular tubes that are filled with an Ar and CO₂ gas mixture. An anode wire is located at the center of each tube, whereas cathode and field-shaping strips are positioned on its borders. They create an electric field that induces an almost uniform drift of ionization electrons produced by charged particles traversing the gas. The charged-particle trajectory is determined from the arrival time of the currents generated on the anode wires of the readout.

The CSCs are made of layers of proportional wire chambers with orthogonal cathode strips and are operated with a gas mixture of Ar, CO₂, and CF₄. Signals are generated on both anode wires and cathode strips. The finely segmented cathode strips and fast readout electronics provide good timing and spatial resolution to trigger on and identify muons.

The RPCs comprise two detecting layers of high-pressure laminate plates that are separated by a thin gap filled with a gas mixture of C₂H₂F₄, i-C₄H₁₀, and SF₆. The electronic readout strips are located between the two layers, and the high voltage is applied to high-conductivity electrodes coated on each plate. The detectors are operated in avalanche mode to cope with the high background rates. Due to their excellent time resolution, they ensure a precise bunch-crossing assignment for muons at the trigger level.

The key feature of the GEM is a foil consisting of a perforated insulating polymer surrounded on the top and bottom by conductors. A voltage difference is applied on the foils producing a strong electric field in the holes. The GEM is operated with a gas mixture of Ar and CO₂. When the gas volume is ionized, electrons are accelerated through the holes and read out on thinly separated strips.

This structure allows for high amplification factors with modest voltages that provide good timing and spatial resolution, and can be operated at high rates.

The precision proton spectrometer (PPS) (section 7) is designed to detect protons scattered at very small angles in interactions where the protons remain intact and only a small fraction of their initial energy goes into the production of particles at small rapidity. In such events, the reconstruction of the kinematic properties of the protons uses their energy loss to determine the invariant mass of the system produced in the quasi-elastic collision. The PPS detector includes tracking and timing stations, which are located inside the LHC tunnel on both ends of the CMS detector about 200 m from the CMS interaction point. Precision tracking and timing is provided by silicon pixel and diamond detectors, respectively. The detectors are enclosed in movable stations, referred to as “roman pots”, within which the detectors can be positioned as close as a few millimeters from the proton beam. The PPS first started in Run 2 as a joint project (CT-PPS) with the TOTEM Collaboration [12]. The initial PPS system consisted of two tracking and one timing station on each side. For Run 3, the PPS was upgraded for improved efficiency and precision with an additional timing station on each side.

The beam radiation instrumentation and luminosity (BRIL) system (section 8) comprises various detectors that measure the instantaneous luminosity and monitor in real time the beam-induced background, beam losses, and timing. Three luminosity detector systems provide robust bunch-by-bunch luminosity measurements in real time. They are: (i) the fast beam condition monitor (BCM1F), which counts hits in silicon pad diodes; (ii) the pixel luminosity telescope (PLT), which counts triple coincidences; and (iii) the HF calorimeter. The HF is instrumented with a dedicated readout for the real-time luminosity measurement and provides hit-tower counting (HFOC) and transverse-energy sums (HFET). The beam condition monitor for losses (BCML), using diamond and sapphire sensors, provides protection against catastrophic beam loss and is part of the LHC beam-abort system. The beam pickup timing device (BPTX) provides logical beam signals to the L1 trigger system. The BRIL system includes the measurement of radiation in the experimental cavern. The measurements are complemented by detailed simulations using the CMS radiation simulation applications.

The data acquisition (DAQ) system (section 9) is responsible for: the readout of all detector data for events accepted by the L1 trigger; the building of complete events from subdetector event fragments; the operation of the filter farm cluster running the HLT; and the transport of event data selected by it to the permanent storage in the Tier 0 computing center. The DAQ consists of: custom-built electronics reading out event fragments; a data-concentrator network transporting the fragments to the surface; a cluster of readout and event-building servers interconnected via the event-building network; the filter-farm cluster of multicore servers connected by the data network running the HLT software; a distributed storage system where event data selected by HLT filtering are buffered; and a transfer system connected to the Tier 0 center via a high-speed network. The DAQ also includes the trigger control and distribution system (TCDS), which distributes timing to the trigger and subdetector electronics, and implements trigger control logic as well as the trigger throttling system (TTS).

The L1 trigger (section 10) consists of electronics responsible for making a fast selection of events based on the presence of high-energy particles in the detector. The L1 trigger receives energy and position information, so-called trigger primitives (TPs), from the calorimeters and the muon detectors. The TPs are evaluated by a trigger processor, which is composed of custom-built electronics and field programmable gate array (FPGA) devices that perform the trigger decision based on a set of predefined trigger algorithms. The L1 trigger operates at trigger rates of about

110 kHz. During LS2, the L1 trigger was upgraded to also process TPs that are designed to select long-lived particles.

The HLT (section 11) is a software-based system in which the full event information is used to select events of interest based on their physics content. The HLT is implemented as a parallel computing system that processes the event data in real time. Since the start of Run 3, the HLT makes use of graphical processing units (GPUs) in the trigger filter farm. The GPUs facilitate the offloading of specific parts of the reconstruction algorithms, e.g., tracking based on the pixel detector, as well as parts of the calorimeter reconstruction. The use of GPUs has led to a substantial reduction of the overall event processing time. With the performance improvements for Run 3, HLT-reconstructed analysis data, referred to as HLT scouting data, are recorded at a rate of about 30 kHz. In parallel, the storage rate of normal triggers was increased to about 2 kHz. Furthermore, the system also stores extra data samples, the so-called parking data sets, at a rate of around 3 kHz. The parking events will only be reconstructed by the offline computing infrastructure at a later time, when the resources will not be needed for the core activities. Other HLT reconstruction improvements in the areas of muon tracking, b jet tagging, and tau lepton reconstruction were also implemented for Run 3.

The offline computing system (section 12) has three key roles: to process the recorded data; to generate sufficiently large Monte Carlo simulation samples based on theoretical models and detector response modeling; and to facilitate physics data analysis performed at the CMS institutes around the world. The CMS data and simulation samples are stored and processed in a globally distributed network of centers, using an ever-growing array of heterogeneous resources. Continuous improvements are made in software and computing performance. Most notably, multithreaded processing and offloading to GPU resources have been introduced.

2 Solenoid magnet

The superconducting solenoid magnet provides a magnetic field of 3.8 T, and forms the center piece of the CMS experiment. A picture of the open CMS detector with visible magnet cryostat is shown in figure 2.

The original plan for the magnet during LS2 had been to turn it off, but to keep it cooled at 4 K. The plan was changed substantially because of a water leak in the experimental cavern inside a diffusion pump of the magnet cryostat discovered during a routine check. To intervene without putting the magnet at risk, it was decided to warm the magnet up to room temperature.

The procedure took place during the Covid-19 lockdown at CERN in April and May 2020. After an outgassing period, the vacuum volume was brought back to atmospheric pressure, and the diffusion pumps were removed and replaced. The vacuum circuit was also fully cleaned, and modifications were implemented for easier access and improved backup capabilities, with new valves and flanges allowing the connection of backup vacuum pumps if needed, as illustrated in the picture in figure 3 (left).

In parallel to this repair, the new free wheel thyristor (FWT) system [14] was installed on the powering circuit of the magnet, visible in figure 3 (right). The FWT bypasses the power converter in a closed loop in case the converter is in a faulty state, e.g., in the event of a power failure or a lack of cooling, thus avoiding a slow discharge to zero current. The FWT contributes to increasing the magnet's lifetime and the operational time at nominal field.




Figure 2. The solenoid magnet cryostat within the open CMS detector. Reproduced with permission from [13].




Figure 3. Photographs of a part of the new vacuum pumping circuit (left) and of the new CMS FWT system installed in the CMS service cavern (right).

As shown in figure 4, a full discharge followed by a ramp up takes about eight hours. A partial ramp down to 9.5 kA, corresponding to 2 T at the pp interaction point, is implemented in case of a cryogenics stop. The idle state of the power converter at constant current, indicated in figure 4 as a dashed line, is devised to provide the time needed for the reconnection of the cold box and refill of the liquid helium dewar, this way avoiding the risk of a fast discharge caused by helium flow fluctuations which may trigger a quench on the power leads and superconducting busbars. Also represented in figure 4 is the “free wheel” mode, when the FWT is triggered, typically after a power glitch, followed by an idle state period after the power converter restart and before the ramp up to nominal field can be resumed.




Figure 4. The CMS magnet current ramp and discharge modes representing the various magnet operation procedures and their duration. A current of 18164 A corresponds to a magnetic field of 3.8 T.

The control systems of the magnet and insulation vacuum were also fully upgraded during LS2. New programmable logic controllers were installed, new control electronics for both the FWT and the renewed vacuum pumping circuit were integrated, and the magnetic measurement system, using Hall probes, flux-loops, and nuclear magnetic resonance devices, was consolidated. The cryogenics system inside the cold box was improved by installation of a large filter with a reduced mesh to limit the recurrent clogging of the turbine filters as much as possible.

The magnet was successfully commissioned in September 2021, just ahead of the LHC pilot beam run, when it was operated at full magnetic field for two weeks with its upgraded powering system and repaired vacuum system. In March 2022, the magnet was ramped up to its nominal field of 3.8 T and declared ready for Run 3.

3 Inner tracking system

3.1 Pixel detector

The silicon pixel detector is the innermost part of the CMS inner tracking system. It provides three-dimensional space points close to the LHC collision point, which allow for high precision tracking and vertex reconstruction.

3.1.1 Detector design

The first CMS pixel detector [1], installed in 2008, consisted of three barrel layers at radii of 44, 73, and 102 mm and two endcap disks on each end at distances of 345 and 465 mm from the detector center. It provided three-point tracking for charged particles and performed very well during Run 1. However, already in Run 1 the instantaneous luminosity delivered by the LHC exceeded the design value of $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, which resulted in a pixel detector readout inefficiency. In order to maintain good tracking performance, this pixel detector was replaced with a more efficient and robust four-point tracking system. In addition, the radius of the beam pipe was reduced in 2014 from 30 to 23 mm, which allowed the innermost pixel layer to be placed closer to the interaction point. The improved pixel detector was installed at the beginning of 2017.

The new detector, referred to as the Phase 1 pixel detector [15], consists of four barrel layers (L1–L4) at radii of 29, 68, 109, and 160 mm, and three disks (D1–D3) on each end at distances of 291, 396, and 516 mm from the center of the detector. The layouts of the two detectors, the original and the upgraded one, are compared in figure 5. The new layout provides four-hit coverage, instead of three, for tracks up to an absolute pseudorapidity of 3.0. The details of the Phase 1 pixel detector design and construction have been published in ref. [16].




Figure 5. Longitudinal view of the Phase 1-upgraded pixel detector compared to the original detector layout. Reproduced from [16]. The Author(s). CC BY 4.0.

In addition to providing more tracking planes, the new detector contains several other improvements: DC-DC power converters are used to supply the necessary current while reusing the existing cables; a two-phase CO₂ cooling system is used; a 400 Mb/s digital readout system is implemented instead of the 40 MHz analog one; and the readout chips are modified to handle higher data rates.

Table 1 shows the main parameters of the Phase 1 pixel detector. The detector consists of two parts, the pixel barrel (BPIX) and the pixel forward disks (FPIX), which are mechanically and electrically independent. The BPIX detector consists of two 540 mm-long half-barrels divided in the y - z plane. The FPIX detector consists of twelve half-disks with radii ranging from 45 to 161 mm. Each half-disk is further divided into two rings of modules. The basic building block of both the BPIX and FPIX is a silicon sensor module comprised of a sensor with 160×416 pixels and a size of $100 \times 150 \mu\text{m}^2$, connected to 16 readout chips (ROCs). In total, 1184 and 672 modules are used in the BPIX and FPIX, respectively. The entire Phase 1 pixel detector comprises a total of 124 million readout channels.

The sensor modules are mounted on light-weight mechanical structures, with thin-walled stainless steel tubes used for the CO₂ evaporative cooling (section 3.1.4). Carbon fiber and graphite materials with high thermal conductivity are incorporated into the detector mechanical structures (section 3.1.3). Both the BPIX and FPIX are connected to four service half-cylinders. They host the auxiliary electronics for readout and powering.

The installation of the Phase 1 pixel detector took place during the extended year-end technical stop of the LHC in 2016–2017. The new detector performed successfully during the Run 2 data-taking period in 2017–2018. However, since the detector installation in 2017 two major interventions

Table 1. Summary of the average radius and z position, as well as the number of modules for the four BPIX layers and six FPIX rings for the Phase 1 pixel detector.

BPIX			
Layer	Radius [mm]	z position [mm]	Number of modules
L1	29	−270 to +270	96
L2	68	−270 to +270	224
L3	109	−270 to +270	352
L4	160	−270 to +270	512

FPIX			
Disk	Radius [mm]	z position [mm]	Number of modules
D1 inner ring	45–110	±338	88
D1 outer ring	96–161	±309	136
D2 inner ring	45–110	±413	88
D2 outer ring	96–161	±384	136
D3 inner ring	45–110	±508	88
D3 outer ring	96–161	±479	136

took place. Due to an unexpected failure (section 3.1.4), all 1216 DC-DC converters had to be replaced in the LHC year-end technical stop 2017–2018. The second intervention was done during LS2 and involved the replacement of the layer-1 modules and, again, all the DC-DC converters. The new layer-1 modules, in addition to the planned replacement of the radiation-damaged silicon sensors, also included new versions of the readout ASICs, which fixed some of the shortcomings observed in the earlier versions, described in section 3.1.2. The LS2 was also used to repair several faulty optical fibers and bad power connections, upgrade the electronic boards in the FPIX system, and fix broken FPIX cooling inlets.

With the innermost layer placed at a radius of 29 mm from the beam line, the modules in this region are exposed to very high radiation doses and hit rates, as shown in table 2: the radiation fluence for L1 with 300 fb^{-1} is $2.2 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$, corresponding to the operational limit of the installed system [16]. To ensure that the inner layer remains fully operational throughout all of Run 3, as planned from the beginning, the innermost BPIX layer was replaced during LS2 in 2019–2021.

3.1.2 Silicon modules

Schematic drawings of the Phase 1 pixel detector modules are shown in figure 6. A module consists of a $18.6 \times 66.6 \text{ mm}^2$ silicon sensor that is bump-bonded to 2×8 ROCs. Each ROC has 80×52 rectangular pixels with a size of $100 \times 150 \mu\text{m}^2$, the same as in the original pixel detector. A high-density interconnect (HDI) flex printed circuit is glued to the sensor and wire-bonded to the 16 ROCs. A token bit manager chip (TBM) is mounted on top of the HDI (two TBMs in the case of L1 modules). The TBM controls the readout of a group of ROCs.

Table 2. Expected hit rate, fluence, and radiation dose for the BPIX layers and FPIX rings. The hit rate corresponds to an instantaneous luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [16]. The fluence and radiation dose are shown for integrated luminosities of 300 fb^{-1} for the BPIX L1 and 500 fb^{-1} for the other BPIX layers and FPIX disks, well beyond the expected integrated luminosities for the detectors at the end of Run 3, of 250 and 370 fb^{-1} , respectively.

	Pixel hit rate [MHz/cm ²]	Fluence [$10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$]	Dose [Mrad]
BPIX L1	580	2.2	100
BPIX L2	120	0.9	47
BPIX L3	58	0.4	22
BPIX L4	32	0.3	13
FPIX inner rings	56–260	0.4–2.0	21–106
FPIX outer rings	30– 75	0.3–0.5	13– 28




Figure 6. Drawings of the pixel detector modules for BPIX L1, BPIX L2–4, and the FPIX detector. Reproduced from [16]. The Author(s). CC BY 4.0.

Sensors. The silicon sensors are of the n-in-n type [16], with strongly n-doped (n^+) pixelated implants on an n-doped silicon bulk and a p-doped back side. The n^+ implants collect electrons, which has the advantage of being less affected by charge trapping caused by high irradiation [17, 18]. It is also advantageous as it allows sensors to be operated under-depleted after the detector is irradiated (when the so-called type inversion [16] is reached).

For the BPIX sensors, the n-side inter-pixel isolation was implemented through the moderated p-spray technique with a punch-through biasing grid. More details can be found in the references of ref. [16]. The sensors are made from approximately $285 \mu\text{m}$ -thick phosphorus-doped 4-inch wafers produced in a float-zone (FZ) process (fabricated at CiS). The FPIX sensors use open p-stops for n-side isolation, where each pixel is surrounded by an individual p-stop structure with an opening on one side. They were produced on $300 \mu\text{m}$ -thick 6-inch FZ-wafers (produced at Sintef).

The radiation resistance of a module is, to a large extent, defined by the possibility of increasing the sensor bias voltage to obtain a sufficiently high signal charge. As already mentioned in section 3.1.1, BPIX L1 modules have the highest radiation exposure and therefore need to be exchanged after approximately 250 fb^{-1} .

During operation in 2017–2018, L1 modules were run with high efficiency at a bias voltage of 450 V, up to an integrated luminosity of almost 120 fb^{-1} . After the replacement of the BPIX L1 during LS2, the new L1 modules must withstand a fluence expected to be about twice as high until

the end of Run 3. In order to maintain a high enough signal charge, the pixel detector power supplies were upgraded to deliver a maximum voltage of 800 V during LS2, as described in section 3.1.4.

Readout chip. The upgraded ROCs, PSI46dig and PROC600 [16], are manufactured in the same 250 nm CMOS technology as the ROCs used in the original pixel detector (PSI46 [19]). Their design requirements are summarized in table 3.

Table 3. Parameters and design requirements for the PSI46dig and PROC600.

	PSI46dig	PROC600
Detector layer	BPIX L2–L4 and FPIX	BPIX L1
ROC size	$10.2 \times 7.9 \text{ mm}^2$	$10.6 \times 7.9 \text{ mm}^2$
Pixel size	$100 \times 150 \mu\text{m}^2$	$100 \times 150 \mu\text{m}^2$
Number of pixels	80×52	80×52
In-time threshold	$<2000 \text{ e}^-$	$<2000 \text{ e}^-$
Pixel hit loss	$<2\%$ at 150 MHz/cm^2	$<3\%$ at 580 MHz/cm^2
Readout speed	160 Mb/s	160 Mb/s
Maximum trigger latency	$6.4 \mu\text{s}$	$6.4 \mu\text{s}$
Radiation tolerance	120 Mrad	120 Mrad

The FPIX detector and layers 2–4 of the BPIX use the PSI46dig. Its design follows very closely the original ROC with the readout architecture based on the column-drain mechanism [20]. The pixel cell remains essentially unchanged except for the implementation of an improved charge discriminator. The improved discriminator reduces cross talk between pixels and the time walk of the signal [19] and thus leads to lower threshold operation (below 2000 e^-). The main modifications were made in the chip periphery in order to overcome the limitations of the PSI46 at high rate. They included a size increase of the data buffers (from 32 to 80 cells) and time-stamp buffers (from 12 to 24 cells) to store the hit information during the trigger latency, the implementation of an additional readout buffer stage to reduce dead time during the column readout, and the adoption of 160 Mb/s digital readout. In contrast to the previous ROC, the data readout is digital, using an 8-bit analog-to-digital converter (ADC) running at 80 MHz. Digitized data are stored in a 64×23 buffer, which is read out serially at 160 MHz.

The PROC600 is used for BPIX layer 1 and has to cope with hit rates of up to 600 MHz/cm^2 . Therefore, the data transfer of pixel hits to the periphery must be much faster than the PSI46dig. This was achieved by a complete redesign of the double column architecture. The pixels within a double column are dynamically grouped into clusters of four and read out simultaneously, which significantly speeds up the readout process.

During operation in 2017 and 2018, the PROC600 delivered high-quality data. However, two shortcomings of the PROC600 were identified, and are discussed in more detail in ref. [16]. The first was cross talk between pixels, which was higher than expected and generated noise at high hit rates. The second was a lower efficiency caused by a rare loss of data synchronization in double columns. Therefore, it was decided to develop for Run 3 a new, revised version of the PROC600. The main change in the new version addresses the rare cases of data synchronization loss. The issue was tracked to a timing error in the time-stamp buffer of the double column, which leads to

inefficiencies at low and high hit rates. It was corrected in the buffer logic of the revised PROC600. The higher-than-expected noise was traced to inappropriate shielding of the circuitry for calibration pulse injection and has been fixed in the revised version of the PROC600. In addition, the routing and shielding of power and address lines were improved. Both changes led to lower noise and lower cross talk between pixels. The revised PROC600 was used to construct the modules for the new L1, which was installed during LS2 for use in Run 3.

Token bit manager chip. The TBM is a custom, mixed-mode, radiation-hard integrated circuit that controls and reads out a group of 8 (L1) or 16 (L2–L4, FPIX) ROCs. To increase the data output bandwidth from a module, two 160 Mb/s ROC signal paths, with one path inverted, are multiplexed into a 320 Mb/s signal, and then encoded into a 400 Mb/s data stream that is optically transmitted to the downstream DAQ system. The TBM has a single output (TBM08) version for L3, L4, and FPIX, and dual output versions, which are used in L2 (TBM09) and L1 (TBM10). The TBM08 version has two independent 160 Mb/s ROC readout paths, and the TBM09 and TBM10 versions have four separate, semi-independent, 160 Mb/s ROC readout paths.

In addition to the increased output bandwidth, several critical features were added to the TBM for Phase 1. As a result of adopting a faster digital readout, finer control over the timing of internal TBM operations and external TBM inputs was needed: delay adjustments were added for the ROC readouts, the token outputs, the data headers and the data trailers, and relative phase adjustments were added between the 40 MHz incoming clock, the 160 MHz clock, and the 400 MHz clock. To prevent very long readouts from blocking the DAQ system, an adjustable token timeout was added that can reset the ROCs and drain buffered data.

Operation of the TBM during collision data taking revealed a vulnerability to a particular single event upset (SEU) that halts the TBM and requires a power cycle of the TBM to recover. An additional iteration of the TBM chips was designed in the spring of 2018 to address this TBM SEU issue and to add an adjustable delay of up to 32 ns to the 40 MHz clock. The delay was added to allow finer adjustment of the relative timing between modules. The new chips are used in the new modules in BPIX L1, incorporated during the consolidation work in LS2.

BPIX module construction. The BPIX detector contains 1184 modules, all having a similar design but coming in three different flavors depending on the requirements of the different layers. The three module types are summarized in table 4.

Table 4. Overview of module types used in the Phase 1 pixel detector.

	ROC	Number of ROCs	TBM	Number of TBMs	Number of 400 Mb/s readout links/module
BPIX L1	PROC600	16	TBM10	2	4
BPIX L2	PSI46dig	16	TBM09	1	2
BPIX L3, L4, & FPIX	PSI46dig	16	TBM08	1	1

A drawing of the detector module cross section for BPIX L2–4 is shown in figure 7. During the module production process, bare modules are made by bump bonding of the 16 ROCs to the sensor. In the next step, the thin four-layer HDI is glued onto the sensor side of the bare module and




Figure 7. Cross sectional view of a pixel detector module for BPIX L2–4, cut along the short side of the module. Reproduced from [16]. The Author(s). CC BY 4.0.

wire-bonded to the ROCs. The TBMs are already mounted, wired-bonded, and tested on the HDI before joining the HDI to the bare module. The BPIX L2–4 modules are mounted using silicon nitride base strips that are glued under the ROCs on the two long sides of the module (figure 6, center). The tight space requirements in the innermost BPIX layer required a different scheme with clamps between two modules (figure 6, left). A cap (not shown in the figure) made from a $75\text{ }\mu\text{m}$ -thick polyimide foil protects the wire bonds of the ROCs and TBMs against mechanical damage from the cables of other modules when mounted on the BPIX mechanics. A single L2 module, excluding the cable, has a mass of about 2.4 g and represents a thickness corresponding to 0.8% of a radiation length at normal incident angle.

The production of the BPIX modules was shared by five different module-assembly centers operated by groups from Germany, Switzerland, Italy, Finland, and Taiwan. The assembly tools and procedures were mostly standardized, however, each center used different bump-bonding techniques, as described in more detail in ref. [16]. The overall yield of the module assembly ranged from 65 to 85% in the different production centers, a bit lower than expected. The main reason was low-yield phases during the production start-up.

FPIX module construction. The 672 FPIX modules installed in the detector all have the same design and are instrumented with PSI46dig ROCs and a TBM08 with one readout link per module. The FPIX modules (2.9 g/module without cable) use a different sensor, HDI, and aluminum/polyimide flat flex cable and are not interchangeable with BPIX modules. The bump-bonding procedure was done at RTI and used an automated Datacon APM2200 bump bonder. Bare modules were sent from RTI to two FPIX module assembly and testing sites in the U.S.A. Modules were assembled using an automated gantry and pick and placement equipment. The wire bonds were encapsulated to protect against humidity and to mechanically support the wire bonds. The average yield for FPIX production modules was 68%, largely limited by bare module quality, especially in the early production.

3.1.3 Mechanics

BPIX mechanics. The detector mechanics have been described in detail in ref. [16], and a brief outline will be given here. The BPIX detector modules are mounted on ladders, with each ladder supporting eight modules. The ladders are mounted on four concentric layers split into half-cylinders. These are staggered by an alternating arrangement of ladders at smaller and larger radii. Such a ladder placement provides between 0.5 and 1.0 mm of sensor overlap in the r - ϕ plane. Ladders are made from carbon-fiber reinforced polymer (CFRP) and have a length of 540 mm and a thickness of $500\text{ }\mu\text{m}$. The end rings, on which the ladders are suspended, consist of a CFRP/Airex/CFRP sandwich structure.

The 1.7 m-long BPIX service half-cylinders contain the DC-DC converters, the opto-electronic components, and the module connections. The modules are connected to the service cylinders with micro-twisted-pair copper cables with a length of about 1 m. These are contained in an additional structure placed in between the BPIX detector mechanics and the service half-cylinder.

The BPIX cooling is provided by complex looping tubes that cool the detector modules and the components on the service half-cylinders. The tubes are made out of stainless steel with an inner diameter of 1.7 mm and a wall thickness of 50 μm . The loops are between 957 and 1225 cm long and dissipate a power of up to 240 W each. On the service half-cylinders the tubes have a wall thickness of 200 μm and an inner diameter of 1.8 and 2.6 mm for supply and return lines, respectively.

FPIX mechanics. The three half-disks forming one FPIX quadrant are supported by a carbon-fiber composite service half-cylinder. The FPIX half-disks consist of two turbine-like mechanical support structures with an inner assembly providing a sensor coverage from radii of 45 to 110 mm with 11 blades, while the outer assembly covers radii from 96 to 161 mm with 17 blades (figure 8). The half-disks serve as the cooling isotherms for the sensor modules. One module is mounted on each side of a blade, such that there is a small overlap in coverage at the outer edge of the blade. The overlap is larger for adjacent modules closer to the beam. The flat panel blades are made of 0.6 mm-thick sheets of thermal pyrolytic graphite (TPG) encapsulated between two 70 μm -thick single-ply carbon fiber face sheets. The blades are suspended between an inner and an outer 2.4 mm thick graphite ring. Each graphite ring houses the stainless steel cooling lines and is reinforced on the side facing away from the blades with a carbon fiber skin.




Figure 8. Drawing of an FPIX half-disk made from two half-rings of modules mounted on blades that are suspended between graphite rings (left), and close up view of a module mounted on a blade (right). Reproduced from [16]. The Author(s). CC BY 4.0.

The FPIX service half-cylinders are 2220 mm long, with an outer radius of 175 mm and an inner radius of 165 mm. The service half-cylinders consist of a double-walled carbon fiber rear section that is 1500 mm long for support elements, a corrugated single-wall carbon fiber front section that houses the half-disks, and an aluminum end flange for the cable and cooling tube pass through. A ruby ball sliding in aluminum bushings on each of the three half-disk support stalks provides a kinematic mount. A detailed description of the FPIX detector mechanics can be found in ref. [16].

3.1.4 Services

Readout architecture and data acquisition system. In order to be efficient, the CMS Phase 1 pixel detector readout architecture was copied from the original detector, with modifications made to handle the 400 Mb/s readout and the increased number of channels. The entirety of the CMS Phase 1 system readout is detailed in ref. [21], and a summary is provided below.

The auxiliary on-detector electronics were modified to include a crystal-driven phase-locked-loop chip (QPLL) [22] that reduced the jitter on the clock signal, a higher bandwidth chip (DLT) [23] to translate the module readout into a level suitable for the laser driver chip, and a new optical transmitter (TOSA from Mitsubishi Electric Corp.) for the higher bandwidth digital signal.

The off-detector VME-based DAQ system used for the original detector was replaced by a μ TCA-based system using a common carrier board (FC7) [24] with conventional 10 Gb/s optical network links for the transfer of the data to the CMS central DAQ system. The DAQ system, used to control and read out the full pixel detector, consists of 108 frontend driver modules (FEDs), which receive and decode the pixel hit information; three (2017–2018) to six (≥ 2021) frontend controller modules (TkFECs), which serve the detector slow control; 16 pixel frontend controller modules (PxFECs) used for module programming and clock and trigger distribution; and 12 AMC13 [25] cards providing the clock and trigger signals. Application-specific mezzanine cards and firmware make an FC7 carrier board into a FED or FEC.

Power system. The Phase 1 pixel detector requires low voltages to supply the readout chips on the pixel modules, a bias voltage to deplete the pixel sensors, and voltages to supply the control electronics components on the supply tubes and service cylinders.

The nominal control voltage of 2.5 V is supplied by A4602 CAEN power supply modules. Power supplies by CAEN of type A4603DH provide the low and high voltages to the pixel modules. The maximum bias voltage that can be delivered was raised from 600 to 800 V during LS2. Each bias voltage channel can provide up to 16 mA, and serves several pixel modules. While the auxiliary and bias voltage supply systems have conceptually not been changed for Phase 1, the low voltage supply system was changed from a direct parallel powering system to a DC-DC conversion powering system. The number of ROCs, and thus the current consumption of the detector, has roughly doubled as compared to the original pixel detector. Thanks to DC-DC conversion factors of 3–4, the currents on the typically 50 m-long supply lines between the detector and the power supplies are reduced with respect to a direct powering scheme.

The DC-DC converter modules are built around the CERN radiation-tolerant FEAST buck converter ASIC [26], and were optimized in terms of dimensions, mass, and performance for the application in the pixel detector [27, 28]. The DC-DC converter modules are installed on the BPIX supply tubes and in the FPIX service cylinders, at a distance of about 1 m from the pixel modules and outside of the sensitive pixel volume. A total of 1216 DC-DC converter modules are used in the pixel detector.

Two types of DC-DC converter modules are needed to supply the pixel modules: one delivering 2.4 V to the analog circuitry of the ROC, and one delivering 3.3 or 3.5 V (depending on the position of the served pixel module in the detector) to the digital circuitry of the ROC and to the TBMs. Each such pair of DC-DC converter modules serves between 1 and 4 pixel modules, depending on the layer and ring of the pixel modules. The DC-DC converter output currents range from 0.4 to 1.7 A

for the analog supply, and from 1.3 to 2.4 A for the digital supply. The power efficiency is 80–84%, depending on the output voltage and load. The DC-DC converters can be disabled and enabled from an already existing chip (CCU) [29] used both in the original and in the Phase 1 pixel detector; when disabled the DC-DC converters do not provide an output voltage. This feature was used during 2017 to power-cycle the TBM chips suffering from a SEU, which could not be recovered otherwise.

Up to seven DC-DC converters (of one variant) are connected to one low voltage power supply channel. The low voltage part of the original A4603 power supplies has been adapted to the DC-DC conversion powering scheme. The maximum output voltage was raised to 12.5 V and the fast remote sensing was abandoned, while a slow-control loop to compensate for voltage drops along the supply cables is still available.

In general, the DC-DC conversion powering system worked very well. However, starting in October 2017, DC-DC converters started to fail, and at the end of the 2017 data-taking period, about 5% of the DC-DC converters were defective. This was traced back to a problem in the FEAST ASIC, namely a radiation-induced leakage current in a transistor [30]. When the chip is disabled, a voltage above the chip specification can build up on a certain node, damaging the chip. For the 2018 data taking, all DC-DC converters were replaced with (almost) identical ones. The input voltage was reduced from about 11 to 9 V, and disabling of the chip was replaced by power-cycling of the power supplies. Due to these operational changes, no DC-DC converter failed in 2018. During LS2, new DC-DC converters were again produced, featuring a new version of the DC-DC ASIC (FEAST2.3). In this new version, the problem is fixed and operational changes are no longer needed.

The use of DC-DC converters meant that the LV and HV modularity of the detector no longer matched. The LV could be switched off by disabling a DC-DC converter, however, the HV stayed on because of the limited number of HV wires. The failure of individual DC-DC converters affected a small number of pixel modules that were kept under bias voltage in 2017 while their LV was off. Having the sensor biased but the pixel readout amplifiers off damaged the amplifiers in the affected pixel modules. Several of those damaged modules were replaced during LS2. In addition, the LV/HV modularity was harmonized in FPIX; for BPIX this was not possible due to the limited number of HV wires.

Cooling. To keep the silicon sensors below 0°C and remove heat from the other detector elements, CO₂ evaporative cooling utilizing the two-phase accumulator controlled loop (2-PACL) approach [31] is used in the Phase 1 pixel detector. Evaporative CO₂ cooling provides low density, low viscosity, and high heat transfer capacity while allowing the use of an all passive, small-diameter, thin-walled stainless steel pipe network inside the detector volume. This results in a lower contribution of cooling to the overall detector material budget. During normal operation (−22°C) the expected power [15] from the BPIX (6 kW) and FPIX detectors (3 kW) is removed using two dedicated 15 kW CO₂ plants, one for each detector, located in the CMS detector cavern. There is enough flexibility and capacity in the system so that the eight cooling loops in each detector can be connected to either cooling plant with no loss of cooling capacity. The typical temperature at the pixel module surface is about 12°C higher than the coolant. A detailed description of the CMS Phase 1 pixel detector cooling system can be found in ref. [16].

During the testing of the removed FPIX in 2019, one of the inlet cooling pipe connections at the detector end flange was broken. It was therefore decided to replace all the FPIX inlet connectors

with a more robust solution that required only a single wrench to make a connection. This sped up the detector installation in 2021 and prevented another occurrence of a broken inlet connector.

3.1.5 Detector operation

Detector live fraction. The detector live fraction, defined as the fraction of working ROCs, was 95.0 and 96.1% for the BPIX and FPIX detectors, respectively, during the first collisions in 2017. The main causes of failures in the BPIX detector were the loss of power due to faulty connectors and modules masked because of readout problems. For the FPIX detector the main problem was an issue with the clock distribution in one sector. Towards the end of 2017, the fraction of nonworking modules was dominated by the failure of the DC-DC converters described in section 3.1.4. The DC-DC problems lowered the working fraction to 90.9 and 85.0% for the BPIX and FPIX detectors, respectively.

During the 2017/2018 LHC year-end technical stop, faulty components were repaired and all DC-DC converters were replaced. A few broken BPIX modules, which were accessible without disassembling the detector layers, were also replaced. These repairs improved the working detector fraction for the 2018 data-taking period. For the BPIX it varied from initially 98 to 93.5% at the year end, with the main drop due to faulty power connectors. The FPIX detector working fraction was stable throughout the entire year at 96.7%.

As already mentioned above, during LS2 a new L1 was installed in the BPIX. Other repairs, in the BPIX and FPIX, involved faulty infrastructure components like bad connectors and broken optical fibers. In addition, eight modules were replaced in BPIX L2. With these improvements, the working fractions for BPIX and FPIX at the start of Run 3 were 99.1 and 98.5%, respectively.

Threshold adjustment. Pixel charge thresholds are an important performance parameter since they directly influence the position resolution. Lower thresholds increase the charge sharing between pixels, resulting in a better resolution. However, too low thresholds result in noise saturating the readout. Therefore, thresholds in all ROCs are adjusted to the lowest possible value, but well above the noise level itself. More details about the threshold adjustment are given in ref. [16].

During the 2017-18 data-taking period, the BPIX L2-4 thresholds were about $1400 e^-$ and similarly for the FPIX detector at about $1500 e^-$. Because of the higher noise in BPIX L1, the thresholds had to be higher, about $2200 e^-$. With these pixel thresholds the number of noise hits was very low, below 10 pixels per bunch crossing per layer, resulting in a per pixel noise hit probability of less than 10^{-6} . Individual pixels that showed a hit probability exceeding 0.1% were masked during operation. The total fraction of masked pixels was less than 0.01%.

The revised PROC600 used in the new L1 installed in 2021 has significantly lower cross-talk noise (as discussed in section 3.1.2). The expected threshold, noise, and time-walk behavior of the revised PROC600 version is similar to that of the PSI46dig. These improvements allow the BPIX L1 to be operated in Run 3 with significantly lower thresholds, similar to the ones for the other layers.

3.1.6 Performance of the pixel tracker

The two most important performance parameters for a pixel detector are hit efficiency and position resolution. Both strongly affect the ability to perform pattern recognition and b tagging, two main roles for a well-functioning pixel detector.




Figure 9. Schematic view of one quadrant in the r - z view of the CMS tracker: single-sided and double-sided strip modules are depicted as red and blue segments, respectively. The pixel detector is shown in green. Reproduced from [32]. CC BY 4.0.

The hit efficiency has greatly improved with the Phase 1 upgrade, mostly due to the ability of the upgraded ROCs to read data at higher speed and operate at lower thresholds, as described in section 3.1.2 and ref. [16]. The hit efficiencies measured at an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ are 97, 98, 99, and 99.5% for BPIX L1–L4, respectively. For the forward disks, the average efficiency is 99%. These numbers present a big improvement with respect to the Phase 0 detector, where the efficiencies would have been much smaller for the same data rates [16].

The position resolution is measured using the “triplet” method [16], where the expected hit position in a detector layer is interpolated from two other layers. For the BPIX L3, this method yields a position resolution of $11.0 \mu\text{m}$ in the r - ϕ direction and $24.3 \mu\text{m}$ in the z direction. For BPIX L1 and L2, the resolutions are somewhat worse due to the higher sensor radiation damage at lower radii. For the forward disks, the resolution is $11.9 \mu\text{m}$ in the r direction and $21.0 \mu\text{m}$ in the z direction. These numbers agree with the design expectations and are consistent with simulations.

Pixel sensors suffer from radiation damage induced by the high density of charge particle tracks. In order to maintain good efficiency and position resolution, it is necessary to keep increasing the sensor’s bias voltage. The nominal voltage after installation was 150 V for the BPIX modules and 300 V for FPIX. Especially for the inner BPIX layers, these voltages had to be increased in a few steps during data taking. Presently, end of 2023, the voltages are 450, 350, 250, and 250 V for the BPIX L1–L4, respectively, and 350 and 300 V for the inner and outer rings of the FPIX disks. The values are expected to increase further throughout Run 3, eventually reaching the maximum of 800 V for BPIX L1 and 600 V for all other layers.

3.2 Strip detector

3.2.1 Detector description

The silicon strip tracker (SST), together with the pixel detector, provides measurements of charged particle trajectories up to a pseudorapidity of $|\eta| < 2.5$. The layout of the SST is shown in figure 9.

The SST has 9.3 million silicon micro-strips and 198 m^2 of active silicon area distributed over 15 148 modules. Single-sided p-on-n micro-strip sensors are used. The detector is 5 m long and has

a diameter of 2.5 m. It has ten layers in the barrel region with four layers in the tracker inner barrel (TIB) and six layers in the tracker outer barrel (TOB). The TIB is supplemented with three tracker inner disks (TID) at each end. In the forward regions, the detector consists of tracker endcaps (TEC). Each TID is composed of three rings of modules and each TEC is composed of up to seven rings. In TIB, TID, and in rings 1–4 of the TECs, sensors with a thickness of $320\ \mu\text{m}$ are used, while in TOB and in rings 5–7 of the TECs, $500\ \mu\text{m}$ thick sensors are used. The modules in the barrel layers measure r and ϕ coordinates, while the modules in the TECs and TIDs are oriented to measure the coordinates in ϕ and z . In four layers in the barrel and three rings in the endcaps, stereo modules are used (figure 9). These modules have a second module mounted back-to-back with a stereo angle of 100 mrad. The stereo modules provide coarse measurements of an additional coordinate (z in the barrel and r in the endcaps).

The analog signals from 128 strips are processed by one APV25 chip. The chip has 128 readout channels, each consisting of a low-noise and charge-sensitive preamplifier, a 50 ns CR-RC type shaper, and a 192-element deep analog pipeline which samples the shaped signals at the LHC frequency of 40 MHz [33]. Signals from two APV25 chips are multiplexed, converted to optical signals by analog opto-hybrids (AOH) [1], and transmitted via optical fibers to front-end drivers (FED), located in the service cavern outside the radiation zone. Pedestal and common mode subtraction, as well as cluster finding, are performed in the FEDs. Clock, trigger information, and control signal are transmitted to the detector by the frontend controllers (FEC), also located in the service cavern. Configuration data for the modules is distributed via the I2C protocol to communication-and-control units (CCU) [29], grouped in token ring networks (control rings). The modules in the SST are grouped in power groups each of which shares one power supply channel. There are 1944 power groups in total. Each power group has two low-voltage channels with 2.5 and 1.25 V regulators and two high-voltage channels that can be regulated up to 600 V [1]. The detector is cooled with C_6F_{14} monophasic coolant by two cooling plants.

The SST has been operated stably and successfully since 2009, and operation is scheduled to continue until the end of Run 3. During Run 1, the SST was operated with its primary cooling at $+4^\circ\text{C}$, significantly above the designed operating temperature, due to insufficient humidity control in the service channels and in the bulkhead region, i.e., the interface region between the detector volume and the outside seal. In 2009, the detector suffered from an over-pressure incident. Both inlet and outlet lines of 90 cooling loops of circulating C_6F_{14} were closed on one of the two cooling plants. After that the detector warmed up. As a result of the over-pressure, some of the cooling lines developed leaks or were detached from modules. Due to this incident there are several regions in the detector that have closed cooling loops or degraded cooling contacts.

During LS1 in 2013–2014, a number of engineering changes were carried out on the detector infrastructure that allowed to lower the operating temperature of the SST below 0°C . Most prominently, a dedicated plant was installed to produce dry air or oxygen-depleted air. The plant has a flow of about $250\ \text{m}^3/\text{h}$, and is able to meet the dew point requirement of the SST (around -60°C). The plant is the primary source of dry gas injection to the detector and its services. In LS1, the insulation of all service channels and the bulkhead of the detector was significantly improved in order to control the humidity conditions in the closed volume of the detector.

All these modifications to the detector infrastructure facilitated the SST operation at -15°C since the beginning of Run 2. However, with increasing irradiation, the leakage currents began to approach the power supply limits (12 mA) in the regions with no cooling or degraded cooling

contacts. Thermal runaway was observed in several power groups of the TIB. As a consequence, since 2018, the SST was operated at -20°C , which sufficiently reduced the leakage currents. By the end of Run 3 it will be necessary to lower the operating temperature to -25°C . During LS2 a test at a temperature of -25°C confirmed that the detector can be operated at this temperature and that there is no degradation of the humidity conditions inside the detector and the service channels.

3.2.2 Performance of the strip tracker

The performance of the SST will be discussed in the following. More details about the results in this section can be found in ref. [34].

Throughout all the years of operation, no SST on-detector components were exchanged because the detector has been inaccessible. The fraction of bad detector components has been largely stable during Run 1 and Run 2. This includes the readout channels that are excluded from the data taking: failing control rings, problems in LV or HV distribution, and individually switched-off modules, single APV25 chips or groups of strips. As can be seen in figure 10, the fraction of bad components was stable throughout Run 2 and amounts to about 4%.




Figure 10. Fraction of bad components for the CMS silicon strip detector as a function of the delivered integrated luminosity. Reproduced with permission from [34].

One of the most important performance characteristics is the signal-to-noise ratio (S/N). The evolution of S/N with accumulated integrated luminosity is shown in figure 11 (left). As expected from irradiation studies [1], the S/N degrades approximately linearly with the integrated luminosity [34]. The decrease observed during Run 2 indicates that the SST will continue to provide high-quality data until its end of life, estimated to be at 500 fb^{-1} , well beyond the expected end of Run 3.

Another important aspect of the SST is the hit efficiency, which is the detection efficiency for a particle traversing a sensor. The measurement of the hit efficiency is performed using tracks that pass the quality criteria as defined in ref. [35]. In order to avoid inactive regions, trajectories that are close to sensor edges or their readout electronics in the studied layer are not considered. The efficiency is determined from the fraction of traversing tracks with a hit in a module anywhere within a range of 15 strips from the expected position. The measured hit efficiency under typical conditions




Figure 11. Left: signal-to-noise ratio as a function of integrated luminosity as accumulated in pp collisions during Run 2, separately for the different detector partitions. Triangles and crosses indicate the results for sensors in the TEC of thickness 320 and 500 μm , respectively. Right: hit efficiency of the silicon strip detector taken from a representative run recorded in 2018 [34] with an average hit efficiency under typical conditions at an instantaneous luminosity of $1.11 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The two gray bands represent regions where the hit efficiency is not measured due to the selection criteria of the analysis [34].

during Run 2, at an average instantaneous luminosity of $1.11 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to about 31 pp interactions per bunch crossing, is shown in figure 11 (right). The average hit efficiency is about 99.5%, depending on the layer. Since the inefficiency mainly depends on the particle flux, the inner layers have a somewhat lower efficiency than the outer ones. Moreover, the inefficiency depends on the sensor thickness and on the pitch.

Radiation effects are also monitored during the operation of the detector, including the increase of the leakage currents in the sensors, the evolution of the full depletion voltage due to the change of the effective sensor doping concentration, and the evolution of the laser-driver performance in the optical readout chain.

During Run 3, due to increasing luminosity, the leakage currents will continue to rise. It can therefore be expected that some modules in regions with closed loops or degraded cooling contact will experience thermal runaway, or that the corresponding HV power-supply channels will reach their limit of 12 mA. Most of the modules are double-sided, and one way to reduce the self-heating effect is to switch off one side of the module. A voltage reduction can also reduce the leakage currents significantly. However, this is possible only if the applied voltage remains above the full depletion voltage. Towards the end of Run 3, it is expected that a lowering of the detector temperature to -25°C will become necessary. It is estimated that this measure will reduce the number of modules experiencing thermal runaway after 500 fb^{-1} of integrated luminosity by roughly a factor of 2.

The sensors of the SST are operated at an applied voltage of 300 V in over-depletion mode, because the sensors are p-on-n type and some of them have undergone type inversion of the bulk material. The full depletion voltage is measured by performing bias voltage scans during pp collisions. A scan of the full detector is done usually twice per year during data taking and once per month

on a selected set of modules. The evolution of the full depletion voltage with integrated luminosity of one module in TIB layer 1 is shown in figure 12. The measurements of the full depletion voltage are also compared with simulations, which describe the change with integrated luminosity well.




Figure 12. Evolution of the full depletion voltage for one TIB layer-1 sensor as a function of the integrated luminosity and fluence until the end of Run 2. The full depletion voltage is measured from the cluster-width variable, shown as black dots, and the predicted curve is based on a model that uses fluence and temperature history as inputs [36]. The hashed area highlights the region at low values of the full depletion voltage where the analysis loses sensitivity [34].

The installation of the pixel detector and the cooling plant maintenance work caused extended periods of time when the silicon detector was not cooled as well as it would have been desirable from the point of view of radiation damage. In figure 12, small increases due to annealing are visible in the simulation around integrated luminosities of 75 and 130 fb⁻¹, corresponding to the winter shutdown periods. As can be seen from measurements and simulation, at around 200 fb⁻¹, the TIB layer-1 sensors are close to the inversion point. The overall situation with the reduction of the full depletion voltage in the SST is shown in figure 13. For each subdetector a decrease of the full depletion voltage is observed that depends on the distance from the interaction point. It is observed that the regions of the detector that are closest to the interaction point, namely TIB layer 1, TID ring 1, and TEC ring 1, are affected the most, as expected.

In summary, the SST has been delivering high quality data for the reconstruction of charged particle tracks since the start of the LHC operation. The performance of the system continues to be excellent also after more than 200 fb⁻¹ of integrated luminosity. Since the beginning of Run 3, the detector has been operated at -20°C. It is expected that the operation temperature will be lowered further to -25°C in order to reduce the leakage current. While radiation effects are visible in all parts of the detector, the margins are large enough for the detector to be operated safely and efficiently, and to provide high-quality data until the end of Run 3.




Figure 13. Decrease of the full depletion voltage for each layer, computed as the difference between the values measured at the time of the tracker construction and the values obtained by the analysis of a bias-voltage scan performed in September 2017 on all the tracker modules. The white (gray) histograms represent modules with 320 (500) μm thick sensors. The average fluence for each layer is shown by the red line.

4 Electromagnetic calorimeter

The electromagnetic calorimeter (ECAL) is placed outside the inner tracking system of CMS. It provides a measurement of the energy of electrons and photons, as well as their impact position and arrival time at the crystals.

4.1 Experimental challenges

The increase in the instantaneous and integrated luminosity, experienced during Run 1 and Run 2 of the LHC and expected to continue in the future, poses operational challenges for the ECAL. The radiation dose deposited in the detector reduces the average light transmission of the PbWO_4 crystals, lowering the signal-to-noise ratio of the electronics readout. The radiation also induces an increase in leakage currents in the barrel photodetectors, which are avalanche photodiodes (APDs), with a corresponding increase in the electronic noise [37, 38]. The instantaneous luminosity reached $2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ during Run 2, compared to $0.75 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ achieved in 2012. The increase in luminosity also poses challenges to the level-1 (L1) trigger system. Specifically, signals from direct energy deposition by particles in the APDs, termed “spikes”, must be rejected. Such spikes occur at a rate that is proportional to the luminosity. The higher radiation has also caused the silicon sensors of the preshower detector to have increased bulk currents, which require regular updates of the HV bias and correspondingly of the calibration of their response.

Additionally, the number of multiple pp interactions in a single bunch crossing (BX), termed pileup, has increased on average from 21 (up to 40) during Run 1 to 34 (up to 80) during Run 2. The bunch spacing in the machine reached its nominal value of 25 ns at the beginning of Run 2, half of what

it was in Run 1. Since the typical signals from the calorimeter, after shaping by the electronics, fall to 10% of their peak value in about 250 ns, the changes in the LHC operation have resulted in an increased number of overlapping signals from neighboring BXs, referred to as out-of-time (OOT) pileup.

These effects will be discussed in more detail in the following sections, along with the improvements in the calibration of the calorimeter and the final performance achieved during Run 2.

4.2 Response monitoring

The PbWO_4 crystals of the ECAL, when subjected to irradiation, undergo transparency changes. This is discussed in greater detail in ref. [39] and can be ascribed to the formation of color centers, which cause absorption bands in the crystal that reduce the light attenuation length. The creation of color centers is a dynamic process depending on the dose rate absorbed by the crystals. Its annealing process spontaneously takes place at room temperature and results in partial recovery of the transmittance. Since the scintillation process remains unaltered, a reference light signal can be used to measure and monitor the transparency and response changes, and corrections can be applied to equalize the crystal-to-crystal response.

To monitor and correct the response of the ECAL, a dedicated laser monitoring system is used that operates primarily at a wavelength of 447 nm, near the peak of the scintillating light spectrum. Additional monitoring wavelengths have been used, in particular a near-infrared one at 796 nm and a green one at 527 nm. These probe the transparency in regions that are much less sensitive to radiation damage (infrared) and more sensitive to the permanent component of the radiation damage (green).

The laser is operated at 100 Hz. To avoid interference with signals from beam collisions, the light is injected into the crystals during the LHC abort gap where there are no bunches in either beam, in intervals of at least $3 \mu\text{s}$. The abort gap is necessary to accommodate the beam abort kicker rise time and is available in all LHC filling schemes. The power of commercial lasers operating at a suitable repetition rate allows the injection of light into a few hundred crystals simultaneously. This is achieved using a system of optical fibers and diffusing spheres acting as homogeneous splitters. Light from a group of 200 fibers is measured by two p-n diodes. The variation in response is obtained by comparing the signal acquired by the APDs with the reference p-n diode. The time-dependent correction factor $LC_i(t)$, derived from the monitoring system for each crystal i , is defined as: $LC_i(t) = [R_i(0)/R_i(t)]^\alpha$, where $R_i(t)$ is the measured response to laser light at time t , and α is a parameter that takes into account the difference in path between the laser and scintillation light. Figure 14 summarizes the long-term evolution of the ECAL response to laser light during Run 1 and Run 2.

4.3 Noise evolution

The electronic noise in the endcaps (EE) is approximately constant. In the barrel (EB), radiation induces damage to the structure of the APD silicon lattice, causing an increase in the leakage current. The evolution of the leakage current is shown in figure 15 (left) as a function of the integrated luminosity for the central rapidity region and the most forward region of the barrel. It is well in line with the expectation from irradiation studies shown in figure 15 (right). The studies were performed using a pair of APDs, equivalent to ones on each of the ECAL barrel crystals. Measurements were done in CMS in situ for the points below $10 \mu\text{A}$, while the points at higher currents are based on laboratory measurements of irradiation with neutrons at different fluences, as indicated in the figure.




Figure 14. Relative response to laser light injected into the ECAL crystals, measured by the laser monitoring system, averaged over all crystals in bins of $|\eta|$. The response change observed in the ECAL channels is up to 13% in the barrel, $|\eta| < 1.5$, and reaches up to 62% at $|\eta| \approx 2.5$, the limit of the CMS inner tracker acceptance. The response change is up to 96% in the region closest to the beam pipe. The recovery of the crystal response during the periods without collisions is visible. These measurements, performed every 40 minutes, are used to correct the physics data. The lower panel shows the LHC instantaneous luminosity as a function of time.

When the signal is corrected for the reduction in average light yield, the electronics noise is effectively amplified. The effective noise, expressed in terms of equivalent energy and equivalent transverse energy, is shown in figure 16. The measurements are extracted from the fluctuation of the signal baseline, and are shown as a function of the pseudorapidity, covering the barrel and endcap regions. The three main data-taking periods of Run 2 are shown, along with the cumulative integrated luminosity since the beginning of Run 1. The shape as a function of $|\eta|$ is the result of the noise increase as a function of rapidity, a consequence of the larger radiation dose received by the forward regions, and of the conversion to equivalent transverse energy, the relevant quantity for physics measurements.

4.4 Signal reconstruction

The signals, after analog processing, are digitized every 25 ns. Upon a trigger, ten consecutive samples are transmitted to the backend electronics [40]. In order to cope with the increased OOT pileup, a novel amplitude reconstruction algorithm was developed for Run 2, based on a template fit, called “multifit” [41]. The multifit algorithm replaced the Run 1 method based on a digital-filtering technique [42], which estimated the energy by weighting five consecutive samples around the pulse maximum and subtracting the pedestals computed on the first three samples before the signal.

The multifit algorithm uses a template fit to extract the amplitude of the in-time pulse and the pulses coming from interactions occurring up to five BXs before and four BXs after, all within




Figure 15. Left: evolution of the APD dark current as a function of the integrated luminosity since the beginning of Run 1. The gray vertical lines represent the ends of the Run 1 and Run 2 data-taking periods. Spontaneous annealing of the radiation-induced defects is visible as vertical steps in the measurements and corresponds to long stops in the data taking, e.g., year-end technical stops. Right: measurement of the electronic noise in the most sensitive amplification range as a function of the measured leakage current of an APD pair. The measurements are explained in the text. The red line is a fit to the data with a square root function. The maximum expected fluence for Run 3 is $4 \times 10^{13} \text{ n}_{\text{eq}}/\text{cm}^2$ at $|\eta| = 1.45$.




Figure 16. Average noise per channel as extracted from data using variations (RMS) of the signal baseline at the end of each year of Run 2 data taking. The amplitude RMS is converted into an energy equivalent (left) and a transverse-energy equivalent (right). The integrated luminosity accumulated since the start of Run 1 is indicated. The spread of the equivalent noise at a given pseudorapidity is indicated by the light-colored bands.

the 10-sample digitization window and potentially contributing to the total signal of one channel. Because of the number of samples and free parameters in the fit, the baseline value of the signal is not computed dynamically but is instead obtained from regular measurements performed during data taking at least once per run. Additional inputs to the fit are the template signal shapes, one per channel, and the noise covariance matrix. Both inputs are regularly measured from data and updated

when found to differ significantly from those in use. This happens typically a few times per year, depending on the luminosity profile of the LHC.

The performance of the multifit algorithm has been measured using events from $\pi^0 \rightarrow \gamma\gamma$ and $Z \rightarrow ee$ decays. The energy resolution is excellent, as is the stability as a function of the OOT pileup. The improvement with respect to a nonoptimized digital-filtering technique is more significant for low-amplitude pulses, where the relative contribution of OOT pileup pulses is larger. The algorithm is sufficiently fast to be used in the HLT, and was adapted for execution on GPUs in the new processor farm used in Run 3.

The arrival time of the signal relative to the digitization window is measured by a digital-filtering technique based on the ratio of consecutive samples [43]. The timing information is subsequently corrected for its dependency on the pulse amplitude, as derived from simulations.

4.5 Trigger

The ECAL provides crystal energy sums, termed trigger primitives (TPs) [44], to the CMS L1 trigger for every BX. The trigger primitives are computed from energy sums of groups of 5×1 crystals, referred to as “strips” [45]. Each strip is served by an individual FENIX chip that performs energy intercalibration, E -to- E_T conversion, amplitude estimation, and BX assignment functions. In the EB, a sixth FENIX chip sums five strip-sums to compute the 5×5 “trigger tower” transverse energy, calculates the “fine-grain” electromagnetic bit based on the compatibility of the deposits with those from an electromagnetic shower, and computes the strip fine-grain bit for the rejection of signals from direct energy deposition in the APDs (“spike killing”) [46]. The strip fine-grain electromagnetic bit is configured to return a 0 for a spike-like energy deposit (a single channel above a configurable transverse energy threshold) or a 1 for a shower-like energy deposit (multiple channels above threshold). In the EE, the five strip sums are transmitted to the off-detector trigger concentrator card (TCC) to complete the formation of the trigger towers.

The TCC is responsible for the transmission of the barrel and endcap TPs to the L1 calorimeter trigger every BX via the optical synchronization and link board (oSLB) mezzanine cards. The TCC also performs the classification of each trigger tower, its transmission to the selective readout processor at each L1 trigger accept signal, and the storage of the trigger primitives for subsequent reading by the data concentrator card.

The ECAL L1 TPs are corrected for the effects of crystal and photodetector response changes due to LHC irradiation. Correction factors are derived using measurements from the laser calibration system, and the same corrections are also applied in the HLT. These corrections were first applied in 2012 only in the endcaps, for 22 individual rings of crystals of the same pseudorapidity, and were updated once per week. During Run 2, because of the higher beam intensities and correspondingly larger response losses, the TP corrections were applied per crystal and extended to the EB. From 2017 onwards, an automated validation procedure was developed to check the impact of the updated conditions on the L1 and HLT trigger rates, and the frequency of the updates was increased to twice per week to better track the response losses versus time.

Radiation-induced changes in the ECAL signal pulse shapes, in particular in the most forward regions of the EE, caused a continually growing probability for the BX to be misassigned in the ECAL TPs. This effect, termed trigger primitive pre-firing, resulted in an inefficiency for recording potentially interesting events of about 0.1% in any given primary data set, and about 1% for events

with two high-energy forward jets of invariant mass around 200 GeV [5]. Following the discovery of this issue in early 2018, η -dependent timing offsets were applied to the ECAL frontend electronics, and throughout 2018, periodic updates to these offsets were made during every LHC technical stop, in order to minimize the level of pre-firing in both the EB and EE.

The ECAL spike-killer algorithm has been retuned for the more challenging beam and detector conditions of Run 2. Spike-like energy deposits are rejected in the formation of the ECAL TPs by exploiting the additional functionality of the FENIX ASICs, the strip fine-grain electromagnetic bit. If the deposit is considered spike-like, and the tower energy is above a second configurable threshold, the tower energy is set to zero and does not contribute to the triggering of the corresponding event.

The spike-killer parameters were updated in 2016 to account for the higher LHC luminosity and the larger single-channel noise observed in the EB during Run 2. These new thresholds reduced the contamination of spikes in the ECAL TPs, corresponding to a transverse energy E_T of more than 30 GeV, by a factor of 2, with negligible impact on the triggering efficiency of electromagnetic signals with $E_T > 20$ GeV.

The spike-killing efficiency is sensitive to drifts in the ECAL signal baseline. Periodic updates in 2018 of up to twice per year of the baseline measurements used in the TP formation were therefore required in order to maintain a stable spike-killing efficiency. By periodically updating the baseline values, the spike contamination for TPs with $E_T > 30$ GeV was maintained below 20% during the 2018 run. These improvements in TP calibration and spike rejection, together with improvements in the L1 trigger system itself, allowed the L1 electron/photon trigger to operate with high efficiency and at the lowest possible E_T thresholds throughout Run 2 [5].

4.6 Channel calibration and synchronization

While the principles and methods of the ECAL calibration have not changed and are described in ref. [47], a brief summary and update is given here to help discuss the results.

The calibration of the calorimeter proceeds in several steps: (i) channels are corrected as a function of time for response changes as measured by the laser monitoring system; (ii) the response of channels at the same pseudorapidity, i.e., within the same ϕ -ring, is intercalibrated using specific physics channels as reference; (iii) ϕ -rings are intercalibrated with each other; and (iv) the absolute energy scale of the detector is fixed. In a separate and independent procedure, channels are synchronized by using the average arrival time of particles in minimum-bias events. Energy selections are applied to ensure an adequate signal-to-noise ratio, and additional criteria remove outliers in the timing distributions and ensure that the pulses have a good shape.

To complete the aforementioned steps (iii) and (iv), $Z \rightarrow ee$ events are used. For step (ii) the combination of a number of independent techniques is employed and is briefly summarized in the following.

The position of the two-photon invariant mass peak from π^0 decays is a good physics standard candle for intercalibration, even if at low energy. At the LHC, π^0 are produced in abundance, and a dedicated trigger and data acquisition stream allow for an efficient collection of a large data set. Events from this stream are saved in a reduced data format that contains only ECAL information in the proximity of the selected photon pair, to optimize the bandwidth at the HLT. Starting from L1 electromagnetic candidates, the reconstruction applies a simplified clustering algorithm that collects energy in a 3×3 crystal matrix centered around an energy deposit, called a “seed”, greater than 0.5

(1.0) GeV in the barrel (endcaps). An offline analysis applies a correction derived from simulation to take into account effects of the readout, e.g., channel zero-suppression, energy lost in the vicinity of the detector boundaries, and dead channels. An iterative fit to the invariant mass of the diphoton pair is performed, varying in each iteration the intercalibration coefficients and recomputing the clustered energy and candidate selection, until the variation from one step to the following is negligible.

The E/p method exploits the distribution of the ratio between the reconstructed calorimeter energy E and the momentum p measured in the tracker of high-energy electrons from W and Z boson decays. In order to obtain a pure electron sample, electron candidates are selected using kinematic, identification, and isolation requirements. The algorithm evaluates the intercalibration in an iterative way. In each iteration, the intercalibration coefficients are updated to constrain the peak of the E/p distribution to equal unity, and the clustered energy is recalculated. A correction is applied to take into account ϕ -dependent biases in the momentum measurement due to the presence of inhomogeneous tracker support structures. The correction is calculated from $Z \rightarrow ee$ events using the tracker momentum measurement in a specific ϕ region for one of the electrons and the ECAL energy measurement in any ϕ region for the other. The nonuniformity in ϕ is on the order of 1%.

Electrons in $Z \rightarrow ee$ events can also be used for the calibration of the detector. Low-bremsstrahlung electrons are selected to minimize the influence of detector material upstream of the ECAL. The definition of low-bremsstrahlung electrons is based on the narrowness of the electromagnetic shower detected in the calorimeter. The well known invariant mass peak and distribution of the dielectron decay provide an almost background-free reference channel. An unbinned likelihood is built for the distribution observed in data, assuming the invariant mass is well described by a classical Breit-Wigner function convolved with a Gaussian function that accounts for detector effects. The resolution and scale of the Gaussian function peak value are the free parameters of the likelihood. The granularity at which the parameters are allowed to vary permits the determination of the crystal intercalibration within ϕ rings, the relative calibration between rings, and the absolute energy scale of the detector. Compared to the other two, this method of intercalibration is particularly effective at large pseudorapidities. The calibration between rings (η -scale) is derived approximately every 5 fb^{-1} , in order to correct for drifts in the detector response.

The azimuthal symmetry of the energy flow in minimum-bias events, which was successfully used during Run 1, became more challenging during Run 2 because of the increased effective noise. While not competitive in precision for intercalibration purposes, it has been used to monitor the single-crystal response over time and provide useful insights into identifying and correcting residual imperfections in the light corrections. An example of such imperfections are the slow regional drifts of approximately 1% over one year of data taking, depending on the integrated luminosity and currently attributed to response changes in the p-n reference diode.

4.6.1 Intercalibration precision

Each of the intercalibration methods described above produces a set of constants with a corresponding statistical and systematic uncertainty. The statistical precision can be evaluated by comparing sets of constants derived from disjoint event samples. In the case of the $Z \rightarrow ee$ method, the precision can be obtained from the fitting procedure.

The π^0 method can be used to provide intercalibration constants for the barrel with a statistical precision that ranges from 0.1 to about 0.3%, slowly increasing with pseudorapidity, for 10 fb^{-1}

of integrated luminosity. In the endcaps the precision is around 1%, due to the larger particle multiplicity, which requires a tighter event selection, and larger detector noise.

The E/p method requires around 50 fb^{-1} (about a calendar year in Run 2) to derive a set of intercalibration constants. The corresponding statistical precision varies from 0.2% for electrons in the inner barrel region to 0.4% for electrons in the outer barrel. The nonuniformity in ϕ of the material in front of the ECAL introduces a systematic uncertainty.

The $Z \rightarrow ee$ method also requires about 50 fb^{-1} to provide a set of constants of precision comparable to the E/p ones in the barrel, while in the endcap it provides a precision far better than the other methods. In Run 2, thanks to a sufficient integrated luminosity, it was possible to intercalibrate crystals in regions of $|\eta| > 2.5$ in the endcaps, where the other methods cannot be used either because of very high pileup contaminations or because the region is outside the tracker coverage.

The calibration constants determined using the three methods are combined using a weight that is proportional to the inverse square of the estimated precision. For the π^0 and E/p methods, the systematic uncertainties are evaluated by studying the impact of the intercalibration constants on the Z boson lineshape, and found to be dominant for π^0 and comparable to the statistical precision for E/p . The intercalibration precision achieved in 2018 with the three methods combined is better than 0.5% for the entire barrel, and between 0.5 and 1% for the endcaps.

4.7 Run 2 operations summary

The ECAL DAQ operated during Run 2 with a luminosity-weighted efficiency larger than 99.6% for the EB and EE, and larger than 99.2% for the preshower detector. The ECAL trigger system also operated with high efficiency and availability during Run 2. The luminosity-weighted efficiency of the trigger system, accounting for trigger downtime and deadtime, i.e., automatic throttling of the readout decisions due to too high input rates, was larger than 99.9%. The fraction of ECAL channels that contributed to the DAQ was larger than 98.6% at the end of Run 2, with a loss of less than 0.2% over the course of the four years of Run 2 operations. The fraction of channels that contributed to the trigger was larger than 99%, and only a few problematic towers, strips, and individual channels were permanently masked.

A number of improvements to the firmware and software of the TCCs [48] were implemented to achieve and maintain these high efficiencies in the more challenging beam conditions of Run 2. These involved the automatic detection and masking of noisy or problematic signals from the frontend readout via configurable thresholds, without the need for manual intervention. The algorithms allowed the setting of individual thresholds per strip in the EE, such that they could be adapted to changing LHC conditions, as well as to increased radiation-induced noise in the forward regions of the EE. As a result of these improvements, which were fully implemented in both the EB and EE before the 2018 run, the number of incidents requiring manual intervention, as well as the deadtime and downtime associated with the ECAL trigger system, were significantly reduced in 2018 compared to 2017 [48].

Additional improvements were made to the data acquisition boards, firmware, and software to be more resilient against and make automatic the recovery from single-event upsets. Despite the luminosity increase, these remained a negligible source of downtime throughout Run 2.

To increase the reliability and ease maintenance, the crates of the high voltage system for the barrel and endcaps have been upgraded from the CAEN SY1525 to the CAEN SY4527, and the low voltage system has been made more redundant.

4.8 Run 2 performance

The performance of the ECAL in terms of energy resolution and stability of the energy scale is evaluated using $Z \rightarrow ee$ events reconstructed using the ECAL information alone. The energy resolution is affected by pileup, noise, and accuracy of the calibration, in relative order of decreasing importance. The resolution as a function of pseudorapidity is shown in figure 17 for electrons with low bremsstrahlung emissions and for an inclusive electron sample.




Figure 17. ECAL resolution for electrons having low bremsstrahlung emissions (left) and for an inclusive selection of electrons (right). The horizontal bars show the bin width.

The stability of the energy scale is monitored as a function of time by measuring the position of the peak in the dielectron invariant mass for $Z \rightarrow ee$ events, and is found to be within 0.1% in the barrel and a few times 0.1% in the endcaps. This figure is valid for electromagnetic deposits that are used in the reconstruction of jets and missing transverse momentum. For precision physics studies involving electrons and photons, the energy scale is further corrected using $Z \rightarrow ee$ events; after corrections, the stability is improved by a factor of 2 to 4, depending mostly on the pseudorapidity.

The ECAL time resolution is measured using $Z \rightarrow ee$ events by comparing the arrival times of the two electrons, defined as the time of the seed crystal in the supercluster. This time is corrected for time-of-flight, determined from the two electron tracks from the primary vertex of the event. Additional reconstruction quality and kinematic criteria ensure a pure sample of electrons. The time resolution resulting from the analysis of 2018 data is displayed in figure 18. It is shown as a function of the effective amplitude normalized to the noise, defined as:

$$\frac{A_{\text{eff}}}{\sigma_n} = \sqrt{\frac{2}{(\sigma_1/A_1)^2 + (\sigma_2/A_2)^2}}, \quad (4.1)$$

where A_1 and A_2 are the amplitudes of the two electron signals, and σ_1 and σ_2 the electronics noise of the corresponding channels. Notable analyses profiting from the ECAL time resolution are those looking for long-lived particles, such as that described in ref. [49].




Figure 18. ECAL timing resolution as measured from $Z \rightarrow ee$ events by comparing the arrival time of the two electrons. The performance has constantly improved over the years due to the frequency at which the synchronization constants have been updated, and is shown for 2018, the last year of Run 2. Updates in the constants are necessary to compensate for pulse shape changes induced by radiation. Vertical bars on the points showing the statistical uncertainties are too small to be seen in the plot. The red line correspond to a fit to the points with the parametrization of the resolution shown in the legend.

4.9 Preparation for Run 3

During LS2, the ECAL activities focused on improving the detector safety and control systems, on the algorithm to determine the trigger primitives, and on the development of a system to automatically compute, validate, and deliver updated calibrations.

4.9.1 Safety and control system

A significant upgrade of the ECAL safety and detector control systems, DSS and DCS, was performed during LS2. The sensor readout system for temperature, humidity, voltage, and current levels, composed of custom readout units, was replaced by industrial analog input (AI) modules. The 12 readout units based on RS-485 interfaces were exchanged by 45 AI modules, which are standard Siemens-certified peripherals (Simatic S7-300 analog input SM 331) connected through Profibus communication buses. This type of connection provides access to extensive diagnostic information and enables the readout of a full sensor at sampling intervals as fast as 0.1 s. This is about one order of magnitude faster than the previous system. Following the update of the readout method, the programmable logic controller (PLC) of the safety system was reprogrammed with completely new software, and is now part of the pool of PLC framework applications in CMS. The action matrix that defines the behavior of the PLC in terms of input and output signals and raises interlocks to protect the detector in case of alarm conditions, did not change, but benefits from the improved hardware. The new system was extensively validated during the regular cosmic ray data-taking campaigns in 2020 and 2021.

4.9.2 Trigger

Several improvements to the ECAL trigger-primitive formation and calibration were implemented for Run 3. These include a further optimization of the spike-killer thresholds (section 4.5) for the

expected Run 3 pileup and noise levels, optimization of the digital-filter weights used to compute the trigger-primitive energies accounting for radiation-induced changes in pulse shapes, and the possible use of a second set of amplitude weights to further improve spike rejection and for the potential tagging of out-of-time signals.

More frequent corrections to account for crystal and photodetector response changes were also implemented for Run 3, and the frequency of the updates was increased from twice a week to once per LHC fill. These corrections are important to maintain stable trigger rates and efficiencies, and improve the energy resolution of the related L1 and HLT objects, particularly electron/photon candidates.

4.10 Calibration

With the aim of reducing, as much as possible, the need for multiple reconstruction of CMS data sets with updated calibrations, a framework has been setup to provide automatic execution and bookkeeping of the workflows necessary to compute and validate updated and refined detector conditions as soon as enough data is available. Not only does this allow to follow closely the prompt reconstruction of CMS data with the best foreseeable conditions, but it also permits to have the best conditions available for the data (re)reconstruction as soon as the data taking is finished. While the techniques and physics standard candles used for calibrating the ECAL were well consolidated during Run 2, novelty and optimization have been introduced at the technical level of the data analysis needed to provide detector conditions. Additionally, with a layer to validate conditions before deployment, key figures of detector performance can be controlled, such as stability, resolution, and projected rates in the HLT. The workflows is synchronized with the online data taking and orchestrated by a Jenkins instance, deployed through Red Hat OpenShift technologies, based on an InfluxDB backend and a Python server.

5 Hadron calorimeter

5.1 The hadron calorimeter in Run 1 and Run 2

The CMS hadron calorimeter [50] (HCAL), shown schematically in figure 19, is composed of four major subdetectors: the hadron barrel (HB) [51], hadron endcap (HE) [50], hadron forward (HF) [52], and hadron outer (HO) calorimeters [53]. The HB and HE cover the pseudorapidity regions $|\eta| < 1.392$ and $1.305 < |\eta| < 3.0$, respectively. The HO provides a measurement of the shower tails in the region $|\eta| < 1.26$, and the HF covers $3.0 < |\eta| < 5.2$.

The HB and HE primarily use brass as the absorber, except for the inner and outer layers of HB, which are constructed from steel. The HB absorber is shown in figure 20 (left). The signals are produced in plastic scintillating tiles (figure 20, right), and the resulting blue light is shifted to green via embedded wavelength-shifting fibers. The towers in HB (HE) have up to 17 (18) scintillator layers, as shown in figure 19. Sequential layers are grouped into “depth” segments: the light from the layers in a given depth segment is optically summed and read out by a single photodetector. Clear plastic fibers send the signal to the hybrid photodetectors (HPDs) in the original design or silicon photomultipliers (SiPMs) after the upgrades. The segmentation is a tower structure in η - ϕ space. The η segmentation is indicated by the black solid lines in figure 19. The towers are referenced using integer indices $i\eta$ and $i\phi$, where the $i\eta$ assignments are given in the figure and $i\phi$ runs from 0 to 71, corresponding to the 72 divisions




Figure 19. Schematic view of the HCAL as of 2016, showing the positions of its four major components: HB, HE, HO, and HF. The layers marked in blue are grouped together as “depth 1,” i.e., the signals from these layers of a given tower are optically summed and read out by a single photodetector. Similarly, the layers shown in yellow and green are combined as depths 2 and 3, respectively, and the layers shown in purple are combined for HO. The notation “FEE” denotes the locations of the HB and HE frontend electronics readout boxes. The solid black lines, roughly projective with the interaction point, denote the η divisions in the tower η - ϕ segmentation, and the numbers at the edge of the tower denote the ieta index. Reproduced from [54]. © 2020 CERN for the benefit of the CMS collaboration. CC BY 3.0.




Figure 20. Left: brass absorber for the hadron barrel calorimeter HB. Right: scintillating tiles with wavelength shifting fibers used as the active media in the barrel, endcap, and outer hadron calorimeters.

in ϕ . Physically, the scintillators are arranged in “megatiles”, which are trays that support an array of scintillator tiles, along with the fibers that route the light to the photodetectors. An HE megatile is shown in figure 21. All channels in a subdetector with the same ϕ are defined as residing in the same wedge.




Figure 21. Physical arrangement of the scintillator and wavelength-shifting fibers into tiles for an HE megatile. Reproduced from [54]. © 2020 CERN for the benefit of the CMS collaboration. CC BY 3.0.

The HF, shown schematically in figure 22, is a 1.65 m-long sampling calorimeter with steel absorber. Plastic-clad quartz fibers with a diameter of 0.6 mm are the active elements, producing Cherenkov light. The fibers are inserted into 1 mm grooves drilled in the steel absorber, 5 mm apart. The fibers are parallel to the beam line. To enable discrimination of electrons and photons from hadrons, the fibers have two lengths: 1.65 m “long fibers” and 1.40 m “short fibers”. The short fibers run from the outside of the HF to 25 cm from the front face of the HF. The HF photodetectors are photomultiplier tubes (PMTs).

The HCAL readout control chain is shown schematically in figure 23. In the frontend readout boxes (RBXs), mounted on the detector, the electronics chain begins with analog signals from the photodetectors. The optical decoder unit (ODU) maps the incoming fibers from the scintillators to the photodetectors. These signals are digitized by the QIE chips (QIE8, QIE10, and QIE11) [56–59] in the readout modules (RMs), which provide both energy information via an analog-to-digital converter (ADC) and, for the QIE10 and QIE11 chips, rising-edge timing information via a time-to-digital converter (TDC). The clock, control, and monitoring (CCM) module provides the LHC clock, synchronous fast commands, slow control, and monitoring for the detector frontend electronics. Each RBX also contains a calibration module (CM), which provides calibration signals from two sources: an in situ LED, which sends light directly to the photosensors, and an off-detector laser, which can send light either to the photosensors or into the scintillator tiles for some of the detector layers. The CM also provides a reference measurement of the LED and laser calibration signal intensities.

5.2 Upgrades

5.2.1 Motivation and overview of the upgrade

The motivation for and the design of the CMS calorimeter system Phase 1 upgrades are given in refs. [60, 61]. The main goals of the HB and HE upgrades were: to replace the HPD photodetectors, which produced anomalous signals [62] and showed signal degradation [63]; to increase the segmentation to that shown in figure 24, to allow for both layer-dependent corrections for the




Figure 22. Schematic view of the CMS hadron forward calorimeter, HF. The yellow region represents the steel absorber with embedded quartz fibers; the grey shaded area to the left represents fibers which deliver the signals to light guides that penetrate the steel plug shielding; the white rectangle to the left of the light guides represents the frontend readout boxes, which house the photomultiplier tubes. The dimensions shown in the diagram are in millimeters.




Figure 23. Schematic diagram of the HCAL readout. The frontend electronics chain begins with analog signals from the HPDs, SiPMs, or PMTs. The HB and HE used HPDs prior to the Phase 1 upgrade and SiPMs after, while the HF uses PMTs. These signals are digitized by the QIE chips in the readout modules (RMs). The “next-generation clock and control module” (ngCCM) is part of the frontend control system. The digitized data are sent to the backend μ HTRs (μ TCA HCAL trigger and readout cards). Data from multiple μ HTRs are concentrated into the AMC13 cards and forwarded to the CMS central data acquisition (DAQ) system. The μ TCA cards also send data to the trigger system. The AMC13 cards also distribute the fast commands arriving from the CMS timing and control distribution system (TCDS) within each μ TCA crate. The detector control system (DCS) software communicates with the “next-generation frontend controller” (ngFEC).

observed radiation damage to the scintillating tiles [64] and better rejection of energy deposits from pileup interactions; to increase the readout bandwidth to allow for the larger number of channels; to add signal arrival time measurements in the HB, HE, and HF; and to standardize the readout electronics across the different calorimeter systems. For the HF, the PMTs were also replaced because they were a source of anomalous signals [62].




Figure 24. The longitudinal and transverse HCAL segmentation for Run 3. Within a tower, layers with the same color are routed to the same SiPM. The location of the frontend electronics is indicated by the letters FEE. Reproduced from [54]. © 2020 CERN for the benefit of the CMS collaboration. CC BY 3.0.

The HCAL plastic scintillators are subject to radiation damage, which results in the reduction of the signal output. The relevant primary characteristics of the LHC operation are the total delivered integrated luminosity, which determines the radiation dose received by the scintillator tiles, and the average instantaneous luminosity, which controls the dose rates. Radiation effects were evaluated by studying the tile performance during the 2017 LHC operation, corresponding to a delivered integrated luminosity of about 50 fb^{-1} . Signal reduction was studied as a function of dose rate. These measurements provide unique information on the radiation damage at dose rates significantly lower than previously studied. The HE tile results were obtained using several complementary methods: a movable radioactive source that can access all the tiles to compare their signal output before and after the 2017 data-taking period; inclusive and isolated-muon energy deposits produced during pp collisions; and a laser calibration system.

The laser system consists of a triggerable excimer laser and light distribution system that delivers UV light (351 nm) to the scintillator tiles in layers 1 and 7 via quartz fibers, as well as directly to the photodetectors. During the 2017 data-taking period, pulses of laser light were injected between LHC fills when there were no collisions. Laser data were collected throughout the 2017 data-taking period. Figure 25 (left) presents the relative signals in layer 1 versus dose for tiles in the ieta range 21–27. The signals show an approximately exponential decrease during periods of stable luminosity, with slopes that depend on the dose rate. Tiles at smaller ieta show more damage per dose than those

at larger $i\eta$, implying that at a fixed dose the damage to the scintillators increases with decreasing dose rate, within the range of our measurements.




Figure 25. Left: relative laser light signal versus the accumulated dose for scintillator tiles in layer 1 and $i\eta$ range 21–27. The average dose rate R for each set of points is given in the legend. The vertical scale is logarithmic and subsequent sets are shifted up by a factor of 1.03 relative to the previous set for better visibility. Each set starts at a dose corresponding to an integrated luminosity of 7 fb^{-1} . The vertical bars give the scaled statistical uncertainties. Right: ratio of the signals from the ^{60}Co source observed before and after the 2017 data-taking period for scintillator tiles in the HE as a function of $i\eta$ and layer number. Tubes in layers 0 and 5 have obstructions and cannot be accessed.

Each individual tile in the HCAL is designed to be serviced by a movable ^{60}Co radioactive source using small tubes that are integrated into the calorimeter. The ^{60}Co source provides photons with energies of 1.17 and 1.33 MeV. The source is attached to a wire that guides it through the tubes. All tiles except those in layers 0 and 5, whose tubes have obstructions, can be accessed. The data were collected during the periods when the LHC did not operate, before both the 2017 and 2018 data-taking periods. Values of the ratio, averaged over $i\phi$ as a function of the scintillating tile layer number and tower index $i\eta$, are shown in figure 25 (right). The signal loss is smaller for tiles at larger radial distance from the beam and for layers that are deeper in the calorimeter.

Because the damage depends on the layer number, increased segmentation allows for more accurate damage corrections to be applied as a function of integrated luminosity, reducing the degradation of the resolution.

The on-detector upgrade components must operate in a radiation field and withstand cumulative effects from both the total ionizing dose (TID) and nonionizing energy loss (NIEL), as well as single-event effects (SEE). The required tolerance for NIELs is quoted as that corresponding to an equivalent fluence (cm^{-2}) of 1 MeV neutrons. Expected SEE rates are characterized by the fluence (cm^{-2}) of hadrons with energy exceeding 20 MeV. Requirements for the HB, HE, and HF are shown in table 5. Since the HE will be replaced after Run 3, the numbers given for the HE are the requirements for Run 3, while the HB and HF values correspond to the full HL-LHC duration.

5.2.2 HB/HE/HO photodetector upgrade

A significant part of the HCAL Phase 1 upgrade was the replacement of the HPD photodetectors with SiPMs for HB, HE, and HO. All three subdetectors used HPDs during Run 1 [65], and were converted to SiPMs at different times. The HO SiPMs were installed prior to the start of Run 2. For

Table 5. Radiation requirements for the Phase 1 upgrade. The HE numbers are for Run 3, while the HB and HF values correspond to the full HL-LHC duration.

Detector	TID [krad]	NIEL [cm ⁻²]	SEE [cm ⁻²]
HB	3.1	1.1×10^{12}	2.0×10^{11}
HE	0.9	9.0×10^{10}	1.6×10^{10}
HF	4.1	7.0×10^{11}	1.8×10^{11}

HE, a ϕ sector of 20° was instrumented with SiPMs during the 2016/2017 year-end technical stop, and the rest of the detector was converted during the 2017/2018 technical stop. The HB used HPDs for all of Run 2, with SiPMs installed prior to the start of Run 3. The grouping of layers in a tower to individual photosensors, i.e., the longitudinal readout segmentation, is indicated by color in figure 19 prior to the upgrade and following the upgrade in figure 24. Prior to Run 3, most of the HB towers had a single readout segment, except for $\eta = 15$ and $\eta = 16$. The HB readout segmentation increased from one to four channels starting in Run 3. For HE, the readout segmentation increased from two to three channels per tower in Run 1 to six to seven channels during Run 2 as the SiPMs were installed.

The SiPM has many advantages over the HPD, including high photon detection efficiency (PDE), high gain, a large linear dynamic range, rapid recovery time, somewhat better radiation tolerance, and insensitivity to magnetic fields. The active area of each SiPM is circular, with diameters of 2.8 mm (sensing up to four fibers) and 3.3 mm (sensing up to seven fibers). In each HE readout unit, 19 HPD channels were replaced by 48 channels of SiPMs. For the HB, each readout unit has 64 channels of SiPMs.

Arrays of eight individual SiPMs were placed into a ceramic carrier known as the “SiPM package”. The packages for the HE and HB are identical. They were designed by CMS and fabricated at Kyocera. Figure 26 (left) shows a top-and-side view of an eight-channel SiPM array in its package. Each SiPM is connected electrically to the package via two wire bonds, one for the signal and one for the bias voltage. The 16 pins at the bottom of the package connect the SiPMs to the downstream electronics. The package precisely locates each SiPM on the ODU, so its mechanical dimensions are important. Before sending them to the SiPM vendor, the mechanical tolerances of each package were measured at the CERN Metrology Laboratory, and more than 30 parameters were compared with the design specifications. Figure 26 (right) shows 48 ceramic SiPM packages prepared for measurement. The yield of good packages was very high with less than one percent rejected for being out of tolerance.




Figure 26. Left: top and side views of an eight-channel SiPM array in its ceramic package. Right: ceramic packages at the CERN Metrology Laboratory for characterization measurements.

Before choosing a vendor and entering final SiPM production for the HE, and then later the HB, careful studies were carried out on preproduction samples from several vendors. The preproduction samples were of order 1500 channels. Key SiPM parameters included PDE, gain, recovery time, uniformity of breakdown voltage, and dark current [66]. A small sample of SiPMs was also irradiated for radiation damage studies [67]. Based on the results, the Hamamatsu SiPMs were chosen for final production.

The quality control program was virtually identical for the HE and HB. We measured the signal response (PDE times gain), breakdown voltage, dark current, forward resistance, and capacitance for each channel. We also looked for spurious noise pulses in every channel. A total of 1400 production HE arrays were tested, and 104 arrays, or 7.43%, were rejected. Since there are eight SiPM channels per array, this corresponded to a yield of good SiPMs better than 99%. For the HB production, a total of 1680 production arrays went through the same quality control testing, with only 82 rejected, again giving a yield of good SiPMs better than 99%. The uniformity was also excellent throughout the two production runs. As an example, figure 27 shows the signal response for 3600 HE SiPMs, with an RMS of about 1.5%. After quality control, a total of 864 (1152) SiPM arrays were installed into the HE (HB).




Figure 27. Distribution of the signal response (photon detection efficiency times gain) for 3600 HE SiPMs.

5.2.3 Readout box

The RBX is an aluminum shell with snake-shaped grooves machined on the surface for insertion of copper cooling pipes. The RBX houses the RMs, control electronics, and the QIE-based frontend electronics, which are connected to the backend in the control room via control and data fibers. The HB and HE frontend electronics, shown schematically in figure 28 are very similar, and primarily differ to reflect the mechanical, structural, and channel multiplicity differences between the two subsystems.

Each HE, HB, and HO RM has 48, 64, and 18 SiPMs, respectively. Each RM houses a single SiPM control card and four frontend readout ADC cards (three cards for the HO). Each RM also houses an ODU. As shown in figure 23, the RBX consists of four RMs installed together with a ngCCM and a CU. In HB, two ngCCMs are installed per RBX, each independently controlling two




Figure 28. Upgraded HCAL endcap (HE) and barrel (HB) DAQ readout chain, including the SiPMs, frontend readout electronics, and backend system. Reproduced from [60]. CC BY 3.0.

RBMs. All RBX components receive the power and communicate via a backplane PCB. The RBX receives its fast and slow control signals through the ngCCM module. This module receives the clock and fast reset signals via optical links from the backend, and distributes these signals to each card in the RBX. Configuration information for the frontend is received over the optical link as well. Finally, the ngCCM is responsible for aggregating and returning all status information about the RBX, which is transmitted to the backend via an optical link.

Each RBX is installed on top of a scintillator megatile in the locations indicated in figure 24. One RBX covers 20° in ϕ . The total number of RBXs is 108, excluding those used for the HF. The RBX is the part of the common CMS infrastructure; therefore, each RBX has individual power, water cooling, and dry gas lines connected to common water and gas manifolds in order to provide proper and stable cooling for the HCAL electronics and decrease the relative humidity around the SiPMs. The maximum temperature of RBX components does not exceed 35°C , and the relative humidity in the SiPM region stays constant within 5%, except on rare occasions when there are cooling or dry gas issues.

The first upgraded RBX for the initial sector of the HE instrumented with SiPMs was installed in 2017. One HE RBX was upgraded at the end of 2016, the rest of the HE in 2017/2018, and the HB RBXs in 2019 after Run 2.

5.2.4 Photodetector control

The use of SiPMs required the design of completely new high-density photodetector control electronics. The new design was done initially for the HO photodetector upgrade, and consisted of an additional board placed next to the QIE cards in the RM, shown schematically in figure 29 (left) and as constructed in figure 29 (right). The HE and HB subdetectors reuse the same design topology to minimize design, testing, and validation time, although small changes were needed to accommodate the different SiPM types. The control board uses a single +8 V power voltage from the RBX backplane and is controlled via an I2C-compatible slow-control connection. The board provides individual bias voltage regulation and leakage current measurements, limiting of bias voltage current for SiPM protection, a Peltier cooler control, and remote temperature/humidity readout.




Figure 29. Left: control board block diagram. Right: HB card pack with one control and four frontend boards. The flex cables provide signal and bias voltage connections to 64 SiPMs.

Bias voltage and temperature regulation strongly affect the SiPM noise and stability. In order to limit the variation in the gain of all the detector channels to 2%, the bias voltage and temperature are strictly controlled, with several key parameters shown in table 6. The measured in situ temperature stability was 0.3°C .

Table 6. HCAL SiPM requirements for the Phase 1 upgrade. The design (measured) parameter values are shown. For the design (measured) value of the BV stability, the RMS (peak-to-peak) is quoted.

Parameter	HB	HE
Channel number per RM	48 (48)	64 (64)
Bias voltage [V]	65–75 (0–80)	65–70 (0–80)
BV resolution [mV]	40 (20)	40 (20)
BV stability [mV]	40 (26)	40 (26)
BV ripple/noise (RMS above 10 kHz) [mV]	7 (3–7)	7 (3–7)
Temperature monitoring [$^{\circ}\text{C}$]	0.25 (0.1)	0.25 (0.1)
Max operating current [μA]	500 (500)	1000 (1000)
BV current resolution [nA]	122 (122)	244 (244)
Total ionizing dose [krad]	3.1 (17)	0.9 (17)
1 MeV-equivalent neutron fluence [cm^{-2}]	9×10^{11} (2×10^{12})	1.1×10^{12} (2×10^{12})
Power for Peltier cooler	(6.5 V $\times$ 1 A)	(6.5 V $\times$ 2 A)

5.2.5 Optical decoder unit

The optical decoder unit (ODU) is the interface between the clear optical fibers bringing light signals from the HCAL scintillator and the SiPMs that convert the light into electrical signals. Each ODU contains a network of short (about 10 cm) clear plastic fibers that remap the planar geometry of the incoming light signals from the megatiles into a tower geometry that is more appropriate and useful for physics analysis. While the size, shape, and some details are different for the HE and HB ODUs, the basic features are the same. Each RM in HB and HE contains one ODU, for a total of 144 ODUs each in HB and HE.

The ODU is essentially a box containing the network of fibers described above. The fibers are fabricated into cables with polished connectors at each end, which mate with the fiber connectors from the detector. The cables are cut in half to create “pigtailed”, which form the basic units for assembly into the ODU. Each HB (HE) pigtail has up to 18 (12) fibers. The routing and connections of the fibers are shown in figure 30 (left). At the top of the picture, the pigtail connectors are attached through open slots in a machined aluminum patch panel, where they connect one-to-one to the fibers coming from the HCAL detector layers. At the left of the picture, the fibers are mapped to their designated holes in the so-called “cookie,” a precisely machined piece of plastic containing 64 (48) holes for the HB (HE) which directs the light to individual SiPMs. The HE cookie is a solid piece of polyether ether ketone plastic, while the HB cookie has two plastic layers with an insulating foam layer in the middle to prevent heat from leaking to the SiPMs. The number of fibers in each hole varies from one to seven, depending on the depth segment being read out. The fibers are glued in place in the cookie, which is optically finished with a diamond flycutter. All the precision-machined parts were fabricated in local industry, while the ODU assembly, inspection, and testing were done by CMS. A completely assembled production HE ODU is shown in figure 30 (right).




Figure 30. Left: view of a spare HE optical decoder unit (ODU), showing its light box and fiber mapping. The fibers route from the patch panel at the top to the “cookie” at the left. The side panels are clear rather than opaque for display purposes. Right: a production HE ODU. The clear fiber pigtail connectors attached to the patch panel are visible at the top. The plastic “cookie” is seen at the front of the ODU.

Extensive quality control testing of both the individual fibers and the fully assembled ODUs was performed to ensure the integrity of the light transmission. Each cable was tested by connecting it to a similar cable of wavelength-shifter fibers that were illuminated individually by LEDs. The light throughput of each fiber in the cable was measured with a photodiode. To make sure there were no problems with the connectors at either end, the cable was reversed, reconnected, and measured again. To test the fully assembled ODUs, light was injected into the ODU through the connectors in

the patch panel, fiber by fiber, and measured by a large photodiode at the face of the cookie. A total of 144 fully qualified ODU's were installed into both the HB and HE.

5.2.6 Frontend readout card

The signals from the SiPMs are integrated and digitized by the QIE11 ASICs [56–59]. The QIE11 integrates negative input charge pulses in 25 ns buckets and digitizes the result with approximately constant resolution over a large dynamic range. The QIE11 achieves an effective 17-bit dynamic range with approximately 1% resolution with only 8 bits through the use of a custom floating point analog-to-digital converter (ADC) with pseudo-logarithmic response. The ADC has a programmable gain via user-configurable input current shunts and a low input impedance ($< 15 \text{ W}$) that is suitable for use with SiPMs. A 6-bit time-to-digital converter (TDC) is also included on the QIE11 chip, consisting of a programmable-threshold discriminator which detects arrival time of the input pulse in 0.5 ns bins. The phasing of the charge integration window relative to the input clock can be adjusted by the user in 0.5 ns steps over the whole 25 ns window. The nominal sensitivity of QIE11 is 3.1 fC per count at the low end (with no programmable input shunt selected). The gain of any particular chip can vary from this nominal value due to process variations, so all chips are calibrated with better than 1% precision. The nominal maximum charge that can be digitized is approximately 350 pC, yielding the nearly 17-bit dynamic range.

Each readout card in the HE (HB) supports 12 (16) SiPM channels and contains a corresponding number of QIE11 ASICs. In addition, each card contains one Microsemi ProASIC3L FPGA that acts as an I2C bridge between the slow-control unit (ngCCM) and the frontend chips, one or two Microsemi IGLOO2 FPGAs that serialize and format the data, and one VTTx [68] module providing two 4.8 Gb/s optical links. The necessary power for operating the frontend readout card is supplied by the CERN FEASTMP DC-DC converters, which are radiation and magnetic field tolerant.

Digital data from the QIE chips are serialized and formatted by two (one) Microsemi IGLOO2 FPGAs in the HB (HE). The flash-based IGLOO2 FPGA achieves sufficient radiation tolerance for the HCal frontend, better than typical SRAM-based FPGA technologies. In the HE, there is sufficient bandwidth to transmit all 8 bits of ADC and 6 bits of TDC data per channel per bunch crossing; however, in the HB, bandwidth constraints require the reduction of TDC information to two bits, which are used to encode four arrival time scenarios: prompt, slightly delayed, and significantly delayed times of arrival, plus the case where no valid pulse is present. After formatting by the IGLOO2, data from each QIE card are transmitted to off-detector backend electronics via 5 Gb/s VTTx optical links.

5.2.7 Slow and fast-control systems

The control systems for the HB, HE, and HF were updated to support the upgraded frontend electronics. For the Phase 1 upgrade of the frontend, the fast-control system, synchronized with the LHC clock, and the slow-control system (also referred to as the DCS), which runs at a much slower frequency than the LHC clock, are handled together within the same hardware. The fast-control system delivers the LHC clock with a maximum jitter of 1 ns, sends the orbit synchronization signal, delivers the so-called “warning test enable” (WTE) signal for the recording of calibration data, provides reset capabilities, and provides fast monitoring.

These requirements are similar for the HB, HE, and HF with a few exceptions: the HF has PMTs instead of SiPMs, and these do not need to receive configuration commands; the HB and HE have a

secondary, redundant control link that can replace the primary control link; and each subsystem has a different number of channels and cards. As a result, a similar architecture is used for the HB, HE, and HF control systems, although the physical realization of the hardware is different. The redundant-link scheme has changed from what was described in the CMS HCAL Phase 1 Upgrade TDR [60]. The redundant link was removed from the HF because it is relatively accessible. In HB and HE, the hardware configuration has been simplified thanks to the installation of more optical fibers. The redundant HE control link was successfully used during the last year of Run 2, when a primary control link failed due to a malfunction of its optical transmitter [69]. For Run 3, the control link with the highest optical power between primary and secondary is used.

The hardware includes the ngCCM modules in the frontend, the ngFEC modules in the backend, and pairs of optical fibers linking the modules (each pair supports bidirectional serial communication), as shown in figure 31. The ngFEC module is a μ TCA baseboard with mezzanine cards and pluggable optical transceivers (SFP+), and is based on the FC7 Kintex 7 FPGA AMC board [24]. Its main functions are to receive TCDS signals from the μ TCA backplane, provide an interface to the DCS computers, merge fast control and slow control over the same bidirectional link used to communicate with the ngCCM, maintain a fixed latency for the fast-control signals across power cycles and ngFEC optical ports, make use of the ngCCM redundant scheme, and support up to twelve bidirectional links.




Figure 31. Block diagram of the HB controls. The ngFEC and ngCCM modules are needed to run and monitor the frontend electronics. All control links connecting the service rooms to the experimental cavern are optical. A secondary link connecting the ngFEC and the ngCCM is available in case of a primary link failure.

The ngCCM is connected electrically to the rest of the frontend over backplanes. The optical links run a modified scheme of the CERN-developed 4.8 Gb/s GBT communication protocol [68, 70]. The GBT protocol was tailored for the GBTx frontend ASIC; however, the GBTx was not available in time for the original HCAL Phase 1 schedule. In the ngCCM, the Microsemi IGLOO2 FPGA was used as a replacement for the GBTx. The performance of the dynamic characteristics of the IGLOO2 FPGA are marginal at 4.8 Gb/s; for CMS purposes, to allow an ample safety factor, it was decided to modify the protocol and run it at half speed. The reduced bandwidth is still sufficient for

all control requirements of the HCAL frontend. The Phase 1 upgrade of the HCAL frontend and backend was successfully completed in October 2019.

5.2.8 Backend electronics: readout

During Run 2, the HCAL backend electronics were upgraded from VME-based to components based on the μ TCA architecture. This replacement was the first step in the Phase 1 upgrade, and the μ TCA-based electronics have operated with a mixture of the original and upgraded frontend throughout Run 2. The design of the backend is described in detail in ref. [60].

In Run 3, the backend and frontend host a much-increased channel count and data volume. With the completion of the Phase 1 upgrade, the number of HB (HE) channels is increased by roughly a factor of 3.5 (2.6), and each channel has a larger number of bits from the upgraded QIE10/11 ADC and TDC. This increase required efficient algorithmic development to calculate the trigger primitives within latency constraints and with the available resources of the μ HTR FPGA. This firmware was achieved, along with the enhancements to the trigger primitive calculations discussed in section 5.3.

The HCAL data acquisition system is managed by an online software suite that controls, configures, and monitors the HCAL electronics. Development of the HCAL online software began in 2005, and it continues to evolve significantly; for example, new components were recently developed to support the upgraded frontend and backend electronics. The software has two main layers that provide complementary functionality. The first layer is composed of custom software written in C++ using the XDAQ [71] framework for distributed data acquisition. For communication with the μ TCA electronics, the IPBus architecture [72] is used within XDAQ. Configuration parameters are provided in human-readable format by a version-controlled system, and larger hardware parameter sets (e.g., SiPM bias voltages, phase delays, etc.) are read from databases. The second software layer allows for coordination between different parts of the HCAL, and with other CMS subsystems. It is written in Java, and based on the RCMS framework [73]. The XDAQ and RCMS frameworks are described in detail in section 9.10.

The frontend electronics are controlled and configured by the ngCCM server, a separate program based on the C++ actor framework [74]. It detects and corrects for single-event upsets automatically, and stabilizes the SiPM temperatures by controlling the Peltier voltages. The ngCCM server communicates via IPBus with special command processors implemented in the ngFEC firmware for each I2C and JTAG channel in the frontend backplanes. The XDAQ-based software communicates with the ngCCM server via JSONRPC over WebSocket [75, 76] to control, configure, and monitor the frontend electronics. The HCAL DCS WinCC-OA communicates with the ngCCM server via a raw socket protocol to obtain voltage, current, temperature, and humidity readings and set the Peltier target temperature. A command-line interface to the ngCCM server additionally allows updates to be uploaded to the frontend FPGAs.

5.2.9 Voltage source upgrades

The Run 1 HCAL frontend electronics power system was based on the CAEN Easy infrastructure. The RBX electronics were fed with two power lines, nominally at 6.5 and 5 V, requiring a total power of around 90 W per RBX. The HPD high voltage (HV) and bias voltage (BV) and the HF PMT HV were provided using power supplies custom made in Bulgaria. The HB, HE, and HO RBXs were powered using CAEN A3016 power supplies, each module having six power channels

capable of powering three RBXs. The HF readout crates were powered using four CAEN A3100 modules for each HF module (positive and negative η). Successive photodetector, frontend, and backend upgrades required some changes to the powering system.

The HO photodetector upgrade, with the replacement of HPD the sensors with SiPMs, required new power supplies, since the new frontend generates the BV internally from the low voltages (LV). No change was required for the LV system, since the A3016 modules were able to provide the extra power demanded by the SiPMs.

The HF LV system for the new frontend electronics also required some changes. Unlike the old electronics, the new electronics only use one supply voltage, in the range 8–10 V, but with more power. During Run 2, the power was provided by eight A3100 units operating at 8 V, four each of the two HF detectors. Each of the sets was in an CAEN EASY crate fed with an individual A3486 power converter. However, a problem was observed during operation, correlated with instantaneous luminosity. The power supplies, located in the HF racks near the detector, were subject to high levels of radiation, causing occasional single-event upsets (SEUs), i.e., incidents in which the control of the modules was lost, leading to the loss of power in the frontend electronics. This resulted in the loss of about 150 pb^{-1} of data in 2017. For 2018, a mitigation protocol was developed, consisting of a fast software detection, reset, and recovery of the power supplies, followed by a reconfiguration of the frontend electronics. The protocol reduced the time of data loss to typically 1–2 minutes, and reduced the total data loss to about 50 pb^{-1} . The rate of SEUs scaled linearly with instantaneous luminosity, with about 20 SEU events in 2018; further, the power cycling was expected to shorten the lifetime of the power supplies, and the high level of irradiation made maintenance prohibitively difficult. Hence a more robust solution was implemented in LS2: the power supplies were moved to an area under CMS where radiation is minimal. Long cables from this location lead the power to the frontends. To compensate for the voltage drop on the cables, the A3100 power supplies have been replaced by A3100HBP units, described below, which can provide a high enough voltage to ensure the frontends receive the nominal operational voltage of 8 V. Following these changes, no SEU incidents were observed during the first year of Run 3.

The HE and HB upgrades, with the increase in the number of SiPM readout channels, led to more changes to the powering system. The HB and HE custom made power supplies were replaced by Keysight N6700C mainframes with N6736B modules. Twenty N6700C units currently provide the BV for the HB and HE SiPMs. The HE A3016 modules were replaced by A3100HBP units. The CAEN A3100HBP model provides one channel with 8–14 V, 50 A, and 600 W output and can power up to two HE RBXs in parallel. These units are also used in the HF, with ten units operating at 10 V used in each half of the HF. A transient voltage spike in one of the HE power supplies, caused by a power cut in June 2018, damaged the inputs of two RBXs. This affected HEM15 and HEM16, which cover 2% of the total HCAL acceptance and 3% of the HB+HE. After identifying the cause of the problem, CAEN introduced a modification to the power supplies to prevent such transients from occurring again. Independently, the HCAL Collaboration also developed an external circuit to suppress such transients. Currently, all the HE RBXs are protected with both the CAEN modification and the external overvoltage protection circuit.

The extra readout channels of the HB SiPMs required more power than what was practical with the A3100HBP units, and Wiener Marathons (PL 508 with five channels 5–15 V/40 A) were chosen as replacements. Eight Marathons are used to power the entire HB frontend electronics. To prevent

damage due to transients, overvoltage protection circuits similar to the ones for the HE are installed on each of the Marathon output channels. New LV cables between the supplies and the RBXs were also installed to keep the voltage drop coming from the additional power within the safe operating range of the frontend electronics.

5.2.10 Photodetector and system calibration instrumentation

Changes in the light yield are expected due to aging of the scintillators, radiation damage, and variations in the photodetector response. The HCAL calibration and monitoring systems are designed to determine the absolute energy scale, to monitor the calorimeter system for changes during the lifetime of the detector, and to derive the energy scale correction factors. In particular, two complementary methods are used for monitoring. The first method consists of a movable radioactive wire source, with source tubes installed in every megatile in such a way that the tubes cross all of its tiles. Measurements for all tiles can be made during long LHC shutdowns to validate the readout-channel mapping. This system was used during the Phase 1 upgrade to provide a relative calibration when the SiPMs replaced the HPDs. The second method is based on light injection from two sources, a UV laser and an LED system. The UV laser light is injected into two layers of each wedge to monitor radiation damage to the scintillator, and both the UV laser and LED light is injected into the optical decoder box that has the photodetectors, to monitor damage to the SiPMs. A calibration module (CM) inside the frontend RBX is responsible for delivering the UV laser and LED light to the photodetectors. The CM was redesigned as part of the Phase 1 upgrade.

Each RBX contains one CM, which includes an LED pulser for the SiPMs, as well as a laser light distribution system. Fibers from the CM are connected to each of the interstitial microfibers in the light mixers, allowing for the LED or laser signals to be delivered to the SiPMs. The CMs also contain pin diodes which directly measure the amplitude of the LED or laser signals, providing a reference value for the calibration. The design of the LED pulser board is the same as that of the original HB/HE calibration module. The controls for the pulser, however, are significantly enhanced. In particular, the CMS global WTE signal is used by the CM to generate LED pulses during the LHC orbit gap. The LED pulse can be positioned between 2 and 65 535 triggering clock synchronization counts (MCLK) from a resetting WTE signal. This pulse can be delayed relative to the MCLK signal from 0 to 25 ns in 0.5 ns increments. The pulse amplitude and width are programmable from about 85 mV to 4.75 V and 0.5 ns to 25 ns, respectively. The pulser board contains six TE Connectivity AMP 147323-1 connectors providing a total current of 0.75 A at 5.5 V DC and a ground for powering the PIN diodes. The pulser board receives communication and MCLK from the ngCCM through the backplane header.

5.2.11 HF upgrade

When shower particles pass through the HF quartz fibers, Cherenkov radiation is produced. However, anomalous signals can also be produced by muons from pp collisions or beam halo interactions whose trajectories pass in the vicinity of the RBXs. Relativistic muons generate Cherenkov radiation when they pass through the glass window of the PMTs. Since the sampling correction for the HF calorimeter is large, these “window events” result in very large signals. Signals produced by muons can be distinguished from shower light through timing: the shower signal is delayed relative to the muon signal due to the longer path length and the lower speed of light in quartz. The PMTs

used in the original HF detector were Hamamatsu Corporation R7525s, with a glass window that is 1.2 mm thick at the center, increasing to 6 mm at the edge. The HF PMT upgrade replaced these with Hamamatsu R7600-M4s. The new PMTs have thinner windows (1 mm of UV glass) and four anodes arranged in a 2×2 grid. The light from the fibers of a given ieta-iphi tower is spread out over the face of the PMT and generates signals on all four anodes; the diagonal anodes in the 2×2 grid are grouped together and read out by a single QIE10 chip. The resulting dual-anode readout allows for the discrimination of signals from muons interacting directly with the PMT, which will typically produce a signal concentrated on a single anode closest to the impact point. The amplitude of signals from muon window events between the old and new PMTs is compared in figure 32.




Figure 32. Cherenkov signals generated in the PMT windows from a muon test beam. Thin windows in the Phase 1 upgrade four-anode PMTs produce smaller signals (red) than those produced in the original PMTs with thick windows (black).

The new PMTs were extensively tested in several beam tests [77, 78] and in CMS during data taking. Algorithms utilizing the signal imbalance between the anodes were proposed and tested during the beam tests. However, a new readout system was needed to take advantage of the additional channels. To be cost-effective, frontend readout cards were redesigned for two-channel readout of the four-anode PMTs (where each readout channel sums two anodes). The backend was also changed from the old VME-based to the new μ TCA-based system.

The new frontend cards included the upgraded QIE10 ASIC, which is similar to the QIE11, except for having a constant $20\,\Omega$ input impedance and no programmable current shunting. As described earlier, the QIE10 chips also have a TDC to measure the arrival time of the signals. Each frontend readout card for HF contains 24 QIE10 ASICs, each digitizing a single dual-anode PMT channel, as well as two Microsemi IGLOO2 FPGAs and three VTTx modules providing six 4.8 Gb/s optical links. Similar to the upgraded electronics in HB and HE, the power is supplied by CERN FEASTMP DC-DC converters. In addition to the replacement of the PMTs and the readout electronics, improvement in data transfer was planned to handle the increased load due to the TDC and two-channel readout.

The HF calorimeter is in a high-radiation area, leading to doses of 1000 Gy in fibers near the beam pipe. To monitor radiation damage, some of the quartz fibers are equipped with a special laser

system that allows both the incident laser light intensity and the intensity of the laser light after traversing the quartz fiber to be measured with the same PMT. In Run 1, the light was provided by a laser system located in the underground service cavern, which is several meters away from the HF calorimeters. The light was then split into several fibers before it reached the HF quartz fibers.

A new laser device was developed that utilized solid state laser diode technology with a wavelength of 450 nm (allowing for the removal of the wavelength shifters used in the old system), and is based on the HF Phase 1 LED calibration unit concept. The laser diode driver and the control circuit reside on a mezzanine card mounted on the QIE board. Optics mix the light and distribute it to four output fibers. The module is installed in the HF frontend crates, which are attached to the HF calorimeters. Figure 33 (left) shows a schematic view of the laser daughter board as part of the QIE frontend card. The upgraded HF light distribution system for monitoring radiation damage is shown in figure 33 (right). The settings of the new system are more flexible, since it is an integral part of the upgraded frontend, allowing for, e.g., a 32 ns delay range adjustable in 0.5 ns steps, and a pulse width adjustable from less than 3 to 30 ns full-width at half-maximum. There is also an on-board PIN diode to independently monitor the performance of the light source.




Figure 33. Left: schematic of the new HF QIE laser board, used in the HF radiation monitoring system. Right: sketch of the upgraded HF radiation damage monitoring system light distribution and electronics chain.

5.3 Trigger

The HCAL trigger is based on the HCAL calorimeter towers, which, in the HB and HE, are composed of projective calorimeter readout channels summed together in depth. While most readout channels have $\Delta\phi = 5^\circ$, the HE readout channels with $|\eta| > 1.74$ have $\Delta\phi = 10^\circ$. To form trigger towers with $\Delta\phi = 5^\circ$, thereby matching the rest of the HCAL, these HE readout channels are split into two in ϕ , with each unit assigned one-half of the measured energy. In the HF, trigger towers are formed from six physical HF towers, with energies of the long and short fibers summed.

The HCAL trigger primitives (TPs) are formed for each trigger tower by combining the information from individual calorimeter towers. Since the HCAL QIE10/11 chips use 8-bit nonlinear ADCs with several ranges, the energy reported by the QIE is first linearized as a 10-bit word via a dedicated look-up table. The linearization depends on several parameters, including the gain, the conversion from the collected charge in fC to ADC counts, and the pedestal. After linearization, the various depth samples are summed, and the amplitude of the pulse is estimated using a pulse filtering scheme that

subtracts out-of-time pileup. These sums define the various trigger channels, and a corresponding TP is produced by the μ HTR. Finally, the trigger tower transverse energies are compressed to the range 0 to 128 GeV with a least significant bit (LSB) of 0.5 GeV, and transmitted to the L1 trigger.

The TP amplitude is reconstructed by taking the sample-of-interest (SOI) bunch crossing as the main contribution to the total signal amplitude. The measurement from the SOI alone works well for HF, where the pulses are shorter than the 25 ns bunch spacing, but for HB and HE, a significant fraction of the pulse leaks into the subsequent bunch crossing. In Run 2, the sum of energies in the SOI and SOI+1 was used to reconstruct the TP energy; this method accounts for the signal leakage into SOI+1, but also incorporates the energy from out-of-time pileup interactions in SOI+1. For Run 3, a new algorithm was developed that instead uses the SOI and SOI-1, i.e., the preceding bunch crossing. The scheme subtracts a weighted amount of the measurement from SOI-1, mitigating the inherent leakage of the out-of-time pileup from SOI-1 in the SOI. To account for the signal leakage into SOI+1, a correction factor is derived from the known pulse shapes rather than using the measurement from SOI+1, thereby avoiding incorporating contributions from out-of-time pileup.

In addition to the TP generation in the HB and HE, six feature bits of information are also generated that can be transmitted to the L1 trigger. These bits facilitate encoding information about: (i) the longitudinal shower profile data for use in calibration, lepton isolation, and identification of minimum ionizing particles, and (ii) the shower time data constructed from the TDC information available in each constituent channel of the trigger tower. Configurable look-up tables determine which time windows within the bunch crossing of interest are represented by the available TDC codes. These time window boundaries have a granularity of 0.5 ns. Starting in Run 3, a subset of the feature bits is used to flag signals characteristic of exotic long-lived particle decays, using either the TDC timing to mark hits with late arrival times or the shower profile data to mark distinctive energy deposits in the various layers of the HCAL. These bits are used by the L1 trigger to select hadronic signatures from long-lived particles with decay lengths of 1–2 m which decay prior to or within the HCAL. The timing precision that can be achieved is estimated to be within 1–2 ns. The resolution is dominated by interchannel synchronization uncertainties and shower-by-shower fluctuations, as determined from studies done using highly energetic hadronic showers in 2018 data.

5.4 System and beam tests

Beam tests were performed with the upgraded Phase 1 frontend at the H2 beam line of the SPS at CERN. The upgrade electronics were installed in the 20°-prototype HE wedge at H2, and were read out by four RMs split between two RBXs. One RBX and two RMs were replaced with the upgraded electronics, while the remaining ones were left unchanged as a reference. The system was tested with 150 GeV muons and with pions having energies ranging from 30 to 300 GeV. Six time samples were recorded for each event, with the beam timed to arrive in the fourth time sample. The total charge for an event was taken as the sum of the last three time samples centered around the beam arrival. The pedestal for each event was estimated from the sum of the first three time samples before the beam trigger and was subtracted from the total charge.

The 2015 beam tests served as the first full test of the entire upgraded readout system and proved the system was functional. The 2017 test beam with production frontend electronics served to quantify the final performance of the upgraded detector. Representative results from this data are shown in figure 34. Using the muon data, the response of the detector instrumented with SiPMs was

measured to be 4.3 photoelectrons per layer. The pion data were used to derive the energy response and resolution of the upgraded detector. For pions, the number of photoelectrons per GeV was measured to be 32.1. The shower energy was measured by taking the sum of all charge collected by each channel in a 3×3 grid of towers around the beam direction. The relative contribution of each detector channel to the energy sum was adjusted using its response to muons to account for differences in the detector response between channels. The beam test was also used to make a unique measurement of the shower profile as a function of depth with a special ODU that allowed the individual readout of each layer in one tower. These studies were particularly important because they measured the maximum energy that can be deposited in any particular layer.




Figure 34. Left: energy resolution of the upgraded prototype detector as a function of the pion beam energy, shown with and without the channel calibrations derived from the response to muons. Right: longitudinal shower profile measured using the special ODU. Bands containing 68, 95, and 99% of all events for each layer are shown.

Irradiation tests of the full suite of frontend electronics were performed at the CERN High Energy Accelerator Mixed-Field (CHARM) facility with the goal of determining the expected error rates from SEEs, as well as the expected lifetime performance related to the cumulative edamage from ionizing and nonionizing radiation. The electronics were subjected to an ionizing dose of 202 Gy, a 1 MeV neutron equivalent fluence of $1.96 \times 10^{12} \text{ cm}^{-2}$, and a > 20 MeV hadron equivalent fluence of $5.88 \times 10^{11} \text{ cm}^{-2}$. At the end of the irradiation, the frontend electronics showed no indication of permanent damage, aside from the expected inability to program the FPGAs. The SEE rates were determined based on bit error counting in the encoded data stream and the pseudo-random bit sequence in the control link. The combined SEE rate was estimated to be less than 30 per fb^{-1} of data. Other upsets requiring a reset or other interventions are estimated to occur at a rate of about one per day.

5.5 Performance

5.5.1 Endcap photodetector performance in Run 2

The photoelectron peaks in the pedestal and low-intensity LED runs, shown in figure 35 (left), can be used to measure the SiPM gains using a single fit function to the observed multipeak charge

spectrum. The gain was found to be stable during the 2017 data taking at the 1% level, corresponding to 42 ± 1 fC per photoelectron, where the uncertainty is the RMS across the 184 channels equipped with SiPMs. The SiPM dark current was also monitored during data taking, as shown in figure 35 (right). The slope of the fitted line is proportional to the SiPM area. The deviations from linearity are due to SiPM annealing when the beam is off and to variations in instantaneous luminosity. With this SiPM dark current growth rate, we expect 110 MeV of noise at the end of Run 3 (500 fb^{-1}).




Figure 35. Left: pedestal distribution for a channel with QIE11 + SiPM readout. The charge is integrated in a time window of 100 ns. The QIE pedestal and photoelectron peaks are visible. Right: dark current increase with the integrated luminosity in 2017, where the slope of the fitted line is proportional to the SiPM area. The deviation from linear behavior is due to SiPM annealing in the absence of beam and variation in the instantaneous luminosity.




Figure 36. Left: energy deposit from muons in pp collision events in the HE tower corresponding to $i\eta = 20$ and depth = 5. The energy spectrum is fitted using the convolution of a Gaussian function with a mean of zero and a Landau distribution. The fitting function has three free parameters: the Landau location parameter, the Landau scale parameter, and the width of the Gaussian. Right: the most probable value of the muon energy deposit per layer in HE towers as a function of depth for different η regions. The vertical bars represent the statistical uncertainty. Muons from collision events are considered when their trajectory is contained within a single HCAL tower. The muon signal peak is fitted with the convolution of a Gaussian and Landau functions. The Landau location parameter is divided by the number of scintillator layers in the considered depth.

In Run 1 and Run 2, the HCAL calibration used E/p from isolated tracks for the energy scale measurement versus η and ϕ -symmetry for equalizing the response along ϕ [54]. After the Phase 1 upgrade, new calibration techniques are needed because of the increased depth segmentation. The scintillator and SiPM aging depends on η and depth, which necessitates depth intercalibration. The signals from isolated, minimum-ionizing muons traversing the HCAL provide an excellent probe for interdepth calibration, corresponding to approximately five photoelectrons per layer. As shown in figure 36 (left), a clear minimum ionizing peak from muons can be observed in collision events when the muon traverses one and only one HCAL tower. The energy deposit per layer is shown in figure 36 (right). The HE detector is homogeneous in depths 2–6, while depth 1 has a thicker and brighter scintillator layer. Because of this, only depths 2–6 are equalized.

5.5.2 HCAL performance in Run 2

For the combined ECAL and HCAL systems, including barrel and endcap, the relative charged pion energy resolution obtained from the test beam can be described as

$$\frac{\sigma}{E} = \frac{84.7\%}{\sqrt{E}} \oplus 7.6\%, \quad (5.1)$$

where E is in GeV. Corrections for the nonlinearity of the calorimetry system due to its noncompensating response to hadronic and electromagnetic energy depositions have been made [79]. The time resolution for energy deposits in the HB and HE, calculated by weighting the QIE digitization times by the associated energies, is 1.2 ns [80]. The HF energy resolution from the test beam [52] is

$$\frac{\sigma}{E} = \frac{280\%}{\sqrt{E}} \oplus 11\%. \quad (5.2)$$

The absolute calibration constants derived from the test beam modules were transferred to the full calorimeter system using an intercalibration from a radioactive ^{60}Co source [64]. The calibration was improved using 13 TeV collision data from 2016, as described in ref. [54]. From the ϕ -symmetry of energy flow in minimum-bias events, the HCAL was intercalibrated to within 3%. An absolute calibration uncertainty of 2% was determined using isolated pions with track momenta between 40 and 60 GeV showering in the HCAL. The energy resolution was found to be 19.4, 18.8, and 23.6% in the HB, HE, and transition region, respectively. Signal loss due to radiation damage is relevant in the HE at high $|\eta|$. It is monitored using a laser calibration system and the response to a ^{60}Co source, and is cross-checked with hadrons and muons from pp collisions [64].

During Run 2, new algorithms for reducing the effect of out-of-time pileup on reconstructed HCAL energies were developed. The HCAL response to incoming particles rises to its maximum within 10 ns, followed by an exponential decay, with 90% of the pulse contained within two 25 ns time samples (TSs). In Run 1, the bunch spacing was 50 ns, and the energy of hits was reconstructed by a simple sum of charges in the SOI and SOI+1 after the contribution from leakage currents was subtracted, with a correction factor applied to account for the tail extending beyond two TSs. In Run 2, the bunch spacing changed to 25 ns, and hence this method was expected to yield a poor energy resolution, since it incorporates contributions from pulses from preceding bunch crossings that overlap with the pulse from the SOI. The new algorithms developed for Run 2 subtract the energy of out-of-time pileup using pulse template fits.

In 2016–2017, two separate pulse-template fitting algorithms, referred to as Method 2 and Method 3, were deployed for offline reconstruction and in the HLT, respectively. The most recent algorithm, called “minimization at HCAL, iteratively” (MAHI), was developed and deployed for data taking in 2018 [81]. Notably, MAHI performs well enough to be executed within the HLT latency requirements, and, for Run 3, has also been ported to run on GPUs for further reduction in processing time. The MAHI algorithm was used for the legacy reprocessing of the CMS Run 2 data. Detailed results are presented in ref. [82].

5.5.3 HF performance in Run 2

The performance of the HF upgrade was evaluated during Run 2. Figure 37 (left) shows the distribution of the signal arrival time as a function of the collected charge [83]. Signals with a time around 7 ns are from showers in the calorimeter, while those at earlier times are due to Cherenkov radiation in the PMT window. The arrival time can thus be used to identify window events. While most signals due to PMT Cherenkov radiation come early, there is a tail to later times, due to imperfect synchronization. Elimination of such background events can be improved by comparing the signals obtained in both channels of the PMT, as shown in figure 37 (right). The charge asymmetry between the PMT channels is calculated as the difference divided by the sum of the signals in the two channels. For genuine events, this value should be close to zero. On the other hand, background events are observed well away from the central peak at zero since they are produced by stray muons hitting one quadrant or a side of the four-anode PMTs. The combination of the arrival time and asymmetry methods improves the elimination of these background events.




Figure 37. Left: HF signal arrival time, as measured in the TDC, versus the collected signal charge. All signals arriving within less than 5 ns are “window events”. The color indicates the number of events using the scale to the right of each plot. The data were taken in early 2017. Right: charge asymmetry between the two channels of a PMT versus the total charge in both channels. The light from genuine collision events is well mixed in the light guides before falling on the PMT, hence similar signals are expected in all four anodes, which are grouped into two channels. The so-called “window events” due to Cherenkov radiation in the PMT window most likely fall on one or two anodes, producing asymmetric signals.

Figure 38 shows the evolution of the missing transverse momentum with improvements to the HF anomalous signal identification based on the arrival time criteria (TDC filters), topological filters, and combined criteria (TDC, charge asymmetry, and topological filters). Topological filters have

been used since the beginning of Run 1 and are based on ratios of energies in the long and short fibers. The new filters, based on the timing and ratios of the PMT channel energies, are as effective as this topological selection. The combination of all the anomalous signal reduction techniques gives the best performance. Additional topological filters based on the shape of jets versus η and ϕ were developed to reject additional noise that escapes these filters [84].




Figure 38. Effect of filters on the HF anomalous energy contributions to the missing transverse momentum measurement. The methods developed based on hardware improvements installed as part of the Phase 1 upgrade are as effective as the topological selections used previously. Including both the new and old filters further reduces the anomalous missing transverse momentum.

6 Muon system

A central feature of the CMS experiment is a powerful system for triggering on and detecting muons. In the previous runs of the LHC, muons have been crucial to many of the physics results of CMS and have contributed to hundreds of published results, including the discovery of the Higgs boson. The importance of muons in the CMS physics program continues to remain high in Run 3 and beyond.

The objectives of the CMS muon system are to identify muons, measure their momenta, and provide signals for triggering on them. These goals are achieved with four complementary detector systems arranged in the steel flux-return yoke of the CMS solenoid. These systems provide efficient detection of muons over a large range of pseudorapidity. The location in the magnetized steel behind the calorimeters and solenoid ensures a low probability of penetration to the muon detectors by particles other than muons and neutrinos.

The physical arrangement of the muon detectors is shown in figure 39. The central section is configured in a barrel geometry with four roughly cylindrical stations at different radii from the beam axis. The endcap section is arranged in four planar stations in z in each endcap.

The drift tube (DT) system in the barrel covers $|\eta| < 1.2$ and is composed of drift chambers with rectangular cells. The DTs provide precise spatial measurements, as well as trigger information. This system is described in more detail in section 6.1.

The cathode strip chamber (CSC) system in the endcap comprises multiwire proportional chambers having cathode strips with an R - ϕ geometry and covering the region $0.9 < |\eta| < 2.4$. The




Figure 39. Schematic view in the r - z plane of a CMS detector quadrant at the start of Run 3. The interaction point is in the lower left corner. The locations of the various muon stations are shown in color: drift tubes (DTs) with labels MB, cathode strip chambers (CSCs) with labels ME, resistive plate chambers (RPCs) with labels RB and RE, and gas electron multipliers (GEMs) with labels GE. The M denotes muon, B stands for barrel, and E for endcap. The magnet yoke is represented by the dark gray areas. Reproduced from [8]. © 2018 CERN for the benefit of the CMS collaboration. CC BY 3.0.

CSC system provides both trigger and precision position information. Because of the higher flux of particles in the endcap region, the CSCs are designed to have a faster response time than the DTs. More details on the CSC system are found in section 6.2.

The resistive plate chambers (RPCs) are double-gap chambers operated in avalanche mode. The RPCs are located in both the barrel and endcap regions, and they complement the DTs and CSCs with a very fast response time that can be used to unambiguously identify the bunch crossing corresponding to a muon trigger candidate. The RPCs are further described in section 6.3.

Finally, GE1/1, a station of gas electron multiplier (GEM) chambers, is located in front of the inner ring of CSC chambers in the first endcap station. The GEMs have both a fast response and good spatial resolution, and they augment the CSCs in a region of high particle flux. More information about the GE1/1 station can be found in section 6.4.

The different groups of muon chambers are labeled with two letters: MB and RB for DTs and RPCs in the barrel region; ME, RE, and GE for CSCs, RPCs, and GEMs in the endcap regions. For MB and RB, these letters are followed by a single number (1–4) indicating the barrel index. Two indices follow ME, RE, and GE, where the first indicates the station in $|z|$ and the second indicates the ring in R . The value of each index increases with increasing distance from the center of the detector.

For each bunch crossing, fast trigger data (“trigger primitives”) are sent from each muon detector system (DT, CSC, RPC, and GEM) to the dedicated level-1 (L1) trigger muon track-finder hardware

in the barrel (BMTF), endcap (EMTF), and overlap (OMTF) regions. The L1 muon trigger is described in more detail in section 10.2.

The muon system is summarized in table 7, which gives the number of chambers for each subsystem, the number of readout channels, and the spatial and time resolution.

Table 7. Properties of the CMS muon system at the beginning of Run 3. The resolutions are quoted for full chambers, and the range indicates the variation over specific chamber types and sizes. The spatial resolution corresponds the precision of the coordinate measurement in bending plane. The time resolution of the RPC of 1.5 ns is currently not fully exploited since the DAQ system records the hit time in steps of 25 ns.

Muon subsystem	Drift tube (DT)	Cathode strip chamber (CSC)	Resistive plate chamber (RPC)	Gas electron multiplier (GEM)
$ \eta $ range	0.0–1.2	0.9–2.4	0.0–1.9	1.55–2.18
Number of chambers	250	540	480 (barrel) 576 (endcap)	72
Number of layers/chamber	8 (R - ϕ) 4 (z , MB1–3)	6	1 2 (RB1, RB2)	2
Surface area of all layers	18 000 m ²	7000 m ²	2300 m ² (barrel) 900 m ² (endcap)	60 m ²
Number of channels	172 000	266 112 (strips) 210 816 (wire groups)	68 136 (barrel) 55 296 (endcap)	442 368
Spatial resolution	100 μ m	50–140 μ m	0.8–1.3 cm	100 μ m
Time resolution	2 ns	3 ns	1.5 ns	<10 ns

The performance of the muon system in Run 1 and the first part of Run 2 is documented in refs. [8, 85]. Much of the muon system is unchanged from that used in Run 1 and described in 2008 [1], but there have been additions and improvements. Notably, there were three major additions to the endcap detector suite: the outer ring of CSC chambers in station four (“ME4/2”), the outer rings of RPC chambers in station four (“RE4/2” and “RE4/3”, collectively “RE4”), and the GEM system in station one (“GE1/1”). In addition, there were important upgrades to the electronics and trigger in many subsystems, which are described in the corresponding sections.

6.1 Drift tubes

6.1.1 General description

Drift tubes (DTs) equip the barrel part of the CMS muon detector, serving as offline tracking devices and providing standalone trigger capabilities. The basic DT detector unit is a rectangular drift cell with a transverse size of 4.2×1.3 cm², whose layout is shown in figure 40 (left). A gold-plated stainless steel anode wire, with a diameter of 50 μ m, is located at the center of the cell, and cathode strips are placed on its side walls. Additionally, electrode strips are located at the top and the bottom of each cell to shape the drift field. The cathodes and electrode strips are set at a voltage of –1200 and 1800 V, respectively, whereas the anode wires operate at applied voltages that vary between 3500 and 3600 V, depending on the individual chambers (more details are discussed in section 6.1.3). Cells are filled




Figure 40. Left: layout of a CMS DT cell showing the drift lines and isochrones. Reproduced from [1]. © 2008 IOP Publishing Ltd and SISSA. All rights reserved. Right: schematic view of a DT chamber. Reproduced from [86]. © 2010 IOP Publishing Ltd and SISSA. All rights reserved.

with a gas mixture of 85% Ar and 15% CO₂, which offers good quenching properties and, under the operational conditions described above, is characterized by a saturated drift velocity around 55 μm/ns.

A schematic layout of a DT chamber is shown in figure 40 (right). Within a chamber, cells are arranged parallel to each other to form layers (L), and groups of four layers, staggered by a half-cell, form superlayers (SL). Each DT chamber is equipped with two SLs that measure the muon trajectory along the bending plane ($R-\phi$). Additionally, chambers from the three innermost detector stations host an SL that measures the position along the longitudinal ($r-z$) plane as well. A total of 250 DT chambers, covering a pseudorapidity range up to $|\eta| < 1.2$, are arranged in five wheels with an identical layout. The wheels are placed parallel to each other along the CMS global z axis, and are labelled W-2, -1, 0, +1, and +2. Within each wheel, chambers are organized in four concentric station rings, labelled from inside-out as MB1 to MB4, and segmented into 12 sectors (S) along the CMS global ϕ coordinate.

The performance of the DT system, measured over Run 1 and Run 2, was found to be remarkably stable and in line with the design expectations. This is documented in refs. [8, 85]. Track segments are typically reconstructed offline with an efficiency above 99%, and they are characterized by spatial and time resolutions around 100 μm and 2 ns, respectively. The efficiency to reconstruct a standalone DT segment in the trigger (also called a trigger primitive) and to correctly identify its bunch crossing (BX) of origin is above 95%. The position (direction) resolution of the DT trigger segments is approximately 1 mm (5 mrad).

Several upgrades, mostly concerning the off-detector electronics, occurred after the end of Run 1 to cope with the twofold increase in LHC instantaneous luminosity during Run 2 with respect to its original design. They are part of the CMS Phase 1 upgrade [87] and are described in section 6.1.2.

Even if a replacement of the DT on-board electronics will be required as part of the Phase 2 CMS muon system upgrade [88], the DT chambers themselves will operate unchanged throughout the HL-LHC period. For this reason, strategies to extend the detector longevity, and maximize the DT performance over a period longer than the one originally expected, were put in place in Run 2 and LS2. They are documented in section 6.1.3.

6.1.2 Phase 1 upgrades of the DT electronics

The increase in instantaneous luminosity during Run 2 required some changes to both the trigger and readout chains, particularly to the electronics hosted in the balconies surrounding the DT wheels

in the experimental cavern, the so-called sector collector. To accommodate the schedule of access opportunities in LS1 and the end-of-year technical stops, and to benefit also from the latest digital technologies, these upgrades took place in multiple stages: first a relocation of the sector collector and, subsequently, the trigger and readout upgrades.

On the other hand, with a single exception, no upgrade of the on-board DT detector electronics [1] was performed or is foreseen until LS3. This holds true for both the frontend (FE) electronics and the HV distribution, which are physically embedded in the chamber gas volume, and for the first level of the DT readout and local trigger electronics, which are hosted in aluminum structures attached to the DT chambers, called minicrates.

The only change that occurred in the minicrates was the replacement of 48 theta trigger boards (TRB), corresponding to the ones located in MB1 of $W_{\pm 2}$, which was performed over LS1. The new TRBs host a Microsemi FPGA that offers the same functionality as the ASIC-based boards originally installed on the detector. This replacement insures the availability of spares for the FE trigger components of the ASIC TRBs until the end of their operation.

A schematic view of the Phase 1 DT detector electronics architecture, as it was in Run 1 and is now in Run 3, is presented in figure 41. Details about this schema are described throughout the following sections.




Figure 41. Schematic views of the Run 1 and Run 3 DT off-detector electronics architectures.

Relocation of the DT sector collector. The first stage of the DT electronics upgrade consisted of a relocation of the so-called sector collector (SC). In its original layout, the SC was composed of ten 9U VME crates, hosting the second level of the DT trigger, the trigger sector collector (TSC) and the readout, the readout server (ROS) electronics, as well as different boards in charge of slow control, monitoring, timing signal distribution, and power distribution. The SC was located on the balconies surrounding the detector in the experimental cavern (UXC).

The relocation project [89] consisted of moving the TSC and ROS boards to the service cavern (USC). As a consequence, the length of the readout and trigger links from the minicrates was increased by approximately 60 m. To cope with the increased length, a dedicated system that converts electrical signals from the chambers into optical ones (CUOF) was installed in the racks previously occupied by the SC. Data were then transmitted by means of optical fibers to the USC, where they were converted back to electrical signals by a second set of dedicated boards (OFCU) before being injected into the ROS and TSC.

The SC relocation brought an increase of operational reliability to the DT system. Moving the TSC and ROS boards to the USC made them always accessible during LHC running periods, allowing the prompt solution of possible problems without the need to wait for technical stops when access to the UXC was possible. It also paved the way for further stages of the Phase 1 DT upgrade, which led to an overall performance improvement. Firstly, the region in which the SC was originally installed was characterized by high radiation levels (up to 0.2 Gy per year in nominal LHC conditions). Without a relocation, these constraints would have led to a limitation in the choice of the electronics components that could be used for the upgrade to radiation-hard units. Secondly, a nonnegligible residual magnetic field (up to 40 mT) is also present on the detector balconies, imposing further restrictions on the use of magnetic components such as inductors and ferrites. The choice of cooling turbines used in the CMS balconies is also constrained by the need to operate in the presence of a magnetic field. If the SC had remained in the UXC, this would have put stringent limits on the power consumption of the upgraded system, since the power dissipation of the original SC had already challenged the cooling capacity of the turbines operating on the balconies. Thanks to the relocation, these constraints were relaxed, leading to more freedom in the design of the upgraded trigger and readout electronics.

The readout and trigger links from the minicrates are based on DS92LV1021 serializers from National Semiconductors, which have an embedded clock. Serial words of 12 bits (10 bits for payload and overhead) are clocked at 20 (40) MHz for the readout (trigger), resulting in a bit rate of 240 (480) Mb/s. Unless an L1 accept trigger is issued, the readout transmits an idle payload that was designed to maximize the DC balancing, resulting in a 40% duty cycle. Instead, in the trigger payload, streams of zeros are preferentially transmitted unless trigger segments are built. This results in a significant DC imbalance in the trigger links, which can be tolerated by operating the CUOF optical transmitters in a low-bias mode, among other improvements upstream, e.g., those to the TwinMux concentrator described below.

A CUOF board consists of a 9U motherboard where four mezzanine cards are typically plugged in. Each mezzanine card carries out the conversion of information received from up to eight links, organized in two FTP cables from either the readout or trigger of one DT chamber. The electrical signals enter the CUOF from RJ45 connectors located at the front of the crate. They are then routed into a line equalizer, restoring levels and compensating for the distortions of the electrical transmission line, and finally are injected into a laser driver that controls the laser diodes. Vertical cavity surface emitting laser (VCSEL) diodes are used. There are eight VCSEL diodes for each mezzanine card, which are connected to a fiber fan-out using LC-type connectors. Each CUOF motherboard also hosts two A3P600L ProASIC3L FPGAs from Microsemi, which control the configuration and the monitoring of the laser drivers. Fine tuning of the drivers' bias and modulation settings is of prime importance to ensure correct transmission of the DC-unbalanced information from the trigger link.

A total of ten CUOF crates (corresponding to two crates per wheel) is installed in the UXC balconies. Each of them hosts thirteen CUOF boards, six (seven) of them dedicated to the transmission of readout (trigger) information, and covering a total of six DT sectors. Two A3050 CAEN modules provide power to the system in each crate. Each of them delivers two independent power supply channels, thus, for a given wheel, four power partitions exist. The power consumption of the CUOF system corresponds roughly to half of that of the previous SCs. A picture of the two CUOF crates instrumented in the balconies surrounding W-1, is shown in figure 42 (left).




Figure 42. Left: front view of two out of the ten DT CUOF crates located in the UXC balconies surrounding the muon barrel (W-1). Center: front view of the five μ TCA crates of the TwinMux in the USC. Right: front view of the three μ TCA crates of the μ ROS in the USC.

Optical communication occurs as 850 nm transmission over OM3 multimode fibers with a 50 (125) μ m core (cladding) diameter. Individual fibers are organized in MTP cords, each containing twelve fibers, which get further assembled in groups of eight to form trunk cables of 96 fibers. Considering the arrangement of fibers into MTP cords, as well as the need for spares, a total of 60 trunk cables is used. This corresponds to a total of 5760 fibers and thus easily covers the minimal need of the system of 3500 fibers. The installation of the UXC-to-USC optical fibers through the trigger tunnels and cable chains to the wheels was labor- and access-intensive, but also leaves a legacy infrastructure that will be exploited by the DT Phase 2 upgrade in LS3. A very important requirement, made at the time of purchasing from the vendor, was that the relative length of different fibers had to be carefully equalized. This was needed to preserve the phase relationship between the signals coming from the different trigger links of a single chamber, which could not be compensated for at the input of the TSC. The propagation delay along different fibers was measured, finding an excellent uniformity, with variations between fibers from the same MTP cord of under 1 ns [90].

The OFCU conversion was performed by dedicated boards that output LVDS signals into RJ45 connectors. These, in turn, were used to transmit information for input to the ROS and TSC boards. Due to different requirements for the readout and trigger electronics, different OFCU boards were developed, but they were both based on commercial parallel optics receivers from AVAGO (HFBR-782BEPZ). At later stages of the Phase 1 DT upgrade, such receivers were re-used as components of the upgraded readout and trigger electronics. The upgrade also included the

design of various slow-control electronics for both the CUOF and OFCUs, and a link to maintain the injection of trigger data into the readout chain.

The entire relocation of the DT SCs took place during LS1, between February 2013 and August 2014. Commissioning with cosmic rays and calibration runs were performed shortly after its completion. It was confirmed that the performance of the upgraded system was consistent with the original one, and the DT system operated very successfully throughout Run 2.

Upgrade of the muon barrel local trigger: the TwinMux concentrator. In the original CMS level-1 (L1) muon trigger architecture, tracks were reconstructed using three different track finders, each one mostly exploiting information from a single muon detector: DT, CSC, or RPC. Candidates from the different track finders were only merged at the last stage of the muon trigger logic. For the Phase 1 L1 trigger upgrade [91], described in detail in section 10, a different layout was chosen. The overall muon trigger chain was designed to exploit information from all the detectors covering the area crossed by a given muon as early as possible in the online reconstruction. This approach was adopted to maximize overall performance and to better control the data acquisition rates. In the case of the muon barrel, the DTs provide excellent position resolution, whereas the RPCs are characterized by excellent time resolution. Hence, in the Phase 1 L1 trigger, information from both detectors is combined by dedicated electronics that provide primitives of superior performance (called super-primitives) already at the input of the barrel muon track finder (BMTF).

To accomplish this, the DT TSC was replaced by a new component, called TwinMux [92], which acts as a concentrator for the data coming from both the DT and barrel RPC chambers. The TwinMux combines information from the two into superprimitives and transmits them to the BMTF and the overlap muon track finder (OMTF) using the 10 Gb/s link protocol exploited by the Phase 1 L1 trigger system.

For the TwinMux, a single slot double-width full-height μ TCA board, called TM7, is used. A TM7 can reach a maximum of 96 optical connections thanks to six front panel Avago optical receivers (72 links limited to 2.7 Gb/s) and two Minipods for high-speed data transmission and reception (up to 13 Gb/s). Figure 43 (left) shows a picture of a TM7 board, where its main components are highlighted. The TM7 board is based on a Xilinx Virtex-7 FPGA that, in the case of the trigger, achieves the merging of several 480 Mb/s links to higher speed serial links and compensates delays to provide BX alignment of the trigger data coming from different inputs. Twelve of the 72 inputs are optionally routable to GTH Gigabit Transceiver inputs [93] in order to handle the GOL-based [94] 1.6 Gb/s links that receive the RPC links. A small mezzanine PCB allows the desired path to be chosen for lines routed to one of the front optical transceiver. Four additional GTHs are used to transfer data on the μ TCA backplane.

From each minicrate, DT trigger information is transmitted as described above. For the RPC detector, five link board masters (LBMs) compress the trigger hit data relative to one muon barrel sector and serialize it through the GOL transmitters. The TwinMux is in charge of forwarding this data to the BMTF and OMTF, by applying a scale-up in the transmission rate (and hence a reduction in the number of links). It is also responsible for duplicating the data to be sent to different track finder processors (up to four times for the sectors of the outer wheels where DT data is shared between the barrel and overlap track finders). Such redundancy is included to reduce the connections between the track finder processors and hence to increase the reliability of the system (which proved




Figure 43. Left: picture of a TM7 board with the main modules highlighted. Right: BX distribution of L1 trigger primitives reconstructed in the muon barrel [95]. Open (red) circles show the performance of trigger primitives reconstructed using information from the DT detector only. Filled (black) circles show the same performance figure for super-primitives built combining information from the DT and RPC detectors.

to be a weak point in the design of the legacy DT track finder). The minimum bandwidth required for forwarding trigger data of one sector is 16 (8) Gb/s for the DT (RPC), implying the need for a total of three 10 Gb/s links. The clock distribution is based on two very-low-jitter PLLs that can broadcast two different clocks to all the FPGA transceivers for performing synchronous or asynchronous data transmission. Finally, a microcontroller is responsible for managing the IPMI interface on the backplane. It handles low-level operations like the hot swap of single boards without the need of switching off a full crate, or monitoring the temperature sensors.

To cover the full barrel, 60 TM7 TwinMux boards are hosted in five μ TCA crates. Each of them is equipped with an AMC13 for clock and slow-control command distribution and for providing a connection to the CMS DAQ. Each crate is also equipped with a commercial μ TCA carrier hub (MCH), a redundant power module, and a JTAG switch used for remotely programming the boards. A picture of the five μ TCA crates hosting the TwinMux is presented in figure 42 (center).

The TwinMux firmware consists of six major blocks. The DT and RPC receiver blocks handle input channels from the two detectors, which correspond to a given muon barrel sector, performing in both cases a link alignment to the same BX. In the case of the DT receiver, a factor 3 oversampling, leading to a frequency of 1.44 Gb/s, is applied to cope with potential data integrity errors caused by the aforementioned DC imbalance of trigger signals. For the RPC, the GOL protocol is decoded. Moreover, clusters are formed of nearby RPC hits and their coordinates are converted to the ones used by the local DT trigger. A so-called super-primitive generation block combines information from the DT and RPC detectors according to algorithms described in more detail later in this section. The output transmitter block formats and sends the trigger output data to the track finders using a 10 Gb/s link protocol. A readout implementation is also included that allows trigger-primitive readout through the CMS DAQ system. Finally, an IPbus-based slow-control block also exists. The TwinMux needs 34.5 BXs of latency for data to go through the FPGA from input to output. Of these, only four BXs are actually used for super-primitive generation, whereas 26.5 are taken for handling of RPC data by the RPC receiver block. Finally, four additional BXs are taken by the serialization of the data sent to the track finders. The DT receiver block works in parallel with, and in the shadows of, the RPC one.

The TwinMux was deployed in production, as part of the Phase 1 L1 trigger upgrade, starting from 2016, after having been tested in a slice of the detector that was sending signals both to the TSC OFCU and TwinMux boards during the 2015 run. The development of the algorithms to build super-primitives, which are used as input to the BMTF, occurred in stages. In 2016, no combination was performed, hence only trigger primitives built by the DT on-board minicrate electronics were fed into the BMTF.

In 2017, an algorithm that combines information from the DT and RPC detectors was deployed to improve the super-primitive BX identification. Within such algorithm, the compatibility of the DT trigger primitives and RPC clusters that are built in nearby chambers is checked by comparing the difference in azimuthal angle ($\Delta\phi$) between them. If a match within a programmable window is found between a DT trigger primitive and (at least) one RPC cluster, and if the difference in terms of BXs between the DT trigger primitive and the RPC cluster is within ± 1 BX, a super-primitive is built using the DT trigger segment position and direction, but the RPC BX. We note that no time correction is attempted if a DT trigger primitive is built exploiting hits from all layers of the ϕ -SLs of a given DT chamber. In any case, a dedicated quality flag, documenting the successful matching with the RPC, is set. This combination better exploits the complementarity of the DT and RPC detectors, relying on the spatial resolution of the former and the time resolution of the latter. The impact, in terms of BX identification performance for trigger primitives caused by muons from pp collisions, is shown in figure 43 (right). The asymmetry in the BX distribution of DT primitives (red open dots), mostly due to the occasional presence of δ rays that can spoil the reconstruction of standalone DT trigger segments, is mitigated when the combination with the RPC is effective (black filled dots). Because of the BX correction, the muon barrel trigger primitive BX assignment efficiency is also increased, on average, by 1.4%.

In 2018, a further improvement was put into production. In the MB1 and MB2 stations, where two RPC layers cover both surfaces of each DT chamber, as described in section 6.3.1, the generation of RPC-only super-primitives is attempted if no DT trigger primitives are present in a given BX. In that case, pseudosegments are built out of RPC cluster pairs that are reconstructed at the same BX in the two different RPC layers. If more pseudosegments are generated at a given BX, only the one whose direction is closest to that of a straight track coming from the CMS interaction point (ϕ_b) is retained. Finally, only pseudosegments for which ϕ_b is below a programmable threshold are accepted. Additionally, also in this case, the RPC-only primitive is marked with a dedicated quality flag. Upon deployment of this algorithm in data taking, a further increase of the super-primitive efficiency around 3% was observed in the MB1 and MB2 stations. Due to the redundancy of the muon system, such local trigger efficiency improvement resulted in a marginal improvement of the BMTF efficiency. Nevertheless, a rate reduction of a few percent was observed for the lowest unscaled L1 muon trigger. This is because, in a few cases, BMTF tracks now get built using a larger number of points along the muon trajectory, improving the measurement of the track transverse momentum.

Given the fact that it uses a different track-building algorithm, no super-primitive generation is attempted on primitives that are sent to the OMTF, since the OMTF receives a full list of DT trigger segments and combines them with RPC information directly at the track-finding step, as also described in section 6.3.2. The super-primitive algorithm that operated throughout 2018 was also deployed at the start of Run 3.

Upgrade of the DT readout: the μ ROS system. Simulation studies have shown that the ROS was the most severe bottleneck in the DT readout chain. The ROS combined information from a total of

25 readout boards (ROB) providing the minicrates with data for a full DT sector. The time needed by a ROS to perform event building from all the input links depended on how hits were distributed within the ROB. A high noise rate from sporadic small groups of channels could be easily sustained. However, high overall background hit rates, which come with a rather uniform distribution within each of the chambers of a sector, take much more time to be processed. In addition, muons crossing the DT can produce up to 44 hits within a single ROS. When both the background and muon hit rates are considered, taking as a proxy an LHC instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the maximum acquisition rate that a ROS can sustain becomes close to the 100 kHz limit imposed by the DAQ and L1 trigger.

For this reason, throughout Run 2, the ROS, as well as the downstream components of the DT readout chain, called device-dependent units (DDU), were replaced with a new system, based on the μ TCA architecture, named μ ROS. For the μ ROS [96], the same TM7 boards designed for the TwinMux are used, but a different firmware is deployed to implement the functionalities needed by the DT readout.

Each TM7 features six 12-fiber MTP receivers, for a total of 72 input links. A total of 25 ROB. s provide the inputs from a given sector. The data from each wheel are thus processed by five μ ROS boards, four of them receiving three sectors each (24 channels per sector, 72 links), and the fifth receiving the 25th channel for each of the twelve sectors. The production system comprises three μ TCA crates (central, positive, and negative wheels) and 25 μ ROS boards. Each crate is equipped with an AMC13 that provides clock and slow-control distribution, as well as a connection to the CMS DAQ. A picture of the μ ROS system is presented in figure 42 (right). With this architecture, no further components of the DT readout system (DDU) are needed. Under the conditions reached over Run 2, the maximum payload bandwidth varied, depending on the wheel, between 0.3 Gb/s in W0 and 0.6 Gb/s in W+2 and -2, remaining well within the AMC13 limits.

In terms of firmware, the components needed to handle the slow control and general board functionality are inherited from the TwinMux. Special care was instead put into the design of the block in charge of data deserialization. This firmware can recover input data with high quality and minimal data losses with respect to the input stream. Carrying data at 240 Mb/s, the receiver samples data at 1.2 Gb/s (a factor of 5 oversampling). Majority filtering is performed on the three central samples of each bit before reassembling the original input word. Bits where a weak majority is found (two instead of three) are marked as transmission error candidates. If the data frame parity error check fails and only one bit is marked as a transmission error candidate, the latter gets corrected. The firmware implements a full verification of the ROB protocol and provides statistics for the different ROB failure cases, which are used for monitoring. Finally, while the legacy ROS system masked channels in case of transmission errors until a resync was issued, the event builder from the μ ROS is capable of recovering from all types of errors as soon as the condition disappears.

The transition to the μ ROS occurred during the 2017/2018 year-end technical stop. Prior to that, in 2017, signals from the ROB. s for a slice of the detector were split and a μ ROS slice-crate was instrumented and integrated into the CMS DAQ as a separate unit. It was used to develop the μ ROS prior to the deployment of the full system, which allowed a smooth transition. In LS2, the FE signals of a sector in the external wheels (W+2 S12) were split to allow the continuation of the same strategy for the Phase 2 upgrade, and another slice-test system will be operated in Run 3. The data of such a Phase 2 slice-test are already read by the CMS DAQ.

The impact, in terms of DT performance, of the transition to the μ ROS is presented in figure 44. The two plots show a chamber-by-chamber map of the DT segment reconstruction efficiency, as measured with a tag-and-probe method. In the measurement, no masking of chambers with hardware or readout issues is applied. The left (right) side of the figure refers to results computed using 2017 (2018) data, collected before (after) the transition to the μ ROS. Bins where the efficiency is significantly lower than 99% are due to possibly sporadic problems that affected chambers or their readout. A better performance was observed after the upgrade to the μ ROS, mostly thanks to a reduction of the number of chambers affected by readout problems.




Figure 44. DT segment reconstruction efficiency measured with the tag-and-probe method using data collected by CMS in 2017 (left) [97] and 2018 (right) [98], before and after the transition to the μ ROS. The efficiency is usually above 99% except for chambers affected by hardware problems, mostly coming from the DT readout. After the deployment of the μ ROS, the overall efficiency improves.

6.1.3 Detector longevity for Run 3 and beyond

The DT system was designed and validated to sustain ten years of LHC operation in nominal conditions, corresponding to approximately 500 fb^{-1} of integrated luminosity [99]. Accelerated aging studies, carried out prior to the installation of the chambers in CMS, indicated that no degradation in the performance of the present detector (and its electronics) is to be expected under those assumptions. Furthermore, at the boundaries of long data-taking periods, data are regularly collected by varying the HV applied to the anode wires (HV scans) and measuring the detector efficiency at each HV point to assess its stability and exclude potential effects due to early aging. Up to the end of Run 2, no efficiency degradation was observed during the HV scans. Additionally, the original longevity studies were complemented with more stringent ones, performed using the CERN gamma irradiation facility (GIF++) [88], which are summarized in the following paragraph. Given the results of all the above studies, and the current projections in terms of integrated luminosity expected for Run 3, the DT performance is foreseen to remain almost constant over the coming run period.

Nonetheless, though several electronics components will be replaced as part of the CMS Phase 2 upgrade program, existing muon detectors will operate throughout the HL-LHC era. For this reason,

further accelerated longevity studies are being performed at GIF++ on two spare DT chambers, as mentioned above. Due to the complexity of the physics and chemistry phenomena driving the aging processes, accelerated studies have large uncertainties, typically taken into account by irradiating with a substantial excess with respect to the needed integrated values (safety factors). The reference targets for instantaneous and integrated luminosities used in the aging studies to derive the DT performance are evaluated assuming safety factors of two, which double both the nominal HL-LHC luminosity ($5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) and the total integrated luminosity (3000 fb^{-1}). Considering these safety factors, by the end of the HL-LHC, the DT hit detection efficiency can possibly drop from its present value of around 96% to approximately 70% in the MB1 stations of $W \pm 2$ (corresponding roughly to 10% of the entire system). Moreover, in a further 20% of the detector (corresponding to the MB4 stations of the DT sectors covering the top half of CMS and to the MB1 stations of $W \pm 1$), the efficiency is projected to range between 85 and 90%. Finally, in the rest of the system, efficiencies above 90% are expected. Given the redundancy in terms of the number of DT layers per chamber and of chamber stations in the muon system, the maximum inefficiency due to DT aging, localized in a narrow region around $|\eta| = 1.0$, is expected to be within 2 (5)% for standalone offline reconstruction (trigger). Though these expected losses are not very large, mitigation strategies, which are described in the following sections, have been put in place during Run 2 and LS2, to maximize the longevity of the DT detector and ensure the highest achievable performance in the long term.

Optimization of the operational working points. For gaseous detectors, deterioration due to aging becomes more significant as the integrated charge released in the gas volume increases. In turn, the integrated charge depends on the rate and type of particles crossing the different regions of the detector, the collected charge per particle, and the total integrated running time.

In the case of the DT detector, background dominates the hit rates and is largest in the MB1 stations of $W \pm 2$ and the MB4 stations for the sectors covering the top half of CMS. The collected charge per particle depends on both the particle type and the gas amplification factor, which is driven by the HV settings of the cell anode wires. Therefore, for a fixed integrated luminosity, reducing the HV working points of the anode wires can result in a significant reduction of the integrated charge. Of course, such an optimization can be performed only within the limits where the impact on detector performance is deemed acceptable.

During the 2017 and 2018 LHC runs, the HV of the DT anode wires was progressively lowered with respect to the default 3600 V used until 2016, in the chambers most exposed to background. In 2017, the wires in MB1 of $W \pm 2$, and the ones in MB4 of sectors 3, 4, and 5 for all wheels were operated at 3550 V. In 2018, a further reduction was applied, reaching the HV values given in table 8. The discrimination threshold values applied in the DT FE electronics were also lowered from 30 to 20 mV for the entire detector.

For the MB1 of the external wheels, lowering the HV of the anode wires to 3500 (3550) V resulted in a relative reduction of 58 (45)% in the drained current with respect to the original 3600 V setting.

The overall performance of the system under the updated operational conditions was thoroughly studied using pp collision data. Firstly, the reduction of the FE thresholds to 20 mV resulted in a marginal increase of overall detector noise, which was handled by masking a few specific noisy wires. Reducing the gain and the FE threshold resulted in shifts of a few ns of the effective drift

Table 8. HV settings for the anode wires of the different DT stations and wheels used for the 2018 LHC run.

	W-2	W-1	W0	W+1	W+2
MB1	3500 V	3550 V	3550 V	3550 V	3500 V
MB2	3550 V	3550 V	3600 V	3550 V	3550 V
MB3	3550 V	3600 V	3600 V	3600 V	3550 V
MB4	3550 V	3550 V	3550 V	3550 V	3550 V

time measured by the electronics caused by opposite-sign effects. These partially cancel but were nevertheless carefully corrected for, by updating the trigger synchronization and offline calibration. Fine-tuned corrections were needed to maintain the optimal performance of the system in terms of BX identification efficiency and time resolution of the offline reconstruction. These remained well in line with what is reported in section 6.1.1. The DT hit detection efficiency was then measured. Results for the ϕ -SLs of the MB1 stations are given as function of the total integrated luminosity in figure 45. Different colors in the plot refer to different DT wheels. The efficiency is stable throughout the different run years, apart from step variations dominated by changes in the FE and HV, which amount at most to 1%. This was proven to be true for the rest of the detector as well and, given the redundancy of layers forming a DT chamber, a negligible impact on the DT local reconstruction efficiency was observed. Finally, the effect of the updated settings on the DT hit resolution was also assessed. In the case of ϕ -SLs, where hit resolution impacts the muon transverse momentum measurement, only a slight worsening was observed. In the chambers where the HV was lowered to 3550 V, the degradation is around 10%. In the MB1 of the external wheels, where the HV was set to 3500 V, the effect is slightly larger, but the hit resolution remains within $250\ \mu\text{m}$, in line with the design expectations.




Figure 45. DT hit detection efficiency, computed as a function of the total CMS integrated luminosity, for the ϕ -SLs of MB1 chambers [98]. Different colors refer to different DT wheels. The plot summarizes how the efficiency evolved during Run 2, mostly as a consequence of the different updates of the detector HV and FE threshold settings.

Given the reduction in terms of expected integrated charge, as well as the very mild impact on the overall performance, the operational working points used in 2018 were chosen as the default configuration for Run 3 as well.

Gas system upgrade and open-loop operation. The DT gas system [1] includes a facility on the surface that stores individual gas components (Ar and CO₂) and combines them to obtain the gas mixture of 85% Ar and 15% CO₂ used by the DT. A second facility, located in the USC, is dedicated to the predistribution of the final gas mixture to the five wheels, as well as to the gas analysis system. Finally, five distribution racks located in the UXC provide the gas flows to the 250 DT chambers, with a flux around 40 l/h to each chamber. The primary goal of the gas analysis system is ensuring the stability of the gas mixture. It consists in the determination of the contamination of O₂ and H₂O, and the drift velocity measurement.

The loss of efficiency caused by aging, described in section 6.1.3, is due to a decrease of gas gain which, in turn, is caused by deposits that form around the anode wires of the DT cells. It is believed that this effect is mostly due to out-gassing of components inside a DT chamber. For this reason, the re-injection and spread from hot detector regions of pollutants contaminating the DT gas should be avoided. Therefore, starting from the 2018 run, the DT gas system operates in the, so-called, open-loop mode, where there is no re-circulation of used gas. Prior to that, the system was operated in a closed-loop mode, where 85% of the used gas was recirculated and only 15% of fresh gas was injected. The closed-loop mode can still safely be used during technical stops or longer periods of inactivity, where no aging due to out-gassing is expected. In order to operate in open-loop mode, an upgrade of the gas system was performed. It allows the intentional introduction of air from a bottle on the surface, to maintain the desired level of O₂, corresponding to 80 ppm. Furthermore, a humidifier with a bypass that ensures the H₂O concentration is kept at 800 ppm was also installed. The presence of small quantities of O₂ and H₂O helps prevent effects that induce chamber aging (such as polymerization) [100], and the target values for such contaminants were derived based on experience from the data taken in Run 1 and Run 2. Because of this upgrade, prompt monitoring of the gas stability has become even more important than in the past.

The O₂ and H₂O measurements are performed using commercial sensors, which have to be recalibrated every year by injecting gas with known components of oxygen and humidity. The drift velocity measurement is done by a more complex system, called VDC [101], which aims to deliver, every ten minutes, a drift velocity measurement. There is one O₂ sensor, one H₂O sensor, and one VDC for each of the five DT wheels, and a separate sampling line for each wheel. The gas sampling permits the selection of the output of any individual chamber, as well as the global output (or the global input) of an entire wheel. The direct measurement of the drift velocity performed by the VDC system is sensitive to possible unknown components or contaminants in the gas. The drift velocity depends on the Ar/CO₂ ratio, as well as on the level of contaminants, such as O₂ and N₂. The drift velocity is a key parameter for the DT local reconstruction, so that any deviation spotted by the VDC system must be investigated and eventually requires immediate intervention on the operation of the DT gas system. An example from the monitoring by the VDC system is presented in figure 46. The first transition from closed-loop to open-loop operation was correctly detected by the drift velocity measurement. In this case, the gas analysis showed very quickly the impact of the injection of air and humidity into the mixture for the case of the open-loop mode, which resulted in a sub-percent variation of the drift velocity.

Other than direct monitoring of the gas mixture and drift velocity, the stability of the gas is also evaluated indirectly with event data, by looking at the stability of the performance of the DT local reconstruction. If a track-segment fit is performed using a sufficient number of hits from a DT




Figure 46. Drift velocity measurement using the fresh gas analyzed by the VDC system. The variation on 8th March 2018 corresponds to the transition between closed-loop and open-loop operation of the DT gas system.

chamber, additional parameters, other than the segment position and direction, can be extracted from the fit itself. In this way, the need for residual corrections on top of the calibration parameters used by the reconstruction, such as the drift velocity, can be evaluated on a segment-by-segment basis. Overall biases in the distribution of such residual corrections are then measured run-by-run for each DT chamber, and their stability across different runs is monitored.

Finally, the DT chambers are operated since Run 1 in a differential pressure mode. The gas control system presently ensures that a positive differential pressure of +3 mbar is applied at the bottom of the wheels. This value, which is well within the mechanical maximum limit (50 mbar), protects against possible contamination. A dedicated system of differential pressure sensors monitors continuously these values, which are transmitted to the DT online monitoring system.

Installation of shields over the outer MB. As mentioned earlier, besides the MB1 stations of the external wheels ($W\pm 2$), the parts of the MB characterized by the highest level of background are the MB4 stations in the sectors covering the top half of CMS (sectors 1 to 7). Studies based on pp collision data, as well as results from simulations, corroborate the hypothesis that background in this region is mostly generated from interactions of low-energy neutrons permeating the cavern. Therefore, a strategy was put in place to effectively protect the top DT chambers by installing proper absorbing shields.

With the aim of designing such shields, absorbing layers were installed on top of the MB4 stations in sector 4 of $W+2$ and $W-2$ during Run 2. Different configurations in terms of material and thickness of absorbers, suggested by simulations of the radiation field in the cavern, were tested each year from 2015 through 2018, and the impact in terms of background reduction was studied. An example of these studies is shown in figure 47 (left), where the magnitude of the linear dependence of the currents in each chamber with respect to the LHC instantaneous luminosity is presented for the 2018 run. The impact of the test shielding configurations installed over the MB4 stations in sector 4 of $W\pm 2$ is clearly visible. Based on the results of these investigations and considering the outcome of simulations of the mechanical stress that different shield setups would induce on the supporting structures, a layout was proposed that is expected to reduce the background by about a factor 2.

The installation of the final DT MB4 shields was performed during LS2 and completed in October 2020. A schematic view of the shielding layout is presented in figure 47 (right). The figure shows the shield installed over $W-2$, -1 , $+1$ and $+2$. In those wheels, shielding cassettes consisting




Figure 47. Left: magnitudes of the linear dependence between the currents from each DT chamber of the MB and the LHC instantaneous luminosity [102]. Results are computed after the optimization of the operational working points. Right: transverse view of an MB wheel highlighting the layout of the MB4 shield, as installed in W-2, -1, +1, and +2. Red (green) lines represent shield layers consisting of thin (thick) cassettes.

of 9 mm of lead and 9 cm of borated polyethylene (thick cassettes) are put on top of sector 4 stations (green lines), whereas cassettes consisting of 9 mm of lead and 3 cm of borated polyethylene (thin cassettes), cover sectors 1, 2, 3, 5, 6, and 7. Due to mechanical restrictions imposed by the presence of services from the vacuum tank and inner detectors, the shields deployed to cover the central wheel have a different layout. In this case, no shielding is applied on top of sector 4, whereas, in the other six sectors, shields consisting of thin cassettes only cover the lowest part of half of each chamber. The deformation on the supporting structures induced by the MB4 shields, which have a weight of approximately 40 tons in total, was measured accurately at the weakest points in the structure and found to be always smaller than 1 mm, in good agreement with finite-element calculations.

In summary, during Run 2, the DT off-detector electronics underwent a set of upgrades that improved the performance of both the readout and trigger. Strategies to increase the detector longevity were also put in place to maximize the DT performance through the end of HL-LHC running. Following a thorough recommissioning over LS2, the DT system successfully entered Run 3, showing offline and online tracking performances that are remarkably consistent with the ones achieved at the end of the previous LHC run.

6.2 Cathode strip chambers

6.2.1 General description

The cathode strip chambers (CSCs) are multiwire proportional chambers with a finely segmented cathode strip readout. The strips run radially in order to measure the muon position in the bending plane, the plane perpendicular to the colliding beams axis, while the anode wires are oriented in azimuth and provide a coarse measurement in the radial direction. A precise measurement of the muon coordinate in the azimuthal direction is obtained from charges induced on the cathode strips. Each CSC module consists of six gas layers, each layer having a plane of radial cathode strips and a




Figure 48. Left: layout of a CSC chamber, with seven trapezoidal panels forming six gas gaps. Only a few wires (lines running from left to right) and strips (gray band running from top to bottom) on the upper right corner are shown for illustration. Reproduced from [103]. CC BY 4.0. Right: installation of the outer CSC chambers (ME4/2) during LS1. Reproduced with permission from [104].

plane of anode wires running perpendicular to the strips. Figure 48 shows the CSC physical layout and a photograph of the installation of some CSCs into the CMS detector.

The gas mixture is 40% Ar, 50% CO₂, and 10% CF₄. The CO₂ component is a nonflammable quencher needed to achieve large gas gains, while the main function of the CF₄ is to prevent anode aging caused by polymerization processes on the wires. Argon is the working gas that is ionized by traversing charged particles. The primary design consideration for the CSC detectors is the ability to provide good spatial and temporal resolution for triggering and identification of a muon. The typical position and time resolutions are 50–140 μm , depending on chamber type, and 3 ns per chamber, respectively.


Figure 39 shows the geometry of the CSC system. At the start of the LHC operation, a total of 468 trapezoidal CSC modules, placed between the steel magnetic flux-return plates, were arranged into four disks (stations). The first station, closest to the interaction point, is further segmented into three rings (ME1/1, ME1/2, and ME1/3), while stations 2 through 4 are separated into just two rings. The rings closest to the beam line (named ME1/1, ME2/1, ME3/1, and ME4/1, or collectively as ME1234/1) are subject to the highest particle rates. Each of the chambers in the ME1/1 ring contains strips split at $|\eta| = 2.1$, which are read out separately and denoted ME1/1a ($2.1 < |\eta| < 2.4$) and ME1/1b ($1.6 < |\eta| < 2.1$). Some key parameters of the CSCs are summarized in table 9.

The ME234/1 chambers each cover 20° in ϕ ; all others cover 10° . All chambers, except for the ME1/3 ring, overlap by five strips at each edge and hence provide ϕ coverage without gaps. The original design included 72 ME4/2 chambers that were descope in the initial construction of the CSC system, and instead were built and installed during LS1 that occurred during 2013–2014. The presence of ME4/2 chambers provides an additional track segment of six hits along the muon trajectory. This allowed for a more robust muon p_T assignment in the L1 trigger described in section 10, so that the rate with a given p_T threshold in the $1.2 < |\eta| < 1.8$ region decreased during Run 2, as shown in figure 49.

As illustrated in figure 39, in the $1.1 < |\eta| < 2.4$ region, almost all possible muon paths cross at least three CSC chambers. In the $0.9 < |\eta| < 1.1$ overlap region, a combination of DT and CSC

Table 9. Key parameters for different types of CSCs.

	ME1/1	ME1/23	ME234/1	ME234/2
Wire diameter [μm]	30	50	50	50
Wire spacing [mm]	2.5	3.2	3.1	3.2
Strip width (narrow) [mm]	3.2	6.6–11.1	6.8–8.6	8.5
Strip width (wide) [mm]	7.6	10.4–14.9	15.6	16.0
Gap between strips [mm]	0.35	0.5	0.5	0.5
Angle subtended by each strip [mrad]	2.96–3.88	2.15–2.32	4.65	2.33
Gas gap [mm]	7	9.5	9.5	9.5
Operating HV [V]	2900	3600	3600	3600


Figure 49. Emulation of the Run 1 algorithms compared to the upgraded Run 2 algorithms, as a function of the L1 muon η . The most common L1 single-muon trigger threshold used in 2017 was $p_T^{\mu, L1} \geq 25$ GeV.

modules also typically provide at least three position measurements along muon paths. Moreover, the four layers of RPCs between the CSCs provide additional muon hits at the trigger level in the $0.9 < |\eta| < 1.6$ region. The extension of the RPC layers and the addition of GEM detectors (GE1/1 in Run 3) further enhance the muon triggering and reconstruction in the $1.6 < |\eta| < 2.2$ region. The high redundancy of the muon system is a central feature of the design and is responsible for the robustness of the muon triggering and reconstruction in CMS.

The present electronics readout system of all CSC chambers, shown schematically in figure 50, share the same architecture as in Run 1. This is briefly summarized below, with a more detailed description available in ref. [1]. On each CSC chamber, an anode local charged track (ALCT) board finds patterns among the six-layer anode hits (sampled by the ALCT board, based on the digitized pulses sent by the on-chamber anode frontend boards) that are consistent with muon tracks, and then assigns a quality rank to each of them. These patterns are called ALCTs. The ALCT board can then send up to two such tracks per each bunch crossing (BX) with the highest quality to the trigger motherboard (TMB). Upon receiving a level-1 accept (L1A) signal from the trigger, coinciding




Figure 50. Schematic of the CSC electronics readout systems for Run 3: ME1/1 (upper), ME234/1 (lower left), all other chambers, ME1234/2 and ME1/3 (lower right). Reproduced with permission from [105].

within 3 BXs with an ALCT, the ALCT board sends all wire hits and ALCTs that have been found in a predefined time window to the data acquisition motherboard (DMB). Both the TMB and DMB are located in VME crates mounted on the edge of the endcap steel yoke.

The cathode signals are processed by another set of on-chamber electronics, the cathode frontend boards (CFEB). Each CFEB amplifies the signals from 16 strips on each of six layers of a CSC, and sends fast trigger information and charged particle hits localized within a half-strip width, as assessed by comparator ASICs (comparator data), to the TMB. On receiving an L1A coinciding within 5 BXs with a cathode local charged track (CLCT), the CFEB digitizes the strip signal waveforms over a 400 ns long period, which is sampled every 50 ns by a 12-bit ADC, and sends the digitized samples to the DMB. The transfer of CFEB data to the TMB and DMB comprises the trigger and data paths of the strip signals, respectively. During Run 1, each CSC had five CFEBs, except for ME1/3, which had four per CSC (since the ME1/3 CSCs have only 64 strips).

The low-voltage power for the on-chamber electronics is distributed at the appropriate voltage levels by the on-chamber low-voltage distribution board (LVDB).

The TMB builds cathode hit patterns into CLCTs and finds coincidences with anode hit patterns that form ALCTs to make local charged tracks (LCTs), also called the CSC trigger primitives. It sends the two LCTs with the highest quality per BX to the muon port card (MPC) and, on receiving

an L1A, to the DMB. The DMB controls the CFEBs on a chamber and collects anode, cathode, and trigger information to send to the detector-dependent unit (DDU), on the arrival of an L1A. The DDU, situated in the frontend driver crates located in the underground service cavern, collects data from 15 DMBs in the CSC system and sends the information through the global DAQ path.

The MPC, situated in the VME crate, collects LCTs from each of up to nine TMBs in a trigger sector, and sends these trigger primitives to the muon track finders described in section 10.2. There is one MPC per peripheral VME crate. The operation of the VME crate is supported by the clock-control board (CCB) and custom VME crate controller (VCC). The CCB serves as the interface between the CSC system and the trigger control and distribution system (TCDS) of CMS. The VCC receives VME commands from the control room and distributes them to the other boards in the peripheral crate via the backplane.

6.2.2 Upgrade of the CSC system since Run 1

While the CSC system operated stably throughout Run 1 and Run 2, some CSC readout electronic boards needed to be upgraded in order to handle the expected longer latency and more stringent trigger requirements at the HL-LHC. Specifically, if the electronics were not altered, the longer latency requirements would fill up and overflow the pipelines of the frontend boards in certain CSC stations, while the higher L1 trigger rates would overwhelm the output bandwidth of various on- and off-chamber electronics boards. To reduce the installation load during LS3, the bulk of the CSC electronics upgrades that required chamber access (dismounting and re-installation) was already performed during LS2.

The CFEBs use switch capacitor arrays to store the charge induced on the cathode strips. These capacitor arrays are capable of storing 96 charge measurements (corresponding to six events worth of data) during the L1 trigger latency. As mentioned above, the digitization of the analog signals and subsequent readout by the DMB only happens when a CFEB receives an L1A that is in coincidence with a CLCT. As a consequence, for CSCs that are closest to the beam (ME1234/1) where the background rate is high, frequent memory overflows and large data losses are expected at the HL-LHC trigger latency. Figure 51 shows the average event loss fraction for ME234/1 rings as a function of the instantaneous luminosity. These curves are based on a statistical model that has been verified by measuring the loss rates in bench tests that emulate the expected background and L1 trigger rates. These results show that data loss would be a severe issue at the HL-LHC with the original CFEBs. The outer CSC chambers will not suffer from these problems because the trigger primitive rates in these rings are lower than those in the inner rings by factors of more than 3.

The CFEBs on ME1/1 have already been vulnerable to data loss at the highest instantaneous luminosities (up to $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) in Run 2. To mitigate such losses, the four CFEBs on the ME1/1b section of the chambers were replaced by four digital cathode frontend boards (DCFEBs) during LS1. The DCFEBs use fast flash 12-bit ADCs that continuously digitize the cathode signals at 20 MHz, as well as having more powerful Virtex-6 FPGAs with large internal memory resources. The resulting digital pipeline can hold up to 700 events, and thus there should be negligible dead time at the HL-LHC. Two optical links running at 3.2 Gb/s are employed per DCFEB to transmit the raw data for DAQ and the fast comparator data for building the L1 trigger to the new optical data acquisition motherboard (ODMB) and optical trigger motherboard (OTMB), respectively, which replaced the original DMB and TMB.




Figure 51. Left: event loss fraction as a function of instantaneous luminosity for different type chambers after different upgrades. The vertical dashed brown line indicates the design HL-LHC luminosity. Right: event loss rate measured in a CFEB under HL-LHC conditions for an ME2/1 chamber, compared to the statistical model [100].




Figure 52. Difference between the position of a reconstructed hit in layer 3 of an ME1/1a chamber and the position obtained by fitting a segment with hits from the other five layers for Run 1 (red) and Run 2 (blue). The spatial resolution is improved by 27% from $\sigma = 64 \mu\text{m}$ in Run 1 to $46 \mu\text{m}$ in Run 2. This is due to the removal of the triple-grouping of strips in ME1/1a, which reduces the capacitance and hence the frontend noise.

In Run 1, due to cost constraints, the 48 strips in the ME1/1a section of the chamber were joined every 16 strips into groups of three, resulting in 16 readout channels served by one CFEB. To remove the ambiguity in triggering and reconstruction, the grouping of the ME1/1a strip readout was removed during LS1, and the single CFEB was replaced by three DCFEBs for strip readout. The ungrouping of the strips in ME1/1a leads to an improvement in spatial resolution of about 20% (figure 52), as well as a decrease in the L1 trigger rate by at least a factor of 3 (figure 49).

Operational experience during Run 2, and additional radiation tests, showed that the programmable read-only memory (PROM) employed in the current DCFEB may not withstand the radiation dose expected at the HL-LHC. Thus, the new DCFEBs (DCFEBv2s) were designed with an option for remote programming of their FPGAs using the CERN-designed GBTx ASIC. The

performance of the two optical transceivers on the original DCFEBs has been occasionally unreliable, particularly due to single-event upsets (SEUs). While mitigation measures in software and firmware have improved the reliability of these transceivers in the DCFEBv2s, they were replaced by the VTTx, a radiation-hard twin optical transmitter designed by CERN. These new DCFEBv2s were installed in the ME1/1 chambers during LS2, thus mitigating the risks associated with the longevity of their PROMs.

The DCFEBs used in the ME1/1 ring during Run 2 were moved to the chambers in the ME2/1, ME3/1, and ME4/1 rings, where radiation levels are lower. In Run 3, although the ME234/1 chambers are fitted with DCFEBs, the new ODMB is not scheduled to be installed before LS3, so these DCFEBs still send data to the DAQ via the original DMBs.

The ALCTs store raw wire-hit information in a pipeline within the FPGA while waiting for an L1A. The ALCTs installed before Run 1 would also suffer significant data loss during HL-LHC operation. This is because they do not have sufficient FPGA memory resources to hold all raw hit information before sending them to the DMB during the HL-LHC L1 trigger latency. Moreover, the output bandwidth for the boards in the inner rings (ME1234/1) would not be capable of handling the expected HL-LHC data rates. Both of these problems were solved by replacing the mezzanine cards in the affected chambers.

During LS1, the ME1/1 chambers were equipped with new ALCT mezzanine cards having more powerful FPGAs (Spartan 6), with 9–12 times more memory than those used previously (Virtex-E). This, in turn, allows the pipeline to be deep enough to satisfy the HL-LHC latency requirements. Similar ALCT mezzanine cards were also installed on the ME4/2 chambers during the same period.

During LS2, as a second phase of the ALCT upgrade, the ALCT mezzanine cards originally installed in ME1/1 during LS1 were moved to ME1/3, and all other mezzanine cards, except for those used in the ME4/2 rings, were replaced. The new ALCT mezzanine cards are largely based on the ALCT design used for cards installed in ME1/1 during LS1, with a few exceptions discussed below. The ALCT mezzanine cards servicing the ME234/1 rings use two new VTTx optical transmitters, each running at 3.2 Gb/s, to increase the output bandwidth for the expected HL-LHC data rates. The new ME1/1 ALCTs, in addition to a different FPGA from those used for the ME234/1 chambers, have a VTRx transceiver instead of two VTTx transmitters on each mezzanine card. This allows for the same remote FPGA programming option (based on the GBTx ASIC) that the DCFEBv2s have, thus mitigating any risks associated with the aging of the PROMs in that ring. From Run 4, the VTRx transceiver (as well as the VTTx transmitters from the ME234/1 ALCTs) will be connected to an updated ODMB board that transmits data to the CMS DAQ over an optical link at 4.8 GB/s. The new mezzanine cards serving the ME234/2 chambers have the same FPGA as those used for the ME1/1 ALCTs, but they will send data to the DMB through copper links even in the future runs due to the lower expected data rate. Since the new design maintains backwards compatibility, the anode raw data sent to the DAQ from the ME1234/1 ALCTs can be transmitted to the DMBs via the copper path during Run 3, and to the new ODMBs via the optical link(s) from Run 4 onwards.

The DCFEBs that were installed on the ME1/1 chambers in LS1 transmitted data to the CMS DAQ and comparator data to the CMS trigger at 3.2 Gb/s via optical fibers. Since the previously installed TMBs and DMBs were designed to work with the CFEBs and have no optical receivers, the TMBs and DMBs for the ME1/1 chambers were upgraded to OTMBs and ODMBs during LS1 as well. They receive optical data from the seven DCFEBs and use a more powerful FPGA. During

LS2, all TMBs used for ME234/1 chambers were upgraded to OTMBs to match the upgrade of their CFEBs to DCFEBs. New OTMBv2s were produced for the ME1/1 and ME2/1 chambers. These are almost identical to the OTMBs installed in LS1, except that obsolete components were replaced and the ability to build trigger primitives including GEM data was added, as described in section 6.4.5. The OTMBs serving the ME1/1 chambers during Run 2 were relocated to the ME34/1 rings.

The new electronics consume more power than the old so the low-voltage (LV) systems were upgraded to provide higher maximum currents at the appropriate voltage levels to ensure reliable operation of the CSCs at the HL-LHC.

For Run 1 and Run 2, the CSC high-voltage (HV) system comprised two independent subsystems: a custom one providing HV to the non-ME1/1 chambers (11 016 channels), and a commercial subsystem supplying the ME1/1 chambers (432 channels) [106]. In both systems, the voltage can be regulated on all channels individually, and the currents drawn are continuously monitored.

The commercial system is no longer supported by the manufacturer and does not allow monitoring of the currents with as fine a precision as the HV system used for the non-ME1/1 chambers, as can be seen in figure 53. For these reasons, during LS2, the commercial HV subsystem was replaced by a custom-made HV one, extending the system used for the non-ME1/1 CSCs. Besides ensuring precise current measurements, such a replacement makes the entire CSC HV system homogeneous and simpler to operate and maintain.




Figure 53. Left: instantaneous luminosity versus time for one of the LHC fills in 2016. Right: current (in nA) in an ME2/1 chamber (the HV segment closest to the beam) for the same fill, as measured with the custom-made HV subsystem used for non-ME1/1 chambers; current (in μA) in an ME1/1 chamber (one plane) for the same fill, as measured with the commercial HV subsystem.

6.2.3 Longevity studies

The CSC chambers will experience a much higher radiation dose than the muon chambers located in the barrel region. The highest neutron fluence and total ionization dose, corresponding to 10 years of HL-LHC running, will reach a level of $1 \times 10^{14} \text{ n}_{\text{eq}}/\text{cm}^2$ and 10 Gy, respectively. The total charge released in the gas volume per unit wire length will be about 0.3 C/cm for the CSCs closest to the beam line, corresponding to an integrated luminosity of 4000 fb^{-1} , as maximally achievable by the end of the HL-LHC.

Longevity tests of CSCs have been underway since 2016 at the CERN GIF++ facility [107]. Typical ME1/1 and ME2/1 chambers have been used. The ME1/1 chambers differ from those of the other rings, of which an ME2/1 chamber is typical, in having thinner wires, smaller gas gaps, and being constructed of different materials. The combination of a photon flux at high rates provided by the GIF++ gamma source and a muon beam periodically provided by the CERN SPS allows the study of not only the longevity of the CSCs but also their performance in an HL-LHC-like environment. The CSCs installed in GIF++ have the same chain of trigger and data acquisition electronics as those installed in CMS during Run 2. The working gas mixture of 40% Ar, 50% CO₂, and 10% CF₄ is supplied with a closed-loop gas system replicating the CSC gas system at CMS. The gas flow and gas refreshing rate are kept at the same level as for nominal CSC operation. While operating at the nominal gas gain under an intense irradiation for more than a year, the chambers accumulated a charge per wire length of 0.33 C/cm; this corresponds to about 1.1 and 1.7 times the amount of charge expected for ten years of HL-LHC running for the ME1/1 and ME2/1 chambers, respectively. Figure 54 shows a modest deterioration of spatial resolution at higher background rates and no deterioration with an accumulated charge up to 0.33 C/cm, the measurements were made with the nominal CSC gas mixture. The average current is proportional to the background intensity, which can be adjusted using a set of filters. The expected typical currents for Run 3 (at $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) and HL-LHC (at $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) background conditions are about 4 and 30 μA for the ME1/1 chamber, and 4.4 and 22 μA for the ME2/1 chamber. No significant deterioration of key chamber parameters such as gas gain, detection efficiency, spurious signal rates, strip-to-strip resistance, or dark currents has been observed. Thus, it is expected that the CSCs themselves will function throughout the planned HL-LHC operation.




Figure 54. Spatial resolutions for an ME1/1b (left) and ME2/1 (right) chamber using a muon beam while being uniformly illuminated by a ^{137}Cs photon source to simulate the background from high luminosity pp collisions, as a function of the average current per layer in one HV segment. The results for four different accumulated charges per unit wire length are shown, along with linear fits to each set.

The chamber gas gain is set by the HV system, which can be individually adjusted in each of up to 30 segments per CSC. Except for the ME1/1 chambers, each wire plane in other CSCs is divided

into three or five independent HV segments, which allows the independent regulation or turning off the HV on any of the individual sections. To maximize the lifetime of a CSC it should be operated at the lowest HV compatible with full efficiency. Until mid-2016, all CSCs operated at the same HV, and the average gas gains in different HV segments varied by up to a factor of 2, as shown by the wide blue distribution in figure 55. In 2016, the gas gains of individual CSCs were modified by tuning each HV channel to reduce the spread, as shown by the narrow red distribution in figure 55. This optimized the CSC gas gains for good efficiency without having unnecessarily high HV on any chamber, thus maximizing chamber longevity. Later, in 2018, the overall HV of all rings except ME234/1 was reduced by about 30 V, which decreased the gas gain by about 20%. The chambers now operate just above the knee of the efficiency plateau in gas gain versus HV and hence remain fully efficient, while the CSC system lifetime is expected to be extended by 20%.




Figure 55. Relative gas gain distribution in CSCs before and after the gas gain equalization campaign in 2016 [100]. Each entry in the histogram presents the mean value of gas gain in each HV channel. The scale of the blue histogram is on the left while the scale of the red histogram is on the right.

The CSC readout electronics can also degrade after exposure to large radiation doses. A systematic program of irradiating the CSC electronic components was performed to identify any of those unable to operate reliably in the HL-LHC radiation environment. All the components used in both the old and upgraded readout boards were found to withstand more than 3 times the expected HL-LHC doses, except for the PROMs used in the frontend boards installed in the ME1/1 chambers, which can only withstand 1–1.5 times the expected dose.

Before the production of the new electronics boards for the Phase 2 upgrade, radiation tests were carried out at the Texas A&M cyclotron and nuclear reactor [108], and at the UC Davis cyclotron. During the tests at the Texas A&M cyclotron, the digital components on the test boards were operated with active data readout while being irradiated with 55 MeV protons. The components tested for SEUs and single-event latch-ups (SEUs) included the FPGAs, PROMs, level adapters, and optical transmitters and receivers. In reactor tests, components were exposed to neutrons with energies up to a few MeV, with exposures equivalent to a TID of 30 kRad. This corresponds to a level of neutron radiation equivalent to about 50 years of that expected at the HL-LHC at the location where

CSC electronics are exposed to the highest radiation flux (the inner portion of the ME1/1 chambers). These tests targeted mostly nondigital components such as voltage regulators and power diodes. The results showed that the components selected for the new electronics will operate reliably in the CMS radiation environment at the HL-LHC.

During Run 2, a new campaign of radiation testing was initiated for the upgraded electronics described above. Particular attention was paid to the PROMs, which are known to have some failures after large radiation exposure [109]. The PROMs in the DCFEBs and ALCs installed during LS1 were tested at the CHARM II mixed radiation facility at CERN, the Texas A&M reactor, and the UC Davis cyclotron. The PROMs performed well up to an exposure of 10 kRad, but both types of PROMs showed some failures at exposures of 15–30 kRad. To mitigate such failures, the option to perform promless programming of the frontend board FPGAs for the ME1/1 chambers, where the radiation is most severe, has been provided. The optical receivers used on the OTMBs upgraded during LS2 were also tested at the UC Davis cyclotron and proved able to sustain the expected HL-LHC radiation dose.

The CSC system has been successfully operating during Run 3 after an enormous upgrade effort in LS1 and LS2. Its longevity has been studied extensively, and it is expected to remain reliable throughout the future running until the end of the HL-LHC era.

6.3 Resistive plate chambers

6.3.1 General description

The CMS resistive plate chambers (RPCs) are gaseous detectors equipped with two gas gaps each having a width of 2 mm and a copper readout plane in between, as shown in figure 56. High voltage is applied to the graphite electrodes, which are coated on the surface of high-pressure (HPL) laminate plates with bulk resistivity in the range of $2\text{--}5 \times 10^{10} \Omega\text{cm}$. The chambers are operated in avalanche mode with a gas mixture of 95.2% $\text{C}_2\text{H}_2\text{F}_4$, 4.5% $i\text{-C}_4\text{H}_{10}$, and 0.3% SF_6 . This allows the chambers to cope with high background rates and ensures an excellent time resolution, as summarized in table 7, facilitating a precise bunch-crossing assignment.




Figure 56. Left: schematic of the double-layer layout of the RPC chambers. Reproduced with permission from [110]. Right: illustration of the RPC technology. Reproduced with permission from [111].

The RPC barrel detector is divided in the direction along the beam axis into five separate wheels labelled $W_{\pm 2}$, $W_{\pm 1}$, and W_0 . The RPCs in each wheel consist of six layers. The first four layers, called RB1in, RB1out, RB2in, and RB2out, are located on the inner and outer sides of the inner two stations of the DT chambers. The other two layers, labelled RB3 and RB4, are located on the inner side of the third and fourth stations of the DTs. There are four disks in each endcap, called $RE_{\pm 4}$,

RE $\pm$ 3, RE $\pm$ 2, and RE $\pm$ 1. Spanning the ϕ direction, each wheel is divided into twelve sectors and each disk into 36 sectors. Due to requirements in the trigger logic, the chambers are divided into two or three pseudorapidity (η) partitions, called rolls. In most of the barrel, there are two rolls: forward and backward. Only RB2in in W $\pm$ 1 and W0 and RB2out in W $\pm$ 2 are divided into three rolls, called forward, middle, and backward. The endcaps are divided into three rolls: A, B, and C. The RPC system consists of 1056 chambers, covering an area of about 3950 m², equipped with 123 688 readout strips. In the barrel, the strips are rectangular in shape with a pitch in the range between 2.28 and 4.10 cm, while in the endcaps they are trapezoidal with a pitch between 1.74 and 3.63 cm.

6.3.2 RPC system upgrades since Run 1

Endcap fourth station upgrade. One of the major upgrades of the CMS experiment, performed during the first long shutdown (LS1) of the LHC in 2013–2014, was to add 144 new RPCs to the fourth stations of the detector endcaps. Adding these stations increased the overall robustness of the muon spectrometer and improved the trigger efficiency in the endcap region over the range $1.2 < |\eta| < 1.6$.

The new RE4 RPCs inherit their design from the previously installed endcap chambers [112]. They are grouped in 72 supermodules, 36 per disk, each consisting of two detectors and covering 10° in azimuthal angle. Prior to their installation in CMS, all RE4 chambers (both RE $\pm$ 4) passed a series of tests at all production stages in order to ensure their quality and performance.




Figure 57. Efficiency distribution of the RE $\pm$ 4 stations in their first year of operation in 2015.

This newest part of the RPC system was successfully commissioned in 2015 at the beginning of the Run 2 data taking. The performance of RE4, shown in figure 57, is in good agreement with expectations [113].

Readout changes and Phase 1 RPC trigger contributions. As mentioned in section 10.2, the Phase 1 level-1 (L1) trigger upgrade [91] moved from a muon detector-based scheme to a geometry-based system. The three muon trigger systems were replaced by three track finders, each covering a specific pseudorapidity region.

The RPC pattern comparator trigger (PACT) has the function of reading out the information from the RPC frontend boards (FEBs). These are installed at chamber level via the link board

system, located in the CMS experimental cavern in the balconies next to the detector. In the link boards, the low-voltage differential signaling (LVDS) signals from the FEBs are synchronized with the LHC clock, converted to optical signals, and sent to the RPC trigger boards in the service cavern. Since the beginning of Run 2, the RPC signals coming from the link boards are split and sent to the muon processors TwinMux, as described in section 6.1.2, and to the concentration pre-processing and fan-out (CPPF) [114]. Both TwinMux and CPPF do the hit clustering and cluster selection and forward the produced clusters to the track finders (TwinMux to the BMTF and CPPF to the EMTF). The OMTF, described in section 10.2, reads the RPC data straight from the link boards, then combines them with the full list of DT trigger segments directly at the track-finding step and performs the clustering and selection.

The contribution of the RPC system to the refactored L1 muon trigger architecture is different for the three muon track finders and can be summarized as follows:

- **BMTF:** RPC timing information is used to improve the DT trigger primitives' bunch crossing assignment. In the first two stations, where two RPC layers are present, a segment is built from the coincidence of hits in the inner and outer chambers. The latter provides redundancy in case of DT inefficiencies.
- **OMTF:** RPC hit position information is used standalone from the eight available chambers (five in the barrel and three in the endcap), per ϕ division (sector).
- **EMTF:** RPC hit position information is used in case there is no corresponding CSC trigger primitive.

The combination of information from the barrel muon detectors at an early stage allows the exploitation of the system redundancy already at the step of building the trigger primitives. As shown in the left plot of figure 58 [95], the combination of DT and RPC leads to an average increase of the station-1 barrel trigger-primitive efficiency of about 1.4%, raising the plateau value from 95% to about 96.5%. However, this is not the only gain from the new architecture. The trigger also benefits from detector complementarity since the use of RPC timing information reduces the number of out-of-time DT trigger primitives, as shown in figure 43.

For the OMTF, the use of the RPC complementarity is even more prominent. Figure 58 (right) shows the efficiency as a function of the reconstructed muon p_T in the η region of the OMTF. If RPC data are not present, the efficiency decreases by about 15%.

6.3.3 RPC system longevity

Detector stability studies. Continuous studies are performed throughout the data-taking periods to ensure correct operation and performance stability of the RPC system.

RPC working point calibration. To ensure the most stable performance possible, the operational high voltage V_{app} of the RPCs is controlled such that the effective high voltage V_{eff} is constant even when environmental conditions change during the data taking [116]. The relation between V_{app} and V_{eff} is as follows [117]:

$$V_{\text{app}} = V_{\text{eff}} \left[1 - \alpha + \alpha \left(\frac{P}{P_0} \right) \left(\frac{T_0}{T} \right) \right], \quad (6.1)$$




Figure 58. Left: station-1 barrel trigger-primitive efficiency as a function of muon pseudorapidity. Right: trigger efficiency as a function of muon p_T for the OMTF, derived from a trigger emulation applied to real data, using (red) and not using (blue) RPC information. Reproduced from [115]. © 2020 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

where P and T are the actual pressure and temperature in the CMS cavern, α is a special pressure correction that equals 0.8, and P_0 and T_0 correspond to values of 965 mbar and 273 K, respectively.

To determine the optimal operating voltage for every chamber, a series of high voltage scans are regularly performed, typically in low and high luminosity conditions. Data are collected for particular detector configurations: the V_{eff} values are equidistantly chosen within a range of 8.8 and 9.8 kV. The data are selected using triggers from the DT and CSC and then reconstructed using the standard CMS muon reconstruction [8]. During HV scan periods, no correction due to pressure or temperature is applied to the voltage. The segment extrapolation method [85] is used to measure the RPC hit efficiency for each HV point. Segments from the nearest DT and CSC chambers are extrapolated to the RPC planes, and the reconstructed RPC clusters (rechits) are matched to the extrapolated points. The efficiency is then calculated as a ratio of the numbers of matched RPC clusters to the extrapolated segments. The resulting efficiency distribution (ε) is fitted for every chamber using a sigmoid function defined as

$$\varepsilon(V_{\text{eff}}) = \frac{\varepsilon_{\text{max}}}{1 + \exp[-\lambda(V_{\text{eff}} - V_{50\%})]}, \quad (6.2)$$

where λ characterizes the slope of the sigmoid at the inflection point, ε_{max} represents the asymptotic efficiency for $V \rightarrow \infty$, and $V_{50\%}$ is the inflection point for which ε_{max} reaches an efficiency of 50%.

The optimal working point (WP) for each chamber is determined by interpolating the efficiency distribution ε using the fitting parameters in eq. (6.2), and calculating the voltage corresponding to 95% of ε_{max} . Then, the WP for each chamber is defined as $V_{95\%} + 100$ V (120 V) for the barrel (endcap) chambers, respectively. The application of individually different offsets leads to similar overall efficiencies in the barrel and endcaps [116, 117].

Since each HV channel supplies two chambers in the endcaps (with a few exceptions), the WP in the endcaps is the average of the WP of the two corresponding chambers for each HV channel. All chambers that operate in single-gap mode or have HV problems are excluded from this method.

Figure 59 shows the evolution of the efficiency as a function of time, as determined from the HV scan data, evaluated at the WP and at $V_{50\%}$ during the LHC Run 1 and Run 2. Despite the changes in the environmental and luminosity conditions, the different calibrations implemented in the detector allowed us to keep the efficiency stable.




Figure 59. Temporal evolution of efficiencies determined from HV scan data at the WP and at $V_{50\%}$, for RE1/2/3 (upper left), RE4 (upper right), and the barrel (lower). The light blue bands show the histograms of the distributions, where the width of the band represents the population of channels having the corresponding efficiency value.

RPC efficiency and cluster size stability. One of the most important measures of the RPC system performance is the efficiency of the rolls, described in section 6.3.1. The efficiency of a roll can be measured as the ratio of the number of reconstructed hits to the number of muons passing through the roll. To obtain the position of the muon trajectory in the RPC layers, the muon track is extrapolated using either the track parameters or the direction of the track segment from the closest other subdetector in the muon system. During the Run 2 data-taking periods, the analysis used both methods: the track [118] and segment [85] extrapolation.

Figure 60 shows the overall efficiency distribution of the RPC rolls in Run 2. The numbers are for all well-performing RPC rolls. Rolls with an efficiency lower than 70% are excluded if the efficiency drop is caused by a known hardware problems, such as a gas leak.




Figure 60. Distributions of the overall RPC efficiencies in the barrel (left) and endcaps (right) during pp data taking in Run 2 [119]. Reproduced from [120]. © 2020 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

The cluster size (CLS) is another important quantity since it affects the RPC spatial resolution. It is defined as the number of adjacent strips firing when a discharge is produced in the RPC. The RPC system has an average cluster size of less than three strips, constant over the years, in agreement with expectations [1]. Keeping the cluster size stable over time is one of the most important prerequisites in the correct operation of the RPC system. During Run 2, the chamber cluster size was monitored run-by-run to guarantee the stability of the system.

The efficiency and cluster-size history for the Run 2 data taking are shown in figure 61. The history follows the changes of the applied HV WPs and changes of the isobutane concentration in the gas mixture. The spread in the cluster size distribution in 2015 was caused by a threshold control problem, which was resolved in October 2015, while the small drop of the RE+4 cluster size in 2016 was due to a temporary LV problem, solved at the end of August 2016. The drop of efficiency in the period 01–19 August 2018 came from a configuration problem. Detailed Run 2 performance results can be found in refs. [119, 121].

One of the main objectives of the RPC analysis is to monitor the system performance with respect to the luminosity. Figure 62 displays the barrel and negative endcap efficiency and cluster size as a function of the instantaneous luminosity during pp data taking in 2016 and 2017. For this comparison, data recorded with the same WP are used. The lower efficiency and cluster size for the barrel in 2016 are caused by the higher isobutane concentration during that year.

The comparison between the 2016 and 2017 results shows a stable efficiency and cluster size. The obtained results can be linearly extrapolated to the expected luminosity at the HL-LHC of $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. A reduction in efficiency of 1.35% is predicted for the barrel and 3.5% for the endcaps. No change is expected for the cluster size under the HL-LHC conditions.

Monitoring of RPC currents. The current drawn by the RPC detectors is one of the base parameters to be considered, both as a measure of the working condition and as an indicator of possible problems. The current should be kept as low as possible to guarantee long-term stability of the RPC gas gaps, and its time evolution is one of the most important input data for longevity studies.

The ohmic current of the RPC system is defined as the current with no beam in the accelerator, at a HV around 7000 V, where the current follows an ohmic law and there is no contribution from




Figure 61. History of the RPC efficiency (upper) and cluster size (lower) during Run 2. Gray areas correspond to the scheduled technical stops.




Figure 62. RPC barrel and endcap efficiency (upper) and cluster size (lower) as a function of the LHC instantaneous luminosity for pp collisions in 2016 and 2017. The linear extrapolation to the instantaneous luminosity expected at the HL-LHC of $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ shows a 1.35% reduction in efficiency for the barrel and 3.5% for the endcap. Reproduced from [121]. © 2019 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

gas amplification. In practice, the ohmic current values are monitored at 6500 V. The “cosmic current” is defined as the current without beam, at the WP voltage, in the region of gas amplification.

For Run 3, the monitoring capability for the RPC detector currents was improved. An especially designed tool was implemented that automatically stores and classifies the currents in a database as a function of different parameters, such as the CMS magnetic field, the instantaneous luminosity, and other parameters. The tool can identify blocks of data of special importance, such as the periods of HV scans, described in section 6.3.3.




Figure 63. Ohmic current history in W+0, RE+1, RE+4, and RE-4. Reproduced from [120]. © 2020 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

Figure 63 shows the currents measured in four RPC stations, W+0 in the barrel, and RE+1, RE+4, and RE-4 in the endcaps. The measured currents are shown as a function of time. An increase of the ohmic currents is observed, in particular for chambers that are more exposed to background, and a linear dependence of the rate on the instantaneous luminosity is seen. In low-background regions such as W+0, the ohmic current increases very slowly. In RE+1 and W+0, the background rate is less than 10 Hz/cm^2 . Both have similar gas flows, corresponding to a volume exchange per hour (v/h) of 0.7 and 0.6, respectively. In contrast, in RE4, the background rate is about 40 Hz/cm^2 and the gas flow is 0.35 v/h , which is lower than the rest of the system, due to a wrong calibration of the flow cells. These rates are the maximum values measured in the upper sectors, i.e., the detector sectors above the beam pipe, at instantaneous luminosities of $1.5\text{--}2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

The RPC currents depend linearly on the instantaneous luminosity [122]. For each LHC fill, the distributions were fit to a linear function to obtain the slope P_1 , also known as the physics current ($i = P_1 L$). Due to the nature of the linear fit, an offset P_0 absorbs the cosmic current (offset + ohmic + gas gain). The slopes P_1 as a function of time for the endcap stations are shown in figure 64. The slope of the RPC current distribution is stable in time. The changes in the middle of August 2018 are due to different HV WPs. Endcap stations that are located at equal distances from the interaction point along the beam pipe have similar slopes. They also have similar rates [122]. No increase due to the integrated luminosity is observed for the slopes over the entire year.

Hydrogen fluoride (HF) is produced in the gas under high electrical discharge. The HF has a high chemical reactivity and electrical conductivity [123], and it is therefore expected to be a source of inner-detector surface damage and relative ohmic current increase that accelerates detector aging.




Figure 64. RPC physics current history in RE±1, RE±2, RE±3 and RE±4. Reproduced from [120]. © 2020 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

In summer 2018, measurements of the HF concentration were performed at the gas exhausts of three regions: W+0 in the barrel, and RE+1 and RE+4 in the positive endcap.

The ohmic currents as a function of HF concentration are shown in figure 65. The RE+1 and W+0 chambers have a similar HF concentration and gas flow (0.7 and 0.6 v/h), and a background of less than 10 Hz/cm². In RE+4, the amount of HF accumulated is higher by a factor 2 at a background of 40 Hz/cm², and the gas flow is lower (0.35 v/h) than in W+0 and RE+1. There is a clear linear dependence between the ohmic current and the HF concentration, which implies that HF trapped in the gas gap may form a thin conductive layer. The HF can be efficiently removed by increasing the gas flow in the chambers, depending on the background rate.




Figure 65. Ohmic current as a function of HF concentration. Reproduced from [120]. © 2020 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

GIF++ longevity studies. The present RPC system was originally certified for ten years of LHC operation at a maximum background rate of 300 Hz/cm² and a total integrated charge of 50 mC/cm² [124, 125]. Based on the data collected in Run 2, assuming a linear dependence of the

background rates as a function of the instantaneous luminosity, and including a safety factor of three, the expected background rates and integrated charge at the HL-LHC will be about 600 Hz/cm^2 and 840 mC/cm^2 , respectively [100]. In such operating conditions, irrecoverable aging effects can appear, as the higher collision rates affect the detector properties and performance.

Therefore, since July 2016, a long-term irradiation test has been carried out at the CERN gamma irradiation facility (GIF++) [126] to study whether the present RPC detectors can survive the difficult background conditions during the HL-LHC running period [127].

Four spare RPC chambers have been irradiated, two each of type RE2/2 and RE4/2 [1, 100]. These are from the endcaps where the backgrounds are expected to be maximal [128]. Two different RPC production types have been tested, reflecting the fact that the RPC endcap production was done in two different periods, 2005 for the RE2 detectors (both RE $\pm$ 2) and 2012–2013 for the RE4 detectors (both RE $\pm$ 4). Two chambers, one from each period, are continuously being irradiated while the other two of the same type are kept as reference and are switched on only from time to time. The detector parameters, such as dark currents, noise rates, currents, and count rates for various background conditions are monitored continuously and compared with the measurements from the reference chambers.

The integrated charge is calculated as the average density current accumulated in time in the three gaps that constitute the detector, since the gamma flux, provided by the $14 \text{ TBq } ^{137}\text{Cs}$ source at the GIF++ [126] is uniformly distributed over the detector surface. The collected integrated charge from the beginning of irradiation until September 2022 is about 813 and 478 mC/cm^2 for the RE2 and RE4 chambers, respectively, which corresponds to approximately 97 and 57% of the expected integrated charge at the HL-LHC.

Dark current and noise rate studies. Dark currents and noise rates are monitored periodically in order to spot aging effects due to irradiation. The dark-current density, i.e., the currents normalized to the surface area, for both the irradiated and reference RE2 chambers are shown in figure 66 as a function of the collected integrated charge. The dark currents were measured at 6.5 kV to determine the ohmic contribution, and at 9.6 kV to include the contribution from gas amplification.

Figure 67 (left) shows the dark-current density monitored as a function of the effective high voltage, i.e., the voltage normalized to the standard temperature of 20°C and pressure of 990 mbar [112] at different values of collected integrated charge. Since the beginning of irradiation, the dark currents have been stable in time, with only small acceptable variations. Figure 67 (right) displays the average noise rate for the irradiated and reference RE2 chambers as a function of the collected integrated charge. The average noise rate is stable with time and less than 1 Hz/cm^2 .

Resistivity and current studies. Other important parameters that are measured periodically are the current in the presence of background radiation and the resistivity of the electrodes. The latter is measured several times per year, since it is a crucial performance parameter. The resistivity is measured by filling the detector with pure argon and operating it in a self-sustaining streamer mode, which occurs when the gas-quenching components such as isobutane are removed. The streamers propagate over the entire detector area, and by measuring the current at a given high voltage, the resistance, and hence the resistivity, can be calculated. The measured resistivity values are normalized to 20°C to allow comparison of the values for different temperatures [130].




Figure 66. Dark-current density for the irradiated (blue squares) and reference (red circles) RE2 chambers as a function of the collected integrated charge at 6.5 (left) and 9.6 kV (right).




Figure 67. Left: dark-current density monitored as a function of the effective high voltage at different values of the collected integrated charge for the irradiated RE2 chamber. Right: average noise rate as a function of the collected integrated charge for the irradiated (blue squares) and reference (red circles) RE2 chambers. Reprinted from [129], Copyright (2023), with permission from Elsevier.

To exclude effects from external parameters, the ratios of resistivity and current between irradiated and reference chambers are taken for a gamma background rate of about $600 \text{ Hz}/\text{cm}^2$, as shown in figure 68. An increase in the resistivity is observed in the irradiated chamber during the first irradiation period, up to $\approx 300 \text{ mC}/\text{cm}^2$, when the detectors are operated in conditions similar to those in CMS: one gas volume exchange per hour and a relative gas humidity of 35–45%. While the RPC system operates in these conditions of humidity and gas volume exchange, they are not the optimal ones for operation at GIF++, where the high background gamma rate is about $600 \text{ Hz}/\text{cm}^2$, causing the HPL plates to dry

and their resistivity to increase. During the longevity test, after an integrated charge of $\approx 300 \text{ mC/cm}^2$ was reached, the relative gas humidity was increased and maintained at about 60%, and the gas flow was increased to three gas volume exchanges per hour. With these changes, the HPL resistivity decreased and the variations were reduced. After $\approx 650 \text{ mC/cm}^2$, an increase in resistivity is observed, which is related to the return to a low gas humidity at about 40% due to sharing the same gas system with other RPCs. The increase of resistivity is confirmed by the decrease of the measured current.




Figure 68. Resistivity ratio (blue squares) and current ratio (red circles) between the irradiated and reference RE2 chambers as a function of the collected integrated charge. Reprinted from [129], Copyright (2023), with permission from Elsevier.

Detector performance monitoring. The detector performance has been measured during test beams before irradiation and at different periods of irradiation. Comparisons between the irradiated RE2 chamber efficiency measured as a function of the effective HV without background radiation and in the presence of a background of 600 Hz/cm^2 , at different values of collected integrated charge, shows that the efficiency is stable over time in the absence of background radiation, and no shift in the WP is observed [127, 131]. In the presence of background, the efficiency is stable at the WP, but a WP shift of 100 V after collecting 260 mC/cm^2 of integrated charge is introduced. This shift in the detector WP is related to the increase in the resistance (R) of the electrodes, which causes an increase of the voltage drop (RI) across them. The latter leads to a difference between the effective voltage V_{eff} applied to the electrodes and the effective voltage across the gas gap V_{gas} , which is compensated by introducing the shift [132]. A similar increase in resistivity is seen in the irradiated RE4 chamber. The quantity V_{gas} is defined as:

$$V_{\text{gas}} = V_{\text{eff}} - RI, \quad (6.3)$$

where R is the HPL resistance and I is the total current.

Since the detector operation regime is invariant with respect to V_{gas} , the efficiency as a function of V_{gas} does not depend on the HPL resistance, as shown in figure 69 (left). After removing the R -increase effect on the electrodes by introducing V_{gas} instead of V_{eff} , all the efficiency curves overlap, and no shift in the WP is observed.




Figure 69. Left: irradiated RE2/2 chamber efficiency as a function of V_{gas} for different background irradiation rates, up to 600 Hz/cm^2 , and different integrated charge values. Different marker shapes of the same color represent different background rates at the same integrated charge values. Right: irradiated RE2 chamber efficiency at the WP as a function of the background rate at different values of the collected integrated charge. Reprinted from [129], Copyright (2023), with permission from Elsevier.

The irradiated RE2 chamber efficiency at the WP is measured at different background rates and at different integrated charge values, as shown in figure 69 (right). The efficiency is stable over time up to the highest background rate expected at the HL-LHC.

Longevity studies on spare RPCs are ongoing at GIF++ under controlled conditions. Preliminary results show no evidence of aging effects. The main detector parameters and performance are stable.

6.3.4 Changes to the RPC system in LS2

Gas leak repairs and green-house gas emission strategy for Run 3. The highest priority during LS2 was the consolidation of the RPC gas system. The aim was to minimize the gas emissions and thus the environmental impact, since the standard gas mixture is composed mainly of F-gases with high global-warming potential. To accomplish this, actions were taken to minimize the number of leaks and to implement a newly developed Freon recuperation system prototype [133].

The RPC gas system is a 13 m^3 closed-loop volume with re-circulation of $7.3 \text{ m}^3/\text{h}$ nominal mixture flow: $5 \text{ m}^3/\text{h}$ for the barrel and $2.3 \text{ m}^3/\text{h}$ for the endcaps. A combination of factors during production of the components and operation of the system can lead to an increase of the leak rate. Environmental conditions in the experimental cavern, such as humidity and temperature, as well as switching from Freon to N_2 and back to Freon, can accelerate the degradation process of the different components of the RPC gas system.

The gas leaks in the RPC system are mainly caused by the T-shaped or L-shaped polycarbonate gas connectors that are broken due to stress applied through the gas pipes, and the low-density polyethylene pipes that are brittle, deteriorated, or cut.

A special gas-leak repair procedure was developed to correctly identify the leaks using an endoscope, which allowed us to determine the exact location and the components that are sources of

leaks. The RPC barrel chambers are coupled with the DT chambers and inserted into the iron yoke. Therefore, in order to reach the leak location, a partial extraction of the entire muon station (RPC and DT) by a distance of 80 cm from the back or front side is required. A surgical cut of the C-aluminum profile is then performed to gain access to the broken gas pipe or T/L connector. In figure 70, an example of the newly developed gas-leak repair procedure is shown. In one scenario, the broken pipe connecting the two chambers is removed, changing the splitting of the gas flow from internal to external (figure 70, upper center) and preserving the parallel gas distribution of the chambers. In another scenario, the repair is accomplished by gluing the T/L connector (figure 70, lower center).



Figure 70. Photographs illustrating the gas-leak repair procedures. Left: scientists working on the repair procedure. Middle left: access to broken component. Middle right: repairing of components. Right: closing and validation.

By November 2020, when all possible repairs were finalized, a total number of 88 gas leaks due to cracked or broken pipes were identified in the barrel chambers. Of those, 50 chambers were successfully repaired and returned to normal operation. This included 17 chambers that did not work during Run 2. The remaining 38 leaking chambers are either not accessible or the partial extraction technique is not applicable, either because the source of the leak is inaccessible or the source of the leak is not identified. In the endcaps, a total of 11 chambers were replaced.

The gas-leak repairs had a large impact on the chamber performance. Figure 71 shows a comparison of the efficiency of repaired chambers using cosmic ray data between the end of Run 2 before the reparation campaign and after all the repairs. This led to an overall gain in efficiency of 1.5%.

To minimize chamber pressure variations, which are considered to be a cause of leaks, automatic regulation valves on pre-distribution racks in the underground service cavern were installed in November 2021. This operation has substantially reduced the development of new leaks in the system.

In the beginning of Run 3 in June 2022, the total number of disconnected barrel chambers was 108, corresponding to about 10% of the entire RPC system. This number includes 14 dysfunctional chambers with gas leaks, 63 operational chambers with gas leaks, and 31 gas-tight chambers that share gas channels with leaking chambers.

Another major activity was testing a Freon-recuperation prototype connected to the RPC gas system exhaust. The aim is to recuperate the $\text{C}_2\text{H}_2\text{F}_4$ from the RPC gas mixture and re-use it [133]. The $\text{C}_2\text{H}_2\text{F}_4$ and $\text{i-C}_4\text{H}_{10}$ gases form an azeotrope, i.e., a mixture of liquids whose proportions can not be altered or changed by simple distillation, because the intramolecular force of same species is




Figure 71. Efficiency comparison between 2018 (red) and 2021 (blue) cosmic ray data for chambers with repaired gas leaks.

much higher than the reciprocal attraction. The separation of the two gases is done by slowly heating the liquefied azeotrope, allowing the enrichment of the liquid R134a and the $i\text{-C}_4\text{H}_{10}$ vapor. The CERN EP-DT gas team is finalizing the R&D on the first custom-built $\text{C}_2\text{H}_2\text{F}_4$ -recuperation system with an expected efficiency of 80% [134]. To maximize the amount of gas in the exhaust line in order for efficient recuperation system operation, the procedure is to turn off and disconnect all leaking chambers. The recuperation system is expected to be ready and operational by summer 2023.

RPC commissioning during LS2. Figure 72 (left) shows the efficiencies per roll, calculated using the segment extrapolation method described in section 6.3.3. Comparing the average efficiencies for 2018 and 2021, an increase of 0.6% is observed. This increase can mainly be attributed to the recovery of chambers from single-gap to double-gap operation mode. The fraction of chambers with more than 70% efficiency rises by 6.1%, considerably raising the redundancy in the measurement of muons. These improvements are a result of the extensive maintenance and repairs during LS2.

In figure 72 (right), the cluster size distribution for barrel wheel W-2 is displayed. The average cluster size is kept around two strips, far below the possible maximum of three, in accordance with the trigger requirements. The comparison between the 2018 and 2021 cosmic ray muon data shows a stable cluster size.

The efficiency measurement using cosmic ray muon data in the endcaps can have large systematic uncertainties, related to the vertical geometry of the CMS muon system in the endcap region. Even so, the current and occupancy values at the WP during the cosmic ray muon data taking show good agreement with the expectation.

After extensive maintenance and repairs, as well as commissioning and longevity studies during LS2, the RPC detector successfully entered the data-taking period of Run 3.

6.4 Gas electron multiplier chambers

6.4.1 Motivation and general description

During the future operation of the High-Luminosity LHC (HL-LHC), the maximum hit rate in the forward region of the muon system is expected to reach 5 kHz/cm^2 in the first muon layer, with an




Figure 72. Left: distribution of the efficiency per roll in the RPC barrel chambers. Only chambers with rolls of efficiency greater than 70% are considered. Right: average cluster size for the RPC barrel chambers in wheel W-2. In both figures, the cosmic ray muon data from 2018 (2021) are indicated in red (blue).

integrated charge per unit area of 100 mC/cm^2 [100]. This increase compared to the present values represents a challenge to the forward muon system, which must be resistant to radiation, have a high rate capability, and maintain an adequate pattern recognition for efficient muon reconstruction, while minimizing the number of misidentified tracks and keeping the level-1 (L1) trigger rate at acceptable levels.

To enhance the track reconstruction and trigger capabilities of the endcap muon spectrometer, large-area triple-layer gas electron multiplier (GEM) detectors [135] were installed in the region covering $1.55 < |\eta| < 2.18$ of the CMS detector for the start of Run 3. This station, denoted GE1/1, is the first of three GEM rings that will be installed for the HL-LHC. Figure 39 shows a quadrant of the CMS detector with the location of the GE1/1 highlighted and outlined in red in the r - z plane. A detailed drawing is shown in figure 73.




Figure 73. Sketch of GE1/1 system of one endcap indicating its location relative to the full endcap muon system, the endcap calorimeter, and the shielding elements. Reprinted from [136], Copyright (2019), with permission from Elsevier.

Muons experience the largest bending in the muon spectrometer at the position of the first muon station as the magnetic field lines bend around the endcap flux return. Due to the increasing background rates at large η and the reduction in the magnetic field in the first muon station, the trigger rate in this region is large and difficult to mitigate. The insertion of the GE1/1 chambers increases the lever arm traversed by muons by a factor of 2.4–3.5, relative to ME1/1 alone, leading to a significant improvement in the muon trigger momentum resolution and a large reduction in the L1 trigger rate. Figure 74 shows the expected trigger rates in the forward muon spectrometer with and without the GE1/1 upgrade. This indicates that the upgrade will lower the trigger rates by a factor of 3–10 depending on the p_T threshold. The muon trigger is described in detail in section 10.2.




Figure 74. L1 muon trigger rate with and without the GE1/1 upgrade, assuming an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, where MS1/1 indicates the first muon station [100].

Each endcap GE1/1 detector consists of 36 double-layered triple GEM chambers located just in front of the first CSC station, labeled ME1/1, each covering a 10° sector in azimuth. The chambers provide full coverage in ϕ and were constructed in two sizes, the odd-numbered GE1/1 are slightly longer in order to maximize the pseudorapidity coverage while fitting in the available space constrained by the support structure, as shown in figure 75. The GEM detector technology can withstand rates up to 1 MHz/cm^2 and has excellent spatial and timing resolution of approximately $250\text{--}500 \mu\text{m}$ and $<10 \text{ ns}$ per layer, respectively. The combined GE1/1 station spatial resolution is on the order of $100 \mu\text{m}$. The GE1/1 detector is placed, in global CMS coordinates, in z between 566 and 574 cm, and at a radius between 145 and 230 cm.

6.4.2 Technical design

The CMS triple GEM detector is a micro-pattern gas detector that comprises four gas gaps separated by three GEM foils, as shown in figure 76. It has an active area of $990 \times (220\text{--}455) \text{ mm}^2$ with a 3/1/2/1 mm wide drift/transfer-1/transfer-2/induction field gap configuration [136]. The bottom of the GEM assembly is a printed circuit board that holds the drift electrode and voltage divider, while the top of the assembly is the readout board consisting of radially oriented readout strips along the long side of the chamber. The strip pitch ranges from 0.6 to 1.2 mm, and the readout board is segmented in up to $10 \times 3 \eta\text{--}\phi$ partitions, each with 128 strips. The triple GEM arrangement allows for a charge amplification factor of up to a factor of several 10^5 , while limiting the probability of electrical breakdown or discharge. The amplified charge induces a measurable signal on the readout




Figure 75. Left: layout of the GE1/1 chambers along the endcap ring, indicating how the short and long chambers fit in the existing volume. Reproduced from [137]. CC BY 4.0. Right: blowup of the trapezoidal detector, GEM foils, and readout planes, indicating the geometry and main elements of the GEM detectors, from ref. [138] John Wiley & Sons. © 2017 Society of Plastics Engineers.




Figure 76. Sketch of a triple GEM detector showing the three foils, cathode, readout PCB and amplification. Reproduced from [100]. CC BY 4.0.

electrode, which is segmented to provide positional information. The gas mixture was chosen to be Ar+CO₂ in a 70:30 proportion.


In the GE1/1 station shown in figure 73, pairs of triple GEM detectors are matched to form a “super-chamber”, providing two measurement planes and maximizing the detection efficiency for the station. Thus, if each chamber operates at 98.0% efficiency, the logical OR of the two signals from the super-chamber provides a 99.9% efficiency. The structure of a GE1/1 chamber is shown in figures 75 (right) and 76, while the dimensions and specifications of the two chamber types, “short” and “long”, are given in table 10.

The main elements of an individual chamber are: the drift board holding the drift electrode, the GEM foils that amplify the ionization signal, the readout board where induced charge is read out on segmented strips, the internal and external frames, and a gas distribution system. A more complete technical description of the layout and assembly of the GEM chambers can be found in ref. [136].

The drift board is a trapezoidally shaped printed circuit board (PCB) inside the active gas volume and coated with a copper layer that serves as the chamber’s drift electrode. Outside the

Table 10. Dimensions and specifications of the GE1/1 short and long chambers, from ref. [136].

Specification	Short	Long
Chamber length [cm]	113.5	128.5
Chamber width [cm]	28–48.4	26.6–51.2
Chamber thickness [cm]	1.42	1.42
Active readout area [cm ²]	3787	4550
Active chamber volume [liters]	2.6	3
Geometric acceptance in η	1.61–2.18	1.55–2.18

**Figure 77.** GE1/1 drift board (left) and a magnified view of the drift board (right) showing the HV pins and the resistor and capacitor network connecting to the chamber HV supply. Reprinted from. [136], Copyright (2019), with permission from Elsevier.

active gas volume, a 100 k Ω resistor and 330 pF capacitor are installed on pads on the PCB to limit the current from the high-voltage (HV) power supply and decouple the signal from the HV. Twelve pins that carry the HV to the GEM foils are mounted as shown in figure 77.

Four internal frames composed of halogen-free glass epoxy with thicknesses of 3/1/2/1 mm are coated with polyurethane varnish and are used to define the spacing between the drift board, the three GEM foils, and the readout board. The GEM foils are attached to the frames by screws that pass through the entire chamber assembly to hold it in place.

The GE1/1 detector uses three identical GEM foils shown in figure 78, which are thin polyimide foil clad with a thin copper layer containing micro-pattern holes etched in a periodic grid. A voltage up to 400 V is applied across the copper-clad surface producing a strong electric field of between 60 and 100 kV/cm. The triplet structure allows for an amplification of around 10^5 when moderately high voltages are applied. The side of the foil facing the readout boards is a continuous conductor, while the strips facing the drift board are segmented into sectors of approximately equal area of 100 cm². The segmentation ensures that, in the extreme case of a large discharge creating a short circuit between the GEM electrodes of a single sector, the affected dead area is limited to the 100 cm² of a single sector instead of deactivating the entire detector. Each sector is connected separately to the HV supply via a 10 M Ω resistor to limit currents from the HV supply and to quench any possible discharge.

The readout board is a printed circuit board with 3072 radially oriented readout strips on the side facing the interior of the chamber. The strips are connected by metalized vias to the outer side of the board, which are routed to readout pads on the exterior-facing side of the board. The board is segmented into 8×3 η - ϕ partitions with three sets of 128 strips in ϕ . Each set of 128 strips is read




Figure 78. Electron scanning microscope image of a GE1/1 foil (left) and a diagram of the multiplication principle (right). Reproduced from [100]. CC BY 4.0.

out by the first stage of the frontend electronics, as described below. The strip pitch varies between 0.6 and 1.2 mm on the shorter and longer ends of the trapezoidally shaped board.

6.4.3 Gas system

During operation, the chambers are run using an Ar:CO₂ gas mixture with a ratio of 70/30. A gas mixer is installed in the surface gas barrack (SGX) where the argon and carbon dioxide gases are mixed to the required percentages. The mixer has input and output lines that go to the underground gas room (UGX). Four racks in the UGX are used to pump and control the gas from the surface using input and output lines to a crate on the periphery of the first muon ring on each endcaps. The rack on the periphery uses twelve gas lines, each of which provides the gas mixture to six consecutive single chambers. Each chamber in turn has a single inlet and outlet through which the gas enters and exits the chamber. In figure 79 a diagram of the periphery crate (left) and the location of the gas lines (center) are shown. The entire system is monitored at different stages (mixer, pre-distribution, and distribution), and the pressure and flow-rate values are supervised by the gas control system (GCS).




Figure 79. Left: periphery GE1/1 gas rack. Center: location of the gas lines feeding the GE1/1 chambers. Right: exploded view of a GE1/1 chamber showing the gas plug attached to the readout board. Reprinted from [136], Copyright (2019), with permission from Elsevier.

Each chamber is connected to the gas system via a single inlet and single outlet attached to the readout board, as shown in figure 79 (right). The gas flows from an inlet placed at one corner of

the readout board across to the outlet placed at the diagonally opposite corner of the trapezoid and through holes in the internal frames.

6.4.4 Electronics

The GEM readout electronics are divided into frontend components that are mounted on the detector itself, and backend electronics that are outside the experimental cavern. A schematic overview of the different components is shown in figure 80. The readout is organized according to 128 readout strips by the VFAT3 chip [139], which is the ASIC on the readout board responsible for the digitization of the induced signals. The signals from individual VFAT3 chips are then sent via traces on a large chamber-sized PCB called the GEM electronics board (GEB) [140]. The VFAT3 and opto-hybrid [141] boards are attached to the GEB, which also routes high and low voltages to the chambers from external off-chamber power supplies. The opto-hybrid board transmits the data via optical links to the off-chamber backend readout and also sends data to the CSC trigger board for processing at L1. The optical links are routed from the opto-hybrid boards on the chambers to patch panels on the periphery crates and then to off-detector electronics in a μ TCA crate [142]. The backend boards are CTP7 μ TCA boards, which are described in ref. [143]. They contain bidirectional links to communicate with the frontend chamber electronics and are connected via AMC13 [25] cards to the CMS L1 trigger and DAQ systems [144], as discussed in section 9.




Figure 80. GE1/1 electronics overview showing the frontend electronics, VFAT, GEB, and opto-hybrid boards, as well as the optical links and backend readout with connections to the L1 trigger and DAQ.

The VFAT3 chip was developed from an existing frontend readout chip, the VFAT2 [145] (very forward ATLAS and TOTEM) ASIC, developed for the ATLAS and TOTEM readout. The VFAT3 design was modified for application in the CMS GEM to provide readout in a high-rate environment, collect the total charge produced by particles crossing the GEM volume, and be radiation hard. A picture and a high-level schematic of a VFAT3 chip are shown in figure 81.

The VFAT3 must collect, amplify, shape, digitize, and buffer the signal produced on each of the 128 channels connected to the readout strips on the GEM readout board. The preamplifier and shaper are charge sensitive and are followed by a constant-fraction discriminator that digitizes the incoming




Figure 81. Picture of the VFAT3 ASIC (left) and a high-level schematic (right). Reproduced from [146]. © CERN 2016. CC BY 3.0.

charge pulses. Following the discriminator, the binary comparator results are synchronized with the LHC clock in the synchronization unit. A programmable threshold allows a channel-by-channel calibration to optimize the signal-to-noise ratio, and is set to obtain approximately 98% efficiency per single hit so that the logical OR of two layers has a 99.9% efficiency.

The data are then split into two paths. A trigger signal, with fixed latency, is sent to the trigger unit, while an asynchronous tracking signal is sent to a buffer for full granularity readout via the communications port. The signal is accompanied by a time stamp, the bunch crossing, for further processing.

The signal can last up to approximately 60 ns depending on the gas mixture, and the VFAT3 chip allows the shaping time to be adjusted to fully integrate the charge and maximize the signal-to-noise. The timing resolution is optimized by the use of a constant-fraction discriminator and a gain with a programmable dynamic range. The buffer for the tracking data is implemented by a SRAM memory that operates as a circular buffer 128 channels wide and 1024 bunch crossings deep. It continuously samples all channels in every cycle of the LHC clock. After 1024 bunch crossings, the buffer overwrites the first entry.

Since the L1 trigger operates with fixed latency, there is a fixed time between the level-1-accept (L1A) trigger signal and the data corresponding to the bunch crossing of interest. Upon receiving an L1A, the appropriate data are transferred from the SRAM1 buffer to the SRAM2 buffer, which is 512 bunch crossings deep and stores the data until they are transmitted off the chip. The chip supports a latency of up to 12.5 μ s with an L1A rate of up to 1 MHz. The data can either be sent in a lossless mode or zero suppressed.

The fixed latency, or trigger path, is used to provide fast hit information synchronous with the LHC clock, which can be put in coincidence with other detectors to decide if the event will generate an L1 trigger. The logical OR of two adjacent channels is sent to the trigger unit.

The slow-control signals are used to send calibration, bias, monitoring, and control information via the communication port and to control and monitor the settings and status of the VFAT3 chip.

The output of each VFAT3 chip is routed along the GEM electronics board (GEB), which is mounted directly on top of the readout board. The GEB is divided into two sections that cover the entire chamber. The GEB routes the signals between the 24 VFAT3s on each chamber and the low voltage from the external power supplies to the chamber and frontend electronics. CAEN power supplies provide voltage to the FEAST boards mounted on the GEB. The FEAST boards perform the DC-DC conversion to provide the proper low voltages for the readout electronics, as well as the HV for the operation of the chamber.

The opto-hybrid board serves as the chamber hub for data communication, transfer, and control. It receives data from all the VFAT3 ASICs, formats it for the entire chamber, and sends the data using the GBT protocol and optical links to the backend electronics for final readout and use in the L1 muon trigger. Additionally, each opto-hybrid board contains three GBTx chipsets [147], one Virtex-6 FPGA [148], three versatile link transceivers (VTRx), and two versatile link transmitters (VTTx) [68]. Trigger data are sent over the VTTx via optical fibers to both the backend readout and the CSC trigger motherboard, while the three VTRx transmit tracking data to the backend readout (CTP7 card). Each GBT can handle up to ten frontend chips at a transfer rate of 320 MHz, and the tracking data are transferred at 4.8 Gb/s through the VTX. The trigger data are transferred at 3.2 Gb/s using 10b/8b encoding. In operation, the trigger data are formatted within the Virtex-6 FPGA, while the tracking data flow directly from the frontend by the GBT VTRx links to the backend. Calibration, and trigger, timing, and control (TTC) signals are sent to the opto-hybrid board and distributed to the frontend via the GEB. Irradiation tests have shown that the single- and double-error rates from single-event upsets (SEUs) are well within acceptable margins for the total ionizing dose (TID) expected over the lifetime of the HL-LHC project. A photo of the opto-hybrid board and the measured error rates are shown in figures 82 and 83 .




Figure 82. Photo of the GE1/1 opto-hybrid board with the Xilinx Virtex-6 FPGA at the center.

The backend electronics are μ TCA cards that were originally developed for a calorimeter trigger Phase 1 upgrade project [143]. The CTP7 has 67 optical receivers, 48 optical transmitters, a Xilinx Virtex-7 FPGA, and a Zync processor. The firmware was adapted for the GEM detector with 36 GBT cores servicing 12 triple GEM detectors with one link to a CMS μ TCA AMC-13 card operating the standard DAQ and trigger links. Using this hardware, all of the GE1/1 can be read out with one μ TCA crate hosting six CTP7 cards and one AMC-13 card.

6.4.5 CSC/GEM trigger for Run 3

Figure 84 shows a view in the global CMS transverse (ϕ - z) plane indicating the location of the GE1/1 and the first CSC station, ME1/1. The diagram indicates the relative position of GE1/1 and ME1/1




Figure 83. Single- and double-error cross sections measured in irradiation tests as a function of the TID.

in the global z coordinate and shows two muon trajectories as they traverse the stations' volume. The deflection of muons with a p_T of 5 and 20 GeV is approximately 3 and 12 mm, respectively, in the local x (global ϕ) direction. By reconstructing the ϕ direction in both the entrance to the GE1/1 station and the exit from the ME1/1 station, an estimate of the track p_T can be made and used for the trigger decision in the L1 muon trigger system. The GE1/1 chambers add two new independent position measurements, which improves the redundancy, lowers the muon misidentification rate, increases the robustness of segments found in the muon spectrometer, and raises the efficiency of the stub-finding algorithm from 90 to 96%. Finally, the joint segments are utilized in the endcap muon track finder (EMTF) to optimize the final track reconstruction and resolution of the L1 trigger, as described more fully in section 10.2.




Figure 84. Side view of the GEM-CSC trigger coincidence (left) and schematic overview of the data flow in the GEM and CSC trigger processor (right). The addition of the GE1/1 station significantly increases (by a factor of 3 to 5) the lever arm of the distance traveled in z by a muon originating from the interaction point. The bending angle between the entrance of the muon to the GE1/1 station and the exit from the ME1/1 CSC station can be used to estimate the momentum of the muon. Reproduced with permission from [149].

The data flow from the GEM chambers to the EMTF via two paths is shown in figure 84. First, trigger data are sent through the opto-hybrid board to the CSC optical trigger motherboard (OTMB) [150] that uses hits in the GE1/1 detector and the CSC ME1/1 to form stubs from the combined stations. These stubs are in turn sent to the CSC muon port card that sorts the stubs

from multiple chambers by sector and sends them to the appropriate EMTF card where the muon momentum is estimated from hits in the CSC, RPC, and GEM using a neural network algorithm [151], as described in section 10.2. In figure 85, simulated distributions of the muon rate in the endcap as a function of p_T and η are shown [100].




Figure 85. Muon rate from simulation as a function of p_T (left) and η (right) with and without the integration of the GEM chambers under various assumptions. Reproduced from [100]. CC BY 4.0.

6.4.6 Detector control system

The CMS detector control system (DCS) [73] is based on the Siemens Simatic WinCC Open Architecture (WinCC OA), previously known as Prozess-Visualisierungs- und Steuerungssystem (PVSS), and is described in ref. [152]. The DCS [153] was designed to control and monitor the high and low voltage and to monitor the gas conditions.

Even though the gas system is controlled by the central CERN gas group, the DCS maintains the function of monitoring the mixture composition (fraction of Ar and CO₂), as well as the pressure and flow rate to the chambers. In particular, the DCS reflects the structure of the gas system, which in CMS is divided into three main elements: mixer, rack, and flowcells.

The DCS system also controls both the high voltage used to operate the GEM foils and the low voltage used to power the frontend electronics on each chamber. The system allows turning on and off either the high or low voltage for every component, as well as monitoring the voltages and currents in the system.

6.4.7 Chamber assembly and installation

To ensure high-quality performance of all the components, a detailed quality control process was put into place to assure the quality of every element. The ten-step process, described in greater detail in ref. [154], can be summarized as follows:

- quality control of the components for production (QC1–2),

- assembly and commissioning of the GE1/1 chambers at the production sites (QC3–5),
- assembly and commissioning at CERN before installation in CMS (QC6–10).

First, all components were cleaned using ultrasonic baths, baked, sand-blasted, and visually inspected for faults. The components were then optically and electrically inspected including verification measurements, specifically measuring the I - V curves and checking for shorts between readout strips on the drift board. The foils were optically inspected, and checked that the leakage currents were less than 30 nA when 500 V was applied between the two sides of the GEM foils. Each chamber was pressurized by dry nitrogen at 30 mbar and checked for gas leakage. The chambers were then flushed for several hours with Ar+CO₂ and monitored for output at moderate HV. The uniformity of the gas gain was verified by placing an X-ray source 1 m away from a chamber enclosed in a copper protective box, which simultaneously illuminated the full chamber.

Once shipped to CERN, the gas leakage and HV tests were repeated for every chamber, and the frontend electronics and chambers were mounted into super-chambers with cooling plates. Electrical conductivity tests, gas leak tests, and electronic noise measurements were repeated with the cooling on. Finally, the super-chambers were mounted on a cosmic ray test stand where efficiency, noise, and tracking studies were done to confirm the correct and uniform operation of each chamber. While in storage and before installation into CMS, the gas-leak and HV stability were tested again and monitored over one month before each GEM was installed and declared ready for commissioning in the CMS detector.

6.4.8 Preliminary commissioning results

Before installation in the CMS cavern, a GE1/1 chamber was tested at the CERN H4 beam, which was extracted from the SPS [155]. A secondary beam of pions was produced when the proton beam struck a beryllium target. Finally, that secondary beam was filtered by collimators to produce a beam of 150 GeV muons from the decay of the pions in the secondary beam.

Several key performance parameters were measured in the test beam. Two are the single-hit efficiency and the time resolution of the GEM chambers. A set of scintillators read out by photomultiplier tubes (PMTs) were placed upstream of the chamber under test. Events were selected requiring a triple coincidence in each of the three layers and the efficiency was estimated by counting the number of hits in the GE1/1 tracking chamber in that region. Time resolution information was obtained by measuring the standard deviation of the distribution of the measured time between the trigger and GE1/1 detector signal. As can be seen in figure 86, an efficiency greater than 98% and a time resolution less than 10 ns were obtained. Other results, such as discharge rate and rate capacity as a function of voltages and gas mixture are given in ref. [155].

During LS2, the GE1/1 chambers were installed in CMS, connected to all services (LV/HV, gas, and optical fiber connection for readout), and the chambers were commissioned in situ. Latency scans were done for all GE1/1 chambers on both endcaps to account for the variable cable and optical fiber lengths, and HV values were optimized for signal-to-noise ratio using cosmic ray muons. It should be noted that, unlike in LHC collisions, cosmic ray muons do not arrive at a fixed and known time. Thus, the latency and other operational parameters cannot be fully optimized for LHC operations until all chambers are exposed to both muons originating from the interaction point and the LHC background conditions.




Figure 86. GEM test beam results showing the efficiency (left) and time resolution (right) as a function of the drift voltage placed on the GEM foils. The chosen gas mixture Ar/CO₂ (70/30) has very similar properties to Ar/CO₂/CF₄ (45/15/40) and was selected as CF₄, a greenhouse gas, is being phased out of industry. Reprinted from [155], Copyright (2020), with permission from Elsevier.

Figure 87 shows the occupancy in both GE1/1 endcap rings. Not all chambers were included in the data taking with cosmic rays due to minor temporary commissioning issues. This explains the missing sectors in the $-z$ endcap. Furthermore, since electronic noise was being addressed, some chambers show higher-than-average occupancy. The GE1/1 system was fully commissioned with cosmic rays and has been successfully operated since the beginning of LHC Run 3 in summer 2022.




Figure 87. GE1/1 occupancy in cosmic ray muon data recorded in 2021.

7 Precision proton spectrometer

The precision proton spectrometer (PPS) [156] is a set of near-beam detectors, located in the LHC tunnel at a distance of about 200 m from the CMS IP, on both sides. Originally conceived and developed as a joint project of the CMS and TOTEM collaborations (CT-PPS), it then evolved into a standard CMS subdetector following a Memorandum of Understanding between CERN, CMS, and TOTEM in 2018. For the sake of clarity, only the PPS acronym is used throughout this section. During Run 2, PPS took part in the CMS data taking, with various configurations between the years 2016 and 2018, and recorded data corresponding to an overall integrated luminosity of about 110 fb^{-1} . Following approval of the extension of the PPS experimental program to the LHC Run 3, a broad upgrade plan has been launched, involving replacement of all PPS detectors.

7.1 Forward protons and the roman pot system

The measurement of very forward (“leading”) protons [156] implies the detection of such protons at large distance from the IP and in proximity of the LHC beam. The main variable characterizing the proton kinematics is its fractional momentum loss, defined as $\xi = (|\vec{p}_i| - |\vec{p}_f|)/|\vec{p}_i|$, where $\vec{p}_i$ and $\vec{p}_f$ are the initial and final proton momentum vectors, respectively. The coverage in ξ at the PPS location is limited by the LHC magnet and collimator lattice in Run 2 and Run 3 to a range below 0.2. In order to extend it as much as possible in the low- ξ region, scattered protons must be detected as close as few (2–3) mm from the beam. This can be accomplished by means of special movable beam pockets, the so-called roman pots (RPs), which host the particle detectors such that they can be moved close to the beam during stable beam operations. The PPS experimental setup includes both tracking detectors, to measure ξ from reconstructed proton tracks, and timing detectors, to suppress the contribution of protons originating from different primary interactions in the same LHC bunch crossing (pileup). Tracking and timing detectors are hosted in separate RPs.




Figure 88. Schematic layout of the beam line between the interaction point and the RP locations in LHC sector 56, corresponding to the negative z direction in the CMS coordinate system. The accelerator magnets are indicated in gray; the collimator system elements in green. The RP units marked in red are those used by PPS during Run 2; the dark gray ones are part of the TOTEM experiment. In Run 3 the “220 near” horizontal unit is used in place of the “210 near”. Reproduced from [12]. The Author(s). CC BY 4.0.

The TOTEM experiment [157] has been using an extensive system of box-shaped roman pots to conduct its physics program since the start of LHC operations in dedicated, low-luminosity runs. The system has been conceived as a two-arm spectrometer, composed of two pairs of stations on each side of the CMS IP, at about 147 and 220 m along the beam line, each station comprising two RPs approaching the beam vertically, and one approaching it horizontally. Within each pair, the station installed closer to the IP is called “near”, and the other “far”. For Run 2, the layout of the spectrometer has been modified as sketched in figure 88 for one of the two arms. The stations at 147 m have been relocated in proximity of those at 220 m and identified as “210 m”. Following these conventions, the four stations are usually named 210-N, 210-F, 220-N, 220-F, where N stands for near, and F for far. In order to allow operations at standard LHC luminosity with PPS, some of the horizontal RPs have been equipped with cylindrical ferrite shields to reduce their radio-frequency impedance [156] and have been reserved for the PPS tracking detectors (section 7.2). Moreover, an additional horizontal RP with modified geometry (cylindrical RP) has been installed in each arm between the near and far stations at 220 m to host timing detectors; it is referred to as 220-C (section 7.3). All vertical RPs remain reserved for the TOTEM physics program. However, their role is crucial for PPS operations at the beginning of each data-taking period, when the tracking detectors they host are used in special runs, in conjunction with those in the horizontal RPs, to determine the global alignment parameters [12].

Figure 89 shows a view of the LHC tunnel, in the region where the RPs are installed in one of the two sectors. The detector housings are made of 2 mm thick stainless steel, with a useful depth to contain the detectors of about 110 mm; the box-shaped pots have a rectangular cross section of $124 \times 50 \text{ mm}^2$, while the cylindrical pots have a diameter of 141 mm. In order for detectors to approach the LHC beam as closely as possible, and to minimize the amount of material traversed by scattered protons, the thickness of the walls is reduced to 200 (300) μm around the active sensor region in the box-shaped (cylindrical) RPs: this section is called “thin window”. Figure 90 shows the structure of a box-shaped and a cylindrical RP, as well as the insertion system installed on a section of the LHC beam pipe.




Figure 89. View of a section of the LHC tunnel in sector 45, with part of the PPS and TOTEM RP stations along the beam line.




Figure 90. Sketches of the horizontal roman pots. Left: box-shaped pot, with the ferrite RF-shield in place. Center: cylindrical pot where the ferrite RF-shield is integrated as a ring in the flange. The thin window shapes, different for the two cases, are visible in the rightmost, vertically centered part of the pots. Right: overview of the insertion system of a box-shaped pot in a section of the LHC beam pipe. Reproduced with permission from [158].

Moderate vacuum conditions have to be kept inside the detector housings, in order to minimize the pressure gradient between the detector volume and the ultra-high vacuum of the LHC beam pipe. The exact value of the pressure depends on the detector technology, but does not exceed 100 mbar. Several RPs are serviced by the same vacuum system.

The solid-state detectors equipping PPS need to operate at low temperatures (-20 to 5°C , depending on the detector) to mitigate the effects of radiation damage. This is accomplished by means of a fluorocarbon evaporative system [159]: the cooling fluid is transported to the RPs, where suitable evaporation circuits, thermally coupled to the detectors, supply the needed cooling power [156, 157].

A more detailed description of the TOTEM and PPS roman pot system and its operations can be found in refs. [156, 157] and references therein.

7.2 Tracking detectors

The energy loss of protons can be measured, through detailed knowledge of the LHC optics parameters and dedicated calibrations [12], from the parameters of their reconstructed tracks. In order to obtain the needed accuracy in ξ , a resolution of few tens of μm is required on the coordinates of the track impact point. The hit rate distribution is highly nonuniform in the region covered by the tracking detectors, with peak particle flux values of about $3 \times 10^9 \text{ p}/(\text{cm}^2 \text{ s})$ and a range spanning over three orders of magnitude. This reflects in the high radiation dose the detectors have to withstand, with reference fluence values in the most irradiated areas of $1\text{--}3 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ for an integrated luminosity of 100 fb^{-1} .

7.2.1 Detector units

Proton tracks are reconstructed in PPS by two tracking stations per arm. During Run 2, the roman pots hosting the trackers and the technology chosen for the detectors have changed over time.

In 2016, in order to advance data collection by one year with respect to the original plan, two of the TOTEM tracking stations, equipped with silicon strip detectors [160, 161], were installed in the horizontal RPs at the 210-N and 210-F locations. The TOTEM strip detectors were originally designed to operate in special LHC runs at very low luminosity, up to about $10^{30} \text{ cm}^{-2} \text{ s}^{-1}$. Within PPS, they have been operated successfully well beyond design specifications, providing excellent position measurements. However, they suffered from severe limitations related to the high hit rate. Firstly, since the hit position in the detector planes is determined by the coincidence of pairs of perpendicular strips, the frequent presence of multiple proton tracks in the same bunch crossing leads to ambiguities that cannot be resolved. Secondly, the radiation dose in the regions with the highest hit rate induced damages that caused an early loss of efficiency: for those regions, the efficiency dropped to zero after $O(10 \text{ fb}^{-1})$ of integrated luminosity.

In 2017, new silicon pixel detectors, based on the 3D technology [162], were installed in the RPs at the 220-F locations, replacing the strip trackers at 210-N. These detectors were specifically developed for PPS, with much improved radiation tolerance and rate capability with respect to the strip detectors, allowing for the reconstruction of multiple tracks per bunch crossing. Finally, in 2018, both stations at 210-F and 220-F were equipped with pixel detectors.

Despite the improved resistance to radiation, the dose accumulated by the pixel sensors and readout chips during Run 2 has been such to degrade significantly the detector performance, as briefly described in the following. This, and the lack of enough replacement parts, called for the construction of new detector modules; with the aim of mitigating the adverse effects of radiation damage, a new design for the support mechanics has been developed for Run 3, which in turn implied a redesign of the frontend electronics. Here, the main differences of the Run 3 design from that used in Run 2 are outlined. A more complete description of the Run 2 setup can be found in ref. [163].

Each pixel tracking station consists of six detector planes, oriented at an angle of 70° with respect to the thin window, i.e., 20° with respect to the plane perpendicular to the beam axis (it was 18.4° in Run 2). This orientation removes the geometrical inefficiency associated to the junction and ohmic columns of the sensor, and increases the probability of charge sharing between nearby pixels for particle signals. The sensor size is driven by that of the readout chip (ROC): for Run 3 it is $16.20 \times 16.65 \text{ mm}^2$, corresponding to a 2×2 ROC matrix; in Run 2 most sensors were longer, $20.40 \times 16.65 \text{ mm}^2$, and read out by a 3×2 ROC matrix. The pixel size is $150 \times 100 \mu\text{m}^2$, also adapted to that of ROC cells; the pixels located at the horizontal or vertical edge of a ROC have double size in the corresponding coordinate, except for those at the sensor edge. The details of the sensor module geometry and the arrangement in the tracking station are shown in figure 91.




Figure 91. Geometry and arrangement of the pixel sensor modules. Upper left: arrangement of the sensor modules in a tracking station, relative to the LHC beam. Upper right: sensor, ROCs, and pixel geometry. The shaded areas with dashed contours refer to the larger, 3×2 modules used in Run 2. Lower: detail of pixels at internal ROC edges.

The 3D sensors have been produced by the Foundation Bruno Kessler (FBK) in Trento, Italy [164]. They are realized with “single-side” technology, where both the junction and ohmic columns are etched from the same side of the silicon wafer; the thickness of the active volume is $150 \mu\text{m}$. A so-called 2E configuration is used, where each pixel has two readout (junction) columns, thus improving the radiation resistance characteristics. In Run 2, the sensors, produced by Centro Nacional de Microelectrónica (CNM) in Barcelona, Spain [165], were realized with “double-side” technology, with the two types of columns etched on opposite sides of the wafer, with $230 \mu\text{m}$ active thickness, in both 1E and 2E configurations. Figure 92 shows schematically a cross section of the column layout for the two productions. Due to the 3D configuration, bulk depletion occurs at inverse polarization

values of a few Volts. Currents below $1\ \mu\text{A}$ are required at the working point, but are generally much lower (a looser requirement was set on Run 2 sensors); the breakdown tension must exceed 60 V.




Figure 92. Schematic layout of the columns of 3D pixels. Left: double-sided technology, used in Run 2. Right: single-sided technology, used in Run 3. Metal bumps on top are used for interconnection with the readout chips; p-type surface implants, such as p-stop or p-spray, provide isolation between contiguous pixels. Reproduced with permission from [166].

7.2.2 Readout electronics

The readout of pixel detectors is based on the frontend electronics developed for the Phase 1 upgrade of the CMS pixel tracker [15]. A schematic view of the readout chain is shown in figure 93.

Signals collected from the sensor electrodes are processed by the PROC600 chips developed for the layer 1 of the barrel pixel detector, described in section 3.1.4: four chips are connected to each sensor via bump bonds. In Run 2, the PSI46dig chips [167] were used. A flexible printed circuit board is glued on top of the sensor; the ROCs are connected to the board via wire bonds. A TBM10d chip (token bit manager), mounted on the board, distributes clock, fast commands and configuration instructions to the ROCs, and collects data through a token ring architecture; only one of the two TBM cores is used, resulting in a single output data stream. In Run 2, the flexible board was mounted around the sensor module, and carried one TBM08c chip. The two versions of the flexible board are shown in figure 94.

The six detector modules of a tracking station are connected, through connectors on the flexible boards, to a concentrator card, called PortCard (figure 95). The connections are made on a section of the PortCard that enters the secondary vacuum of the RP through a vacuum-sealed feed-through. The rest of the board, operating at atmospheric pressure, contains all the active electronics needed to perform the board tasks:

- manage the configuration of the integrated circuits equipping the board itself via I²C commands sent on the slow-control ring;
- send the clock signal, fast commands, and configuration instructions to the frontend readout (via the TBM) through an optical link connection;
- collect and send, through individual optical links, the data received from the detector modules;




Figure 93. Schematic diagram of the PPS pixel tracker readout chain. The various components shown are described in ref. [15]. Reproduced from [15]. CC BY 3.0.




Figure 94. The detector module of the silicon pixel tracker, in its Run 2 (upper) and Run 3 (lower) versions.




Figure 95. The pixel PortCard for Run 3; the small flange near the bottom acts as a feed-through for the board. The section below is housed inside the secondary vacuum volume of the pot.

- route lines needed to read out environmental parameters in the detector housing and control the internal movement system (described later).

The PortCard can also control, through a parallel port register, the selective powering of the detector planes.

7.2.3 Support structure and internal motion system

The sensor modules and the flexible boards are glued, using a flowable silicon-based sealant, to aluminum support plates, ensuring thermal connection to the cooling circuit on the sides of the detector package (figure 95). Precision holes in the plates are used to control the position of the sensors with respect to the thin window of the detector housing. In Run 2, larger support plates were employed, realized in thermal pyrolytic graphite (TPG) enclosed in thin aluminum sheets.

An aluminum support structure carries the six detector modules, ensuring their precise positioning and coupling them to the evaporators of the cooling system. Six slits on the sides of the support structure host the aluminum plates of the modules, and pins on both sides keep them in place. Three adjustable transfer-ball feet provide contact with the floor of the detector housing; on the opposite side, two spring-loaded frames allow the correct sealing of the upper flange of the RP and the transverse movement of the detector package, as described in the following. In Run 2, planes were arranged in pairs in separate sections of the support structure, eventually assembled together and to the rest of the package mechanics. Feet were static and no internal movement was allowed. Figure 96 shows the assembled detector packages, as well as the connections to the readout and services outside the secondary vacuum volume.

Both the sensors and the readout chips used in Run 2 were designed to withstand fluence values up to $2\text{--}3 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ [162, 168]. However, the highly nonuniform distribution of the hit rate (figure 97) implied localized radiation damage in the readout chip structures that could not be compensated by changing global configuration registers of the chip itself. In particular, the position of the time window to accept hits (WBC register) could not be optimized for all pixels simultaneously: after a data-taking period corresponding to about 20 fb^{-1} of integrated luminosity, the regions with the highest rate could not be read out.




Figure 96. Details of the pixel detector package. Left: Run 2 version. Center: Run 3 version. Right: mechanical, cooling, and electrical connections to the upper part of the RP in Run 3: the vacuum section of the PortCard is visible, as well as the cooling circuit and the connections to the environmental sensors.




Figure 97. Proton fluence in two planes (one per arm) of the PPS pixel tracker in a sample run in 2017, normalized to an integrated luminosity of 1 fb^{-1} . The horizontal and vertical lines in the distributions are an artifact due to the different pixel size of the sensors in those regions.

With the aim of extending the lifetime of the detectors, in Run 2 the support of the corresponding RP was raised with respect to the floor of the LHC tunnel, in two steps of 0.5 mm, during technical stops. Because the inefficient area was so localized, these small shifts were sufficient to move it to a region with much smaller hit rate, and thus almost restore the overall initial efficiency. However, the range for this kind of movements cannot be further extended; moreover, the operation is onerous and not without risk.

In order to perform the same kind of movement in a more practical and safe fashion, an internal (vertical) movement system for the pixel detector package has been developed. The very design of the new detector modules has been driven by this requirement. Their width has been minimized (by

placing the flexible board on top of the sensor) so as to exploit all the available space in front of the RP thin window. The width of the thin window is 29 mm, while that of the sensor module is just below 22 mm. A movement range of 5 mm for the detectors has been foreseen, including some safety margin. This will allow to shift vertically the sensor modules, during the data-taking period, in 11, 500 μm -wide steps, thus distributing the radiation damage.

The task is accomplished by means of a miniaturized stepping-motor linear actuator (Zaber LAC10A-T4A), mounted inside the detector package (figure 98). The actuator has an excursion of 10 mm and a step size of about 23.8 nm. Its body is fastened to the lower frame of the spring-loaded structure; the head of the actuator is screwed into one of the side support parts. These two parts can slide on each other thanks to low-friction elements realized in polyether ether ketone (PEEK). Their relative position is measured by a simple linear motion potentiometer (Bourns 3048L-5-103), mounted in a similar fashion as the actuator, whose resistance can be read out externally. Extensive tests have been performed on the fully assembled system, in particular performing repeated movements inside an exact reproduction of the RP and in conditions similar to the working ones ($T \approx -20^\circ\text{C}$, $P \approx 15\text{ mbar}$). Some results are shown in figure 98. The tests have demonstrated that the measured variation in resistance is consistent with the expected shift, showing the adequacy of the position measurement, where only a moderate precision is required. Some hysteresis effect is observed, most likely due to the mechanics of the potentiometer, when the shift direction is inverted; however, this effect is reproducible over several movement cycles.




Figure 98. The system for internal motion of the pixel detector package. Left: detail of a partially assembled detector package: the stepper motor is the black object on top; the blue object below is the potentiometer used to monitor the position; both of them have their body tied up to the sliding slit on top, while their mobile tip is screwed to the support structure for the modules. Right: results of a motion test inside a RP at standard working conditions (-20°C and about 5 mbar), with measured versus nominal position. Two sets of points can be identified for forward and backward movements, revealing a hysteresis effect.

The linear actuators are controlled by dedicated two-channel controllers (Zaber X-MCC2). One controller per arm is installed inside the RR53 and RR57 shielded areas (“alcoves”) along the tunnel. They are connected to Raspberry PI microcomputers via USB link, which also reads out the detector position through an external ADC board; the microcomputers can be accessed from outside via the 4G network in the tunnel. This kind of setup is adequate in view of the noncontinuous operation of the motors. In fact, only a limited number of detector shifts is performed over the whole data-taking time, in interfill periods.

7.3 Timing detectors

In the presence of multiple proton-proton interactions in the same bunch crossing (pileup), the separation of overlapping events in CMS relies on the reconstruction of multiple primary interaction vertices. However, because the proton tracks reconstructed by PPS have scattering angles very close to zero, the tracking system described in the previous section cannot associate them to CMS primary vertices. For events with two protons detected on opposite sides, this limitation can be overcome if a precise measurement of the arrival time of protons is available: from the difference in time $\Delta t = t_+ - t_-$, where t_+ and t_- are the time measurements in the positive and negative arm of PPS, respectively, the position in z of the pp vertex can be inferred from $z_{pp} = c\Delta t/2$. Dedicated studies [156] have shown that a time resolution of $O(10\text{ ps})$ is needed to achieve the correct association of forward protons to the events reconstructed centrally by CMS. For pileup conditions corresponding to $\mu = 50$, a time measurement with 10 (30) ps resolution could reject the combinatorial background from random combination of uncorrelated protons by about a factor 30 (20) while keeping about 60 (50)% of the signal. However, larger time resolution values, within 100 ps, can still help in reducing the background. As in the case of the tracking system, timing detectors must be able to operate with large and highly nonuniform particle fluxes, to tolerate high radiation doses and to work in a vacuum. Moreover, the material thickness must be limited.

During Run 2, one roman pot on each arm was used for timing, with four planes per RP. For Run 3, instrumenting an additional station at the 220-N location will allow a second timing RP on each arm, giving a total of eight planes per arm.

7.3.1 Detector modules

The PPS timing detectors [169–171] are based on synthetic single crystal chemical vapor deposit (scCVD) diamonds. The detectors combine good time resolution, extreme hardness against large and nonuniform radiation, low material budget, and fine segmentation near the beam.




Figure 99. Left: details of the four-pad and two-pad segmentation of the diamond sensors used in the Run 3 modules. Right: arrangement of the four crystals in a Run 3 module, where the position of the beam is indicated by the spot on the left. Reproduced with permission from [172].

At the conclusion of Run 2, half of the planes were based on an improved double-diamond architecture, in which two diamonds are connected to the same amplification channel. This nearly doubles the signal, while keeping a similar noise level, resulting in a significant improvement of the timing resolution compared to a single-diamond design. In Run 3, all timing planes are of the double-diamond design. The crystals have dimensions of $4.5 \times 4.5\text{ mm}^2$, with a total active surface area of approximately $20 \times 4.5\text{ mm}^2$ per plane. The final segmentation is achieved

during the metalization process, resulting in a total of 12 channels per plane. Due to the highly nonuniform flux of particles, the segmentation varies with the distance from the beam in the x direction. Several configurations of pad dimensions and sensor layout have been used in Run 2; in Run 3, sensors with two-pad and four-pad segmentation are used, with dimensions as detailed in figure 99, where the layout of four of these sensors in a detector module is also shown. Pads close to the beam position have smaller size (as small as 0.55 mm), while a coarser segmentation is used farther from the beam. This results in a more uniform occupancy, with low inefficiency due to multiple hits in the same channel. Because of the horizontal crossing of the LHC beams at the IP, hits are more widely distributed along the x axis. For this reason, the detectors are only segmented horizontally.

The diamonds for Run 3 include newly produced crystals, and crystals that were previously used in Run 2, cleaned and remetalized. Samples of the latter were first studied in beam tests after dismounting, and found to maintain high efficiency and achieve single plane time resolutions of $\sim 80\text{--}95$ ps, after being exposed to fluences as large as 5×10^{15} p/cm² [173]. Compared to the nominal resolution of about 50 ps per plane for a new double-diamond detector, the decrease in resolution was found to be largely consistent with radiation damage to the preamplification electronics.




Figure 100. Left: the hybrid board for the Run 3 timing detectors' readout; the lower, wider section hosts the sensors and is housed inside the secondary vacuum volume of the pot. Right: detail of the diamond sensors on one side, connected via wire bonds to the frontend electronics.

The diamonds are glued to a hybrid board (figure 100), containing both the sensors and a multistage amplification chain for each of the 12 channels. As mentioned above, radiation damage to the amplification chain was identified as a limiting factor for the timing performance in Run 2. Therefore, a revision of the hybrid board was designed for Run 3, with a modified layout to mitigate the exposure of the preamplifiers to radiation. Additional modifications to the design were made to improve the high-voltage isolation and the stability against RF noise pickup. Finally, remote control of the amplifier gains was implemented, giving the opportunity to better fine tune the settings and compensate for any degradation during data taking.

7.3.2 Readout electronics

The signals from the hybrid boards are transmitted by individual coaxial cables to custom “NINO boards”, mounted in a mini-crate about 1 m above the LHC beam pipe (figure 101). The main data path for reading all channels at the full trigger rate is based on the fast, low-power NINO ASIC [174], with four chips per board. The NINO performs discrimination of the input signals above an adjustable threshold. The width of the output signal is proportional to the input charge, and is then stretched by a constant value for compatibility with the next piece in the readout chain. This serves as a proxy for the signal amplitude, allowing for time-walk corrections to be derived from the data. In Run 3, a revision of the NINO board also allows for signals from a subset of channels to be split off to record the full waveform information, at reduced rate.




Figure 101. Left: the digital readout unit. In this example both an HPTDC mezzanine (upper left) and a SAMPIC mezzanine (upper right) are mounted on the motherboard. Right: location of PPS mini-crates installed above the LHC beam pipe, for both the cylindrical timing RP already used in Run 2, and the rectangular 220-N RP newly equipped for timing in Run 3. The readout boxes (labeled “RB”) contain the NINO boards and DRUs as described in the text. The box labeled “CB” holds components for the precise reference clock.

From the NINO boards, the signals are transmitted to the digital readout units (DRUs), mounted in the same mini-crate above the beam pipe. Each DRU consists of a digitizer motherboard supporting two types of mezzanine for timing measurements. The first, used in Run 2 and for all channels in Run 3, includes four high-performance time-to-digital converter (HPTDC) chips [175]. The HPDTC for PPS is used in very high resolution mode, where an on-chip four-point sampling and interpolation of a delay-locked loop is performed to improve the time resolution. This results in a TDC with a binning of about 25 ps, at the cost of limiting the number of output channels per chip to eight. The HPTDC records the time of both the leading and trailing edges of the signal, allowing for time-walk corrections to be performed as a function of the reconstructed input charge value. In Run 3, additional DRUs are also equipped with a second type of mezzanine, based on the SAMPIC [176] fast waveform sampler. This allows the readout of the full analog signal shape, for calibration purposes, on a fraction of the data, for trigger rates up to 100 kHz. Each digitizer motherboard includes a radiation-hard MicroSemi SmartFusion2 FPGA, which formats the data from the HPTDC (or SAMPIC) mezzanine into the final data frames. The data from each board are then transmitted over optical fibers to the backend using two GOH opto hybrid mezzanines.

In addition to its role in the data acquisition, the DRU is also responsible for receiving and distributing slow-control commands, clock, and fast commands. It receives both the standard clock

and the dedicated precise reference clock [169] that is used for the timing measurement. Along with the configuring of components on the DRU itself, it can propagate commands via I²C interface to set the thresholds on the NINO board, and, for Run 3, the low voltage settings on the diamond hybrid. The ability to remotely set the low voltage will allow for better fine tuning of the operational settings, and provide the ability to compensate for possible radiation damage effects.

A schematic view of the timing control and readout chain components installed in the LHC tunnel is shown in figure 102.




Figure 102. Schematic diagram of the full readout chain for the PPS timing detectors in Run 3, also showing the configuration used in Run 2. The arrows on the right represent electrical connections to the slow-control ring and optical connections to components in the underground service cavern of CMS, described in section 7.4. The central set of DRUs (3 units) was already present in Run 2, but was employed for the diamond detectors equipping the TOTEM experiment at that time, and for one layer of UFSD detectors used in 2017 for the PPS timing.

7.3.3 Reference clock distribution

For precise timing measurements, a clock distribution system that provides time information at points separated by large distances, with a picosecond precision, is needed. This requires a system capable of the highest precision and of the utmost time stability.

PPS employs the clock distribution system developed by the TOTEM Collaboration [177], an adaptation of the universal picosecond timing system [178, 179], developed for FAIR (Facility for Antiproton and Ion Research) at GSI, Darmstadt. The system has been installed and used during the TOTEM and PPS runs since 2017.

The system can be logically subdivided into four major blocks: the transmission unit, the distribution unit, the measurement unit, and the receiving unit. Receiving units are installed in the LHC tunnel, as close as possible to each timing detector, while the transmission, distribution, and measurement units are located in the IP5 counting room. A block diagram of the entire system is shown in figure 103. The system leverages the use of a dense wavelength division multiplex (DWDM) technique, exploiting the transmission of multiple signals of different wavelengths over a common single-mode fiber. This allows the use of standard telecommunication modules compliant with ITU (International Telecommunications Union) standards.

The transmission unit optically modulates two reference clock signals using two different DWDM wavelength carriers: λ_1 and λ_2 . These optical signals are multiplexed into a single fiber and transmitted to the distribution unit, where they are split to be distributed to all receiving units located in the tunnel on both sides. The multiplexed signal is then optically amplified via an erbium-doped fiber amplifier (EDFA) to compensate for the attenuation due to the multiplexing and splitting steps.

These signals are further multiplexed with a third one, of wavelength λ_M , with the goal of measuring the transmission delay over each fiber. This is done in the measurement unit, where a network analyzer drives the optical modulation of this third DWDM signal, which is then sent to and reflected back by the receiving unit. In this way, the signal delays can be determined and possible drifts can be monitored.

In the tunnel, receiving units separate and convert the multiplexed optical signals generated in the transmission unit back to electrical signals, and reflect, via a fiber Bragg grating reflector, the one generated in the measurement unit. The electric signals are then routed to frontend and readout electronics.

Measurements performed with the installed system have shown an additional contribution to jitter of slightly less than 1 ps, mainly due to the inherent jitter of the clock source signal, the noise added by the optical components, and the bandwidth of the transmission system itself.

7.4 Data acquisition and detector control

The backend data acquisition systems for the various detectors used in PPS have been developed independently and are based on different hardware. The strip trackers have maintained the original structure developed for TOTEM, and are integrated into the CMS data acquisition (DAQ) system, described in section 9. A similar architecture is used for the timing detectors. The pixel tracker system is based on the DAQ scheme employed by the Phase 1 upgrade of the CMS pixel detector, described in section 3.1.4.

7.4.1 Pixel detectors

The pixel DAQ is based on μ TCA FC7 carrier boards, following the design illustrated in figure 93, and described in more detail in ref. [21]. The boards are employed as either FEDs or FECs, depending on the type of attached mezzanine card and the firmware deployed. Two FC7 boards are equipped as FEDs, and receive data transmitted by POHs over optical fibers from the pixel detectors on the two arms of the PPS spectrometer. The typical data volume in Run 2 was about 0.5 kB/event or less, depending on the instantaneous luminosity and other LHC conditions, and is similar in Run 3. Additional FC7s are equipped with FEC mezzanines, transmitting signals over fibers to the frontend boards in the LHC tunnel. Depending on the firmware, the FECs are used as either PxFECs, sending




Figure 103. A block scheme of the clock distribution system for the PPS timing detectors. The four receiver units correspond each to a timing station in the tunnel; the remaining elements of the system are all located in the counting room. Different colors of the fiber lines represent the different wavelength carriers used by the system, λ_1 , λ_2 , λ_M (black, green, and yellow, respectively), and the multiplexed signals (red).

the clock and trigger, or TkFECs, responsible for sending slow-control commands. The entire system, shown in figure 104 (left), is housed in a single μ TCA crate, containing one AMC13 for clock and trigger distribution, and one MCH crate controller.




Figure 104. Left: the PPS pixel DAQ crate. On the left are two FC7 FECs, sending clock, trigger and slow-control commands; in the center-right are two FC7 FEDs, receiving data from the two arms of the PPS spectrometer; on the far left is the MCH crate controller, on the far right is the AMC13. Right: the timing and strips DAQ crate. The SLinks are placed in the backplane of the OptoRx boards delivering data to upstream CMS DAQ.

7.4.2 Strip and timing detectors

The DAQ for the PPS timing detectors, and for the strip detectors used in special alignment fills, was developed by the TOTEM Collaboration, based on a scalable readout system (SRS) [180]. Optical receiver mezzanines (OptoRx) are connected to a custom designed board, which further connects to a frontend concentrator board. The OptoRx receives data transmitted over fibers from GOHs on the timing digital readout units. A scalable readout unit (SRU) is responsible for receiving and distributing the clock, trigger, and fast commands within the SRS crate (figure 104 right). The system was designed for full compatibility with the CMS DAQ. Data are transmitted via SLink to the FRLs of the central CMS DAQ, where they are treated in the same way as those from other FEDs. During Run 2, a total of two VME-based FEDs were used for the timing detectors in normal data taking, while four additional such FEDs were used to read out the strip detectors in the vertical RPs for special alignment runs. In Run 3, a total of six timing FEDs are used in standard runs, to support the addition of a second timing detector station, and the readout of a limited amount of data with a SAMPIC waveform digitizer.

The SRS firmware was originally designed to read data from the VFAT chip [145] used by the TOTEM strip detectors. In order to use this system with minimal changes, the timing data from the HPTDCs are packed into a VFAT-like frame, with no zero-suppression. In Run 2, the data volume from the full timing system was around 2 kB/event. In Run 3, this will more than double with the addition of a second timing station and the SAMPIC readout option. The system is capable of sustaining the CMS trigger rate of more than 100 kHz.

The distribution of (non-reference) clock, trigger, and slow-control commands for the timing electronics is handled by a VME FEC, of the same type and configuration as used in the electromagnetic calorimeter preshower (section 4). The commands are transmitted over fibers to the frontend boards of the timing system. One control loop is dedicated to the strip detectors, while another

is dedicated to the timing detectors. In addition, firmware for the FPGA of the timing digitizer motherboard can be deployed over the slow-control loop, using a new software development that takes advantage of the JTAG master of the CCU25 communication and control unit [29]. This allows firmware updates to be made remotely, without requiring access to the LHC tunnel.

7.4.3 Detector control system

The PPS detector control system (DCS) is a legacy from the TOTEM experiment built with the industrial WinCC OA software (PVSS) used to control the low and high voltage of detector packages, to monitor the radiation dose, pressure, and temperature sensors, as well as the roman pots movements and the status of interlocks of the detector safety system (DSS). The TOTEM DCS system includes the automatic matrix actions to manage the power-supply state of the roman pots according to the LHC beam status. The system for Run 3 is integrated within the CMS DCS framework, and runs with WinCC OA version 3.16.

7.5 Roman pot insertion and running scenarios

The position of the roman pots with respect to the LHC beam is controlled by the LHC operators through standardized procedures. During the injection, acceleration and luminosity tuning stages, the RPs are kept in a retracted (“garage”) position, at about 40 mm from the proton beam. When “Stable Beams” are declared, the RPs are moved closer to the protons, at a distance depending on the beam width at that location, $\sigma_{x, \text{XRP}}$. This distance, d_{XRP} , was given in Run 2 by the following rule from machine protection:

$$d_{\text{XRP}} = \max[(n_{\text{TCT}} + 3)\sigma_{x, \text{XRP}} + 0.3 \text{ mm}, 1.5 \text{ mm}], \quad (7.1)$$

where n_{TCT} is the distance of the tertiary collimator (TCT) from the beam center in units of the beam width $\sigma_{x, \text{TCT}}$ at its location. The 3σ retraction ensures that the RPs stay in the shadow of the TCT, while the additional 0.3 mm margin protects against accidental beam orbit deviations. From arguments of mechanical rigidity of the RP thin window, which could bulge towards the beam in case of a secondary vacuum loss, an absolute lower limit of 1.5 mm was imposed. In Run 2, with $n_{\text{TCT}} = 8.5$, the resulting d_{XRP} ranged between 2.2 mm (210-F) and 1.5 mm (220-F).

Operations in Run 3 are characterized by a far more complex luminosity-leveling scheme with concurrent changes in the crossing angle, called α in this section, and the beta function at the IP, β^* . The most recent concept of the leveling scheme can be represented by the trajectories shown in figure 105 for the years 2022 and 2023.

Both α and β^* , as well as the collimation scheme, have a decisive impact on operation and performance of the RP spectrometer.

7.5.1 TCT collimator and roman pot insertion scheme

The RP approach distance d_{XRP} , as given in eq. (7.1), depends on β^* through the beam width, $\sigma_{x, \text{XRP}}$, and through the TCT distance, $n_{\text{TCT}} = d_{\text{TCT}}/\sigma_{x, \text{TCT}}$. The function $d_{\text{TCT}}(\beta^*)$ characterizes the TCT collimation scheme.

In the old standard collimation scheme, used until the end of 2022, the TCT did not move during stable beams, i.e., $d_{\text{TCT}}(\beta^*) = \text{const}$, so the nominal distance d_{XRP} was entirely determined by the




Figure 105. Luminosity-leveling trajectories for typical LHC fills in the $(\alpha/2, \beta^*)$ plane for 2022 (dashed line) and 2023 (continuous line).

evolution of the beam widths, $\sigma_{x,\text{XRP}}$ and $\sigma_{x,\text{TCT}}$, with β^* . With β^* -leveling, d_{XRP} changes during the fill. However, an automated RP movement synchronized with the β^* evolution during the fill is not foreseen in the control system and would be difficult to implement, in particular because also the position limits in the interlock system would have to follow this movement. In Run 2, this was not a problem, since the range of β^* (25–30 cm in 2018) was so small that the variation of d_{XRP} was negligible. The RPs were kept fixed at the maximum distances along their nominal trajectories. In Run 3, the wider β^* range complicates the situation. The evolution of the RP distance with β^* , for the old collimation scheme with fixed TCT positions, is shown in figure 106, for two example RPs.

In 2022, where the β^* leveling covered the range from 0.6 to 0.3 m, the RP with the widest nominal distance range was 210-F. Ideally, it would have had to move by 0.22 mm. The RP 220-F, on the other hand, could stay fixed at the limit of 1.5 mm throughout the fill. The situation of the two other pots, 220-N and 220-C, was between these extremes, with ideal movement ranges smaller than 0.2 mm.

In 2023, the much wider β^* range (1.2 to 0.3 m) would have led to movement amplitudes up to 1 mm. To judge whether it is acceptable to keep the RPs fixed at the most distant points along their trajectories like in Run 2, the impact of the RP distance on the mass acceptance has to be considered. The minimum accepted mass M_{min} of centrally produced states X from the process $pp \rightarrow pXp$ with double proton detection is given by:

$$M_{\text{min}} = \sqrt{s} \frac{d_{\text{XRP}} + \delta}{D_{x,\text{XRP}}}, \quad (7.2)$$

where $\delta \approx 0.5$ mm is the insensitive margin from the outer RP window to the point of full efficiency in the detector, and $\sqrt{s} = 13.6$ TeV. The horizontal dispersion $D_{x,\text{XRP}}$ of the LHC from IP5 to the RP depends linearly on the crossing angle, according to:

$$D_{x,\text{XRP}}(\alpha) = D_{x,\text{XRP}}(0) - D'_x \alpha, \quad (7.3)$$




Figure 106. RP insertion distance d_{XRP} for the leveling trajectories of 2022 (dashed lines) and 2023 (continuous lines) from figure 105, evaluated for the pots 210-F (left) and 220-F (right) in the case of the TCT fixed to $d_{\text{TCT}} = 8.5\sigma_{x,\text{TCT}}(\beta^* = 30 \text{ cm})$. The magenta line shows the nominal distance according to eq. (7.1). The blue lines show the most conservative constant RP distance, i.e., never closer to the beam than the nominal distance. If this blue distance is smaller than the limit of 1.5 mm (red line), which is the case for 220-F, the pot has to stay at that limit. The fill evolves from the right to the left.

with a constant $D'_x > 0$, implying that smaller crossing angles yield a larger dispersion and hence a better low-mass acceptance.

Figure 107 shows the evolution of M_{min} during β^* leveling in 2022 and 2023 for the old collimation scheme with fixed TCTs. The differences in M_{min} between 2022 (dashed) and 2023 (continuous) for a given β^* are caused by the concurrent crossing-angle leveling in 2023. The magenta lines represent the case of the RPs moving along their nominal d_{XRP} . The blue lines show the effect of keeping the RPs fixed at the most distant position along their trajectories. If a blue line crosses the limiting red line (1.5 mm distance), the RP has to be fixed at 1.5 mm, and its M_{min} follows the red line.

For 2022, the loss in M_{min} at the end of the fill, at a β^* of 30 cm, by not following the nominal trajectory, at most 35 GeV for 210-F, was considered to be still acceptable. Therefore, the run in 2022 proceeded along the old scenario. Despite the slight loss from omitting the movement, the acceptance for a given $(\alpha/2, \beta^*)$ point was better than in Run 2.

For 2023, on the other hand, the M_{min} acceptance loss would have amounted to 150 GeV. To mitigate this performance deterioration, while maintaining fixed RP positions throughout the fill, the collimation working group has elaborated an alternative scheme where the TCTs, having a more sophisticated movement control system, adapt their positions to the β^* evolution. In this scheme, operationally implemented in 2023, the constant RP positions are all close to the 1.5 mm limit and thus the M_{min} acceptance limits (figure 108) near the safely achievable optimum.

A quantitative mass acceptance overview for 2018, 2022, and 2023 is given in table 11, anticipating the upper mass limits, M_{max} , discussed in section 7.5.2. The higher dispersion in Run 3 leads to generally lower values for both M_{min} and M_{max} than in Run 2.




Figure 107. Minimum accepted central mass $M_{\min}$ in the RPs 210-F (left) and 220-F (right) for the old collimation scheme with the TCTs fixed at $d_{\text{TCT}} = 8.5\sigma_{x,\text{TCT}}(\beta^* = 30 \text{ cm})$ and two cases for the RP positions. Magenta lines: RPs moving according to eq. (7.1) and figure 106. Blue lines: RP positions fixed on the most distant point of the nominal trajectory. The red lines correspond to the 1.5 mm distance limit. The fill evolves from the right to the left.




Figure 108. Minimum accepted central mass $M_{\min}$ in the RPs 210-F (left) and 220-F (right) for the new collimation scheme with moving TCTs and fixed RP positions. The red lines correspond to the 1.5 mm distance limit. The fill evolves from the right to the left.

7.5.2 TCL collimator insertion scheme

The debris collimators, TCL4 and TCL5, in the outgoing beams upstream of the RPs determine the high-mass acceptance limit $M_{\max}$ of PPS. A TCL collimator inserted to the distance d_{TCL} creates a

Table 11. Accepted central mass range, $[M_{\min}, M_{\max}]$ (in GeV), for each RP at the beginning and at the end of the levelling trajectory in 2018 ($\sqrt{s} = 13$ TeV), and 2022 and 2023 ($\sqrt{s} = 13.6$ TeV), for the collimation schemes laid out in the text. Due to the coincidence requirement, the RP with the highest $M_{\min}$ (typeset in bold face) defines the spectrometer acceptance.

Year	2018	2022	2023	
$\alpha/2$ [μrad]	160	160	135	160
β^* [cm]	30–25	60–30	120	30
210-F	[449, 2485]	[377, 2086]	[329, 1968]	[360, 2086]
220-N	[407, 2485]	[355, 2086]	[324, 1968]	[355, 2086]
220-C	[397, 2485]	[350, 2086]	[323, 1968]	[354, 2086]
220-F	[347, 2485]	[342, 2086]	[321, 1968]	[352, 2086]

mass cut-off:

$$M_{\max} = \sqrt{s} \frac{d_{\text{TCL}}}{D_{x,\text{TCL}}}, \quad (7.4)$$

where $D_{x,\text{TCL}}$ is the horizontal dispersion at the location of the TCL. Throughout Run 2, TCL4 was inserted to a fixed distance of typically $15\sigma_{x,\text{TCL4}}(\beta^* = 30 \text{ cm})$ from the beam center. The distance of TCL5 was chosen such that it did not introduce any tighter cut: $d_{\text{TCL5}} = 35\sigma_{x,\text{TCL5}}(\beta^* = 30 \text{ cm})$.

In Run 3, the distance of TCL4 is unchanged with respect to Run 2. Due to a slightly different optics, it now corresponds to $17\sigma_{x,\text{TCL4}}(\beta^* = 30 \text{ cm})$. The TCL5 collimator is placed at $42\sigma_{x,\text{TCL5}}(\beta^* = 30 \text{ cm})$, again a position mostly in the shadow of TCL4. Because of the crossing-angle dependence of the dispersion, $M_{\max}$ varies throughout the fill as shown in figure 109. The $M_{\max}$ values for 2018, 2022, and 2023 are given in table 11.




Figure 109. Upper mass cut-off in Run 3 caused by the debris collimators TCL4 and TCL5 for the settings explained in the text as a function of the crossing angle. The fill evolves from the left to the right between the magenta lines.

8 Luminosity and beam conditions

The main deliverables of the beam radiation instrumentation and luminosity (BRIL) system are the measurements of the online and offline luminosity, the beam-induced background, the beam losses and timing, and the radiation products in the CMS experimental cavern. The latter are used to compare with the values from corresponding simulations. The BRIL system delivers luminosity and beam-condition data in real time to the LHC and CMS control rooms. In this section, the instrumentation and technologies used to carry out these measurements are described. The software, computing, and monitoring aspects of the BRIL systems are also discussed. The precision measurement of the integrated luminosity for the different data-taking periods of Run 2 has been described in refs. [181–187]. For the 2016 data, an uncertainty of 1.2% was achieved [181]. For 2017, 2018, and Run 3, improved instrumentation and refined analysis techniques are expected to lead to further reduced uncertainties. Key ingredients of the precision are the availability of the various independently calibrated luminometers with different instrumental systematics and the in situ monitoring of their performance in terms of efficiency and linearity using short vdM-like transverse beam separation scans in physics fills.

For the online bunch-by-bunch luminosity measurement in Run 3, we rely on two fully dedicated luminosity detectors (luminometers): the pixel luminosity telescope (PLT) and the fast beam conditions monitor (BCM1F). In addition, the forward hadron calorimeter (HF) is used for the online luminosity measurement. The HF has a separate trigger-level readout dedicated to the luminosity measurement.

Both PLT and BCM1F were already operated successfully throughout Run 2. For operation during Run 3, improved versions of these detectors were constructed and installed during LS2. Other subsystems relevant to the luminosity measurement are the pixel detector and the muon drift tube system. These various luminometers provide redundancy and complementarity for the precise measurement of the online and offline luminosity and the online monitoring of the backgrounds.

The three independent online luminometers, PLT, BCM1F, and HF, are based on fundamentally different technologies, and each provide precise bunch-by-bunch measurements of the instantaneous luminosity and the beam-induced backgrounds (BIB). The BRIL online measurements are available at all times, in particular whenever there is beam in the LHC, independently of the status of other CMS subsystems and the central CMS DAQ system. The BRIL data-acquisition system (BRILDAQ) delivers data in real time every 1.458 s, corresponding to four so-called “lumi nibbles” (NB). The duration of one NB is 0.3645 s, corresponding to 2^{12} orbits, and 64 NB make up one luminosity section (LS) of about 23.3 s.

Data from the pixel detector and the muon systems contain important additional information that is used to ensure a precise calibration and long-term monitoring of the stability and linearity of the luminosity detectors. Tracker data are collected via the central CMS DAQ for offline analysis, and provide pixel cluster and vertex information per event during calibration runs, and bunch-by-bunch cluster counts per LS for data-taking periods with collisions. Muon data in Run 3 are collected via the BRILDAQ from three sources: (i) orbit-integrated muon track stubs from the muon barrel track finder; (ii) trigger primitives from the DT chambers that are equipped with prototype Phase 2 readout; and (iii) muon counts via the 40 MHz scouting system. The first source is the same as that used during Run 2, while the last two are new and were introduced for Run 3 in preparation for the Phase 2 upgrade.

To monitor the beam conditions, in addition to the dedicated luminometers, we also use the beam-halo monitor (BHM) and beam-condition monitor for beam losses (BCML).

A simulation and monitoring strategy is in place for the radiation background to evaluate and, if necessary, mitigate its effects.

8.1 Real-time bunch-by-bunch luminometers

8.1.1 Pixel luminosity telescope (PLT)

The PLT detector [188–192] is an independent system for the luminosity measurement. It consists of 48 silicon pixel sensors arranged on 16 “telescopes”, eight on each side at 1.75 m from the interaction point (IP), close to the beam pipe, at a pseudorapidity $|\eta| \approx 4.2$. Each telescope is 7.5 cm long and contains three planes with individual silicon sensors. Their inner edges are as close as 4.7 cm from the beam line. Figure 110 shows a sketch of the layout of one end of the PLT (left) and a photograph of the actual detector showing the two half-frames that each side consists of (right).




Figure 110. Left: sketch of the general PLT geometry. The sensors are indicated by the purple squares. Right: the actual PLT detector at one end of CMS, showing the arrangement of the eight telescopes around the beam pipe. Reproduced from [193]. CC BY 4.0.

The silicon sensors [194, 195] are n-in-n (45 sensors) and n-in-p (3 sensors) type and consist of 80 rows and 52 columns of pixels (or 26 double columns of 160 pixels each). The pixels are $150\ \mu\text{m}$ wide and $100\ \mu\text{m}$ high, and cover a total active area of $8 \times 8\ \text{mm}^2$ with an active thickness of 285 and $150\ \mu\text{m}$, respectively. During data taking, it is possible to set a smaller active area to decrease the contribution from accidental hits, which are triple coincidences of signals that do not actually come from a single particle track originating from the IP. The sensors are bump-bonded to a PSI46v2 readout chip (ROC) [196, 197]. Both the n-in-n silicon sensors and the PSI46v2 ROCs are inherited from the CMS Phase 0 pixel detector, while the telescope with n-in-p sensors of $150\ \mu\text{m}$ thickness uses prototype CMS Phase 2 silicon sensors.

The PLT detector is usually operated in a “fast-or” mode in which the data are read out at the full bunch-crossing frequency of 40 MHz. The fast-or mode does not contain full information on hit

positions. In fast-or mode, if any pixels in a sensor register a hit over threshold, a single signal is produced. The ROCs also have a separate data path to read out the full pixel data (hit location and pulse height) at a lower rate (up to about 3.3 kHz) upon receipt of a level-1 accept (L1A) trigger signal. The full pixel information, read out at a lower rate, is used to measure calibration constants, corrections, and systematic uncertainties for the online and offline measurements.

The three ROCs for a single telescope are connected to a high-density interconnect (HDI) card, which contains a token bit manager (TBM) chip [198]. The TBM chip coordinates the readout of the three individual ROCs and produces a single readout for each telescope. Four telescopes are connected to a port card, which manages the communication and control signals for a single half-frame of the detector. The port card is in turn connected to the opto-motherboard (OMB), which converts the electrical signals into optical signals that are then sent by fibers to the CMS service cavern, where the backend readout electronics are located. Of the different kinds of application-specific integrated circuits (ASICs) used in the PLT, all but two (SlowHub and PLTdriver) were developed for the Phase 0 CMS pixel detector [1]. The SlowHub chip, on the OMB, uses blocks from the pixel TBM chip to extract the slow I²C signals from the fast 40 MHz commands needed for the pixel readout, thus significantly reducing the total number of fibers needed. The PLTdriver chip manages the fast-or signal outputs from the individual ROCs.

The backend readout electronics comprise a single frontend controller (FEC) card, which issues commands to the ROCs, TBMs, and OMB, and three frontend driver (FED) cards. One FED is used to read out and decode the full pixel data from the ROCs, and is identical to the FEDs used by the Phase 0 pixel detector [199]. This data are then read out over an Slink [200] connection and saved to a dedicated PC. The other two are the fast-or FEDs, one for each side of the PLT. These read out the fast-or data from the ROCs and histogram the number of triple coincidences per channel and per bunch crossing (BX). These data are read out via a CAEN VME optical bridge to a separate dedicated PC. A dedicated trigger is used, in which events from colliding BXs are read out without a requirement for specific activity in the event, this way equally sampling all colliding BXs in the LHC orbit.

The readout hardware counts the number of triple coincidences, events where all three planes in a telescope register a signal, to determine the instantaneous luminosity. This fast-or readout allows the PLT to provide online per-bunch luminosity with excellent statistical precision, with the triple coincidence requirement providing a strong suppression of background from noise, residual radioactivity of the detector material (afterglow), and BIB sources from beam-gas interactions and beam halo. The luminosity measurement is obtained using a zero-counting technique, described, e.g., in ref. [181], in which the actual rate is inferred from the measured rate of no hits, thereby correcting for the overlap of signals.

8.1.2 Fast beam conditions monitor (BCM1F)

The BCM1F detector is a fast particle counter installed around the beam pipe, ± 1.8 m away from the IP, at a radius of about 6 cm, corresponding to $|\eta| \approx 4.1$. The BCM1F provides a real-time measurement of both luminosity and beam-induced backgrounds.

The experience gained from BCM1F operation in Run 2 [201, 202] led to an update of the design for Run 3 [203]. For Run 3, the mix of silicon and poly-crystalline (pCVD) diamond sensors used in Run 2 was changed to an all-silicon sensor configuration with active cooling that has an

improved stability with increasing integrated fluence and an improved linear response of the sensors as a function of the instantaneous luminosity.

As in Run 2, the Run 3 detector consists of four “C-shaped” printed circuit boards (PCB), referred to as C-shapes (figure 111). A total of four C-shapes form two rings, one on each side of the IP. Each C-shape has six double-diode silicon sensors that are used to measure particle hits. The entire system comprises 48 identical channels. The C-shapes for Run 3 were designed for improved robustness and with additional space for active cooling contacts.




Figure 111. Photographs of the BCM1F detector used in Run 3. Reproduced with permission from [204]. Left: one full BCM1F C-shape printed circuit board. Right: a closeup of the frontend module, which includes the processing chip with two silicon double-diode sensors on each side.

Performance studies led to the choice of a sensor per-channel with an area of $1.7 \times 1.7 \text{ mm}^2$, balancing the occupancy at high interaction rates with the statistical precision at low luminosity. The sensors were designed as alternating current (AC)-coupled double-diodes and produced on the CMS Phase 2 outer tracker strip wafers [32] with a full depletion thickness of $290 \text{ }\mu\text{m}$.

Each sensor is directly connected via bond wires to one of the three fast JK16 [205] CMS ASIC chips. This preamplifier chip was designed to have a fast rise time of 7 ns , narrow pulses with 10 ns full width at half maximum, and a return to the baseline within 25 ns . The connections are established via AC coupling to prevent baseline drifts due to radiation-damage-induced leakage currents that would not be compensated by the preamplifier. The coupling capacitance is part of the sensor design, while a resistor on the PCB is used as the bias resistor.

Copper pads were added on the C-shapes for Run 3 to make the cooling contact with a new titanium 3D-printed cooling pipe, which is placed in series with the PLT cooling loop. This improves the thermal contact of the sensors to the cooling system and prevents thermal runaway induced by radiation damage, thereby ensuring the expected lifetime of the BCM1F silicon sensors exceeds the duration of Run 3.

Using analog opto-hybrids [206], the AC signals from the C-shapes are propagated to the backend electronics that includes CAEN V895 discriminators, real-time histogramming units (RHU) [207], and a CAEN V1721 flash ADC, all of which use a VME bus. Parallel μTCA electronics consist of an 8-bit ADC mezzanine card with a 1.25 GB/s sampling rate on a gigabit link interface board (GLIB) carrier board [208, 209], using an FPGA-based derivative peak-finder algorithm [210]. At the RHU, which was already used for Run 2, the data are filled into histograms with 6.25 ns granularity. This

time interval corresponds to the time of flight for relativistic particles arriving from the IP that is a quarter of the time between two collisions.

The BCM1F detector, situated in this location and delivering a fast response with zero dead time, allows for the bunch-by-bunch separation of collision signals from backgrounds. The luminosity is measured bunch-by-bunch using the zero-counting method mentioned above, followed by a correction for out-of-time hits and the application of a calibration constant. The relative contribution from background can be determined using the first bunch in the train or a noncolliding bunch in the beginning of the orbit, where the background from out-of-time hits due to activation from previous collisions is small. The background measurement is of crucial importance for the operation of the CMS detector since it is used by the silicon strip and pixel detectors as an automatic switch-on semaphore.

8.1.3 The forward hadron calorimeter (HF)

A comprehensive description of the HF is provided in section 5. Here, we highlight the aspects that are relevant for the luminosity measurement.

Like the PLT and BCM1F, the HF reports luminosity measurements at all times, independently of whether the main CMS DAQ system is in operation. The outputs from the circuits used to digitize the signals from the HF photomultipliers (PMTs) are routed to FPGAs that are part of the HF readout [60]. The FPGAs tap into the primary readout path in a noninvasive way and collect channel-occupancy (OC) and transverse-energy-sum data in histograms that have one bin for each of the 3564 BX time windows of one beam orbit. The histograms are periodically transmitted to a central BRILDAQ processor node. With this configuration, the HF delivers a continuous and precise measurement of the bunch-by-bunch luminosity at all times.

While all HF channels are available for use in the BRIL backend electronics, Monte Carlo simulation studies indicate that the best linearity of the luminosity measurement is obtained when using only the η rings 31 and 32, corresponding to a pseudorapidity range $3.15 < |\eta| < 3.5$, where the long-term radiation damage is predicted to be the smallest.

The HF luminosity system allows the extraction of a real-time instantaneous luminosity relying on two algorithms. The first (HFOC) is based on zero counting, in which the average fraction of below threshold towers is used to infer the mean number of interactions per BX. The second (HFET) exploits the linear relationship between the average transverse energy per tower and the luminosity. Both the HFOC and HFET methods exhibit excellent linearity, based on the experience acquired during Run 2 [181–183, 183, 186] and in the beginning of Run 3 [211].

8.2 Additional luminometers

8.2.1 Tracker

The CMS silicon tracker, described in section 3, is characterized by a low occupancy in the silicon pixel detector. This feature is exploited by the pixel cluster counting (PCC) method to measure the luminosity offline with excellent precision. The most recent PCC luminosity measurement during Run 2 [183] was based on data collected by zero-bias (colliding bunches only) and random triggers at recording rates of approximately 2 kHz and 400 Hz, respectively. The zero-bias data were used for the luminosity calculation, while random-triggered events served for the determination and subtraction of the afterglow backgrounds present in the zero-bias data. The method achieved a

statistical precision of about 0.2% per luminosity section during physics runs, with excellent stability after removing problematic sensor modules.

We applied a similar strategy for Run 3, with improvements in the data processing and recording to reduce and streamline the data. The cluster reconstruction and counting is implemented in the HLT software, which reduces the data size by three orders of magnitude. Potentially unstable modules are removed using an iterative method in which modules with large fluctuations relative to the average cluster count are removed. The implementation of a new processing path is in progress during Run 3, in preparation for the Phase 2 upgrade of the luminosity instrumentation for the HL-LHC [212]. In this path, the data from the HLT are transferred directly to the BRILDAQ.

An alternative method using information from the silicon tracker involves counting the number of primary-interaction vertices [181]. This method is used during van-der-Meer (vdM) calibration fills as a tool in the measurement of beam-dependent parameters. Vertices are also analyzed to derive the bunch density distributions of the beams, and to calibrate the distance by which the steering magnets displace the beams in the transverse direction.

8.2.2 Muon system

The BRIL system also makes use of level-1 trigger information from the muon barrel (MB) drift tube (DT) detectors, which deliver counts of the number of orbit-integrated muon stubs. While the rate of muons in the barrel is low, the stability and robustness of the muon system makes the DT a valuable source for luminosity monitoring. The muon detectors and trigger are described in sections 6 and 10.2, respectively. The muon information used by BRIL is aggregated by the barrel muon track-finder algorithm (BMTF) [213] of the level-1 muon trigger. The BRILDAQ system receives the data from the BMTF via a dedicated readout system.

In addition to the orbit-integrated data, during Run 3 a new histogramming firmware module was added to the backend of the DT slice test that uses Phase 2 readout technology to test how trigger primitives (muon segments per DT chamber) can be used for bunch-by-bunch luminosity measurements. This system serves as a demonstration of how a luminosity measurement can be made with per-bunch granularity for the CMS Phase 2 upgrade.

Figure 112 illustrates the readout diagram of the DT slice test. A detailed description is provided in ref. [88]. The MB1 and MB2 chambers are read out using the upgraded electronics. Data duplication was established for the MB3 and MB4 layers, thus allowing us to evaluate the performance of the upgraded electronics with respect to the older system. The BRIL histogramming module is placed on the so-called AB7 backend boards. Each AB7 board produces a single histogram. Each of the MB1, MB2, and MB3 chambers are read out by a single AB7, whereas two AB7 boards are needed to read out the MB4 chamber.

In addition, the 40 MHz scouting system in Run 3, described in detail in section 10.6, provides an alternative means to use the bunch-by-bunch muon rate from the level-1 muon trigger for the luminosity measurements. The concept of the 40 MHz scouting system is to capture part or all of the trigger data streams from different trigger layers or their subsystems using spare optical outputs. A demonstration of the scouting system was already tested in Run 2 based on the output from the global muon trigger algorithm μ GMT described in section 10.2. Due to the inclusion of the layer-2 calorimeter trigger and the BMTF data, the Run 3 scouting demonstration offers additional objects with more sophisticated customizable selection possibilities. The histogramming for measuring




Figure 112. Diagram comparing the DT readout (RO) and trigger chains in the Phase 1 system used for data taking in Run 3 and in the DT slice test that acts as a demonstrator of the Phase 2 system in Run 3. Reproduced from [88]. CC BY 4.0. The central part of the figure indicates the number of Phase 2 on-board DT (OBDT) boards installed in each DT chamber. The left part of the figure shows how the Phase 1 readout and trigger-primitive generation are performed by the legacy on-board DT electronics (minicrates). Information, transmitted by optical fibers from the detector balconies to the counting room, is further processed independently for the readout (μ ROS) and trigger (TwinMux). The right part of the figure illustrates how the slice test TDC data are streamed by each OBDT to AB7 boards hosted in the counting room, which are used for event building and trigger-primitive generation.

the luminosity is performed per object type using the already mentioned generic histogramming firmware loaded on the scouting board.

The histogramming firmware is a generic module developed by BRIL that targets luminosity measurements using various, primarily Phase 2, CMS subsystems. The DT muon detector and the 40 MHz scouting system are the first to use this new module, and multiple instances of the module are placed on backend FPGAs. The module comes in two flavors: a synchronous type is adapted to the bunch-by-bunch luminosity measurements, when the counted objects arrive synchronously with the bunch clock, as is the case of the DT system. An asynchronous flavor is used for the 40 MHz scouting system in which preprocessing of the trigger primitives is performed at a rate of 250 MHz, providing histogram input approximately every six clock periods.

The module is implemented for the Xilinx 7-series and newer FPGAs. The readout of the histograms can be performed using any suitable data transfer protocol. In particular, the IPbus and DMA readout interfaces have been implemented for the DT and 40 MHz scouting systems, respectively.

8.2.3 Z boson counting

Due to a clean signature and a relatively high cross section, the $Z \rightarrow \mu^+ \mu^-$ process is of particular interest for the luminosity measurement. It provides an alternative and complementary method to transfer integrated luminosity measurements between data sets, as discussed in refs. [214, 215]. In Run 3, Z boson counting is implemented as part of the data reproduction chain with the capability of providing precise results within about one week of the data taking. A full initial analysis of the 2022 data has been performed [216] and used to confirm the early calibration of the luminometers

by comparison with the predicted cross section [217]. The Z boson rates are also used as an additional method to directly compare the integrated luminosities delivered to the ATLAS and CMS experiments [218, 219].

8.3 Beam monitoring instrumentation

8.3.1 Beam-halo monitor (BHM)

To monitor the BIB at high radius, an efficient detection system was installed during LS1 and operated in 2015–2016. It was then recommissioned for Run 3. The BHM monitors [220] are composed of a total of forty cylindrical quartz Cherenkov detectors, twenty on each end of the CMS detector, mounted on the rotating shield. The rotating shield is situated at a distance of about 20 m from the interaction point. The exact location was chosen to maximize the time difference between the arrival of BIB and particles from the pp collisions. The BHMs are mounted on the outer radius of the shield, at a distance of 1.83 m from the beam line. The units are uniformly distributed in ϕ , starting from $\phi = -30^\circ$ and continuing through $\phi = 210^\circ$.




Figure 113. Left: photograph of the active element of one BHM detector unit, which is a 100 mm-long by 52 mm-diameter cylindrical quartz radiator, connected to a Hamamatsu R2059 photomultiplier. Right: shielding of a BHM detector unit, consisting of a low-carbon steel tube shown in blue, a mu-metal tube in gray, and permalloy in yellow. The quartz radiator, photomultiplier tube, and socket are shown in white, light green, and dark green, respectively. Reproduced from [220]. © CERN 2015. CC BY 3.0.

The quartz radiators, shown in figure 113 (left), are optically coupled to a fast ultraviolet-sensitive PMT on one end and painted black on the other. These detector units use the directional nature and timing of the BIB background and the pp collision products. In order to use the directional property, the PMT side of the quartz detector faces the center of the CMS detector. When a BIB muon arrives with the incoming beam and produces Cherenkov radiation in the quartz, the light propagates forward and is collected by the PMT. In contrast, pp collision products arrive from the opposite direction, and the produced light is absorbed by the black paint.

The shielding of each BHM unit consists of three coaxially arranged metallic tubes, as illustrated in figure 113 (right). The outer tube is a 10 mm-thick, 38 cm-long tube made of mild steel with endcaps on both sides composed of soft iron (ARMCO) with the same thickness. The intermediate layer is a 30 cm-long, 1.5 mm-thick mu-metal tube, and the innermost tube is 27 cm long and 0.8 mm thick made of permalloy. The shielding is designed to protect the PMT tubes from the 17 mT residual magnetic field and from radiation that mainly originates from low-energy electrons and positrons produced in the pp collisions, which otherwise would form the dominant contribution to the BIB.

The readout of the BHM detector makes use of many components developed for the HCAL Phase 1 electronics upgrade [60], with dedicated firmware and readout adapted to the beam-monitoring requirements [221]. The PMT signal is digitized by a charge-integrating ASIC (QIE10),

providing both the signal rise time and charge integrated over one bunch crossing. It ensures dead-timeless readout of the signal amplitude and edge time information with 500 ps resolution. As a part of the readout from the backend electronics, the occupancy histograms are obtained in time bins of a configurable duration. The histogram bins are typically 6.25 ns wide and are read out every 4 NB (about 1.458 s). In special running periods, detailed self-triggered data containing event-level information about amplitude and time-of-arrival of each hit can be also collected.

A calibration and monitoring system was also installed for the BHMs [222] to evaluate possible changes in performance of the PMTs and quartz radiator due to aging and radiation damage. The system uses a light signal produced by UV-emitting pulsed LEDs, which is sent to each detector unit through quartz optical fibers and optical splitters. The monitoring system can be set to send a laser pulse periodically in an empty orbit or in periods when beams are not present.

8.3.2 Beam-condition monitor for beam losses (BCML)

The BCML detector [223–225] is linked to the abort system of the LHC, and protects the CMS silicon tracker from beam-loss events. The detectors simultaneously measure the integrated currents over 12 different time intervals and produce an unmaskable beam-abort trigger if one of them exceeds a threshold, which is configured to be at least three orders of magnitude less than the amount estimated to cause damage. The BCML system is situated in four different locations within the CMS detector and includes the BCML1 detector located at $z = \pm 1.8$ m and the BCML2 detector located at $z = \pm 14.4$ m from the interaction point. A total of 16 channels of the BCML detector are actively used in the beam-abort system, with four channels in each location, providing CMS with monitoring redundancy.

The BCML system primarily uses pCVD diamond sensors, most of which were replaced in LS2. In addition, BCML2 has sapphire sensors installed, as described below. Diamond detectors are typically used as robust beam monitors in locations where the radiation levels are very high. Diamond is radiation hard and does not require active cooling.

There are four pCVD diamond sensors installed at each end of both the BCML1 (figure 114) and the BCML2 (figure 115) detectors, 16 sensors in total, each having an active volume of $10 \times 10 \times 0.4$ mm³. The current created by ionization in the sensor is proportional to the ionizing particle flux through the active detector material, and it thus provides a good observable to determine the amount of radiation that a sensor receives. The BCML sensor signals are read out using the LHC beam-loss monitor (BLM) electronics [226–228].

If a beam loss signal, i.e., an integrated current over threshold, is detected in BCML1, the beam can be dumped within two to three LHC turns. The threshold is defined for two different integration time intervals. A short-duration (40 μ s) threshold protects the silicon tracker from potentially damaging amounts of energy deposition in the electronics, and a longer-term (≈ 83.3 s) integration threshold protects against high beam background conditions that can result in problematic conditions for data taking and reconstruction, as well as an increased dose rate to the inner detector region. A machine beam interlock system [229] ensures that the BCML system is fully operational before beam injection into the LHC can begin. The BCML systems are powered independently of CMS from an uninterruptible power system, so that the LHC can run even if the local CMS power distribution is not functional.

New optical grade single-crystal sapphire sensors were installed in BCML2 during LS2 in six stacks. These are produced industrially via the Czochralski process at low cost [230]. Similar to




Figure 114. Mechanics for the BCML1 detector system. Left: a BCML1 mounted on the BCM1F C-shape PCB, attached to the PLT support structure. Right: a BCML1 mounted on a C-shape PCB (upper right), and a single sensor in a Faraday cage (lower right).




Figure 115. Mechanics for the BCML2 detector system. Left: the support structure of the BCML2 detectors, which also serves as mechanical protection. This structure is compatible with the Phase 2 beam pipe. Right: single-sensor base plate PCB for the large sensors used in BCML2 (upper right), and a complete sensor box assembly, with a stack of five sensors (lower right).

diamond, sapphire, chemically aluminum oxide Al_2O_3 , behaves as a wide-band-gap semiconductor, with a band gap of 9.9 eV. First measurements indicated a very high tolerance against radiation damage [225, 231], although the charge collection efficiency is low ($\approx 1\text{--}10\%$) [231]. The sapphire Phase 2 prototype sensors used in BCML2 have an active volume of $25 \times 25 \times 0.5 \text{ mm}^3$ and can be installed in units ranging from a single sensor to a stack of up to five sensors, as shown in figure 115. In the stack option, the sensors are electrically mounted in parallel such that the observed current is the sum of the currents in each sensor.

To normalize the BCML2 readings and thereby monitor nonlinearities and instabilities in the response from the diamond and sapphire sensors, in addition to the solid-state detectors, two ionization chambers (one per end) of the LHC BLM type [226] were placed on top of the table originally designed to support the Castor subdetector. This is located at about $z = 1450 \text{ cm}$, directly behind BCML2, and the chambers are connected to the same readout electronics as the BCML2 system. These detectors provide a good dynamic range, demonstrate linearity with a deviation less than 1% [232], and no radiation degradation is expected. The proximity to the BCML2 detectors, and the fact that these ionization chambers are in the same readout system as BCML2, makes them ideal for normalizing the BCML2 readings and thereby monitoring the nonlinearities and instabilities in the response from the diamond and sapphire sensors. Since the response of these ionization chambers is also well known for a mixed radiation field, the detectors can also be used to benchmark the radiation simulations.

8.4 Radiation instrumentation and simulation

Radiation background can cause unwanted signals in detectors, activation of materials, damage to electronics, and detector degradation. A radiation simulation and monitoring strategy is in place such that background radiation can be estimated, detected, analyzed, and lowered where necessary.

8.4.1 Radiation monitoring

The radiation monitoring strategy for Run 3 includes the use of the following systems:

- The HF RadMons [233] are gas-filled proportional counters situated inside the HF regions. There are eight in use during Run 3, four on either end of the CMS detector. The detection principle is based on neutron capture by boron with the emission of an α particle in the argon-filled proportional counter. The incident neutrons are slowed by a surrounding polyethylene moderator with a 3-inch radius. The HF RadMons provide information on the total neutron flux for benchmarking the simulation estimates.
- As the part of a wider radiation and environment monitoring unified supervision (REMUS) system under the responsibility of the CERN radiation protection group (HSE/RP), during Run 3 there are 12 inducted activity monitors (IAM) in the CMS cavern, read out by the radiation monitoring system for the environment and safety (RAMSES) [234, 235]. Ten of these are in the same place as in Run 2, and two additional ones were installed in the same location as the existing monitors on the HF, to support luminosity stability studies. The primary function of the IAMs is to measure the ambient-dose equivalent rates, $H^*(10)$, associated with the residual radiation field when the LHC is not operational. However, their

large dynamic range means that they can also provide reliable measurements of the prompt radiation field during collisions. They are used to monitor any upgrades to shielding in the rotating shield region and provide benchmark data for simulations. This system is also used by BRIL for luminosity stability and linearity studies.

- The LHC RadMon system [236, 237], which is maintained by the CERN EN/STI group, is designed to monitor the radiation that is related to electronics damage. There are currently 29 LHC RadMon units situated inside the experimental cavern, and it is foreseen to add ten more during Run 3. Each unit, upgraded to “Type v6”, has 13 detectors, which includes two radiation-sensitive field effect transistors (RadFETs) to measure the total ionizing dose, three Pin Diodes (in series) to measure the 1 MeV-equivalent fluence in silicon, and eight static random access memory (SRAM) modules. The latter measure the cumulative fluence of high-energy hadrons via single-event effects and thermal neutrons using a different voltage setting. The BRIL group also uses the LHC RadMon data for radiation simulation benchmarking studies.

In addition to the monitoring systems listed above, several established methods are used to provide information about the radiation field. These include activation samples that are placed in the outer cavern with subsequent measurements performed by the HSE/RP group; monitoring of CMS detector degradation; and the two LHC beam-loss monitors that are used as part of the BCML setup, described in the previous section. Radiation monitoring is complemented by dedicated simulation studies, used to predict and understand in more detail the background radiation in CMS.

8.4.2 Radiation simulation

Monte Carlo simulations to predict radiation levels in the experimental cavern and the CMS detector are typically performed with the CERN FLUKA radiation transport code [238, 238]. The maintenance of the FLUKA geometry models is the responsibility of the BRIL radiation simulation team, as well as the dissemination of the simulation results to the CMS subdetector teams via a web-based tool. The baseline pp collision simulation using a geometry model that reflects the current CMS Run 3 configuration is tagged v.5.0.0.0. Relative to the Run 2 FLUKA geometry, this includes the implementation of the LS2 upgrades, the replacement of the central beam pipe from the IP to $z \approx 16.7$ m, modifications of material within the rotating shield regions including reinforcements to fill existing gaps, as well as modifications to the installation of the LHC vacuum assembly for experimental area (VAX) equipment in LS3. The radiation levels per integrated luminosity (or collision) in the central detectors for Run 3 are expected to be higher than in Run 2 by up to a factor of 1.5, depending on the location and type of particle, owing to changes in the shape of the central beam pipe. Changes in the predicted radiation levels in the outer detectors for Run 3 are influenced by the increased showering in the central regions (resulting in a reduction downstream) and modifications inside and outside the rotating shield region where a lot of the secondary radiation is generated and leaks into the cavern. The lighter aluminum beam pipe results in significantly lower activation levels for a given irradiation pattern. The FLUKA predictions of the radiation environment in Run 3 are illustrated in figure 116 for the CMS cavern and detector. The effects of an additional new forward shield, expected to be installed in 2024 around the existing rotating shield and included in geometry model v.5.1.0.2, are shown in figure 116 (lower).




Figure 116. The FLUKA predictions for the expected Run 3 fluence of hadrons with energies greater than 20 MeV (upper left) and neutron fluence (upper right) normalized to an integrated luminosity of 200 fb⁻¹ at 7 TeV per beam are shown for the CMS cavern and detector. For the central part of CMS, 1 MeV-neutron-equivalent fluence (middle left), and absorbed-dose (middle right) are also presented. The lower two plots show the expected effect of the new forward shield as the ratio of hadron (left) and neutron fluences (right) in the CMS cavern comparing the Run 3 FLUKA simulation results of v5.1.0.2 with the Run 3 baseline of v5.0.0.0.

8.5 The BRIL online data acquisition and monitoring

The BRIL data acquisition (BRILDAQ) system facilitates the real-time recording, processing, and monitoring of luminosity and beam background data at a latency of only a few seconds. Essential BRIL data are delivered in real-time to the CMS and LHC control rooms, allowing for prompt and efficient operation of the LHC machine and CMS detector. BRILDAQ is completely separate from the central CMS DAQ system, such that BRIL can operate continuously and independently of the CMS data-taking status even when all other CMS detector components are switched off.

Similar to the central DAQ system of CMS, BRILDAQ manages distributed and heterogeneous subsystems using a homogeneous distributed software architecture, where reliability and redundancy are important to prevent down time. The BRIL data from different subsystems can be viewed together and correlated in real time, and new subsystems can easily be integrated.

The core software is based on the CMS online framework XDAQ [239]. This event-driven and message-oriented architecture enables large numbers of loosely coupled software components

and services to exchange information in near real time. Messages are exchanged via the publisher/subscriber model (eventing) where the publishers send messages that can be received by any number of subscribers. Communication between BRILDAQ and the LHC is based on the data interchange protocol (DIP) [240], which also uses a publisher/subscriber model.

To fully exploit the redundant online luminosity and beam background monitoring systems, the histogramming of the BRIL subsystem frontends is synchronized using common timing signals distributed by the CMS TCDS system (section 9), defining the hit count integration interval boundaries. Long command counters are also distributed for additional intersystem synchronization. Using such a technique facilitates uniform accounting of the delivered luminosity and the downstream data handling by BRILDAQ.

Typically, a subsystem provides one or several source instances, which read the raw histograms from the hardware, and one single processor instance for data aggregation and calibration. Both kinds of applications are stateless to ensure the availability of the data, regardless of the status of the LHC beams and the running state of CMS. In Run 3, the participating subsystems are PLT, including the fast-or and the Slink data, BCM1F, with both RHU and μ TCA readouts, HF, DT, and REMUS for luminosity measurements, and other BRIL measurements from subsystems such as the beam-pickup system for timing measurements (BPTX) [212], HF-RadMon gas-filled proportional chambers, BHM, and BCML. High-level components include the storage manager, the vdM monitor providing the luminometer calibration constants in real time after a beam-separation scan, the luminosity monitor, and the best-luminosity selector. The latter selects the value of the instantaneous luminosity that is delivered to the LHC. The DIP-related components use the service provided centrally by the CMS DAQ team. In addition, it is foreseen to process PCC data out of the HLT in near-real-time fashion, as mentioned in section 3.

The run control system, responsible for controlling and tracking the configurations of the BRIL applications, is based on the CMS run control framework (RCMS) [73] and its configuration database. Each BRIL process is controlled by a function manager (FM) that manages a simple finite-state machine.

The run control web frontend allows the operator to control the life cycle of BRILDAQ applications and to manage the configurations, which are versioned and can be retrieved and stored in the database via the web. A set of other web tools completes the system by providing useful functionalities such as displaying logs of the processes and showing results of the data analysis.

The BRIL online web monitor is a single-page application built on the Angular framework [241] where the website interacts with the user by dynamically rewriting the current web page with new data instead of reloading the entire page. It displays real-time and historical charts of online quantities. Monitoring data are published by BRILDAQ applications, then stored in the ElasticSearch [242] backend database, which responds to requests from the web client. The monitoring system is nonintrusive and flexible for easy integration of new data and charts.

The raw and calibrated data are stored on local disks, then transferred to an offline storage area, and eventually moved to tape. Summary data are loaded into the BRIL database. All types of data can be reprocessed, and multiple versions can coexist on disk and in the databases.

To provide high-quality luminosity measurements for use in physics analysis, an application toolkit is provided to the physics community, centered around the so-called “brilcalc” tool for luminosity calculation. It provides delivered and recorded luminosity with different calibration sets,

using a “normtag” that defines the best detectors and calibrations to use for each time period. This strategy has proven to be successful and remains unchanged in Run 3.

9 Data acquisition system

This section describes the scope, design choices, and implementation of the experiment’s central data acquisition (DAQ) system. The present implementation for Run 3 is described in detail along with a discussion of its performance. We also give an overview of the evolution of the system over the lifetime of the CMS experiment.

9.1 Scope

The CMS online event selection is performed using two trigger levels: the level-one (L1) trigger, described in section 10, implemented in custom electronics, which selects approximately 100 kHz of events based on coarse information from the calorimeters and the muon detectors; and the high-level trigger (HLT), described in section 11, which runs on a farm of commercial computer nodes integrated with the DAQ data flow. The HLT processes fully assembled events, applying algorithms similar to those used in offline reconstruction, and selects a few kHz of events for storage on disk.

While this two-level approach with full event building after the first level greatly simplifies the overall system design compared to approaches with more trigger levels and/or partial event building, the resulting requirements on the data acquisition system are demanding: the DAQ system needs to read out approximately 700 detector backend boards at a rate of ≈ 100 kHz and perform event building and distribution with a throughput of about 100 GB/s— a challenging task when the DAQ system was first implemented [1] for Run 1 with the hardware available in the late 2000s.

The DAQ system is also responsible for collecting events selected by the HLT, buffering them at the experiment site, and transferring them to the Tier 0 computing center at CERN. The general structure of the DAQ system has remained fairly constant from its original implementation for Run 1 up to its current implementation (illustrated in figure 117). Across implementations, the main components of the system are:

- A trigger throttling system (TTS), consisting of custom electronics modules collecting fast status information from all the backend boards in order to throttle the trigger to avoid buffer overflows.
- Since Run 2, the trigger control and distribution system (TCDS), implementing the trigger control logic, reacting to fast status signals from the TTS or collected directly from upgraded FEDs, and distributing timing and trigger signals to the entire experiment.
- Custom electronics modules, the frontend readout (optical) link boards that receive data over a custom link from the detector backend boards (in CMS called frontend driver or FED boards), and send data over a commercial link.
- A data concentrator network based on commercial network technology.




Figure 117. Diagram of the Run 3 DAQ system. The total numbers of cabled elements are given, including elements used in the MiniDAQ systems described in section 9.11 and elements installed as hot spares or for contingency. In typical global data-taking configurations, a subset is used, as described in the text.

- Readout unit (RU) servers that aggregate data from $O(10)$ FEDs into super-fragments and buffer the data.
- An event builder network based on commercial network technology.
- Builder unit (BU) servers that build full events from super-fragments received from all the RUs. In Run 3, the RUs and BUs run on the same set of servers (RU/BU nodes).
- An event backbone network, based on commercial network technology, connecting the BUs to the filter unit (FU) servers.
- The FU servers that execute the HLT algorithms.
- A storage and transfer system (STS), merging events selected by the HLT into larger files, buffering them locally, and transferring them to the Tier 0 at CERN via the central data-recording (CDR) network.

The DAQ system comprises software for data handling and controlling custom electronics, a hierarchical control system to manage most of the aforementioned components, a monitoring system that collects metrics, and several monitoring clients that interpret these metrics.

9.2 Evolution

Since its first implementation for Run 1, the CMS DAQ system has undergone multiple upgrades to keep pace with the evolving needs of the experiment. Towards the end of Run 1, it became clear that during Run 2 the LHC would reach higher instantaneous luminosities than originally anticipated, resulting in increased pileup and larger event sizes. In response, several subdetectors implemented upgrades and replaced VME-based backend electronics with new systems based on the μ TCA standard. A new optical readout link with a higher bandwidth of up to 10 GB/s was developed for these new backend systems. The custom electronics of the DAQ system were upgraded to support this optical link and to support a new network technology for the data-concentrator network, from Run 2 onwards.

The bulk of the DAQ system downstream of the custom readout electronics consists of commercial computing and networking equipment, for which a regular replacement cycle must be observed. This circumstance provided the opportunity to reimplement large parts of the DAQ system for Run 2, and again for Run 3, benefiting from advances in technology to achieve a much more compact design while doubling the event building bandwidth from Run 2 onwards. The choice of networking technology has been adapted following trends in the supercomputing industry. An overview of the three generations of the CMS DAQ system is provided in table 12.

The HLT computing capacity was scaled up according to the experiment's needs on a yearly basis over Run 1 and Run 2. The Run 3 system includes general purpose graphics processing units (GPUs), providing cost-effective computing acceleration. This has required a major effort in porting the HLT code. Major changes in the software enabled better decoupling of the DAQ and HLT processes since Run 2. Various automated features and automatic diagnostics have been added to the control system to maximize the uptime of the DAQ system. Driven by the need for more detector partitions than the original system could support, and to host μ TCA backends, a new trigger control and distribution system (TCDS) was developed for Run 2, absorbing the functionality of the original trigger control system (TCS), and TTC system, and providing optical inputs for the TTS signals from the upgraded backends. In the following sections, the upgrades to the DAQ system are described in more detail with an emphasis on the implementation for Run 3.

9.3 Subdetector readout interface

During Run 1, the CMS subdetectors were read out exclusively through the S-LINK64 [200] DAQ readout link, an LVDS-based copper link capable of transferring up to 400 MB/s. One or two such links are received by the frontend readout link (FRL) [246], a custom Compact-PCI card. The FRL forwarded the data to a commercial Myrinet network interface card (NIC) via an internal PCI-64 bus at 66 MHz. Super-fragments were built on the Myrinet NICs using custom firmware [246]. Assembled super-fragments were then transferred into the memory of a readout unit (RU) server via DMA.

At the start of Run 2, the Myrinet NIC on the FRL was replaced by a custom-developed PCI-X card, the frontend readout optical link (FEROL) [247]. The FEROL acts as a 10 Gb/s Ethernet NIC, sending data to the event builder. It receives data either via the PCI-X interface from the FRL (from the bulk of the subsystems that continue to use S-LINK64) or via optical SlinkExpress inputs from new or upgraded subsystems. Up to two SlinkExpress inputs at 6 Gb/s or one input at 10 Gb/s are supported. While S-LINK64 senders are mezzanine cards plugged onto the backend boards of the subdetectors, requiring a significant footprint, the SlinkExpress sender is a firmware IP-core that

Table 12. Key parameters of the CMS DAQ system in Run 1 [1, 243], Run 2 [244, 245], and Run 3.

	Run 1	Run 2	Run 3
Event building rate pp ^a	100 kHz	100 kHz	100 kHz
Event size pp ^a	1 MB	2 MB	2 MB
Read-out links S-LINK64 (copper) 400 MB/s ^b	636 ^c	575 ^c –532 ^d	528 ^c
Read-out links optical ^e 6 Gb/s ^b	—	55 ^{d,f}	55 ^{c,f}
Read-out links optical ^e 10 Gb/s	—	60 ^c –167 ^d	176 ^c
FED builder network technology	Myrinet	Ethernet	Ethernet
FED builder network speed	2 rails of 2.5 Gb/s	10 & 40 Gb/s	10 & 100 Gb/s
Event builder # of readout units	640	108 ^d	57 ^g
Event builder network technology	Ethernet	Infiniband	Ethernet RoCE v2 ^h
Event builder link speed	1–3 rails of 1 Gb/s	56 Gb/s	100 Gb/s
Event builder parallel slices	8	1	1
Event builder network throughput	1.0 Tb/s	1.6 Tb/s	1.6 Tb/s
Event builder # of builder units	1260 ⁱ	73 ^d	57 ^g
BU RAM disk buffer	none	16 TB	10 TB
HLT # of filter units	720 ^{c,i} –1260 ^{d,i}	900 ^c –1084 ^d	200
HLT # of cores	5.8k ^c –13k ^d	16k ^c –31k ^d	26k ^j
HLT computing power (MHS06)	0.05 ^c –0.20 ^d	0.34 ^c –0.72 ^d	0.65 ^j
HLT # of NVIDIA T4 GPUs	—	—	400
Storage system technology	16 SAN ^k systems	1 cluster file system	1 cluster file system
Storage system bandwidth write + read	2 GB/s	9 GB/s	30 GB/s
Storage system capacity	300 TB	500 TB	1.2 PB
Transfer system to Tier 0 speed	2 × 10 Gb/s	4 × 40 Gb/s	4 × 100 Gb/s

^aDesign value. ^bMain data-taking configuration, excluding links from partition managers used for partitioned running. ^cAt the beginning of the run. ^dAt the end of the run. ^eSlinkExpress. ^f54 links from mezzanine cards with optical SlinkExpress. ^gReadout and builder unit running on the same server (“folded event builder”). ^hRemote DMA over Converged Ethernet. ⁱFilter and builder units running on the same server. ^jNot including the GPU computing power. ^kStorage-area network.

can be included in the FPGAs of the backend boards. Only the footprint of an optical transceiver cage is needed, facilitating the move to smaller form factors such as μ TCA. The SlinkExpress works with 8b/10b encoding at up to 6.3 Gb/s or 64b/66b encoding at 10.3 Gb/s, resulting in an effective bandwidth of up to 5.0 or 10.0 Gb/s, respectively. The data format (definition of headers and trailers) is identical to that of S-LINK64. The SlinkExpress is packet-based and supports re-transmission at the packet level. Packets have a variable size of at most 4096 bytes. Fragments up to 4096 bytes are transferred into individual packets, while larger fragments are split across multiple packets.

At the output side, the FEROL sends data via TCP/IP, employing a custom TCP/IP engine implemented in the FPGA logic. This was achieved through a simplification of the TCP/IP protocol for unidirectional use, which reduced the number of states from 11 to 3 [248]. The TCP/IP streams (one per input) are sent via an Ethernet network to a standard Ethernet NIC in the readout unit (RU) server where they can be received with the standard Linux TCP/IP stack. Using performance tuning as described in ref. [247], a sustained point-to-point throughput of 9.7 Gb/s has been achieved for fragments larger than 1 kB.

For subdetectors that were potentially limited by the copper S-LINK64's bandwidth, but did not upgrade their backend electronics, a new type of mezzanine card with optical transmission was developed. It plugs onto the subdetector readout electronics in place of the original mezzanine card and transmits data using the 6 Gb/s SlinkExpress, increasing the bandwidth to 625 MB/s. A version using 10 Gb/s SlinkExpress is also available. This mezzanine card has been deployed for the ECAL subdetector with 6 Gb/s SlinkExpress links.

In 2017, a new pixel detector was installed with new backend electronics requiring readout through 108 links at 10 Gb/s. Instead of producing more FEROL boards that are powered and controlled through PCI-X because of the legacy interface to the FRL, a new readout board was implemented, based on the μ TCA standard, incorporating the features of four FEROL boards. The new FEROL-40 board [249] receives up to four channels at 10 Gb/s using the SlinkExpress and sends the data on four links of 10 Gb/s Ethernet using TCP/IP.

Table 13 shows the readout parameters of the CMS experiment at the end of Run 1 and Run 2, and at the start of Run 3. Starting in Run 3, the HCAL barrel partitions are read out with higher segmentation, as described in section 5.2.8, resulting in a significant increase in data size. Nine additional readout links were added for the HCAL. The dependence of the data size on the pileup (PU) and the resulting number of vertices has been studied for each data source. Data sizes observed at peak PU in Run 2 resulted in a total event size of 1.4 MB. The event size in Run 3 at the planned peak luminosity of $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ ($\langle \text{PU} \rangle = 56$) is 1.6 MB. An extrapolation, using a polynomial fit, to $\mathcal{L} = 3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, $\langle \text{PU} \rangle = 85$, results in an estimated event size of 2.0 MB. This is still within the design value of the DAQ system, which is therefore considered capable of handling the readout bandwidth at conditions beyond the planned peak luminosity in Run 3.

For heavy ion runs, where the L1 trigger rate will be 50 kHz, the expected event size based on the latest heavy ion run in 2018, taking into account the upgrade of the HCAL readout, is 3.3 MB. The overall throughput at the input to the DAQ system will thus be similar to that in pp runs. The FED sizes (FED can indicate both the subdetector electronic interface and, in this context, the data payload that is read out from that interface) is, however, distributed in a different way. In many cases, a special optimization of zero-suppression and/or selective readout algorithms at the level of the FEDs is applied in order not to be limited by the bandwidth of individual FEDs.

9.4 Event builder

The CMS experiment uses a two-stage event builder (EVB) system responsible for assembling event fragments retrieved from around 760 detector backend boards into a single event payload and delivering the built events to the HLT. The first stage, called the FED builder, reads fragments from the underground FEROL and FEROL-40 boards and, using the switched network, aggregates them in the RU/BU nodes. In the second stage, the core EVB, all the RU/BU nodes transfer fragments from each event into one destination node, assigned on a per-event basis. With the Run 3 DAQ system, each node in this way handles around 2 kHz of fully-built events. On each node, events are written into a large 200 GB RAM buffer and made available to the HLT via a dedicated data network.

Network interconnect technologies are a main driver for the EVB design and scaling. The Run 1 FED builder system [243] was based on the Myrinet 2.5 Gb/s network, comprising 640 RU nodes and split into 8 parallel slices for performance. Prior to Run 2, Infiniband and 40 Gb/s Ethernet became established mainstream standards in high-performance computing (HPC). They

Table 13. Subdetector readout configuration.

Subdetector	End of Run 1		End of Run 2 (start of Run 3)					Data size [kB]	
	# FED	# FRL	# FED	# FRL/FEROL copper	6 Gb/s	10 Gb/s	# F40 ^a	Run 2 ^b ⟨PU⟩ = 56	Run 3 ^c ⟨PU⟩ = 56
Tracker pixel	40	40	108	—	—	—	32	259	
Tracker strips	440	250	440	250	—	—	—	731	
Preshower	40	28	40	26	—	—	—	54	
ECAL	54	54	54	—	54 ^d	—	—	74	
HCAL	32	32	32 (41)	—	—	32 (41)	—	170	+221
Muons CSC	8	8	36	18	—	—	—	40	
Muons RPC	3	3	3	3	—	—	—	0.3	
Muons DT	10	10	9	—	—	9	—	22	
Trigger	5	5	14	—	—	14	—	41	
CASTOR	4	4	4 (0)	4 (0)	—	—	—	—	
TCDS	—	—	1	—	1	—	—	1.0	
CTPPS	—	—	11	9	—	2	—	2.2	
Muons GEM	—	—	2	—	—	2	—	— ^e	+22
Total	636	434	754 (759)	310 (306)	55	59 (68)	32	1.39 MB	1.63 MB

^aFEROL-40 (4 × 10 Gb/s). ^bObserved at Run 2 peak luminosity. ^cOnly the increment in data size with respect to Run 2 at the same peak luminosity is shown. ^dMezzanine card with optical SlinkExpress. ^eNot included for regular data taking.

were evaluated [244, 250] and ultimately adopted [245, 251] as technologies of choice for the Run 2 DAQ upgrade. This allowed the system to be scaled down by an order of magnitude and implemented in a single slice, while increasing the overall EVB throughput [252].

For Run 3, the Run 2 equipment, which had reached the end of its vendor support, had to be retired. The design choice was taken to implement the EVB with 100 Gb/s Ethernet for the Run 3 system. This system, described in the following sections, satisfies nearly identical performance requirements as the Run 2 DAQ, and, owing to hardware evolution, is reduced in scale by more than a factor of two. Some design details of the preceding Run 2 system are also outlined.

9.4.1 FED builder

The FED builder transports data fragments over 10 Gb/s links from the underground cavern, over a distance of around 200 m, to data concentration network Ethernet switches located on the surface. A total of 557 links from the FEROL and FEROL-40 boards are used in a typical Run 3 data-taking configuration, as detailed in table 13.

The data concentration network is implemented using a Juniper QFX10016 [253] chassis-based 100 Gb/s Ethernet switch, providing a flat network, which allows full flexibility in routing the fragment traffic between readout nodes. For inbound traffic from the FEROLs, seven QFX10000-36Q line-cards are installed, capable of supporting a total of 672 10 Gb/s links, split from 24 × 40 Gb/s ports per line-card. Three QFX10000-30C cards installed in the switch provide a capacity of up to ninety

100 Gb/s ports for connections to the RU/BU nodes. A pause frame [254] flow control mechanism is used to achieve a lossless low-latency network and ensure optimal buffer occupancy on the FEROLs. The system is configured to support Ethernet jumbo frames of up to 9 kB to improve performance.

9.4.2 Core event builder

The core EVB has three distinct functional units, the RU, BU, and event manager (EVM). They are implemented as software applications and run on a single set of RU/BU computers, interconnected using a dedicated event-building data network. These applications are implemented within the XDAQ framework described in section 9.10.

The RUs unpack and merge fragments received via TCP streams from the FEROLs over the data concentration network. Each fragment checksum is verified before being merged with others into a super-fragment, i.e., a collection of fragments from the same event aggregated by a particular RU, and buffered in memory. The EVM, which runs on a single EVB node, orchestrates the event building process by performing a destination assignment, i.e., allocating a given event to a specific BU and sending a message to the RUs to send super-fragments to that destination. When the BUs receive a super-fragment from each RU the event is completely built. Completed events are written into ramdisk, a memory buffer structured as a file system, as described in section 9.5. This file system is exported through the event backbone network to the FUs. An FU group of typically three or four FUs is assigned to read the event data from each RU/BU node.

Starting in Run 2, EVB senders and receivers were implemented using the Infiniband Verbs-API, described in section 9.10, for communication over the network using remote direct memory access (RDMA) capabilities of network-interface cards and switches. This technology facilitates the bypass of the operating system networking stack and delivers the payloads directly into the application-accessible memory, avoiding CPU and memory overheads associated with the handling of high bandwidth and packet rate networking in software.

For the Run 3 EVB, the cost-effective 100 Gb/s Ethernet technology was chosen over the native Infiniband used in Run 2. A second Juniper QFX10016 chassis-based switch is used as the core EVB network backbone (equipped with nine QFX10000-30C line-cards having a total of 270 100 Gb/s ports). The switch supports a lossless Ethernet network, a prerequisite for RDMA over Converged Ethernet (RoCE) v2 protocol, providing encapsulated Infiniband protocol over Ethernet hardware. Importantly, like native Infiniband, it facilitates direct memory access between communicating nodes using offloading by NICs. A benefit of using RoCE v2 is that the Verbs-API-based EVB applications developed in Run 2 can be reused with minor adaptations.

In addition to the core EVB network, the switch runs the event backbone network, which handles the traffic between the EVB and HLT farm. The same chassis switch supports also the storage and transfer system (including transfers to Tier 0), traffic from the DAQ to the online data quality monitoring system, and running the online cloud on some of the legacy Run 2 HLT nodes. The 62 RU/BU nodes are connected to the EVB network over 100 Gb/s optical links, including 57 nodes as part of the nominal data-taking configuration and five as hot spares. In addition, there are seven cold spare nodes without network interface cards.

As pointed out above, the Run 3 system also introduced a folded core EVB setup, where the RU/BU nodes serve as both RU and BU functional units (RU/BU). This allows the bidirectional utilization of the network links, nearly halving the number of needed RU/BU nodes and the network

bandwidth compared to a nonfolded setup, thus reducing the overall size and cost of the system for the same throughput requirement. On the other hand, this design is demanding on the I/O and memory performance of the nodes, which in such a system need to receive and merge fragments, exchange super-fragments with other RU/BU nodes, build and serve full events, and pass-through the HLT output to the STS, all over multiple 100 Gb/s interfaces. Therefore, a modern server architecture was required for the task. Dell R7515 servers, with the 32-core AMD EPYC Rome 7502P CPU running at 2.5 GHz, equipped with 512 GB of DDR4 RAM, were selected, with all memory channels populated for the maximum memory performance. Each server is equipped with two Mellanox ConnectX6 dual-100 Gb/s PCIe Gen4 NICs. One card has both links used for the respective connections to the FED builder and event backbone networks, while the second card uses one link for connection to the core EVB network.

9.4.3 Performance

Detailed studies were done comparing the EVB performance of both the Intel Xeon Skylake dual-socket servers [144] and comparable AMD EPYC Rome single-socket platforms. The AMD platform was ultimately chosen for RU/BU nodes for Run 3 due to better performance and a simplified memory architecture. The AMD 7502P CPU is internally assembled from multiple 4-core silicon dies (CCDs), each having a separate L3 CPU cache and an interconnect fabric providing eight memory controllers, as well as multiple PCIe Gen4 interfaces. Due to these characteristics, the CPU internally resembles, to an extent, the nonuniform memory access (NUMA) memory architecture. It can be configured in a 4-, 2-, and single-node NUMA mode with respect to the CPU die and memory controller topology.

The single-node NUMA mode uses interleaved access to the RAM controllers with improved maximum bandwidth, with the drawback of potentially higher memory and I/O latency. In evaluations, this mode was found to be a well-balanced setup for folded EVB requirements, avoiding the delicate tuning of thread affinities to CPU cores, which would have been needed to optimize performance. The setup maintains CPU affinity settings for groups of threads running similar tasks to the same or adjacent CPU cores and placing them closer to the corresponding NIC PCIe lines in the interconnect topology.

A memory-based file system based on Linux tmpfs [255] is used as the EVB event output ramdisk. To achieve high read/write performance, 2 MB huge-page support for tmpfs is used, available in recent Linux kernel versions. It significantly improves the ramdisk throughput since it avoids bottlenecks when the CPU is handling a large number of 4 kB memory pages.

Measurement of the DAQ event building performance with emulated data generated by the RUs and discarded after event building, is shown in figure 118. This demonstrates that the EVB system is capable of handling the nominal 100 kHz event rate from the L1 trigger with event sizes up to ≈ 2.5 MB. A throughput of approximately 10 GB/s is achieved per RU/BU node at the plateau, which is 80% of the available network bandwidth.

The performance for input generated in the FEROLs and either discarded after the EVB or passed through the full-chain DAQ system (including the HLT and STS) is shown in figure 119. Two setups of a smaller scale than the data-taking setup were used for this measurement, comprising either four or 19 RU/BU servers in a folded setup, with an additional machine running an EVM application. Each RU application was concentrating 20 fragments from the FEROLs (two fragment streams per 10 Gb/s FEROL link, or a total of 100 Gb/s of the input network bandwidth per RU). A total of either 16 or 76 FUs, four assigned to each RU/BU, were used. These servers were Run 2




Figure 118. Core event building throughput (left) and event rate (right) of the full-scale RU/BU EVB setup shown for a range of built event sizes. Emulated input data are generated in the RU/BUs and discarded after the event building stage.

FU nodes equipped with 10 Gb/s Ethernet NICs. Thus, the total available bandwidth was 40 Gb/s over the event backbone network to the FU group. With each system, a plateau throughput of about 9 GB/s was achieved per RU (BU), with discarding the data after the EVB, amounting to over 70% of the available network bandwidth from the FEROLs. About 4 GB/s per server was achieved with the full-chain DAQ, which is about 80% of the available network bandwidth to the FUs.




Figure 119. The DAQ throughput (left) and event rate (right) per RU/BU node for a range of uniform fragment sizes using a mode with the data discarded after the EVB and with the traffic flow through the HLT and STS. Emulated input data are generated at the FEROLs with 20 fragments concentrated per RU/BU.

This result demonstrates that the system is capable of handling traffic for fragment sizes larger than 2000 bytes at rates above 100 kHz. When extrapolated to the larger system of 50 RU/BUs (nearly the scale of the system used in Run 3), a fragment range around 2000 bytes corresponds to the total event size of over 2 MB, which is above the size expected in Run 3. With the fixed number of FUs per RU/BU, the achieved bandwidth approximately scales from the small- to the medium-size

system and gives confidence that the system comprising over 50 RU/BU nodes can robustly scale and handle the Run 3 requirements. Furthermore, while tests were performed using a 10 Gb/s network bandwidth per FU, the Run 3 FUs have more bandwidth available from the 25 Gb/s Ethernet links, as mentioned in section 9.5.5, and, therefore, higher full-chain throughput capacity.

9.4.4 Event builder load balancing

The EVB protocol implements a global dynamic re-balancing mechanism of the event building throughput on the RU/BU nodes. The assignment algorithm, running on the EVM, periodically determines the load status of each node, depending on factors such as the local output buffer occupancy, as well as the number and type of HLT computers attached to a particular RU/BU. From these, it determines the rate at which each node can build events to fairly distribute the load and equalize the processing latency at the EVB and HLT across the cluster. A local throttling mechanisms on the RU/BUs can also get triggered, based on occupancy thresholds and latency information reported by the HLT infrastructure software. This results in the BU application temporarily blocking or throttling the event building at the corresponding RU/BU.

9.5 Event filter

The HLT performs the second stage of event filtering, analyzing at a rate of approximately 100 kHz, the events selected by the L1 trigger, and accepting events at a rate on the order of a few kHz. It runs on a large cluster of multicore servers (FUs) that perform the reconstruction and filtering in software, implemented within the CMS software framework (CMSSW) [256].

By using the same framework and code base for the online and offline reconstruction, most of the algorithms developed in CMSSW for the offline reconstruction can also be employed in the HLT. This enables the reuse of the same data structures, supports a unified approach to the use of detector condition data, and allows for rapid deployment of both newly developed and established offline algorithms.

Starting from the first year of Run 2, the CMSSW and HLT implemented task-based multithreaded event processing. This feature was instrumental in reducing the memory requirements of the HLT: by running a smaller number of processes on the FUs, each using multiple threads, common data such as detector conditions and calibrations can be shared among threads, decreasing the overall memory footprint on the servers. As a result, this made it possible to fully exploit all the logical cores of the CPUs, leading to increased overall computing efficiency. At the beginning of Run 3, support for offloading to accelerators, such as GPUs, was also added. In addition, reducing the number of running processes per server was found to be beneficial in reducing the memory footprint on the GPU.

The framework architecture is described in more detail in section 12. The HLT algorithms and offloading are detailed in section 11. A dedicated software infrastructure couples the CMSSW workflow with the EVB and the storage and transfer system. It handles the data input and output for all these processes, as well as the bookkeeping of event processing and monitoring.

9.5.1 File-based filter farm

The HLT interface to the DAQ, running in the same process as the HLT algorithms, is implemented in ordinary CMSSW modules and services, which facilitate the communication and reading of input data from the RU/BU nodes. The data transfer into the HLT starts from a ramdisk file system serving

as a data buffer on the RU/BU [257]. An NFS (network file system) v4 server provides access to the file system over the network. The ramdisk is sized to buffer the data during the HLT startup time of 60–90 s, dominated by the loading of the software libraries, calibrations, and conditions, as well as to absorb fluctuations in the HLT event processing time.

The BU application writes fully built events into the ramdisk, using a custom binary file format with a header prepended to each event describing its total length and providing information such as an identification number and a checksum. The files created by the BU application, each containing around 100 events, are distributed among CMSSW processes running on several FUs assigned to the RU/BU node. A dedicated application running on the RU/BU nodes, called the file broker, is responsible for the unique assignment of files to the processes. The requests for files are placed with the file broker using an HTTP-based protocol, while the standard file system access API is used to read files from the NFS. All input-related communication and file access is implemented in the CMSSW input source module written in C++. This module operates several dedicated threads for pre-buffering of the input data to minimize stalling the CPU cores while waiting for input. The input source parses the assigned input files, extracts and verifies the event payload against stored fragment checksums, and hands it over to the framework, which schedules the HLT modules performing the event reconstruction, analysis, and filtering.

9.5.2 The HLT software infrastructure

The HLT daemon is a system service orchestrating the HLT operation in the RU/BUs and FUs. Implemented as a Python application, it relies on a file notification API of the Linux kernel, `inotify` [258], which provides low-overhead monitoring of file system events. The start of a run is triggered by the appearance of raw event data files in the ramdisk. Similarly, a run completes when all files have been processed by the HLT. The service runs on all RU/BUs, where it is responsible for contacting the attached FU nodes to process a run, as well as the FUs, where it is responsible for managing the life-cycle of the locally running CMSSW processes. The HLT daemon starts the CMSSW applications as forked child processes, monitors their status, and can restart them in case of unexpected process termination. Each child CMSSW process is allocated a number of CPU cores or hyperthreads (in the case of using simultaneous multithreading (SMT), a technique allowing separate tasks to run on the same CPU core), also taking into account the number of parallel threads configured in the CMSSW processes, as also described in section 11. By using such an allocation strategy, servers can be used at their maximum computing capacity.

9.5.3 The HLT menu and output data streams

The HLT configuration, also referred to as the “HLT menu” and described in section 11, defines a set of paths in which physics objects are reconstructed and events are filtered based on specific physics requirements. Events that are accepted through HLT paths for similar physics processes are stored in primary data sets, defined also in the HLT menu. The primary data sets are defined such that the total event rate of each data set is kept within the limits imposed by the offline data processing, while minimizing the overlap among different data sets. In turn, the primary data sets are grouped into streams, corresponding to the actual files that are output by the HLT processes and transferred to the Tier 0 for offline processing. In addition to those used to collect events for physics analyses, the HLT defines dedicated paths and streams to collect data for detector calibrations and for online

data quality monitoring (DQM). Finally, special paths and streams are used to collect statistical information about the HLT processes themselves, such as the individual trigger rates and CPU usage.

Events selected by the HLT are serialized using ROOT [259], compressed using gzip [260], LZMA [261], or Zstandard [262], and written to files on the local partition. Lossless compression used for the HLT output, applied to the entire serialized event payload, reduces the total event size to about 70% of the input event size for proton runs (specified in table 13), as observed in Run 2 and the first year of Run 3. For proton runs, it is preferable to minimize the CPU usage of the compression in order to preserve resources for the HLT reconstruction. The gzip tool, which was used in Run 2 and in 2022, typically used up to on the order of 1% of the CPU in the HLT during Run 2 and 2022.

The application of Zstandard was evaluated after the first year of Run 3 and shown to have improved performance for a compression factor comparable to gzip, with an estimated use of less than 0.5% of the CPU. The HLT began using this algorithm for compression in 2023 for proton runs. For the heavy ion run, the compressed size is estimated at about 3 MB for centrally colliding collision events and 1.2 MB for minimum-bias events, based on the 2018 heavy ion run. The LZMA algorithm, which can typically improve the compression factor by an additional 10–15%, but is an order of magnitude more CPU demanding than gzip or Zstandard, is being considered in order to maximize the HLT output rate within the DAQ bandwidth limits. The approaches to size reduction of heavy-ion events are discussed in section 11.6.

Each output stream is written into a separate file and includes the event raw data, a selection or a summary of the HLT objects, and the trigger decision for each path specific to a stream. The streamer file format used for output allows files belonging to the same stream, but produced by different processes and nodes, to be trivially concatenated. This step, which is performed separately for each stream, is called merging. In the first merging stage, streams are merged from local CMSSW processes and copied to the NFS-mounted output partition on the RU/BU. Subsequent merging stages are handled by the storage and transfer system (STS), as described in section 9.6.

The CMSSW input, output, and merging operate on sets of data and metadata files created separately for each stream at a regular time interval called the luminosity section (LS), spanning 2^{18} LHC orbits, which takes approximately 23.3 s as explained in section 8. Metadata, written by the BU, by the CMSSW processes, and by the HLT output merging at LS intervals, specify information such as event counts, checksums, and total output size. Bookkeeping (completion) and integrity checks are performed at the granularity of an LS by comparing input and output metadata. In the case of a failed integrity check, alerts are raised to notify the shift crew, and the affected data are discarded and marked as missing in the bookkeeping.

The HLT daemon services handle all the local output collection and merging tasks on the FUs. They also report the latency in event processing and merging to the BU application. The BU uses this feedback to throttle event building in case of high delays and can activate throttling in case of high ramdisk occupancy.

9.5.4 Monitoring

The file-based filter farm (F3) monitoring system was designed from the ground up in Run 2 [242] around the Elasticsearch NoSQL database [263]. Structured JSON-formatted [264] monitoring data are injected into an Elasticsearch cluster from the HLT daemon service, STS, and other sources. This includes bookkeeping such as input and output event counts, event file size, bandwidth, CPU

usage, HLT information, and the hardware state at the granularity of each FU node. Textual logs and error reporting from the HLT processes are also handled. Information is stored for an extended time of up to several years, allowing real-time as well as follow-up analysis of performance and faults in the system. Several web-based tools have been developed as frontends to the system, visualizing the information for experts and the shift crew. Monitoring of the HLT performance is also collected using non-event streams, such as histograms, which are merged and shipped to the DQM system, and trigger statistics, which are merged and stored in a database.

9.5.5 Evolution of the HLT farm

The F3 was gradually expanded almost every year during Run 2, as the HLT processing requirements increased due to evolving LHC and detector conditions. The FU nodes are typically replaced with a new generation of hardware after the end of their 5-year warranty periods, with the old nodes being assigned to the online cloud, as described in section 9.9.2. Table 14 summarizes the composition of the HLT farm at the end of Run 2 and at the beginning of Run 3, along with the computing power estimates based on the HEPSPec 2006 [265] (HS06) benchmark measurements. Prior to the start of Run 3 data taking, the entire farm was replaced with new AMD CPU nodes, each equipped with two NVIDIA T4 GPUs [266]. In this configuration the average processing time of the HLT can be up to 500 ms per event, at the nominal L1 rate of 100 kHz.

Table 14. Summary of the HLT filter farm unit specifications, thermal design power, and performance based on HS06 in the final year of Run 2 and first year of Run 3.

Run	Run 2 (2018)			Run 3 (2022)
Architecture	Intel Haswell	Intel Broadwell	Intel Skylake	AMD Milan
CPU model	dual E5-2680v3	dual E5-2680v4	dual Gold 6130	dual 7763
CPU cores	2×12	2×14	2×16	2×64
Nominal freq. [GHz]	2.5	2.4	2.1	2.45
Turbo freq. [GHz]	3.3	3.3	3.7	3.5
TDP [W]	120	120	125	280
Memory [GB]	64	64	96	256
Nodes	360	324	400	200
CPU cores (total)	8640	9072	12800	25600
HS ^a /node	538	659	773	3224
TDP [W]/kHS ^a	223	182	162	87
kHS ^a	194	214	309	645
GPU card	—	—	—	2×NVIDIA T4

^aHS06 measurements only take into account the CPU performance. The precision is around 1%.

Specific HLT reconstruction algorithms, such as those used in data scouting, described in section 11.4, can be offloaded to these GPUs, reducing the average HLT event processing time by over 40%, as shown in figure 134 and discussed in section 11. In comparison to an HLT farm

equipped only with CPUs, this corresponds to a reduction in the farm’s overall cost by approximately 15% and its power consumption by 30%.

On F3 nodes, CMSSW jobs are divided into two groups, and, for each group, the CPU and memory affinity is pinned to a single CPU socket NUMA domain on dual-socket AMD nodes. Each GPU is attached to PCIe lanes on an individual CPU socket and assigned to the group of processes running on the same socket. The NVIDIA multiprocess daemon is used to schedule GPU access between processes, providing a small performance enhancement over direct access. A total of eight jobs with 32 threads and 24 streams each (i.e., parallel event processing pipelines in CMSSW) is used per node, since this configuration was found to fit within the memory capacity of the nodes and GPUs. It was determined that increasing the number of threads per process in CMSSW adds a small additional overhead, and thus it was not further increased to keep the optimal processing capacity of the F3.

To support NFS data transfers in Run 3, a Juniper QFX5120 top-of-the-rack (ToR) switch is employed, which uses 8×100 Gb/s up-links to the event backbone network from each rack. The network contains approximately 40 FU nodes with 25 Gb/s Ethernet connections to the ToR switch. A flat interconnect network is provided, enabling any FU to communicate with any RU/BU node.

9.6 Storage and transfer system

The storage and transfer system is the last stage of the DAQ data flow. It collects the HLT output from each RU/BU and associated FU group and writes it to a cluster file system. It later transfers the data to Tier 0 for repacking from the streamer format into the ROOT format and to permanent storage on disk and tape servers.

The HLT daemon services on the FUs within the FU group concatenate output streamer files, analyze metadata, and copy the corresponding files into a dedicated RU/BU output partition. This was a spinning disk RAID array in Run 2, and is a 200 GB ramdisk partition in Run 3 to support higher throughput. The merger service running on the RU/BU nodes periodically polls the streamer and JSON metadata files in this area, written by each FU for a particular stream and luminosity section. The JSON metadata files are used to verify the completion of per-LS output. Once all FUs have copied a complete set of files, the service distributes tasks to several worker threads to read the output of all streamer files and append them into a single file location in the distributed file system. For most streams and a majority of the output bandwidth, this step is done by simultaneously writing into a single file at a different offset. This technique is used to efficiently merge data into the final file object, one per stream and LS. This file can be transferred by a single copy operation to Tier 0, instead of requiring an additional read and write operation to perform such a concatenation.

The final stage of merging is performed on a set of dedicated STS nodes, with each handling a subset of streams. For most streams this amounts to verification of the metadata and checking the completion of files written by the merger service from each RU/BU. Fully merged streamer files, residing in a cluster file system, are handled and transferred to their destination by the transfer service, also running on the same nodes as the merger service. For data destined for the Tier 0, a pool of threads starts the file copy jobs to transfer data to the EOS disk system [267] using a high-speed link to the central data recording (CDR). The transferred file size is kept below 16 GB for optimal transfer throughput performance. The network infrastructure used to transfer to the Tier 0 is detailed in section 9.9. Streams are also delivered to several other destinations, such as the DQM and the

calibration cluster, or are parsed to extract the HLT and L1 trigger rate monitoring information and inject it into the relational database.

Extensive bookkeeping is required to track all the files passing through the system. This information is injected and visualized in the F3 monitoring system. Metadata relevant for the transfers are also written to an SQL database for bookkeeping and to provide transfer metadata to the Tier 0.

For Run 2, Lustre [268] was selected as the cluster file system, after initial evaluation in the DAQ integration system. The production system consisted of two disk servers (OSS) and a metadata server (MDS), providing on the order of 9 GB/s of total read and write throughput, with around 500 TB of storage space provided by a redundant hard drive setup.

In Run 3, the storage system requirements are driven mainly by the heavy ion running, which, apart from centrally colliding lead ion events, aims to collect a large amount of minimum-bias events. An estimated throughput of 17 GB/s is required in a scenario that includes a trigger selection of around 1 kHz of centrally colliding heavy ion events and a minimum-bias rate of 10 kHz at the HLT output, as described in section 9.5.3.

To cover these requirements, a hardware refresh was pursued in Run 3, retaining the same file system technology while significantly improving the bandwidth capability compared to Run 2. A new system comprising two DDN EXAScaler [269] SFA7990X data servers and a single SFA400NVX metadata server was acquired for the task. Each data server consists of 124 SAS 8 TB 7.2k RPM hard disks, organized in a RAID6 array. The metadata server comprises 23 SSDs, each of 1.8 TB capacity, in a RAID6 array. The system is connected via the chassis-based Ethernet switch also used for the EVB and HLT. Seven STS nodes were additionally installed and are connected to the chassis-based Ethernet switch. Together with the RU/BU nodes, these nodes are set up as Lustre clients and the file system made accessible through mount points. The Run 3 Lustre system storage space was scaled to provide several days of storage space for proton LHC runs or up to a day for heavy ion runs should the Tier 0 connection fail. The usable disk capacity is 1.2 PB, and the system is capable of simultaneously writing 24 GB/s and reading 11 GB/s using standard file system benchmarking tools. The TCP/IP protocol was used for communication in these tests. The system is capable of temporarily prioritizing the write bandwidth at the expense of the read bandwidth, up to the limit of the available disk space. This is particularly useful for runs with high peak bandwidth, such as heavy ion runs, where the read rate recovers and allows draining the accumulated files towards the end of the LHC fill and in interfill periods.

9.7 Trigger throttling system

The trigger throttling system (TTS) collects fast readiness signals from all FEDs, merges them per TTC partition with a priority logic, and makes them available to the trigger control logic to avoid buffer overflows by inhibiting triggers or drive recovery actions. While in Run 1 this trigger control logic was implemented in the trigger control system (TCS), since Run 2 it is part of the trigger control and distribution system (TCDS), described in section 9.8. A FED may signal the following main TTS states: Ready (to accept triggers), Warning/Busy (buffer fill level above high-water mark), Out-Of-Sync (synchronization loss), and Error (other error situation), in ascending order of priority. The TCDS reacts to the highest-priority TTS state by inhibiting triggers, executing a re-synchronization sequence, or executing a reset-sequence.

For the legacy FEDs, TTS signals are sent using four LVDS pairs, which are merged using compact-PCI based fast merging modules (FMMs) [246] that combine the TTS states of up to 20 FEDs using a priority logic. For partitions with more than 20 FEDs, the FMMs are arranged in a tree structure. For upgraded FEDs, TTS signals, transmitted over optical fiber, are merged with a similar logic by μ TCA-based TCDS partition interface (PI) modules, as described in section 9.8.

9.8 Trigger control and distribution system

At the start of Run 2, a new TCDS [270] was introduced, replacing the Run 1 trigger control system that was integrated into the L1 trigger system, and the Run 1 TTC system. The TCDS provides support for an enlarged range of detector partitions and for detector backends sending their trigger throttling signals over optical fiber. These were both needed to integrate the additional and upgraded subdetectors with μ TCA backends during Run 2. The TCDS distributes timing and control (synchronization) data that are flowing to the detector backends and frontends and receives back status information related to the readiness of the detector systems to handle more triggers. The clock reference that is distributed along with the fast control information is synchronous with the beams in the LHC, which is required to keep the data taking in step across the various detector systems. The TCDS is implemented in the μ TCA architecture.

As illustrated in figure 120, a central crate contains a central partition manager (CPM) board and up to twelve local partition manager (LPM) boards. The CPM receives the LHC clock from the TTC machine interface and the L1 accept signal from the global trigger, as described in section 10.3). Each of the LPM boards contains eight independent integrated CMS interface (iCI) logic blocks that are able to control a detector partition for local running. The iCI blocks translate the generic TCDS synchronization commands to subdetector specific commands. The iCI block also contains a partition-specific emulator of the APV25 tracker readout chip buffer levels that inhibits triggers that would lead to overflows of these buffers. Each LPM board also contains two partition manager (PM) blocks that can orchestrate combined runs with sets of partitions in the same LPM. The CPM contains one PM block orchestrating global runs with all partitions. The PM blocks provide the following functionality:

- Trigger throttling, taking into account:
 - the TTS state of partitions;
 - trigger rules, suppressing bursts of triggers by limiting the number of triggers in certain windows of bunch crossings, as required by the subdetector frontend electronics;
 - protection against overflows in the pre-shower frontend ASIC (in a similar way to the protection against overflows in the APV25);
 - resonant trigger protection, i.e., protection against triggers arriving at regular intervals over prolonged periods, which could give rise to resonant vibrations that might damage systems (such as wire bonds),
 - the DAQ back pressure to the PM's readout link.
- Bunch mask trigger veto that can be used to inhibit triggers (such as prefiring triggers) in LHC bunch crossings that are not filled.

- Generation of random triggers.
- Generation of (periodic) calibration triggers to trigger the readout of calibration events (interleaved with physics events).
- Definition of sequences of synchronization commands such as run start, run stop, and detector backend recovery.
- Automatic triggering of sequences based on the trigger throttling state, for example, for global re-synchronization or subsystem-specific recovery actions.
- Monitoring of trigger rates, sequences, and dead times.

A secondary set of local and central partition manager hardware, connected in parallel to the primary partition managers, is available for final validation of firmware and/or software updates, or as hot spares.




Figure 120. Overview of the CMS trigger control and distribution system (TCDS).

Each of the iCI blocks is connected to a partition interface (PI) board, fanning out triggers and synchronization commands to a single detector partition and merging the throttling signals from the partition. These boards are located in separate μ TCA crates. For the bulk of the backend boards, the PI board receives the TTS signal from an FMM, as described in section 9.7, while triggers and synchronization commands are sent via backwards-compatible TTC fibers and the legacy TTC hardware. For upgraded subsystems, the PI sends triggers and synchronization commands to up to

10 backend boards via bidirectional optical links. In the reverse direction, it receives TTS signals from the backends and merges these signals with a priority logic similar to the FMM.

9.9 Networking and computing infrastructure

The operation of the CMS experiment is supported by the CMS service network, a high-performance distributed network connecting all the computers directly related to the operation of the experiment. While isolated from the CERN general purpose network (GPN), connections to the CDR and LHC technical network (LHC-TN), as well as the GPN, are allowed for use in specific cases. The network follows a redundant design established in the early years of the experiment. Currently, routers provide 1 Gb/s Ethernet connections to the servers. Available bandwidth is commonly used for user access, control, and monitoring. The majority of the DAQ computers are located in racks, to which the service network is provided using ToR switches.

The DAQ data networks have been significantly redesigned and upgraded throughout the lifetime of the experiment, as described previously in this section. To facilitate sending data taken by the experiment in Run 3 to the Tier 0, four 100 Gb/s links from the event backbone network to the CDR are employed. This facility is also used by the online cloud for accessing external data stores and services, as described in section 9.9.2.

A network-attached storage (NAS) system is used by CMS for core storage needs. This includes home directories and subsystem storage, exposed via standard remote file system protocols to nodes in the CMS network. The system acquired for Run 3 provides 1.3 PB of storage space. The CMS online computing system runs the CERN-supported version of Linux, CERN CentOS 7 (CC7) [271], as well as Red Hat Enterprise Linux 8 (RHEL8) [272]. Both of these operating systems are also used for the DAQ operation. Machines are initially installed (or reinstalled) using a network preboot execution environment (PXE) boot installation service. After Run 1, the Quattor [273] system was replaced with the Puppet [274] software configuration management tool, which handles the OS installation, configuration, and deployment of the online software in distributed and reproducible fashion on thousands of DAQ and subsystem computers.

9.9.1 Virtualization

Virtualization allows the reuse of physical computers by sharing them for multiple services that are running in virtual-machine (VM) instances, while simultaneously providing an isolated execution environment for each service. A virtual-machine infrastructure has been set up and used to run the CMS online services based on the oVirt [275] open-source virtualization management platform on the CentOS 7 OS. The VMs running on this infrastructure are installed from OS boot images provided on the network and using the Puppet configuration management system described above. The VMs are accessible on the service network via dedicated network names assigned by the DNS.

9.9.2 Online cloud

The HLT farm consists of a large number of multicore nodes, representing significant computing power. During interfill periods, week-long technical stops, and longer LHC and detector upgrade periods, this capacity is mostly unused for the HLT tasks. Thus, starting in Run 2, infrastructure was developed [276] to run an Openstack-based cloud overlay to use the computing resources for

offline production jobs. In order to not overlap with the HLT-based workflows, the FUs are able to run virtual machines (VM) that provide the necessary software environment for the worldwide LHC computing grid (WLCG) CMS jobs. These jobs run tasks such as simulation and reconstruction in full isolation from the HLT. In section 12, more details on the WLCG and workflows are provided. An API is implemented in the HLT daemon to facilitate the automatic suspension and startup of the cloud mode or resumption of the HLT. The cloud images can be suspended to disk and later resumed, allowing the quick save of unfinished jobs when, for example, the cloud needs to be suspended based on the LHC status. The system can also perform the switch automatically, reacting to the LHC state or, in another mode, monitoring the HLT CPU usage and dynamically re-allocating unused fractions of the HLT to the cloud. About 5% of the farm remains in HLT mode at all times to provide an operable HLT for cosmic ray data taking, commissioning, and tests.

Openstack VMs use a virtual local area network (VLAN) for access to the CERN services such as EOS, where the job input data, as well as the destination of job results, are located. The VLAN is implemented in the event backbone network and routed through CDR links to the CERN IT infrastructure.

The overlay cloud has been extensively used since Run 2 and successfully expands the CMS computing resources through opportunistic re-purposing of the hardware. Furthermore, it consists not only of the active HLT cluster, but also computing nodes that were retired from the HLT. They are kept in cloud operation as long as the computing infrastructure support is possible. Overall, the online cloud, which is classified as a Tier 2 CMS site, as described in section 12, is one of the major contributors to CMS offline computing, comparable to the largest Tier 1 sites in the amount of CMS production workload.

9.10 Software, control, and monitoring of the DAQ

Two software frameworks have been developed within CMS to implement the bulk of the experiment's online software: the C++ based XDAQ framework [239], used to implement hardware access and data transport, and the Java based run control and monitoring system (RCMS) [73, 277], used to implement the hierarchical control structure and main user interface. These two frameworks, which have been adopted by the central DAQ system and by all subdetectors, are maintained and enhanced according to the evolving requirements of the experiment and continue to be used in Run 3.

The XDAQ software is a platform designed specifically for the implementation of distributed DAQ systems. It has a layered middleware structure, providing support for communication, a web user interface, high-speed networking, hardware access, multithreading, performance tuning, monitoring, error reporting, and logging. The XDAQ system builds upon industry standards, open protocols, and libraries, e.g., TCP, HTTP, XML, and Apache Xerces. Notable enhancements to XDAQ include the development of new “peer transport” plug-ins to support new network technologies, such as RDMA using Infiniband Verbs [278, 279], and a new service-based approach to the configuration of the built-in monitoring and alarming infrastructure [280]. New XDAQ-based applications have been added to control new types of custom hardware and existing applications, such as the event builder enhanced with features such as load balancing and fault tolerance.

The RCMS is a framework based on web applications running inside container instances (Apache Tomcat), which provides the building blocks to compose a distributed hierarchy of nodes to control and monitor the state of XDAQ applications and other online applications used during data

taking. Control nodes are based on state machines, with system-specific control logic implemented in Java. They are steered and monitored through web user interfaces. The RCMS includes database schemas to hold the configuration of all software components, define hardware configurations, and manage the complex interconnects required for the two-stage event builder. Extensive tools for configuration management are available. In the area of the RCMS, the development of guidance systems for operators and automation have helped to make the operation of the experiment less error-prone and less reliant on the knowledge of experts. The overall configuration of the experiment can be automatically selected based on the state of the LHC, actions needed in response to certain LHC state changes, and actions like high-voltage ramping in the detectors that are performed automatically [277]. Recovery from regular single-event upsets and from other typical data-taking problems is fully automated [281, 282]. Configuration management tools have been enhanced according to new requirements in Run 2 and Run 3 supporting, for example, fine-grained data-flow optimization according to the network topology.

The filter farm and the storage and transfer system are controlled and monitored by online software based on Python and Elasticsearch, as described in sections 9.5.2 and 9.5.4.

A number of monitoring clients transform the raw monitoring data from both the XDAQ monitoring system and the Elasticsearch-based monitoring system into web-based graphical and textual monitoring displays used by the shift crew. The monitoring clients typically display instant data with a latency of a few seconds and can also be used to browse historic data to facilitate post-mortem analysis. The Java-based DAQ expert tool [283–285] detects all common data-taking problems by evaluating rules encapsulated in logic modules using snapshots of monitoring data. This helps the shift-crew with the sometimes difficult task of pin-pointing the cause of data-flow problems. With problems for which a recovery is known, the tool can drive completely automatic recovery actions. It consists of several micro-services responsible for reasoning, notification, and control of the recovery.

A switch monitoring system with a web-based graphical representation was developed to monitor the link status and performance metrics of the DAQ data networks and assist experts in diagnosing network-related failures.

In addition to specific aggregation and presentation tools for the DAQ and HLT, a general service is provided to the subsystems and collaboration at large to aggregate and present online monitoring data stored in the different databases. The online monitoring service (OMS) is a new Run 3 software tool replacing a set of web-based monitoring tools (WBM [286]) used in Run 1 and Run 2 to provide unified remote access to the monitoring data. The OMS uses a generic relational database model (Data Warehouse) and interface, and a web-based presentation framework by which information across heterogeneous data sources and formats is aggregated and presented. The presentation is organized in a structure of folders and pages that contain portlets typically displaying information in the form of tables and graphs, showing, for example, run and fill details, trigger rates, or subdetector monitoring.

9.11 MiniDAQ

In addition to the global DAQ system, self-service DAQ systems, called MiniDAQs, are provided for most of the CMS subdetectors. These setups can be used at any time by the subdetector groups for calibration runs, tests, and debugging using detector partitions that are not participating in the

global data taking. These setups have proven extremely useful since they allow independent testing under almost the same conditions as in global data taking. Trigger control for MiniDAQ systems is provided by one of the PM blocks in a TCDS LPM. During Run 1 and Run 2, these MiniDAQ systems ran on dedicated RU, BU, and FU servers and provided limited bandwidth with respect to the global DAQ system. In Run 3, MiniDAQ systems share the RU/BU and FU servers with the global system and provide a full bandwidth to each subdetector. The configuration of the MiniDAQ systems is dynamically updated to follow any changes in the global system.

10 Level-1 trigger

The level-1 (L1) trigger is implemented in custom hardware processors. It comprises calorimeter and muon trigger systems that provide jets, e/γ , hadronic τ , and muon candidates, along with calculations of energy sums, to the global trigger (GT). At the GT, the trigger decision is generated, based on the multiplicity and kinematic information of the various candidate trigger objects. The trigger configuration is implemented in a trigger “menu” comprised of several hundred “seed” algorithms. Upon a positive GT decision, the full detector data are read out for further filtering in the higher-level trigger (HLT). During LS1, in 2013–2014, the L1 trigger hardware was entirely upgraded, and has subsequently been operated successfully since 2016. A detailed report on this Phase 1 L1 trigger upgrade and performance with Run 2 data is given in ref. [5].

For Run 3, although no major trigger hardware upgrade was performed, new capabilities have become available already through new algorithmic approaches, some of which are based on machine learning (ML) techniques. Software such as HLS4ML [287] facilitates the use of ML techniques in FPGAs. Developments for Run 3 within the L1 trigger mostly focus on broadening the physics reach of CMS through the addition of dedicated triggers for long-lived particle (LLP) signatures, improving object measurement and calibration, utilizing the upgraded calorimeter trigger primitives (TPs) and additional muon TPs from the new GEM muon detector, and implementing additional calculations in the global trigger to provide greater flexibility in the design of L1 trigger algorithms.

The addition of a 40 MHz scouting system, commissioned in the early stages of Run 3, that receives data from both the calorimeter and muon L1 trigger subsystems, has the potential to further broaden the physics reach of CMS. It enables the readout of unfiltered data, reconstructed in situ at limited precision but at full bunch-crossing rate, and provides unprecedented monitoring capabilities. The following sections describe the Run 3 developments specific to each of the L1 trigger subsystems.

10.1 Calorimeter trigger

10.1.1 Calorimeter layer 1 trigger

The calorimeter layer 1 trigger receives TPs from ECAL, HCAL, and HF, calibrates them, combines the ECAL and HCAL TPs into single trigger towers (TTs), and transmits the TTs to layer 2 for further processing. Calorimeter TPs for triggered events are readout to DAQ, and used in the data quality monitoring (DQM) system, where they provide the input to the software emulator, such that online and emulated data can be compared in real time for monitoring purposes. For Run 3, layer 1 receives updated HCAL TPs (section 5.3), which improves the mitigation of out-of-time pileup, and updated ECAL TPs (section 4.5) with improved rejection of spikes caused by particles striking the avalanche photodiodes.

The TT energies are calibrated to account for energy losses due to inactive material in front of the calorimeters. The calibration is performed separately for each calorimeter. Since the inactive-material map is symmetric in ϕ , the calibration is performed as a function of η and transverse energy, E_T , only. Whereas for the ECAL TP calibrations, the scale factors vary by less than 20% across the E_T range, the HCAL and HF calibration scale factors are much more E_T dependent, varying by about 50%. Both ECAL and HCAL scale factors vary by about 20% across the η range.

Due to the large volume of data produced by the TPs and limited DAQ bandwidth, only TPs for triggered events are read out by the DAQ system. To study possible trigger bias, for approximately every 100th event that passes the full L1 trigger selection, validation data are read out, that also contain the TT information. In Run 3, validation events contain the ECAL TP data for five bunch crossings, including the two bunch crossings before and after the triggered bunch crossing. This information is useful for studying unexpected detector effects, and can be used to study ECAL prefiring (section 4.5), both using the DQM and offline analysis. This additional ECAL TP data can also be used for various optimization studies, such as monitoring the timing of the ECAL TPs through the L1 trigger path during commissioning phases.

10.1.2 Calorimeter layer 2 trigger

Calorimeter layer 2 receives calibrated TTs from layer 1, reconstructs jet, e/γ , and τ candidates, and computes energy sums. The energies of jet, e/γ , and τ candidates are calibrated as a function of p_T and η , and isolation and ID criteria are applied to e/γ and τ candidates. Pileup mitigation is applied to all objects to reduce the rates while maintaining high efficiencies. The layer 2 hardware remains the same as for Run 2. Ten main processor cards each process data from the entire calorimeter for a single bunch crossing in a time-multiplexed configuration. A single demultiplexer processor receives data from the main processors, performs the final calculation of the energy sums, and forwards the object collections to the global trigger. More details can be found in ref. [5].

A range of improvements to the layer 2 algorithms are being investigated for Run 3, most of which involve utilizing the updated ECAL and HCAL calorimeter TPs as discussed in sections 4.5 and 5.3, respectively. In particular, the ability to trigger on LLP signatures by identifying displaced jets using the additional HCAL timing and depth information available for Run 3 is being pursued to help broaden the LLP physics program of CMS.

Jets are reconstructed by summing the energies of a 9×9 window of TTs centered on a jet seed that must have an energy greater than 4 GeV, which corresponds to approximately the same jet size as jets reconstructed offline with $\Delta R = 0.4$ within the barrel calorimeter. The energy contribution due to pileup is estimated by summing the three lowest energy out of the four 3×9 regions on the boundaries of the jet and subtracting this pileup estimate from the jet energy, which is then calibrated.

For Run 3, an LLP jet identification algorithm is implemented that uses the HCAL timing and depth information. Each TT has an HCAL feature bit set in layer 1, which compresses six feature bits received from the HCAL backend, containing timing and depth information, into one feature bit. When the jets are reconstructed from TTs at layer 2, an LLP jet ID bit is set if the jet contains more than a configurable number of TTs with the HCAL feature bit set. Additional LLP jet algorithms have been added to the GT menu that require the LLP jet ID bit to be true. In addition to tagging LLPs, the ability to tag boosted jets with substructure using a pattern-matching technique, and the use of ML techniques to calibrate the jet energy and perform pileup subtraction are being investigated for Run 3.

The missing transverse momentum (p_T^{miss}) is calculated as the vector sum of the TT energies across the full calorimeter. During Run 2, a significant increase in the instantaneous luminosity and thus the pileup per bunch crossing revealed a nonlinear relation between pileup and L1 p_T^{miss} trigger rate, leading to a large increase of rate at large pileup. To maintain the L1 p_T^{miss} trigger thresholds while keeping the rates manageable, a pileup mitigation procedure was implemented using a lookup table (LUT) to exclude TTs below a configurable energy threshold from the p_T^{miss} calculation, based on the η of the TT, and an estimate of the pileup of the event. The TT energy thresholds were rederived for Run 3, using Monte Carlo (MC) simulation samples containing inclusive pp events with a pileup distribution reflecting that of Run 3.

A comparison of the L1 p_T^{miss} trigger efficiencies between 2018 and Run 3, for different L1 p_T^{miss} trigger thresholds that provide the same L1 trigger rate, is shown in figure 121. Compared to the 2018 LUT, the updated Run 3 LUTs show a significant improvement of the efficiency relative to the true p_T^{miss} in the event. As the LHC conditions evolve, the pileup mitigation will be updated to ensure optimal performance. In the current tuning, the energy threshold below which TTs are excluded is set such that 99.5% of TTs are excluded for inclusive pp events.




Figure 121. Comparison of the L1 p_T^{miss} trigger efficiency using pileup mitigation in 2018 (circles) and in Run 3 (squares) for thresholds that provide a rate of 4.3 kHz, for $Z \rightarrow \mu\mu$ events.

The e/γ and τ candidates are constructed by clustering TTs containing energy deposits greater than 1 GeV around a cluster seed tower with an energy deposit of at least 2 GeV. This clustering is done dynamically using the available tower-level information. The cluster is trimmed by removing towers, and a veto based on the trimmed cluster shape is applied, to reject pileup and reduce background rates. A fine-grain veto is applied in the barrel calorimeter that quantifies the compactness of the electromagnetic shower within the seed tower to reject hadron-induced showers. A veto is also applied that requires a low HCAL-to-ECAL energy ratio (H/E) in the seed tower, with different thresholds used in the barrel and the endcap regions. Isolation requirements are applied to set an isolation bit and provide isolated candidates to the GT. Merged τ candidates are constructed by merging nearby clusters that pass a set of proximity conditions, to capture multiprong hadronic τ decays. Energy calibrations for e/γ candidates use the cluster shape and those for τ use H/E and the

presence of merged clusters, in addition to the p_T and η . While no significant changes to either the e/γ or τ algorithms have been implemented for Run 3, the existing calibrations and isolation working points have been and will be rederived throughout Run 3 to reflect updated detector conditions and calorimeter TP algorithms. Methods to improve the isolation working points of e/γ and τ candidates utilizing ML techniques are being implemented.

10.2 Muon trigger

The L1 muon trigger for Run 3 receives TPs from four partially overlapping muon subdetectors: DT, CSC, RPC, and GEM. As described in detail in section 6, three of these subdetectors, DT, CSC, and RPC, were operated during Run 2, while the GEM detector was added as part of the Phase 1 upgrade and is used for the first time in Run 3. The L1 muon trigger system reconstructs muon tracks and provides measurements of muon track parameters using TPs which provide position, timing, and quality information from detector hits. In the barrel, accurate directional information is also provided. The geometrical arrangement of the muon subdetectors, including the new GE1/1 detector in front of ME1/1, is shown in figure 39. In this section, the changes to the muon track finders are discussed in detail.

The L1 muon trigger system in Run 3 comprises the same overall design as in Run 2. Three muon track finders (TFs) reconstruct muon tracks in three distinct pseudorapidity regions using TPs from muon detectors. The barrel muon track finder (BMTF) receives inputs from DT and RPC in the barrel ($|\eta| < 0.83$), the overlap muon track finder (OMTF) uses DT, CSC, and RPC in the overlap between barrel and endcap ($0.83 < |\eta| < 1.2$), while the endcap muon track finder (EMTF) takes inputs from CSC, RPC, and GEM in the endcap ($1.2 < |\eta| < 2.4$). All three muon track finders transmit up to 36 muons each per bunch crossing to the global muon trigger (μ GMT), which resolves duplicates and transmits a maximum of eight muon tracks per bunch crossing to the GT, similar to Run 2.

In Run 3, all three muon track finders additionally provide measurements of parameters for muon tracks that are displaced from the primary interaction point. The beamspot constraint requiring the track to originate from the interaction point is removed. The newly available track parameters are used in the GT to provide L1 muon trigger seeds targeting displaced-muon signatures that could originate from LLPs. Additionally, in Run 3, the EMTF receives TPs also from the GEM detectors, and this information can be used to improve both prompt and displaced muon triggering.

The new displaced-muon algorithms provide a p_T , measured without the beamspot constraint, and a transverse displacement, d_{xy} , from the beam line, for muon tracks obtained from propagating back to just the first muon station. Algorithms optimized for prompt muons typically underestimate the p_T of highly displaced muons, as the displacement is mistaken as increased track curvature due to the beamspot constraint. The new displaced-muon algorithms improve the p_T estimation for displaced muons, hence improving the efficiencies for these triggers. The displaced TF algorithms in general do not affect the muon track building, but provide additional measurements of unconstrained quantities. Due to this approach, the trigger efficiency is increased significantly for muon $p_T > 10$ GeV when the displacements are larger than 20 cm, while in the case of lower p_T muons, the displaced muon algorithms perform similarly to the prompt algorithms. This is due to the fact that all the prompt TF algorithms have a minimum p_T assignment value (2–4 GeV depending on the TF), and the underestimation of p_T for displaced muons becomes less important at low p_T values.

10.2.1 Barrel muon track finder (BMTF)

In the BMTF, the kBMTF algorithm reconstructs muons in the barrel region using TPs received from the DTs and RPCs via the TwinMux concentrator cards, and has been used online since 2018 [5]. It is the successor of the original BMTF algorithm, used between 2016 and 2018, and based on an approximate Kalman filter algorithm in which muon tracks are reconstructed from detector hits starting from the outermost muon station and propagating inwards while updating the track parameters.

In Run 3, the updated version of the kBMTF algorithm mainly improves the displaced-muon triggering performance. The new algorithm retains high efficiency for prompt muon tracks up to $d_{xy} \approx 50$ cm, while the displaced-kBMTF algorithm increases the efficiency of muon tracks with displacements above 50 cm. The expected performance for Run 3 was evaluated using a Run 2 cosmic ray muon data sample in which the displaced algorithm shows efficiencies around 80% for $L1$ $p_T > 10$ GeV, up to displacements of 100 cm (figure 122).




Figure 122. Displaced (blue) and prompt (black) kBMTF trigger efficiencies compared to the prompt BMTF (red) trigger efficiency with respect to the muon track d_{xy} , obtained using a sample of cosmic ray muons from 2018 data. The efficiencies are measured using muon candidates with $p_T > 10$ GeV. The prompt kBMTF improves BMTF efficiencies up to about 90% for up to 50 cm displacements, while displaced kBMTF retains efficiencies above 80% for up to 90 cm displacements.

10.2.2 Overlap muon track finder (OMTF)

The OMTF builds muon tracks using the TPs from the DTs and RPCs in the barrel and CSCs and RPCs in the endcap. The algorithm for Run 3, which is mostly identical to that of Run 2, uses a Bayes classifier algorithm based on precomputed patterns generated from simulated events to associate hits in each station with the reference hit in the pattern. The patterns contain information about muon track propagation and the probability density function of the hit distribution in ϕ with respect to the reference hit. There are 26 patterns for each muon charge, corresponding to p_T values between 2 and 140 GeV, that are then used to estimate the p_T of the muon track based on

the likelihood that the track matches a pattern. In Run 3, the OMTF includes additional patterns to improve displaced-muon triggering. Although the general structure of the algorithm remains the same, the updated Run 3 algorithm now finds the best matching prompt and displaced patterns for a given muon track. There are 22 displaced patterns for each muon charge corresponding to different d_{xy} values which are valid for high-momentum tracks ($p_T > 30$ GeV).

In Run 3, the OMTF uses prompt and displaced patterns in parallel for each track. The algorithm identifies prompt-muon tracks and estimates their p_T using prompt patterns, and at the same time it estimates the d_{xy} of the displaced muon using the displaced patterns. The expected performance for Run 3 was evaluated using a displaced-muon gun simulation sample, in which just muons are simulated with zero pileup, in order to compare ideal efficiencies, in which the displaced-OMTF algorithm shows efficiencies around 80% up to displacements of 200 cm (figure 123).




Figure 123. The OMTF trigger efficiencies for displaced and prompt algorithms with respect to muon track d_{xy} obtained using a displaced-muon gun sample. The efficiency curves are plotted for different values of the p_T estimate from the prompt algorithm (red, yellow, and blue), for the displaced algorithm (green), and for the combination (black). The prompt algorithm underestimates the p_T of displaced tracks, causing most of the tracks to have $p_T < 10$ GeV. The displaced algorithm can recover these tracks and improve the efficiencies to be around 80% for up to 200 cm displacements.

10.2.3 Endcap muon track finder (EMTF)

The EMTF builds muon tracks using TPs from CSCs, RPCs, and GEMs in both endcaps. The Run 3 EMTF algorithm includes GEM inputs for the first time. Already in Run 2, the algorithm used one TP per endcap muon station, with CSCs having priority over RPCs, to build muon tracks based on a set of predefined patterns. After the patterns were found, the differences in ϕ and θ angles between hits in different stations were used to estimate the p_T of the muon track using a boosted decision tree (BDT). The BDT was trained using simulated single-muon events and the output values of the BDT were stored in LUTs for fast evaluation. Since the beginning of Run 3, the EMTF algorithm also receives the new GEM TPs in station 1, as well as the updated CSC TPs in all stations, which provide better position and bending resolution. These TPs can be used both at the track building stage, as

well as the p_T estimation stage to improve the performance of the EMTF algorithm. Additionally, the EMTF uses a neural network (NN), implemented in FPGA logic, to estimate the p_T and d_{xy} of displaced-muon tracks, and forwards data to a hadronic shower trigger that uses the multiplicity of hits in the endcap CSCs to trigger on LLPs producing showers as they enter the endcap muon systems.

The first of the GEM detectors (GE1/1) is included in the EMTF algorithm in station 1 of the CMS endcap muon system. The Run 3 EMTF algorithm can use GE1/1 hits in conjunction with ME1/1 hits to improve prompt and displaced trigger efficiencies and reduce rates caused by mismeasured muons. Due to the placement of the GE1/1 and ME1/1 chambers in the CMS endcaps, the strong magnetic field in this region causes a larger bending of the muon track. This bending information between GE1/1 and ME1/1 is foreseen to be used to improve the p_T assignment for both prompt and displaced-muon tracks.

Similar to the other track finders, the EMTF also includes a new algorithm to improve displaced-muon triggering. The EMTF for Run 3 includes a NN-based p_T and d_{xy} assignment algorithm, which runs in parallel to the prompt algorithm. The NN has been directly incorporated into the EMTF firmware and estimates the p_T and d_{xy} of muon tracks that are built by the EMTF track building algorithm. The EMTF performance for prompt muons originating from the primary vertex remains identical. As shown in figure 124, the displaced-EMTF algorithm (NN-EMTF) shows an improved efficiency up to a d_{xy} of about 100 cm, while the prompt EMTF algorithm retains a high efficiency for prompt muon tracks up to about 25 cm. The expected performance for Run 3 was evaluated using a displaced-muon gun simulation sample with zero pileup. Zero-pileup samples were used since they are more useful in optimizing the NN and comparing ideal efficiencies. For L1 $p_T > 10$ GeV, the NN-EMTF algorithm shows efficiencies above 80% for $1.2 < |\eta| < 1.6$ and up to 100 cm displacements, while efficiencies for $1.6 < |\eta| < 2.1$ and $2.1 < |\eta| < 2.5$ at 60 cm displacement are around 20 and 5%, respectively (figure 124).

Finally, the EMTF for Run 3 forwards CSC hit information to provide a standalone method for triggering on hadronic showers occurring in the CSC detectors. An LLP decaying to hadronic particles within or slightly before the endcap muon systems can cause a shower of charged particles hitting the muon detectors, which are then recognized through a high hit multiplicity. In Run 3, the CSC detector sends information on whether a high multiplicity is found in any given chamber by comparing the measured multiplicity to a set of predetermined thresholds, individually for each CSC station and ring combination. The EMTF processes this information to decide whether there was a hadronic shower of a given quality in at least one of the CSC chambers in any given sector. The expected performance of this algorithm for Run 3 was evaluated using multiple physics simulation samples containing LLPs and found to provide efficiencies around 30%.

10.3 Global trigger

The hardware for the present global trigger system, the μ GT, was installed as part of the Phase 1 upgrade, and was used for most of LHC Run 2 and is used for Run 3. The flexibility of the L1 trigger system has allowed for the addition of a μ GT test crate containing the same hardware to be added to further extend the global trigger capabilities. It is used for testing and development purposes, for example, to test new or experimental trigger menus, or to test Phase 2 algorithms using ML or autoencoding, as described in ref. [88]. It receives the same optical inputs as the production crate from a passive optical splitter panel. For Run 3, this crate has been included as an optional component




Figure 124. The EMTF trigger efficiencies for prompt and displaced-muon algorithms for L1 $p_T > 10$ GeV with respect to muon track d_{xy} obtained using a displaced-muon gun sample. The solid stars show displaced NN performance while hollow squares show the prompt BDT performance. The different colors show different η regions: $1.2 < |\eta| < 1.6$ (black), $1.6 < |\eta| < 2.1$ (red), and $2.1 < |\eta| < 2.5$ (blue).

into the CMS data acquisition system. While the test crate does not issue triggers itself, for events triggered by the production crate, the full test crate data (inputs and results of calculations) is available in the CMS data stream for offline analysis. The production and test crates are shown in figure 125.

By running the test crate with the same firmware, and thus the same trigger algorithms, as the production crate, it is possible to carry out consistency and hardware checks. More importantly, one can also run with different firmware in the test crate. This allows trigger developers to test new trigger menus with actual data. During normal data taking, a stream of zero-bias data is included, where the only trigger requirement is that the bunches for that crossing are filled. The zero-bias trigger has a high prescale applied, such that only 1 in N events are recorded, where N is the prescale value. This makes it possible to investigate the performance of new triggers in the test crate menu. A test crate algorithm that is more restrictive than a specific algorithm in the production system can also be studied by using the latter algorithm as a reference, thus benefiting from a larger data sample. The fact that the test crate can be included in data acquisition during normal running at no additional cost means that large amounts of realistic data can be collected, thus allowing for accurate tests even of very restrictive algorithms that only rarely provide a trigger. By normalizing the offline data to the number of zero-bias triggers taken or by recording the online monitoring data, one can also check the total trigger rate of a new menu. This is very useful, since it is not trivial to accurately estimate the total menu rate offline using the rate of individual algorithms, since multiple algorithms can fire the same event simultaneously. This is important in order to guarantee that when running with the new menu the L1 trigger rate remains within the total bandwidth, so that deadtime is minimized.

The upgraded monitoring backend is configured to collect and store data from the test crate in order to provide integration with the central monitoring services. Trigger menu rates are stored in the cen-




Figure 125. Production (upper) and the new test crate (lower) of the μ GT.

tral Prometheus monitoring database, described in section 10.5, together with those from the production crate for online monitoring and prompt offline studies. Data are also expected to be sent to offline computing systems for long-term storage. Alerting services are configured to warn shifters and experts with relaxed severity compared to the production system, as the system is not critical for data taking.

10.4 Trigger menu

Expanding the CMS physics reach significantly beyond that explored already during Run 2 [5] requires important changes to the Run 3 trigger menu, and relies mostly on improvements to the L1 trigger algorithms. The LHC beam conditions and filling scheme, as well as other aspects like luminosity leveling, described in more detail in section 7.5, should be taken into account in order for the menu to be efficient and robust within the L1 bandwidth limit of about 100 kHz.

The updates to the calorimeter and muon trigger systems discussed earlier aim to retain the Run 2 physics coverage described in ref. [5], while providing additional access to signatures that do not originate at the primary vertex, such as displaced muons or displaced jets. Furthermore, the GT is able to perform two new kinematic computations that can be utilized to increase the physics coverage: the three-body invariant mass and the di-object ratio between invariant mass and the $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ between the objects. Special L1 trigger menus are available in addition to

the standard menus to account for various data-taking scenarios, as was done during Run 2, e.g., targeting signatures related to the physics of bottom quarks using the so-called “B parking” [288], recording additional data for B physics and other studies, as further described in section 11.5.

10.4.1 Trigger seeds for displaced muons

Physics signatures involving LLPs were not well represented within the L1 trigger menu during Run 2, since specific algorithms capable of triggering on displaced objects within the detector were not yet implemented. The changes for Run 3 provide the ability to trigger on displaced muons with higher efficiency compared to Run 2. Updates to the muon algorithms in all three muon track finders provide unconstrained p_T and d_{xy} measurements for displaced muons, as discussed in section 10.2.

To benefit from the displaced muon algorithms, the L1 trigger menu of Run 3 extends that of Run 2 by offering seeds that use the unconstrained p_T and d_{xy} measurements. These new seed features are available for muon objects with $p_T > 10$ GeV, and possible additional selections are defined according to the planned physics program. The chosen p_T threshold is based on the expected gain in the trigger efficiency described in section 10.2, concerning only muons with $p_T > 10$ GeV and displacements larger than approximately 20 cm. Since the displaced muon algorithms do not perform better compared to the prompt algorithms in the case of low p_T muons, no new seed is created if the leading muon in the considered seed has a p_T below 10 GeV.

10.4.2 Trigger seeds using new kinematic variables

Low-mass resonance searches, for example those used to study the physics of bottom quarks, are based on targeting a final state with low p_T objects. While the final states targeted in B-physics searches are often expected in the barrel region, many other interesting physics scenarios predict a wide pseudorapidity distribution for the final state objects. Ideal seeds with a minimum object selection with low object thresholds would result in a high trigger rate that is unsustainable, exceeding the available trigger bandwidth under any LHC data taking conditions, and provide low purity.

During Run 2, the most used unscaled seeds relying on a single object had relatively high p_T selections to keep trigger rates manageable, e.g., the single-muon trigger with $p_T > 22$ GeV. Special B parking triggers, used for events that are written to disk storage without full online event reconstruction at HLT, were developed to lower the L1 trigger thresholds while keeping within the HLT trigger bandwidth. To provide reasonable rates, typical B parking seeds in general restrict pseudorapidity to the barrel and overlap regions, e.g., the double-muon trigger with no p_T selection but with $|\eta| < 1.5$ and $\Delta R_{\mu\mu} < 1.4$. However, the HLT thresholds used for the B parking seeds, generally above 5 GeV, would often significantly reduce the total acceptance of interesting physics signals.

The substantial L1 trigger rates related to low p_T trigger thresholds can alternatively be managed by using kinematic variables that are optimized to increase the acceptance of predicted signal events. During Run 3, the GT is able to perform two new kinematic variable computations: the three-body invariant mass, and the di-object ratio between invariant mass and the ΔR . These allow the rate of seeds with low p_T thresholds to be reduced. For example, the three-body invariant mass can be harnessed to target the decays of a τ into three muons, a final state with low- p_T final objects with wide pseudorapidity distributions. Alternatively, the di-object ratio between invariant mass and the ΔR could benefit a low-mass dimuon resonance search providing a dark photon interpretation.

10.4.3 Run 3 trigger rates

The baseline L1 trigger menu in Run 3 is identical to that of Run 2 detailed in ref. [5]. Since the beam and detector conditions, and upgraded TPs are different from those during Run 2, the baseline trigger menu has been reviewed using Run 3 simulation samples to understand if slight modifications to the existing seeds are necessary to respect the total rate budget of 100 kHz. Two preliminary rate studies have been performed by reweighting the pileup distribution in Run 3 simulation samples, which has an unrealistic flat shape between 30 and 80. Firstly, by reweighting the Run 3 MC simulation samples to the Run 2 pileup distribution, it was confirmed that the rates using Run 3 simulations with updated TPs provide the same rate as obtained from Run 2 data, for the same pileup conditions. Secondly, the expected rates for realistic Run 3 pileup conditions were estimated by reweighting using a pileup distribution determined from the expected instantaneous luminosity during luminosity leveling. This approach takes into account various beam conditions, such as the bunch-to-bunch variation in pileup, which can have a significant impact on the rates. The rate allocation for single- and multi-object triggers and cross triggers is shown in figure 126. These results can be compared to a similar figure for the L1 trigger rate allocation under the Run 2 conditions, provided in ref. [5]. Additional seeds, relying on the aforementioned new features, are implemented on top of the baseline L1 trigger menu. This menu is then tuned according to the desired physics program of CMS for Run 3.




Figure 126. Fractions of the 100 kHz rate allocation for single- and multi-object triggers and cross triggers in the baseline Run 3 menu, calculated using Run 3 Monte Carlo simulation samples of inclusive pp events with appropriate pileup.

10.5 Online software and monitoring

During LS2, a number of improvements to the development workflow for the L1 trigger online software have been used. Practices from the well-established methodologies DevOps [289] and Agile [290] have been introduced in the Run 3 online software to simplify the development and deployment of upgrades and enable continuous improvement. The online software is built in centrally-provided Docker images that provide a replicable development and test environment. Every online software project employs continuous integration (CI) pipelines in GitLab to automatically build and test new commits. In order to minimise maintenance efforts and facilitate the development

of CI pipelines, an in-house Auto-DevOps platform was developed, taking inspiration from the GitLab Auto-DevOps project. Auto-DevOps is employed to centrally provide a set of configurable and generic CI jobs that represent typical tasks in the online software development, such as building RPM package manager packages and Docker images. Feedback from introducing standard CI pipelines in all software projects was excellent: given the large number of software components and teams that the online software consists of, this procedure established standard workflows that greatly increased the safety of new deployments. New automated tests are constantly developed and integrated into the CI whenever new incidents occur, reinforcing the importance of feedback from operations to minimise disruptions of service. Work is ongoing to develop a testing cluster managed by RedHat OpenShift running on CERN computing services where test deployments can be performed from CI pipelines. The goal is to provide a complete environment where software checks on the entire infrastructure can be run before deploying on the CMS computing resources.

The Run 3 monitoring backend is based on industry-standard tools and is centered around a Prometheus monitoring database instance collecting monitoring data from all trigger subsystems [291]. The database is expected to collect around 300 000 metrics every 20 s and correlate monitoring information coming from different sources to provide a powerful monitoring and alerting system. Alerts are handled by the Prometheus Alertmanager service [291]. The software enables binding alerts to external software based on conditions to page shifters and escalate to experts if necessary. Actions to automatically mitigate or solve problems can also be configured to minimise downtime. Monitoring dashboards are built using Grafana and can provide powerful inspection tools for online and prompt post-run analysis (figure 127) [292]. Prometheus is not designed for providing long-term storage of metrics, therefore integration with offline computing systems is under consideration.




Figure 127. Screenshot of the Grafana L1 trigger monitoring dashboard.

The Run 3 trigger configuration editor and main monitoring page are designed in the Vue.js JavaScript framework, which superseded the outdated and slower Polymer framework. The Run 3 main monitoring page has also been updated to present data taken from the updated monitoring backend.

The significantly improved L1 online software and monitoring enables users to spot problems much more efficiently, and react more quickly compared to Run 2. It also facilitates more efficient diagnosis of problems after they have occurred, since more historical information is available and is easy to access.

10.6 The L1 scouting system

For Run 3, the L1 trigger project added dedicated hardware for the triggerless recording of objects reconstructed in the level-1 trigger, at a rate of 40 MHz without trigger or filtering. The data are received by dedicated FPGA-based processing cards, either housed in powerful servers (I/O nodes) or operating standalone in custom enclosures. Once received and pre-processed, the information is provided to a computing farm via Ethernet. Triggerless data recording, referred to in CMS as “40 MHz scouting” or “L1 scouting”, facilitates improved precision of measurements such as the luminosity, as well as unprecedented levels of trigger monitoring. The term “scouting” is also used in the context of HLT, described in section 11.4, where it refers to the selection and high-rate recording of small-size analysis data sets.

The L1 scouting system deployed for Run 3 is meant as the first large-scale demonstration of a larger system planned for the Phase 2 L1 trigger upgrade [88]. Besides providing a testing ground for the technical implementation of the readout, it provides the opportunity to study solutions for large-scale distributed processing of high-rate data, the correlation of multi-bunch-crossing signals, as well as first studies of possible physics applications using limited resolution trigger data. The measurement of a range of physics processes, such as LLPs with displaced muons and flavor anomalies in τ physics, could potentially benefit from L1 scouting.

10.6.1 Architecture of the L1 scouting system

The L1 scouting system receives data from the level-1 trigger via spare output links and processes them quasi-online in a dedicated computing farm. The system operates largely independently from the standard CMS trigger and data acquisition chain. In the first test system in Run 2, the data sent from the μ GMT to the μ GT was duplicated, and final muon objects, as well as intermediate muon candidates derived from the BMTF inputs, were transmitted over eight 10 Gb/s optical links to the L1 scouting system.

For Run 3, an additional set of duplicated μ GMT outputs supply a L1 scouting processor dedicated to luminosity monitoring, as described in section 10.6.2. Each of the twelve BMTF processors dedicates two 10 Gb/s links to L1 scouting information. L1 scouting data from the calorimeter layer 2 system also mirrors the trigger objects provided to the μ GT. Each μ GT processor transmits 512 bits per bunch crossing over three 10 Gb/s links, indicating which of the trigger algorithms has fired in a given bunch crossing.

The architecture of the initial Run 3 L1 scouting system is a scaled up version of the Run 2 system, and consists of I/O nodes housing FPGA-based input boards that receive up to eight 10 Gb/s links using the L1 trigger link protocol. Data are transmitted unidirectionally with no back-pressure to the trigger. The L1 scouting system therefore does not interfere with the standard trigger system in any way. The FPGA logic performs both zero suppression and preprocessing of the data such as reformatting or recalibration. The use of fast NN algorithms, implemented within FPGA resources to improve recalibration performance has been demonstrated (see section 10.6.2). Data are transferred

by direct memory access (DMA) via a Gen-3 PCIe x16 bus into the memory of the I/O nodes, from where they are sent to dedicated processing units in the surface data center via 100 Gb/s Ethernet over coarse wavelength division multiplexing CWDM4 single mode optical infrastructure. Even after full zero-suppression, the long-term storage of the huge amount of raw data produced by the trigger processors, in view of a subsequent “classic” multitiered offline analysis and reduction, does not represent a viable approach. Data taken in the early months of LHC Run 3 are being used to investigate various methods for a real-time analysis.

In the Run 2 demonstrator system, a KCU1500 Xilinx development kit, equipped with a KU15P FPGA, was used to capture the μ GMT inputs. Additional I/O nodes for Run 3 are equipped with more powerful boards, such as the Micron SB-852, using a large Xilinx Ultrascale+ VU9P FPGA and also providing access to the Micron Deep Learning Accelerator (MDLA) [293], which is a proprietary compiler that translates pre-trained ML networks into instructions for an FPGA-based hardware implementation. Monitoring of the readout board is being developed using the AXI-lite interface provided by the Xilinx xDMA core. Custom software exposes access to monitoring and control registers via a RESTful interface and monitoring and diagnostic data are exported to the Prometheus server described in section 10.5.

In addition to the I/O nodes with dedicated receiver boards, Xilinx VCU128 development kits [294] are used as standalone receivers for L1 scouting data. These boards, which are housed in a custom enclosure with PCIe extender buses for control and monitoring, are equipped with a VU37P FPGA with 8 GB of high bandwidth memory. The VU35P, with a very similar architecture, is planned for use on the DAQ-800 board currently being designed for the CMS Phase 2 data acquisition [295], and is the anticipated readout board for the L1 scouting system of the Phase 2 upgrade. The VCU128 boards in Run 3 thus allow a realistic test of the Phase 2 design, albeit with reduced input bandwidth. Equipped with an additional FMC mezzanine to provide up to 32 10 Gb/s input links, they transmit L1 scouting data directly over TCP/IP to the surface computing farm, providing more efficient utilization of the bandwidth available on the 100 Gb/s Ethernet links. The Run 3 scouting architecture is illustrated in figure 128.

10.6.2 Applications of the L1 scouting system

The use of ML algorithms to improve the physics potential of L1 trigger objects captured by the triggerless recording of trigger data at 40 MHz has been investigated. An example of this is the use of deep NNs, optimized for a throughput of around 1 MHz, to recalibrate L1 muon objects in real time, such that the accuracy of the p_T , η , and ϕ parameters of the muons can be improved over the standard L1 trigger reconstruction, which is optimized for efficiency at threshold, not for a full physics analysis. Neural networks have been trained on zero-bias triggered data from LHC Run 2, to predict the offline fully reconstructed muon parameters of matched L1 muon objects. This recalibration has been shown to improve the accuracy of the L1 muon parameters when compared to the standard L1 reconstruction, both for μ GMT and BMTF L1 muon parameters [296, 297]. Additionally, neural networks that are trained to reject pairs of L1 muon candidates that do not correspond to matched fully reconstructed muons, and networks for anomaly detection have also been implemented in the Xilinx VU9P FPGA with a throughput of about 1 MHz, using the MDLA.

A copy of the μ GMT scouting system will be used to provide direct per-bunch measurements of the L1 muon multiplicity to the CMS BRIL system described in section 8. FPGA firmware




Figure 128. Architecture of the Run 3 L1 scouting prototype system. The input system, located in the experiment service cavern, consists of a combination of different types of FPGA receiver boards, hosted in the PCIe bus of I/O server nodes or extender. The boards receive and pre-process data from the different trigger systems. Two of the boards (KCU1500 and SB852) use DMA to move their data to host memory, from where they are transmitted to the surface data center over 100 Gb/s Ethernet links. The VCU128 implements a TCP/IP core in the FPGA and directly transmits data to the surface system. In the surface data center, links from the L1 scouting input system are connected to a switched network. Data streams are received through the said network by L1 scouting buffer servers (DSBU) and buffered in files on large RAMdisks. The L1 scouting processing units (DSPU) access buffered data from the DSBU to perform data reduction and analysis. The processed data are finally moved to a Lustre cluster file system for long-term storage.

has been developed and implemented in both the KCU1500 and SB-852 boards to histogram the muon multiplicity per bunch over a period of four lumi nibbles, corresponding to 1.458 s. These histograms are read out to the host PC in real time via the PCIe DMA implementation.

Cosmic muon data taken with the L1 scouting system in special runs during LS2 have been used to analyse the performance of the new kBMTF algorithm. The L1 track finding algorithms assume that muons always originate from the interaction region. Therefore, a cosmic muon traversing the full detector appears as two back-to-back muons, usually separated in time by one or two bunch crossings. The data were taken when the magnet was off. In the absence of a magnetic field, muons traverse the detector in a straight line, leaving two muon tracks in the barrel drift tubes, typically two bunch-crossings apart ($\Delta BX = 2$). Figures 129–130 show the relationship between the L1 reconstructed impact parameters, d_{xy} , and the angle ($\phi_{in} - \phi_{out}$) of the two corresponding L1 tracks.

The Run 3 demonstrator system was used to systematically collect and store data during collision runs throughout 2022 from the μ GMT and calorimeter layer 2. The offline analysis of these data provided invaluable insight into the reliability of the acquisition system and in particular of the hardware,




Figure 129. Illustration of the correlation between the impact parameter d_{xy} , highlighted in red, and the difference between the angles measured for the incoming and outgoing legs of a cosmic muon.




Figure 130. The impact parameter d_{xy} of the incoming (left) and outgoing (right) cosmic ray muon tracks as measured by the BM TF as a function of the difference of azimuthal coordinates of the incoming and outgoing legs. The BM TF firmware encodes the impact parameter in two bits, hence the range of values on the left y axis. The orange curves model this dependence of the actual impact parameter as $R_{MS} \cos((\phi_{in} - \phi_{out})/2)$, where R_{MS} is the radius at which the BM TF measures the ϕ coordinate of the track. The right hand side y axis shows the d_{xy} values (in cm) as predicted by this model, which exhibits remarkable consistency with the measurement (for the values within the range), if one assumes one unit of the left axis to correspond to an impact parameter of about 100 cm. Reprinted from [298], Copyright (2023), with permission from Elsevier.

firmware and software. The results of this work have been and are being used to guide the design of further extensions of the demonstrator and the design of the Phase 2 systems. The quality of the data was studied in view of assessing their usability for trigger diagnostics, luminosity measurement, and physics studies. This work will continue throughout the rest of Run 3, along with the completion of the demonstrator with the commissioning of the VCU128 acquisition boards, collecting data from the uGT and BM TF systems, which represents a major step forward in validating the Phase 2 design.

11 High-level trigger

11.1 Overview

To select events of potential physics interest, the CMS trigger system divides the processing into two levels: the level-1 (L1) trigger, implemented in custom hardware as described earlier in section 10, and the high-level trigger (HLT), implemented in software, running on a farm of commercial computers. The HLT refines the purity of the physics objects selected by the L1 trigger, with a maximum input rate increased to 110 kHz in 2023 from 100 kHz earlier. In 2022, for standard pp collisions at an average instantaneous luminosity of $1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the average rate to offline storage for promptly reconstructed physics events was approximately 1.7 kHz. In 2023, for a peak luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the rate of promptly reconstructed physics events was about 2.6 kHz. Additional data streams for calibration purposes or “HLT data scouting” are stored at higher rates with a smaller event content. The HLT data scouting differs from the L1 scouting (described in section 10.6) in that events must still satisfy a subset of L1 triggers before being written to disk at a high rate.

The HLT runs on a cluster of 200 nodes, each equipped with two AMD EPYC “Milan” 7763 CPUs, two NVIDIA T4 GPUs, and 256 GB of memory, running Red Hat Enterprise Linux 8. The HLT farm is described in section 9.5, with the details specific to Run 3 highlighted in section 9.5.5.

The HLT data processing uses the concept of “paths” to structure its workflow. These paths are sequences of algorithmic steps designed to reconstruct physics objects and make selections based on specific physics requirements. Steps within a path are typically organized in ascending order of complexity, reconstruction refinement, and physics sophistication. For example, the resource-intensive track reconstruction process is usually carried out after completing a series of initial reconstruction and selection steps involving the data from the calorimeters and muon detectors.

The reconstruction modules and selection filters of the HLT use the same software framework that is also used for offline simulation, reconstruction, and analysis (CMSSW [256]). As noted in section 9.5.3, HLT paths selecting similar physics object topologies are grouped into primary data sets for subsequent offline processing, and collections of primary data sets are organized into streams for efficient handling.

In preparation for Run 3, the HLT software was adapted to make use of heterogeneous computing architectures, and several reconstruction modules were developed to take advantage of that to meet the challenges of processing data at ever increasing luminosity and pileup. Algorithms implemented to run on both CPUs and GPUs, are automatically directed to run on a GPU if a GPU is available; otherwise, the CPU-based version of the algorithm is executed. The Patatrack project [299] has created parallelized versions of pixel track and vertex reconstruction algorithms that can run on an NVIDIA GPU and were written using the NVIDIA CUDA language. The data structures are optimized for GPU, and the entire reconstruction chain is executed on the GPU to minimize time-consuming data transformations and transfers. A subset of ECAL and HCAL local reconstruction algorithms have also been ported to GPU, also using CUDA. Based on these efforts a reduction in overall event processing time of about 40% has been achieved. More details can be found in section 11.3. To make full use of the gain, CMS deployed and commissioned a filter farm composed of nodes comprising two GPUs in addition to two CPUs, as noted in section 9.5. More information about multithreaded processing and GPU offloading is given in section 12.5.

The HLT selects data for storage through the application of a trigger “menu”, in which the collection of individual HLT paths is configured. The trigger path definitions, physics object thresholds, and rate allocations are set to meet the physics objectives of the experiment. In 2022, the HLT menus for pp data taking typically contained around 600 paths. This includes the primary HLT paths for analysis as well as paths for calibration and efficiency measurements that are typically looser than the primary paths. These latter HLT paths are often “prescaled”, i.e., only a fraction of the events that pass the requirements are actually accepted, to limit processing time and storage rate. Different trigger menus are used for the recording of heavy-ion collision data. The rates, physics breakdown, and CPU timing of the pp menu are described further in section 11.3.

11.2 HLT reconstruction

The HLT paths in the menu depend on the modules that produce the physics objects from the all-silicon inner tracker and from the crystal electromagnetic and brass-scintillator hadron calorimeters, operating inside a 3.8 T superconducting solenoid, together with data from the gas-ionization muon detectors embedded in the flux-return yoke outside the solenoid. A foundation to many of the specific object reconstructions is the particle-flow algorithm [300], which uses information from these systems to identify candidates for charged and neutral hadrons, electrons, photons, and muons. The main features of the HLT physics object reconstruction and improvements for Run 3 are described in the following subsections.

11.2.1 Tracking

Tracking using the hits recorded by the pixel and strip trackers is generally performed iteratively using a combinatorial Kalman filter, starting with tight requirements for the track seeds that become looser for each subsequent iteration. Hits in the tracking detectors that have already been used in a track are removed at the beginning of the next iteration. For Run 2, initially, the track reconstruction in the HLT consisted of three iterations. The first two iterations required four consecutive hits in the pixel detector to seed the tracking. These iterations target first higher p_T tracks and then lower p_T ones, and use the full volume of the pixel detector. The third iteration relaxes the requirement on the number of hits in the pixel detector to three, and is restricted to the vicinity of jet candidates identified from calorimeter information and the tracks reconstructed in the two previous iterations. The track reconstruction is limited to tracks that are consistent with the leading vertices reconstructed with the pixel detector (those vertices with the largest summed p_T^2 of pixel tracks).

In 2017, several issues with the Phase 1 pixel detector were identified that led to a nonnegligible fraction of inactive pixel modules (section 3.1.1). During the 2017–2018 year-end technical stop, the performance of the pixel detector was restored. Nevertheless, an additional recovery iteration was added to the tracking, to safeguard against a recurrence of this or possible other detector failures. In particular, track seeds consisting of just two pixel hits were allowed to be reconstructed in regions of the detector where two inactive modules overlap.

For Run 3 the seeding and tracking were significantly revised. Tracking is now performed using only a single global iteration, and is seeded by a loose selection of the pixel tracks reconstructed by the Patatrack algorithm, which offers improved performance over the four-hit pixel tracking used for data taking in 2018 [299]. To be used as track seeds, Patatrack pixel tracks must satisfy a loose selection requiring three or more pixel hits, $p_T > 0.3$ GeV, and be compatible with a primary

vertex candidate. Despite fewer iterations and less CPU time required, the Run 3 tracking has improved performance over that used for Run 2. The tracking efficiency and fake rate are measured in simulated $t\bar{t}$ events with an average pileup of 63. Figure 131 shows the Run 3 tracking efficiency and fake rate as determined from Monte Carlo (MC) simulation. The tracking efficiency is defined as the fraction of simulated particles from the signal interaction within the considered p_T and η regions, and longitudinal (transverse) impact parameters <35 (70) cm that are matched to a reconstructed track. The fake rate is defined as the fraction of reconstructed tracks that could not be matched to a simulated particle. The Run 3 tracking efficiency is higher than that of Run 2 for $p_T > 0.7$ GeV, mostly in the central tracking region and the overall fake rate is lower, particularly around the transition region between the barrel and the endcap pixel detectors ($0.9 < |\eta| < 2.1$).

11.2.2 Muons

Tracking algorithms are also deployed to identify and reconstruct muons measured in the muon detectors and in combination with the pixel and strip trackers. While the algorithms used during Run 2 are described in more detail in ref. [301], a brief summary is given here along with the changes implemented for Run 3.

Muon track reconstruction at the HLT takes place in two steps: first using hits only in the muon system (L2 reconstruction), followed by a combination with hits in the inner tracking system (L3 reconstruction). The L2 reconstruction is equivalent to the offline standalone muon reconstruction. The L3 reconstruction is seeded by an L2 muon and follows an iterative track reconstruction similar to that described in the previous section in a region around the seed starting from the outer tracking layers and working inward (“outside-in”) or from the inner tracking layers working out (“inside-out”). The inside-out approach can also be seeded directly by L1 trigger muons. The L3 reconstruction is essentially 100% efficient with respect to that in the L1 trigger, while consuming only about 30% of the overall HLT CPU time. After the muon track reconstruction, identification criteria are applied, as well as isolation criteria for the isolated muon category. The isolation is based on the sum of p_T from additional tracks associated with the primary vertex and calorimeter energy deposits that are clustered using a particle-flow algorithm in a cone of radius $\Delta R = 0.3$ around the muon. The estimated contribution from pileup to the energy deposits in the calorimeter is subtracted.

For Run 3, several modifications to the muon reconstruction were made to improve the HLT performance, in particular with respect to CPU timing. The L2 muon reconstruction was extended to include hits from the GEM detectors, described in section 6.4. At L3, the tracking was adapted to make use of the Patatrack pixel track seeds, followed by a single iteration with the full tracker. The efficiency for low p_T muons was improved by optimizing the search regions around seeds in which muon tracks are reconstructed in the tracker, achieving an efficiency above 80% for $p_T > 2$ GeV. Additionally, machine-learning algorithms were incorporated: a boosted decision tree (BDT) classifier for the search algorithm and a neural network (NN) algorithm for seeding strategy. The inside-out algorithm uses the BDT classifier to consider only the seeds with high quality. The NN algorithm was developed to choose the best seeding strategy for the outside-in reconstruction from the L2 muon and thus limit the total number of seeds considered. With these changes, the L3 efficiency remains unchanged, while the CPU time was reduced by 15%.




Figure 131. Tracking efficiency for the Run 2 HLT tracking (blue) and the Run 3 HLT single-iteration tracking (red) as a function of the simulated track p_T (upper left) and track η (upper right). Only simulated tracks with $|\eta| < 3.0$ are considered in the efficiency measurement, with $p_T > 0.4$ (0.9) GeV required for the upper left (right) plots. The tracking fake rate (lower) is shown as a function of the reconstructed track η for the Run 2 HLT tracking (blue) and the Run 3 HLT single-iteration tracking (red).

11.2.3 Electrons and photons

The HLT electron and photon identification uses the L1 calorimeter trigger candidates as a starting point to perform a “regional” reconstruction of the energies deposited in the ECAL crystals around them. Subsequently, superclusters (clusters of ECAL deposits within a certain geometric area around the seed cluster) are built using the same reconstruction algorithm as used offline [7]. The energy correction applied to HLT superclusters is simpler than the one used offline in that it uses ECAL information only. Requirements are then imposed on the minimal energy, as well as other properties of the energy deposits in the ECAL and HCAL subdetectors.

For electrons, the ECAL supercluster is associated to a reconstructed track with a direction compatible with its location. The first step is a match with pixel detector hits. Since 2017, the pixel matching algorithm requires three pixel hits, to maximize early background rejection, while a hit doublet is accepted only if the trajectory passes through a maximum of three active modules. If the supercluster is successfully matched with the pixel hit seeds, the electron track is reconstructed using a Gaussian sum filter (GSF) tracking algorithm [302].

Variables to enhance the identification of true electrons and photons are applied based on the shower shape in the ECAL, the energy deposition in the HCAL, and, in the case of electrons, the matching between the track and the ECAL supercluster, as well as the quality of the GSF track. As with the other leptons, isolation criteria are generally applied, except for some specific paths, to electrons and photons based on the calorimeter energy deposits in a cone of radius $\Delta R = 0.3$ around the electron or photon and the sum of p_T from additional tracks associated with the primary vertex. Several HLT paths with different isolation criteria and p_T thresholds are defined to provide a range of efficiencies and rates for specific physics analyses.

11.2.4 Tau leptons

The reconstruction of hadronic tau-lepton decays (τ_h) at the HLT is also of crucial importance for the physics program. During Run 2, it was performed in three steps. The first step, the L2 reconstruction, is seeded by L1 τ_h candidates. The energy depositions in the calorimeter towers around the candidates within a cone of radius 0.8 are clustered, and L2 τ_h candidates are reconstructed by using the anti- k_T algorithm [303, 304] with a distance parameter of 0.2.

In the second step, known as L2.5, a charged particle isolation criterion, based on pixel detector information, is implemented. Pixel tracks are reconstructed around L2 τ_h candidates with $p_T > 20$ GeV and $|\eta| < 2.5$ in a region of $\Delta\eta \times \Delta\phi = 0.5 \times 0.5$. Tracks originating from the primary vertex with a transverse impact parameter $d_{xy} < 0.2$ cm, at least three hits, and a trajectory in an isolation cone of $0.15 < \Delta R < 0.4$ around an L2 τ_h candidate, are considered for the isolation sum. An L2 τ_h candidate is considered isolated if the scalar sum of the p_T of the associated pixel tracks is less than 4.5 GeV.

The final step, the L3 reconstruction, includes track reconstruction using the full tracker. For Run 2, tracking used a reduced number of iterations to fit into CPU time budget. Moreover, the track reconstruction was performed regionally around the L2 τ_h candidates. Until mid 2018, the L3 reconstruction was performed using a cone-based algorithm. It was then upgraded to the hadrons-plus-strips (HPS) algorithm [305] that is also used in offline reconstruction. Both algorithms start with jets reconstructed by the anti- k_T algorithm with a distance parameter of 0.4.

In Run 3, the L2 and L2.5 sequences were replaced by a convolutional neural network, in which the pixel tracks from the Patatrack algorithm and the calorimeter candidates are used as input. This improved the rejection rate at L2 by about a factor of 2 for similar efficiency. Additionally, the efficiency of the L3 reconstruction was increased by introducing a neural network, DeepTau [306], adapted from the offline reconstruction such that it matched HLT requirements for speed and performance.

11.2.5 Jets and global energy sums

Jets are reconstructed at the HLT using the anti- k_T clustering algorithm [303, 304] with a nominal distance parameter of 0.4, or of 0.8 in the case of wide jets used in boosted topologies and multijet triggers. Inputs to the jet algorithm are usually particle-flow candidates, or alternatively calorimeter

towers. Depending on detector and beam conditions, corrections are applied to the jet energy scale, as well as the measured particle-flow hadron and average pileup energies.

Triggers using jets and global energy sums are used in CMS across a wide spectrum of physics analyses. Multijet trigger paths, for example, are key to select vector boson fusion event candidates which contain two very forward jets in opposite endcaps, with a large angular separation and a large dijet invariant mass. Final states with boosted multijet signatures can also be identified using dedicated jet substructure techniques such as the soft drop approach [307]. Signatures with many jets in the final state can also be triggered using H_T , the transverse energy sum of all jets, or S_T that combines jets with leptons.

Paths based on missing transverse momentum $\vec{p}_T^{\text{miss}}$, defined as the negative vector sum of the p_T of input objects, also exist. For these paths, the accounting for noise and beam-induced backgrounds is especially important in order to keep rates and resolutions under control. Trigger paths based on p_T^{miss} alone, or in combination with jets, leptons, or photons in the event, are also used, e.g., to search for weakly interacting particles.

11.2.6 b jet tagging

The identification of b jets at the HLT is essential in order to enhance the fraction of events containing heavy flavor jets from processes like vector-boson associated Higgs boson production where the Higgs boson decays into a pair of b quarks. Such processes would otherwise be unlikely to pass the standard thresholds for leptons, jets, or missing transverse momentum.

Since b tagging relies on the measurement of tracks that are displaced with respect to the primary vertex, both the pixel and the silicon strip trackers are used to improve the spatial and momentum resolutions of such tracks. In 2016, the combined secondary vertex algorithm CSVv2 was used. Subsequently, in 2017 and 2018, the multiclassifier neural network DeepCSV was implemented [308]. With DeepCSV, the b jet tagging efficiency was improved by 5–15% at constant gluon or light-quark misidentification rates.

For Run 3, two new neural network taggers, DeepJet [309] and ParticleNet [310], were deployed in 2022, with further improved performance. In addition to tracks, the DeepJet algorithm also uses information from neutral and charged particle-flow jet constituents. The ParticleNet algorithm provides multiclass jet-flavor classification for categories of b, c, and light quarks, gluons, and hadronically decaying tau leptons. For use in HLT, the Run 3 tagging algorithms were trained on dedicated HLT-reconstructed simulation samples.

Figure 132 shows the light-flavor jet misidentification rate versus the b jet efficiency for the different tagging algorithms, evaluated on simulated top-quark pair production events with an HLT jet selection of $p_T > 30$ GeV and $|\eta| < 2.5$. Compared to the performance of the DeepCSV algorithm used during Run 2, the DeepJet algorithm trained using HLT quantities has a light-flavor jet misidentification rate that is lower by about a factor of 3 (up to efficiencies of about 75%). The ParticleNet algorithm reduces the misidentification rate by another factor of 2.5.

11.2.7 New HLT paths for long-lived particles

In addition to the improved reconstruction algorithms discussed in section 11.2, the Run 3 HLT menu has been significantly expanded to explore new and unconventional physics signatures. The Run 3 HLT menu includes new dedicated triggers targeting long-lived particle (LLP) signatures, such as




Figure 132. Light-flavor jet misidentification rate versus the b jet efficiency for the various b tagging algorithms. The solid curves show the performance of the DeepCSV (blue), DeepJet (red), and ParticleNet (magenta) algorithms in the HLT. The dashed curves show the corresponding offline performance for DeepJet (red) and ParticleNet (magenta) taggers using offline reconstruction and training.

displaced jets, displaced leptons and photons, and delayed jets using timing information from ECAL and HCAL, as well as HCAL depth information, as detailed in section 5.3. Furthermore, it features high-multiplicity trigger (HMT) paths that trigger on hadronic showers in the muon system, and new triggers for B physics, as described in section 11.5. In 2023, further resources were deployed to enhance the sensitivity for demanding event topologies, such as the production of Higgs boson pairs decaying with at least two b jets, vector boson fusion (VBF) processes, and events with LLP signatures. Several of these new HLT paths are also seeded by new dedicated L1 trigger candidates (section 10.4), further improving the signal acceptance for their target searches. The new triggers for B and VBF physics will be described in more detail in section 11.5.

There are now several flavors of displaced and delayed jet paths available at the HLT in Run 3. First, there is a suite of displaced dijet paths that were already available to a certain extent in Run 2, but that have gone through major improvements over time. These displaced dijet paths are either inclusive, in that they select events with calorimeter $H_T > 650$ GeV and two jets with less than two prompt tracks each, or they are more exclusive and require events with calorimeter $H_T > 430$ GeV and two jets with less than two prompt tracks and at least one displaced track. For Run 3, the displaced jet selections have been improved at the HLT. In particular, the selection on the number of prompt tracks was tightened in order to reduce the rate, while at the same time the definitions of prompt and displaced tracks were loosened in order to improve the signal efficiency for low-mass LLPs. Furthermore, additional L1 trigger seeds have been added. An L1 seed that requires a single muon in addition to a small amount of H_T helps to reduce the displaced dijet H_T threshold at the HLT. In addition, the new HCAL timing and depth seeds, described in section 5.3, bring improved efficiency for LLPs with $c\tau > 0.5$ m. All of these improvements to the displaced dijet paths provide better efficiency to trigger on low-mass LLPs, especially those with heavy-flavor decays.

In addition to delayed jet paths that use HCAL timing, there are also new HLT paths that exploit the ECAL timing. For LLPs that produce jets with delays of about 1 ns or more, the signal efficiency is improved by an order of magnitude, with respect to the MET triggers that were available for this analysis in Run 2. In particular, there are two different kinds of delayed jet triggers that use ECAL timing at the HLT; there are paths that are seeded by H_T , and there are paths that are seeded by L1 τ objects. For both seeds, different requirements are made at the HLT, namely one or two jets, independently of whether those jets are trackless or not, and the amount of timing delay. The paths seeded by H_T improve the sensitivity to low-mass LLPs with respect to the MET paths, and the paths that are seeded by L1 τ objects increase the efficiency to trigger on Higgs boson decays to long-lived scalars that decay to four b jets as well as to four τ leptons.

Neutral LLPs with particularly long lifetimes could decay hadronically beyond the calorimeters, creating a high-multiplicity shower in the muon system. Such showers are expected to consist of hundreds of hits, but no tracks or jets reconstructed in the inner detectors. Essentially, the CMS muon system would act as a sampling calorimeter. As mentioned in section 10.2.3, new L1 seeds have been developed to collect these high-multiplicity events in the CSCs. The high-multiplicity triggers (HMTs) at L1 are used to seed several HLT paths. At the HLT, a clustering of hits in the muon system is performed using the Cambridge-Aachen algorithm [311, 312]. The first HMT HLT path reconstructs a single CSC cluster, with stricter cluster requirements than at L1 in order to control the rate. The second HMT HLT path reconstructs a CSC cluster as at L1, and then additionally requires a cluster in the DTs with at least 50 hits. The last available HMT HLT path reconstructs a single DT cluster with 50 hits, makes no requirements on CSC clusters, and is seeded by MET triggers at the L1 trigger. As compared with the MET triggers that were available in Run 2, the trigger efficiency for these unique signals is improved by factors of 3 to 20 depending on the path.

Paths for displaced muons at the HLT have also been improved in Run 3. Displaced dimuon paths are seeded by two L1 muons with low p_T thresholds and by new displaced kBMFTF double muon seeds with unconstrained p_T and d_{xy} , as described in section 10.2.1. These seeds feed into several types of displaced dimuon paths at the HLT. There are L2 double-muon paths that require an in-time collision based on the beam pickup timing device with a veto on prompt muons, complemented by paths that make use of a seed developed for cosmic ray muons. These two paths require displacements of at least 1 cm. Lastly, there are L3 double-muon paths that require displacements of at least $100\ \mu\text{m}$. This suite of HLT paths covers a wide range of displacements and improves the signal efficiency over that of Run 2. The efficiency, measured in a cosmic ray muon sample recorded in 2022, was measured to be 100% for displacements $d_{xy} < 100\ \text{cm}$ for the L2 pp seed + prompt-veto path, and 90% for $d_{xy} < 350\ \text{cm}$ for the L2 cosmic seed + prompt veto path. The efficiency of the L3 displaced dimuon path is measured to be 85% in the data sample of non-prompt J/ψ events. At the same time, the background efficiency is small: it is measured to be $< 1\%$ in Drell-Yan data events for all three displaced dimuon HLT paths.

11.3 Run 3 HLT menu composition, rates, and timing

While the complete list of HLT paths in the Run 3 HLT menu is too long to be listed here, a representative sample of some standard triggers with their HLT thresholds and rates is provided in table 15.

Table 15. HLT thresholds and rates of some generic triggers in the Run 3 HLT menu. The rates were obtained from measurements during an LHC fill in November 2022 and have been scaled to a luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

HLT algorithm	Rate
Isolated muon with $p_T > 24 \text{ GeV}$	250 Hz
Isolated electron with $E_T > 32 \text{ GeV}$	182 Hz
Particle-flow based $p_T^{\text{miss}} > 110 \text{ GeV}$	81 Hz
4 PF jets with $p_T > 70, 50, 40, \text{ and } 35 \text{ GeV}$ with two b tags	57 Hz
Two isolated tau leptons with $p_T > 35 \text{ GeV}$	54 Hz
Muon with $p_T > 50 \text{ GeV}$	51 Hz
Two electrons with $E_T > 25 \text{ GeV}$	21 Hz
AK4 PF jet with $p_T > 500 \text{ GeV}$	16 Hz
Two same-sign muons with $p_T > 18 \text{ and } 9 \text{ GeV}$	10 Hz




Figure 133. The HLT rate allocation by physics group for the Run 3 menu deployed in November 2022, scaled to a luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The total rate (blue bar) is the inclusive rate of all triggers used by a physics group and the pure rate (green bar) is the exclusive rate of all triggers unique to that group. The shared rate (orange bar) is the rate calculated by dividing the rate of each trigger equally among all physics groups that use it, before summing the total group rate.

Figure 133 shows the “standard physics” HLT rates consumed by each CMS physics group, estimated from a fraction of events recorded by CMS from a fill taken in November 2022. “Standard physics” refers to the collection of triggers included in the primary data sets whose reconstruction starts within 48 hours after the data were recorded. This does not include data scouting at the HLT or parking triggers, as detailed in sections 11.4–11.5. The average delivered instantaneous luminosity during this fill was $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, but the measured rates have been scaled to correspond to a luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. For the rate measurement, events are assigned to a physics group if the group uses at least one of the HLT algorithms that triggered the event. The group categories correspond to the physics analysis working groups: Higgs boson physics (HIG), searches for new physics in boosted signatures (B2G), searches for new physics in final states with imbalanced transverse momentum (SUS), top quark physics (TOP), standard model physics (SMP), B physics (BPH), and searches for exotica (EXO). The calibration category corresponds to all HLT algorithms used for subdetector alignment and calibration purposes. The “objects” category corresponds to HLT algorithms used for monitoring and calibration by the so-called physics object groups.

What can be seen from figure 133 is that roughly one third of the menu rate has been devoted to standard model physics processes (HIG, SMP, TOP), one third to searches for physics processes beyond the standard model (EXO, SUS, B2G), and the remaining one third to B physics processes, physics objects (for monitoring and calibration purposes for example), and subdetector calibration. The trigger selection for B physics and LLP (included in EXO) are largely unique to those groups.

The distribution of HLT time spent processing the data is shown in figure 134. The processing time running only on CPUs (left) is compared to that when part of the reconstruction is offloaded to GPUs (right). These results were obtained for an HLT configuration representative of the 2022 conditions, running over a sample of 64 000 pp collision events with an average pileup of 56 collisions. The measurements were performed on a machine identical to those used in the HLT farm, as described in section 9.4, equipped with $2 \times$ AMD EPYC Milan 7763 CPUs and $2 \times$ NVIDIA T4 GPUs. The node was configured identically to the HLT farm, with simultaneous multithreading (SMT) enabled, NVIDIA multiprocess server (MPS) enabled, and running eight jobs in parallel with 32 CPU threads and 24 concurrent events each. The average processing time per event is 690 ms when running only on CPUs and 384 ms when offloading part of the reconstruction to GPUs, corresponding to a speedup of over 40%. The maximum processing time per event for the initial Run 3 event filter farm configuration is 500 ms, as noted in section 9.5.5.

11.4 Data scouting at the HLT

A limiting factor for the data-acquisition rate is the bandwidth of the data to record on disk (a few GB/s), not the event rate per se. Thus, if the size of the data per event is reduced, a higher rate of events can be recorded, i.e., using significantly lower trigger thresholds. In the so-called “HLT data scouting”, only the most relevant physics information is stored, as reconstructed by the HLT, and not the complete set of raw data. Scouting was first implemented in Run 1 [313], and was developed further during Run 2, for selected physics objects, such as jets [314] and dimuons [315]. For instance, the H_T threshold in the HLT scouting data was reduced from 800 to 410 GeV, and this data was used for a search for three-jet resonances [316].

A further benefit of HLT data scouting is that events are reconstructed only once, using the resources of the HLT, and thus do not require further computing resources to perform the offline




Figure 134. Pie chart distributions of the processing time for the HLT reconstruction running only on CPUs (left) and offloading part of the reconstruction to GPUs (right). The slices represent the time spent in different physics object or detector reconstruction modules. The empty slice indicates the time spent outside of the individual algorithms.

reconstruction step, which is described in section 12. On the other hand, a complete reprocessing of HLT scouting data is not possible. The quality of the scouting data depends on the calibrations and alignments used in the HLT, but these must anyway be precise to maintain good trigger performance. Furthermore, the CPU processing time budget per event is limited at the HLT (less than about 300 ms per CPU core during Run 2) and scouting requires extra processing, which must fit in the constraints.

During the beginning of Run 3, HLT scouting data was recorded with a rate of up to 30 (22) kHz for the 2022 (2023) data-taking periods, respectively, and an event size of about 13 kB, compared to the full raw data event size of about 1 MB. Events are reconstructed in the HLT scouting scheme if they are accepted by an array of L1 triggers targeting one or two electrons, muons or jets or a moderate amount of H_T with thresholds lower than for other HLT paths. For Run 3, a special version of the particle-flow reconstruction algorithm using pixel tracks reconstructed with Patatrack (section 11.2.1) was deployed. This special version allowed offloading to GPUs at the expense of a slightly degraded parameter resolution in comparison to tracks reconstructed using the full tracker information. The reduced reconstruction time made it possible to use particle-flow reconstruction on a larger fraction of input events.

In addition to the objects from the particle-flow reconstruction [300] that were already stored during Run 2 (muons, jets, and particle-flow candidates), HLT scouting in Run 3 includes the reconstruction and storage of electrons, photons, and tracks. At the beginning of the 2023 data-taking period, the electron reconstruction was further optimized by loosening the L1 seeding requirement, resulting in an increased efficiency for low p_T electron reconstruction. No event selection is applied after the reconstruction. Reconstructed objects are stored if they fulfill relatively loose criteria, such as $p_T > 20$ GeV and $|\eta| < 3$ for jets, and $p_T > 0.6$ GeV and $|\eta| < 3$ for particle-flow candidates. In addition to kinematic quantities like p_T and η , other information to facilitate offline analyses is also

stored. In the case of charged particle-flow candidates, this includes parameters of associated tracks to facilitate, e.g., the training of jet tagging using machine-learning algorithms. Muons are stored not only with track parameters but also with information about the hits in the tracker, allowing the possibility to refit the dimuon vertices.

Since HLT scouting is able to access low-momentum objects with higher rates than conventional HLT trigger paths, it is well suited for analyses targeting low momenta and low-mass particles. Current studies include, for example, analyses of low-mass dimuons, diphotons, and dielectrons as well as a $H \rightarrow b\bar{b}$ analysis that benefits from the decreased threshold on H_T .

11.5 Data parking

One limiting factor for the HLT output rate is the bandwidth of the prompt event reconstruction at Tier 0. An alternative approach to increase the amount of data available for physics analysis is to increase the storage rate on disk, while delaying the reconstruction of the data until a later time, when the necessary computing resources are available. The reconstruction can be scheduled during a year-end technical stop or a long shutdown, for instance. This concept, known as “data parking”, was already implemented in Run 1 [6] and Run 2 to record additional data for B physics and other studies. During an LHC fill, as the luminosity decreases, the bandwidth to trigger additional events increases, and the trigger thresholds for data-parking events are gradually relaxed to record parking data.

In 2018, for example, the collection of $b\bar{b}$ events was enhanced by tagging and storing events containing at least one displaced muon, e.g., from a semileptonic B decay. The p_T threshold at the HLT for a single isolated muon was 24 GeV for the standard physics menu. For the parked B data the p_T threshold was as low as 7 GeV and the HLT output rate reached 5 kHz at the end of fills, enabling CMS to accumulate about 10^{10} b hadrons [288]. In Run 3, data parking still targets B physics, but it also includes a rich set of other physics data. As of the end of 2022, the parking streams record events with at least one muon candidate with $p_T > 12$ GeV and a transverse impact parameter significance larger than 6; events with two muons with $p_T > 4$ and 3 GeV with an invariant mass less than 8.5 GeV; and events with two electrons with $|\eta| < 1.22$, $p_T < 4$ GeV, and an invariant mass less than 6 GeV. With the addition of the parking streams, the total HLT reaches peak output rates of 6 kHz. This can be seen in figure 135, which shows (separately) the HLT output rates for promptly reconstructed events and for parked data for an LHC fill recorded in 2023 with a peak levelled luminosity of about $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

In 2023 the parking strategy was extended to improve the signal acceptance for critical Higgs boson measurements and searches. By dropping the single-muon parking approach, which was limited to instantaneous luminosities lower than $1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, and by improving the purity of the dielectron triggers, CMS is now able to dedicate bandwidth for events targeting final states with two b-tagged jets, the VBF production mechanism, and LLP signatures. The two b-tagged jet criterion, mainly designed for the production and decay of HH into four b quarks, relies on the presence of four jets with $p_T > 30$ GeV, two loose b-tagged jets using the ParticleNet tagger, and an aggressive threshold $H_T > 280$ GeV. Figure 136 (left) shows the trigger efficiencies for prompt and parking data. A clear improvement of more than 20% is observed with respect to the trigger efficiency of Run 2. The measured efficiency in a single-muon dataset recorded in 2023 confirms a plateau efficiency of more than 90% (figure 136, right).




Figure 135. The HLT rates for promptly reconstructed data streams (blue) and parked data (black) as a function of time during an LHC fill in 2023.




Figure 136. Left: comparison of the trigger efficiency of the $HH \rightarrow b\bar{b}$ trigger among the three different strategies used in Run 2 (black), 2022 (blue), and 2023 (orange) using the signal MC sample. Right: trigger efficiency of the $HH \rightarrow b\bar{b}$ trigger using events collected by the single muon trigger in 2023.

For VBF event candidates a twofold strategy, one inclusive and one exclusive, is pursued. The inclusive approach relies only on the invariant mass of two forward jets and applies a tight selection of $m_{jj} > 1000 \text{ GeV}$, while in the exclusive approach a looser threshold on the invariant mass is affordable because there are additional requirements on the central objects in the events. Table 16 summarizes the details of the two complementary approaches.

11.6 Heavy ion physics

Heavy-ion collisions impose unique challenges on the DAQ and HLT systems, as also noted in section 9. A completely custom HLT menu was developed with almost no overlap in content

Table 16. HLT thresholds and rates of the VBF triggers, as obtained from measurements during an LHC fill in June 2023 at an average instantaneous luminosity of $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to a pileup of 61.

HLT algorithm	Rate [Hz]
2 jets with $p_T > 105/40 \text{ GeV}$, $m_{jj} > 1000 \text{ GeV}$, $\Delta\eta_{jj} > 3.5$	720
2 jets with $p_T > 70/40 \text{ GeV}$, $m_{jj} > 600 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, 2 central jets $p_T > 60 \text{ GeV}$	430
2 jets with $p_T > 105/40 \text{ GeV}$, $m_{jj} > 1000 \text{ GeV}$, $\Delta\eta_{jj} > 3.5$, 3 central jets	120
2 jets with $p_T > 90/40 \text{ GeV}$, $m_{jj} > 600 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, isolated muon $p_T > 3 \text{ GeV}$	110
2 jets with $p_T > 75/40 \text{ GeV}$, $m_{jj} > 500 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, $p_T^{\text{miss}} > 85 \text{ GeV}$	110
2 jets with $p_T > 45 \text{ GeV}$, $m_{jj} > 500 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, tau $p_T > 45 \text{ GeV}$	40
2 jets with $p_T > 90/40 \text{ GeV}$, $m_{jj} > 600 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, isol. μ $p_T > 3 \text{ GeV}$, 3 jets	12
2 jets with $p_T > 45 \text{ GeV}$, $m_{jj} > 500 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, electron $p_T > 12 \text{ GeV}$	5
2 jets with $p_T > 75/40 \text{ GeV}$, $m_{jj} > 500 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, $p_T^{\text{miss}} > 85 \text{ GeV}$, 3 jets	5
2 jets with $p_T > 70/40 \text{ GeV}$, $m_{jj} > 600 \text{ GeV}$, $\Delta\eta_{jj} > 2.5$, 2 jets $p_T > 60 \text{ GeV}$	3

and paths with the pp menu. Additionally, because of the dense environment, the physics object reconstruction algorithms used by the HLT are generally customized for heavy-ion running. For example, while jet reconstruction remains based on the anti- k_T algorithm, the underlying event energy subtraction differs from that used to handle pileup in pp running.

In preparation for heavy-ion data taking in Run 3, significant effort is devoted to increasing the available L1 bandwidth. In 2018 up to 30 kHz was achieved, whereas we plan up to 50 kHz in Run 3 conditions. The most critical components to study include the ECAL, pixel, and tracker detectors. During a heavy ion test run in 2022, a scan of different ECAL readout settings was performed to study the readout size and potential impact on physics objects. A size reduction is crucial to achieve the target L1 rate. In order to fully utilize the L1 bandwidth, fractional prescale factors will be used to operate the trigger at the optimal point between rate and dead time.

The target for the PbPb collision run in 2023 is to record about 5×10^9 events. The HLT will be operated at an output rate above 10 kHz. About 10 kHz are “minimum bias” events, collision events with only a minimal L1 selection. The budget also contains 1–2 kHz of triggers of selected physics objects, such as muons, electrons, photons, jets, and track multiplicity conditions in central collisions. A significant fraction of the latter events will also be contained in the minimum bias data set.

To use the bandwidth between the HLT and the Tier 0 center in an optimal way, all triggers are collected into a single “HIPhysics” data set. Only later, during the offline processing step, which is described in section 12, will they be split into secondary data sets. Tools have been developed to automatically configure the splitting of large data sets into multiple outputs.

In addition to hadronic collision events, a suite of ultra-peripheral collision triggers will also be deployed. These target, for example, processes where the two ultrarelativistic nuclei do not collide directly, but their electromagnetic fields interact with each other. Such events are typically very clean, with only a few physics objects in the detector and no other visible activity. The size of the event data is minimal and contributes little to the bandwidth.

With the aging of the detector, the efficiency to trigger on some of the more peripheral collisions drops. Therefore a task force was formed to understand how the decrease in efficiency can be

avoided, including, for example, combining information with other detector parts like the Zero Degree Calorimeter. The heavy ion event size for hadronic collisions is comparable to that for pp collisions, and a sustained throughput of 17 GB/s is anticipated. Even though the bandwidth is much higher than that in Run 2, given the event size, the target number of recorded minimum bias events will be hard to achieve. To increase the event rate further, a new approach has been developed in which one of the most significant components of the event record, the tracker information, is removed. Instead, for each strip cluster, only summary information is written out. This approach is expected to reduce the event size by about 30%. The algorithm was commissioned during the 2022 heavy-ion test run, where both event format contents (the reduced format, “RawPrime”, and the complete information) were written out.

12 Offline software and computing

12.1 Overview

CMS offline computing has evolved over the past 15 years to support the ever-growing needs to trigger, filter, store, transfer, calibrate, reconstruct, and analyze the recorded and simulated data of the experiment. The offline system receives a subset of the real-time detector information from the data acquisition system, once it is filtered by the high-level trigger at the experimental site, ensures safe curation of the raw data, and produces data for physics analysis. Major activities are also the production and distribution of Monte Carlo (MC) simulation data, as well as the processing of conditions and calibration information, and other nonevent data. The data input and output layer of the CMS data processing software is provided by the ROOT framework [259].

Key components of the offline computing system, described in the sections that follow, include an event data model and corresponding application framework, the processing chain and data tiers, computing centers, referred to as Tier 1, Tier 2, and Tier 3, which provide storage and processing resources all connected through a distributed world-wide computing grid, and a set of computing services that provide tools to transfer, locate, and process the data.

A timeline of the major data processing and computing software improvements put into production over the past decade or so is shown in figure 137. CMS developed a highly flexible computing model, with the goal of using efficiently all of the resources available to the experiment while minimizing both hardware and personnel needs and maximizing overall throughput. The innovations led to an increase in efficiency in the use of computing resources, and enabled the experiment to use additional new and diverse types of computing resources that were not available a decade ago.

12.2 Detector simulation

The CMS detector simulation is based on the GEANT4 [317] toolkit. It is augmented with computationally efficient techniques, such as shower libraries for the forward calorimeters, and specific identification criteria per particle type and detector region for neutrons [318], which guarantee high fidelity of the simulation. The “Full MC” simulation chain includes execution of the standard reconstruction chain. A subset of physics analyses, such as scans of new-physics signatures over large parameter space, use the “Fast MC” chain or “Fast Simulation” application. This does not rely on GEANT4, but rather a parameterized approach [319] with a simplified CMS geometry. The




Figure 137. Timeline of the major data processing and computing software improvements put in production since 2010.

Fast MC chain also uses a faster version of the reconstruction step. The performance of the CMS simulation is detailed in ref. [320].

12.3 Event reconstruction

Event reconstruction is the processing step that transforms the event information contained in the raw data, i.e., packed detector readout data, into high-level physics objects, such as electrons, muons, photons, or jets, which are used in physics data analysis. In CMS, the event reconstruction is logically partitioned into several steps, starting from the local reconstruction in which the data are processed individually by a single detector component, and ending with global particle-flow reconstruction [300] and object identification. The particle-flow algorithm aims to reconstruct and identify each individual particle in an event, with an optimized combination of information from the various elements of the CMS detector.

A great deal of attention is dedicated to the computing performance of the event reconstruction, both for pp and heavy-ion collision data. The run-time optimization is a priority for CMS, and is achieved through algorithmic and technical improvements. Examples of technical improvements are the adoption of a more recent compiler version, utilization of special compilation flags and other code optimizations, without impact on physics performance. The timing is benchmarked continuously and in great detail in order to identify and remove performance degrading patterns as early as possible in the release integration process [321].

Run-time performance improvements must be paired with precise physics validation of the software, delivered swiftly to the code authors. Physics validation happens at multiple levels during the software integration process, mainly through automatic comparison of data quality monitoring (DQM) histograms of relevant distributions and physics quantities. Before new reconstruction code is merged into the central software repository, a limited number of events are processed and the

quality of the reconstruction is assessed. During twice-daily integration builds of the CMS software, a growing level of detail is added to the automatic validation, and a larger number of events are processed to obtain the DQM histograms.

Every major release, which may be intended for data processing or mass production of MC simulation samples, is preceded by weekly or biweekly pre-releases. The validation of pre-releases involves confirmation by detector and physics object experts who assess the quality of the results obtained with the software compared to actual Monte Carlo and data processing campaigns of a few million events. Final validation, referred to as release validation [322], is required before new CMS software releases are put into production.

12.4 Computer architectures and platforms

CMS strives to use efficiently all of the computing resources at the experiment’s disposal. To this end, the data processing software supports a generic mechanism, described in section 12.5, to offload work onto accelerators such as GPUs, described in section 12.5.2.

Our builds support three different CPU architectures, x86_64, ARM, and IBM POWER, driven by the opportunities to obtain allocations at HPC centers where these architectures are available (section 12.7.1). We perform integration and unit testing of the CMS code, and create regular releases of the whole CMS software suite CMSSW [256] for these architectures [323]. ARM support began in 2012 and POWER in 2016. Installations of non-x86_64 releases are performed since 2014 on the file system used as a vector for CMS software, CVMFS, an aggressively cached distributed read-only file system [324, 325].

CMS not only builds, tests, and runs code regularly on different CPU architectures, but also employs two different compilers, GCC and Clang, as well as different operating systems. The combination of a CPU architecture, a compiler, and an operating system is commonly referred to as a “platform”. The ability to build, test, and run CMSSW and perform integration tests on several platforms significantly contributes to achieving top code quality, a capability which CMS plans to preserve for the entire Run 3 and beyond. There is an increased risk of bugs and unstable algorithms silently altering the results of computations without being noticed if the code is tested and executed only on one platform.

12.5 Application framework

12.5.1 Multithreading

For LHC Run 2, the application memory usage was foreseen to increase beyond the 2 GB-per-CPU core limit of the computing grid worker nodes, due to increased pileup compared with Run 1. In order to reduce memory usage per CPU core, the CMS application framework was enhanced to support parallelism through a multithreading paradigm [256, 326]. The multithreaded framework follows a task-parallel paradigm implemented with Intel’s oneTBB [327] library, which expresses concurrent units of work as tasks and passes them to oneTBB’s task scheduler to run. The multithreaded framework was introduced for production jobs in 2015 for Tier 1 computing resources, and during 2016 for the majority of other resources. Currently, production jobs use eight threads by default, which is a good compromise between application memory usage and CPU efficiency. In addition, 8-core slots have been agreed as a standard job size on the shared worldwide LHC computing grid (WLCG) [328] resources.

The initial version of the multithreaded framework processed only separate events concurrently. The framework had to synchronize between the worker threads at specific stages of the data processing, and combined with large variability between the processing times of different collision events, this led to noticeable inefficiencies in the CPU utilization. The threading efficiency has been gradually improved [329], and currently many levels of concurrency are exploited. Within one event, independent modules are run simultaneously whenever possible according to their data dependencies. Modules can use oneTBB's parallel constructs in their internal work including so-called event setup modules that process conditions data. Events from multiple luminosity sections (section 9) and interval-of-validity ranges of the conditions data can be processed in parallel. The framework supports physics modules with different levels of thread-safety guarantee, which are also associated with their threading efficiency or memory usage.

12.5.2 Offloading to accelerators

The framework has generic support for offloading computations from the CPU worker threads, which allows those worker threads to continue to work on other computations [330]. A module that offloads computations has its event processing function split into two stages, where the first function is to offload the computations, and the second function is called when the offloaded computations have been completed.

Support for specific offloading technologies is implemented on top of the framework's generic mechanisms. The Nvidia GPUs on the same computing node are supported with the CUDA API [330]. Utility classes help with asynchronous execution, sharing resources between modules via the event, and minimizing data movements. At the time of writing, pixel local reconstruction, pixel track and vertex reconstruction [299], ECAL unpacking and local reconstruction, and HCAL local reconstruction have GPU implementations. Offloading them to a GPU reduces the HLT CPU usage by about 40% [295].

Historically, the accelerator vendors have provided their own APIs. However, developing and maintaining separate versions of algorithms for each platform is unsustainable and, therefore, CMS investigated ways to achieve performance portability with a single code base. A performance portability framework makes it possible to have one single code base and to build libraries for different classes of hardware, such as CPUs or different vendors of GPUs and accelerators. The suitability of Alpaka [331] and Kokkos [332] libraries for the CMS data processing model and software were explored in detail [333, 334], with the conclusion that Alpaka was better suited for CMS for LHC Run 3.

The ability to use accelerators on a remote computing node, or in separate processes on the same node, could allow more flexible use of accelerators. Using remote accelerators for machine learning (ML) inference has already been demonstrated from the CMS application framework [335–337] using inference servers such as Nvidia Triton [338]. In general, ML algorithms are expected to be easily portable between various accelerators.

12.5.3 Geometry

Geometry information is fundamental for several CMS applications, such as simulation, reconstruction, and event visualization. For the start of Run 3, the in-house developed geometry description tool used during Run 1 and Run 2, DDD [339], was replaced with the community tool, DD4Hep [340],

Table 17. Description of the data tiers regularly produced by centrally managed workflows. The ROOT framework is used to write and read the data.

Name	Description
GEN	Intermediate and outgoing stable ($c\tau \gtrsim 1$ cm) particles from the collision simulation. May include Les Houches accord event (LHE) data from the matrix-element generator, if applicable.
SIM	Detailed description of energy deposits left by stable outgoing particles in the detector material. Two options are available: a highly-accurate GEANT4-based application (Full MC); and a parametric fast simulation application (Fast MC), which trades accuracy for a 100-fold decrease in detector simulation time or 10-fold decrease in total CPU time per simulated event. The level of inaccuracy introduced by Fast MC is typically a difference of less than 10% in final analysis observables.
DIGI	Digitized detector readout or simulation thereof. In simulation, the effect of additional collision events (pileup) is folded into the event description in this step. In Run 2, a “premixing” technique was introduced, where the additional events are summed in a separate processing step and then applied to the simulated primary event.
RAW	Packed detector readout data.
RECO	Detailed description of calibrated detector hits and low-level physics objects.
AOD	Reduced description of calibrated detector hits and low-level physics objects, uncalibrated high-level physics objects.
MiniAOD	Reduced low-level physics objects and calibrated high-level physics objects. A truncated floating-point representation is used for most object attributes. Introduced for Run 2 to reduce the number of analyses requiring AOD inputs.
NanoAOD	Compact data format containing only high-level physics object attributes stored as (arrays of) primitive data types. Introduced during Run 2 to reduce the number of analyses requiring MiniAOD inputs.

which is also used by the LHCb experiment and others. The selection of DD4Hep was made because its library is well behaved in multithreaded environments and because the replacement of an in-house solution with a community-supported tool improves the sustainability of our software stack. The migration of the geometry for Run 3 took place during LS2 and was an opportunity to review the entire description of the CMS detector, even improving it in some respects. The migration to DD4Hep of the Phase 2 geometry was completed during 2022.

12.6 Data formats and processing

Data and simulation processing workflows are broken into several steps, each defined by the output data structures per event it produces, referred to as a “data tier”. The data tiers in use for centrally produced simulation and reconstruction workflows are listed in table 17. No changes in the data formats were made between Run 2 and Run 3. A single executable process may produce multiple outputs corresponding to different data tiers, which reduces I/O operations when the necessary data structures are already in process memory.

The smaller-size analysis formats are key to reducing both the overall amount of data stored by CMS and analysis processing time per event. In Run 2, approximate event sizes in each format are 400 kB for AOD (Analysis Object Data), 40 kB for MiniAOD, and 1–2 kB for NanoAOD. The goal is for 50% of CMS analyses to use only NanoAOD data sets as input before the end of Run 3. Although this format is not suitable for every analysis, the increase in usability and speed should appeal to many users.

Nonevent data are used to interpret and reconstruct events [1]. Four types of nonevent data remain in use: construction data, generated during the construction of the detector; equipment management data; configuration data, comprising programmable parameters related to detector operation; and conditions data, including calibrations, alignments, and detector status information. A procedure for deriving a selection of online calibration constants has been in place since Run 1. This has since been consolidated and the ability to create calibration constants from data collected across multiple runs was developed along with new calibration workflows.

A typical Run 3 simulation workflow will be a so-called “step-chain” job composing GEN-SIM-DIGI/PU MIX-RAW-RECO-AOD steps and the reduction to MiniAOD and NanoAOD formats. For detector data, only the RECO-AOD step is performed, followed by the reduction to Mini and NanoAOD. For a given data-taking period, AOD is produced 1–3 times in large-scale processing campaigns, while MiniAOD [341] and NanoAOD [342] data tiers are reproduced more frequently as high-level physics object calibrations are updated.

In the CMS workflow management system, a processing step applied to a given set of inputs (or requested number of events in the case of GEN) forms a task. Tasks are chained together to form a complete workflow, with intermediate tasks writing their output to site-local storage, and optionally registering the output in the data management system, as discussed in section 12.8.1. To reduce data transfer, task chains are converted to step-chains when possible, with intermediate output kept only on job-local scratch disk and all subsequent steps executed in a single job. As discussed in section 12.5, the steps are typically executed multithreaded, using up to eight cores per executable, which reduces the workflow management overhead. The GEN step uses software tools from the HEP theory community, and although some tools may force a step to be executed single-threaded, work is ongoing to improve per-executable parallelism in this context.

12.6.1 Premixing

In Run 2 the increase in pileup events resulted in a more I/O and computing-intensive pileup simulation. A number of individual minimum-bias events comparable to the pileup level had to be read from local or remote disk pools and superimposed to the hard scatter event. For this reason, a “premixing” simulation method [318, 343] was introduced to drastically reduce the I/O by summarising all the parasitic pileup collisions in one single “pileup-only” event. Hard-scatter events are generated and simulated without pileup; separately, a sample of MC pileup-only events is also produced, using the pileup distribution for a certain running period and considering in-time and out-of-time interactions in all subdetectors. A selected premixed set of pileup events is then overlaid on the hard scatter events. This approach reduces I/O operations by 90% compared with the previous method of overlaying individual minimum-bias events. On average, digitization and reconstruction of simulated events with pileup is twice as fast [343]. The size of the premix library is proportional to the number of simulated events and contained about 200 million events in Run 2. For Run 3, the integrated luminosity is expected to be higher, and larger premixed samples will be generated.

12.7 Computing centers

The organization of the distributed computing infrastructure used by CMS was initially based on a model described by MONARC [344], where the sites were organized in “Tiers”, pledging CPU, disk storage, and tape resources, proportional to the commitment of each funding agency within the CMS Collaboration. The roles of the various Tiers were established as follows:

- A Tier 0 center close to the experiment (CERN) to execute a first calibration and reconstruction pass, the so-called “prompt reconstruction”, as well as to maintain a custodial copy of all RAW data on tape; 24h/7 support is guaranteed;
- Tier 1 regional centers (six in use by CMS), which maintain a second distributed custodial copy of the RAW data on tape, and provide CPU for re-reconstruction and MC simulation; 24h/7 support is guaranteed;
- Tier 2 local centers (about 50 used by CMS), providing support for analysis activity and MC simulation; guaranteed support only during working hours.

The Tier 0 is the largest of the CMS sites. It provides computing capacity, disk and tape storage and is hosted by CERN. At the Tier 0, prompt reconstruction starts 48 hours after the data are acquired. This delay is necessary for an initial set of detector calibrations to be bootstrapped, starting from the execution of alignment and calibration processing sequences on the newly acquired data — the so-called “express” reconstruction. One such workflow is described in ref. [345]. Once the calibrations are derived, they are consolidated in a payload written to a database that is then read by prompt reconstruction at a later stage. The flexibility of the CMS central job submission infrastructure makes it possible to produce substantial amounts of simulated samples at the Tier 0.

France, Germany, Italy, the US, the U.K., and Spain provide the six Tier 1 sites of CMS. These sites have a primary role in the computing model, offering the precious combination of computing capacity, disk storage and tape archival space, all together with high bandwidth connectivity to CERN and Tier 1 s through the LHCOPN private IP network, as well as to other sites through the LHConet network. Tier 1 s are used to store the active copy of the RAW data on tape, to perform re-reconstruction passes when data processing algorithms and calibrations improve so much to require it. The Tier 1 centers have been adapted to support the Tier 0 in the task of prompt data processing, whenever needed. In the current computing model a job can now run, in principle, wherever free CPU is available while accessing input data through the WAN (figure 138). The total amount of computing resources pledged by the largest and smallest of the Tier 1 centers of CMS differs by an order of magnitude, still these special sites are equally critical for the support of the CMS physics program.

The computing model used in Run 2 and Run 3 is a significant evolution of the MONARC hierarchical model. The roles of the Tier 1/2s have become more similar in order to optimize the usage of the resources with high efficiency, and following collaboration-wide priorities. The presence of high-speed wide-area network (WAN) connections, the development of advanced data federation [346, 347] and caching [348] technologies and the creation of optimized data tiers such as MiniAOD and NanoAOD, allow distribution of the tasks among the Tier 1’s and Tier 2’s following CMS priorities, with minimal net distinction between analysis, MC production, and data (re-)processing. However, Tier 1 sites still tend to be assigned a larger proportion of central production jobs.




Figure 138. The evolution of the CMS computing model from a hierarchical (left) to fully connected structure (right).

Tier 3 sites are those which do not pledge resources, but nevertheless provide CPU and storage in varying amounts. These could be sites with batch farms for mainly local use, or those that predominantly support another experiment, but allow CMS to take CPU slots opportunistically when they are available.

An important aspect of the computing model in CMS is the continuous monitoring of the distributed data centers, both via recording of CPU and data metrics of the workloads, and with automatic tests checking periodically the health and reliability of the various services. These tests allow the production system to temporarily disable sites with transient issues and to re-enable them when working again.

In order to easily run computing jobs on diverse resources, the environment was standardized during Run 2 via the use of virtualization and containerization in Singularity [349].

12.7.1 High performance computing (HPC)

National and supranational bodies are investing substantial resources in supercomputers. These machines are part of our computing infrastructure and, given the advantage they represent for many scientific and industrial applications, they are here to stay. In the exascale era, HEP might be compelled to count on these resources for the bulk of its computing capacity. CMS has already started to take advantage of HPCs, adapting its software and computing tools. However, the integration of individually highly optimized HPC systems with the CMS experiment computing environment is not a trivial task. From a technical perspective, HPC centers differ in a variety of specialized hardware setups and strict usage policies can apply, particularly related to security. As a consequence, a variety of conditions can be encountered when attempting to integrate these centers, spanning from less common operating systems, absence of local scratch disk space on computing nodes, low memory availability per core, limited or even absent outbound network connectivity, and, of course, distinct architectures. Given the unique nature of HPC systems, CMS has invested a substantial amount of effort to develop solutions case by case.

CMS has deployed several solutions which can be classified broadly into two categories: a transparent site-extension model where HPC computing resources are seamlessly integrated into an existing CMS grid site, and the HEPCloud model in which an additional layer of job submission infrastructure sits between the experiment's workflow management system and the HPC computing resources. In both cases, the storage of existing CMS grid sites is used through remote access from the HPC, complemented by on-site caches or local storage at some HPC centers.

Integration of HPCs that are transparent to central computing operations have been performed in several countries in Europe, e.g., Italy, Switzerland, Germany, and Spain, for example expanding existing Tier 1 or Tier 2 sites elastically into these machines [350–352], or by deploying so-called overlay batch systems [353]. It is also through this model that a validation of the physics performance of CMS software on the PowerPC architecture could be performed [354].

The HPC resources located in the US are integrated into the CMS systems via HEPCloud [355], a portal to an ecosystem of diverse computing resources, commercial or academic, hosted at FNAL. Since 2020, HEPCloud has provisioned resources from seven different HPC centers located for example at the NERSC, PSC, TACC, or SDSC computing centers.

12.7.2 Data archive at CERN

CMS migrated to the CERN tape archive (CTA) [356] system in December 2020. The CTA system provides the tape backend to the CERN EOS disk system [267], and together EOS and CTA replace CASTOR. This upgrade was necessary to prepare for the higher data rates of Run 3 and beyond, as well as to provide a uniform API for disk and tape operations, with support for newer protocols and authentication methods. It is difficult to achieve high throughput using a traditional buffer made up of spinning disks, due to contention on the drives from multiple simultaneous streams. Therefore a new approach was taken, with a small, fast SSD buffer in front of the physical tapes. This allows the tape drives to operate close to their nominal speed of 400 MB/s.

12.8 Computing services

12.8.1 Data management

During Run 1 and Run 2, CMS used PhEDEx [357] and Dynamo [358] as data management tools. However, for Run 3 and beyond it was necessary to adopt a more scalable, flexible, and powerful system to increasingly automate the data management, and include the possibility to scale up transfers to around 100 petabytes per day by the late 2020s for the start of the High-Luminosity LHC (HL-LHC) [359]. The new software should also support future technologies such as token authorization and non-FTS (File Transfer Service) [360] transfers. The Rucio system [361], a data management project for scientific communities, was adopted by CMS at the end of 2020. It can perform all of the functions of PhEDEx and Dynamo, and also make higher-level decisions about data placement. Rucio is run centrally, without an agent at every grid site as required by PhEDEx.

The CMS Rucio infrastructure is based on Helm [362], Kubernetes [363], and Docker [364], which are industry standards. All of the Rucio services are built into a single Kubernetes cluster which can be brought up from scratch in under an hour. Rucio removes data as additional space is needed at a site. Only data that is not held in place by one or more rules is eligible to be removed. To make this decision, Rucio uses the last access time of the data. This information is taken from job reports, CMSSW file reads, and monitoring of the AAA (Any data, Anywhere, Any time) system, which is described in the following subsection.

CMS developed a mechanism to check the consistency of the Rucio database with the actual state of files on disk. Remote disks are scanned using XRootD tools, comparing the results with the state of the database, and generating lists of missing and unneeded files. These lists are passed to additional Rucio components for file removal or re-transfer.

12.8.2 Data transfer protocols

During Run 1 and Run 2, CMS transferred data among sites via protocols based on the grid security infrastructure (GSI) [365, 366], with GSIFTP the most common protocol used and GridFTP its most common implementation. When the end of support for GSI was announced, the LHC community started looking for a replacement both for the protocol and an authentication mechanism. The WebDAV protocol [367], an extension of HTTP which supports third party copy (TPC) transfers and tokens for authentication, was selected. The adoption of the WebDAV protocol started early in 2020 and was completed in 2022.

The “Any Data, Anywhere, Anytime” (AAA) data federation [346] was introduced in 2014–2015 during LS1. This is a model for effective federation of distributed storage resources via an XRootD cluster at each computing center, allowing for remote access to any file within the CMS namespace. The XRootD framework supports partial reads of files and is commonly used by analysis jobs. The combination of AAA and caching [348] technologies allows nonlocality between data and CPU, and is expected to lead to entirely storage-free computing centers in Run 3. These are the first steps towards a “Data Lakes” architecture [368], a centralized data repository, as envisioned for the HL-LHC era.

12.8.3 Central processing and production

Large-scale MC sample production and data event reconstruction activities are performed in a distributed computing infrastructure, coupled to a specialized workload management system (WMS), described in section 12.8.4. A global batch queue manages the distribution of production and analysis jobs to the CMS distributed computing system in an optimized and flexible way. The submission infrastructure (SI) employs GlideinWMS [369] and HTCondor [370] software suites in order to build and manage a “Global Pool” [371] of computing resources where the majority of the CMS tasks are executed.

The SI comprises multiple interconnected HTCondor pools [372], as shown in figure 139, redundantly deployed at CERN and FNAL in order to ensure a high-availability service. The main component of the SI, the Global Pool, obtains the majority of its resources via the submission of pilot jobs to WLCG [328] and open science grid (OSG) sites. However, locally instantiated processing nodes via, for example, DODAS [373] or BOINC [374], as well as opportunistic resources, such as the HLT filter farm [276] when not in use for data taking, can also be employed. As described in section 12.7.1, the SI computing capacity has recently expanded into HPC facilities which are integrated as part of the Global or HEPCloud pools. CERN on-site CMS resources, along with opportunistic local (BEER [375]) and cloud [376, 377] computing slots, are organized into a third HTCondor pool, built on a dedicated set of hosts to isolate it from potential issues in the main Global Pool, given its critical role in supporting Tier 0 tasks during data-taking periods. Specialized nodes known as “schedds” control workload submission. While being primarily attached to one “pool”, these schedds can interact with other federated pools, requesting additional resources when demands are not covered in the primary pool, as indicated in figure 139.

The SI can leverage any grid, cloud, HPC, or opportunistic resource available to CMS, which in total have tripled in size over the last five years, currently amounting to around 400k CPU cores. The Global Pool matches diverse CMS workloads, which include single-core as well as multicore




Figure 139. Schematic diagram of the submission infrastructure, including multiple distributed central processing and production (WMAgent) and analysis (CRAB) job submission agents (schedds). Reproduced from [372]. CC BY 4.0. Computing resources allocated from diverse origins (green boxes) are grouped into HTCondor pools (gray boxes), federated via workload flocking. The collector and negotiator agents (yellow boxes) keep the state of each pool and perform the workload-to-resource matchmaking.

requests [378] to these resources. A successful scheduling is achieved by simultaneously ensuring that all available resources are efficiently used [379], a fair share of resources between users is reached, and the completion of CMS tasks follows their prioritization, minimizing job failures and manual intervention. While the SI typically manages 100k to 150k simultaneously executing tasks, recent scalability tests [380] have demonstrated the capacity of the infrastructure to sustain in excess of half a million concurrently running jobs.

12.8.4 Workload management system

The CMS production WMS, based on the WMCore framework, comprises many services, including WMAgent, the most important and complex system in this architecture. It is a stateful software, responsible for the work splitting, event and data bookkeeping, job management and submission, while interfacing with the data management tools for global data bookkeeping and grid data transfers. The WMAgent system is tightly coupled to the schedds in the SI, enabling it to make use of the Global Pool resources, handling thousands of requests in parallel and enforcing the workflow prioritization. Other components in the production WMS include the WMStats system, responsible for the workflow and job monitoring, and providing logging and recovery features, and the request manager, which manages and stores workflow descriptions, submitting workloads for processing on grid resources. A suite of “Microservices” has been developed to perform very specific tasks such as automatic input and output data placement, monitoring, and cleanup of the data used in central production. The introduction of Microservices has resulted in multiple features of the system being automated, thus reducing the required operational effort.

12.8.5 Distributed analysis

For CMS scientists, a dedicated tool known as CMS remote analysis builder (CRAB) has been available since the start of Run 1, to manage the submission on the grid of data analysis applications or small-scale productions of data samples that are not centrally managed. CRAB offers the same functionalities as originally planned [1]. Users interact with a thin client running inside the CMSSW environment, which packs the user analysis environment into a sandbox and uploads it together with a configuration file to a central, database-centric service hosted in the CMSWeb services framework described in section 12.8.7. Centralized components run either as Docker containers, or under the control of HTCondor [370] or DAGMan [381]. These components handle preparation, execution, bookkeeping, error recovery, monitoring, and output delivery for the user application. CRAB has been continuously improved to provide increased automation and scaling, leaner operations, better error recovery, access to data archived on tape and integration with the Rucio data management system, while keeping the same basic interface for the user. For users needing a simpler, more interactive service to access resources from the Global Pool to make histograms, for example, the CMS Connect [382] service is available as an alternative to CRAB.

12.8.6 DBS database

The data bookkeeping service (DBS) is a catalog holding event metadata for all simulated and experimental data processed and stored by the CMS experiment. DBS contains all necessary information for tracking the data sets, including provenance information, processing history and parentage relationships between files and data sets. DBS is used by all computing jobs in CMS. The current version (version 3) was completely redesigned and reimplemented following a project review in 2009, in order to better fit the evolved CMS data processing model, to better integrate with the evolved CMS data management and workflow management (DMWM) projects, and to ensure a higher scalability. DBS3 is implemented as a Python-based web service using a standardized architecture provided by the CMS DMWM project. Client-server communication in combination with thin client APIs using the JSON format as a lightweight replacement of XML-RPC led to better scalability. In addition, the database schema has been streamlined and fine-tuned in order to avoid both excessive table joins, reduce query latency, and improve server stability, which enhance data accessibility.

The data aggregation system (DAS) [383] provides access to distributed CMS metadata via a common query language (QL). Users may access DBS, Rucio and other metadata sources by placing simple QL queries without prior knowledge of service APIs, database schemas, and implementation or location of the services.

12.8.7 Web services and security

In order to guarantee successful data taking, CMS uses a series of web applications, which perform various tasks related to job submission, job monitoring and bookkeeping, and the location of the data sets. These services are provided and maintained by several teams and centrally deployed under the umbrella CMSWeb.

To achieve optimal resource utilization and high availability, the services have been containerized and are now being deployed in a Kubernetes cluster [384]. Detailed monitoring, for all the services and the nodes they run on, has been built on top, using the central CMS computing monitoring

tools. Authorization and authentication (A&A) is also being offered as a service in the CMSWeb cluster. Currently, A&A is using X509 certificates, but is rapidly migrating to authentication tokens, following the shift to the OAuth standard initiated by the WLCG. This migration should take place by 2024, according to current plans [385].

Security is a very important aspect of web services. A great effort is required to stay up-to-date with the latest security recommendations. The web services and security teams work closely with the CERN IT security team, following up vulnerabilities. In a joint effort, proactive penetration tests are organized several times a year to discover potential issues ahead of time and to ensure that the infrastructure remains safe. As part of improving security, the logs of all the CMSWeb services are now pushed to the security operations center (SOC), which analyses them and identifies possible threats.

12.8.8 Monitoring and analytics

A scalable and reliable monitoring system is required to ensure efficient operation of the CMS distributed computing services, and to provide comprehensive measurements of the system performance. Metrics such as wall-time consumption of computing resources, memory, CPU, storage usage, and data access patterns are monitored to study the evolution of the performance over time and allow for in-depth analyses of the main system parameters.

The main components of the CMS monitoring infrastructure are presented in ref. [386]. CMS relies on central CERN [387] and dedicated CMS monitoring infrastructures. The former, supported by the CERN IT department, is used extensively to store data from CMS computing subsystems, such as HTCondor, submission infrastructure, CMSWeb user activities, analysis and production workflows coming from the WMAgent and CRAB job submission tools, and data-transfer and storage information from Rucio. The status of all CMS computing systems is monitored in real time using predefined views and dashboards.

The dedicated CMS monitoring infrastructure is composed of several Kubernetes clusters. It is mostly used to monitor computing nodes and services, and to provide additional features such as a sophisticated alert and notification system, data aggregation, annotations, and automation of several workflows. Several hundred nodes and services are monitored, and a fault-tolerant infrastructure is provided with minimal operational and maintenance effort.

13 Summary

Since the beginning of the LHC operation in 2009, the CMS detector has undergone a number of changes and upgrades, adapting the experiment to operating conditions at luminosities well beyond the original design. In 2022, the LHC Run 3 began and CMS successfully recorded its first 40 fb^{-1} of proton-proton data at a center-of-mass energy of 13.6 TeV with an operation efficiency of 92%. This paper describes the modifications, as installed and commissioned for LHC Run 3.

The upgraded pixel tracking detector was installed in early 2017. In the new detector, the number of barrel layers was increased from three to four, and the number of disks in each endcap from two to three, whereas the material budget was reduced, leading to a better tracking performance up to an absolute pseudorapidity of 3.0 for pixel tracks. The upgrade also involved a new readout chip enabling increased hit detection efficiencies at higher occupancy.

In the electromagnetic calorimeter, measures were taken to improve the monitoring and calibration of effects caused by irradiation, leading to a loss in the PbWO_4 crystal transparency and an increase of the leakage current in the avalanche photodiodes. In the second long shutdown (LS2) the calibration algorithms were refined to identify and remove spurious spike signals and to determine time-dependent correction factors for the laser monitoring system.

The upgrade of the hadron calorimeter included new readout electronics with finer granularity, leading to an increase in the number of channels and longitudinal segmentation. The previous generation of photosensors was replaced by silicon photomultipliers, which measure the scintillator light output with a better signal-to-noise ratio.

In the muon system, a gas electron multiplier (GEM) detector, consisting of four gas gaps separated by three GEM foils, was added in the endcaps. The other subsystems, drift tubes (DT), cathode strip chambers (CSC), and resistive-plate chambers (RPC), underwent several upgrades. In the DT, the muon trigger logic was replaced by a new data concentrator based on μTCA architecture. The top of CMS was covered with a neutron shield to reduce the background in the top external DT chambers. An outer ring of CSCs (ME4/2) was added in LS1, and in view of the High-Luminosity LHC, the bulk of the CSC electronics upgrades that required chamber access were performed already during LS2. An outer rings of the RPC chambers in station four (RE4/2 and RE4/3) were added as well. The endcap muon track finder of the L1 trigger was upgraded to utilize GEM-CSC joint track segments to optimize the final track reconstruction and resolution at the trigger level.

The precision proton spectrometer was upgraded significantly. Its tracker radiation-damaged sensors and chips were replaced. The mechanics of the detector, as well as the front-end electronics, were completely redesigned to add a novel internal motion system designed to mitigate the effects of radiation damage. In the timing system a second station was installed in each arm. All detector modules were replaced by new double-diamond modules with the aim of further improving the timing resolution.

In the beam radiation instrumentation and luminosity system, new versions of the pixel luminosity telescope (PLT), the fast beam conditions monitor (BCM1F), and the beam conditions monitor for losses (BCML) were installed for Run 3.

To cope with increasing instantaneous luminosities, the CMS data acquisition (DAQ) system underwent multiple upgrades. The backend technology was gradually moved to the more powerful μTCA standard. A new optical readout link with a higher bandwidth of 10 Gb/s was developed. The bulk of the DAQ system downstream from the custom readout benefited from advances in technology to achieve a much more compact design, while doubling the event building bandwidth.

The first level (L1) trigger, composed of custom hardware processors, uses information from the calorimeters and muon detectors to select events at a rate of up to 110 kHz within a fixed latency of about $4\ \mu\text{s}$. The developments in the L1 trigger mostly focused on the addition of dedicated triggers that became possible due to enhanced capabilities of the global trigger logic and increased trigger information delivered by the calorimeters and muon systems. Among other applications, new triggers for long-lived particle signatures were implemented. The addition of a 40 MHz scouting system that receives data from both the calorimeter and muon systems, further broadens the physics reach of CMS.

The high-level trigger (HLT) performs the second stage of event filtering and accepts events at a sustained rate of the order of 5 kHz. Since Run 3 began, an additional 30 kHz of HLT scouting data is recorded. Since 2016, the HLT has been operated using multithreaded event processing software,

minimizing memory requirements through reduction of the number of processes concurrently running. For Run 3, GPUs were successfully deployed in the HLT.

Substantial improvements were achieved in the physics performance and speed of the software, as well as in the computing infrastructure. Some of the major changes are: support for multithreaded processes and utilization of GPUs; direct remote data access; and usage of high-performance computing centers. New tools such as Rucio for data management were adopted with future data rates in mind. Considerable effort was put into the automation of the workflows and the validation of the software. Physics analyses have been moved to smaller and smaller formats for centrally produced and experiment-wide shared data samples, the most recent of which is the NanoAOD.

The development of the CMS detector, as described in this paper, constitutes a solid basis for future data taking.

Acknowledgments

We congratulate our colleagues in the CERN accelerator departments for the excellent performance of the LHC and thank the technical and administrative staffs at CERN and at other CMS institutes for their contributions to the success of the CMS effort. In addition, we gratefully acknowledge the computing centers and personnel of the Worldwide LHC Computing Grid and other centers for delivering so effectively the computing infrastructure essential to our analyses. Finally, we acknowledge the enduring support for the construction and operation of the LHC, the CMS detector, and the supporting computing infrastructure provided by the following funding agencies: the Armenian Science Committee, project no. 22rl-037; the Austrian Federal Ministry of Education, Science and Research and the Austrian Science Fund; the Belgian Fonds de la Recherche Scientifique, and Fonds voor Wetenschappelijk Onderzoek; the Brazilian Funding Agencies (CNPq, CAPES, FAPERJ, FAPERGS, and FAPESP); the Bulgarian Ministry of Education and Science, and the Bulgarian National Science Fund; CERN; the Chinese Academy of Sciences, Ministry of Science and Technology, the National Natural Science Foundation of China, and Fundamental Research Funds for the Central Universities; the Ministerio de Ciencia Tecnología e Innovación (Minciencias), Colombia; the Croatian Ministry of Science, Education and Sport, and the Croatian Science Foundation; the Research and Innovation Foundation, Cyprus; the Secretariat for Higher Education, Science, Technology and Innovation, Ecuador; the Ministry of Education and Research, Estonian Research Council via PRG780, PRG803 and PRG445 and European Regional Development Fund, Estonia; the Academy of Finland, Finnish Ministry of Education and Culture, and Helsinki Institute of Physics; the Institut National de Physique Nucléaire et de Physique des Particules / CNRS, and Commissariat à l'Énergie Atomique et aux Énergies Alternatives / CEA, France; the Bundesministerium für Bildung und Forschung, the Deutsche Forschungsgemeinschaft (DFG), under Germany's Excellence Strategy — EXC 2121 “Quantum Universe” — 390833306, and under project number 400140256 - GRK2497, and Helmholtz-Gemeinschaft Deutscher Forschungszentren, Germany; the General Secretariat for Research and Innovation and the Hellenic Foundation for Research and Innovation (HFRI), Project Number 2288, Greece; the National Research, Development and Innovation Office (NKFIH), Hungary; the Department of Atomic Energy and the Department of Science and Technology, India; the Institute for Studies in Theoretical Physics and Mathematics, Iran; the Science Foundation, Ireland; the Istituto Nazionale di Fisica Nucleare, Italy; the Ministry of Science, ICT and Future

Planning, and National Research Foundation (NRF), Republic of Korea; the Ministry of Education and Science of the Republic of Latvia; the Lithuanian Academy of Sciences; the Ministry of Education, and University of Malaya (Malaysia); the Ministry of Science of Montenegro; the Mexican Funding Agencies (BUAP, CINVESTAV, CONACYT, LNS, SEP, and UASLP-FAI); the Ministry of Business, Innovation and Employment, New Zealand; the Pakistan Atomic Energy Commission; the Ministry of Education and Science and the National Science Center, Poland; the Fundação para a Ciência e a Tecnologia, grants CERN/FIS-PAR/0025/2019 and CERN/FIS-INS/0032/2019, Portugal; the Ministry of Education, Science and Technological Development of Serbia; MCIN/AEI/10.13039/501100011033, ERDF “a way of making Europe”, Programa Estatal de Fomento de la Investigación Científica y Técnica de Excelencia María de Maeztu, grant MDM-2017-0765, projects PID2020-113705RB, PID2020-113304RB, PID2020-116262RB and PID2020-113341RB-I00, and Plan de Ciencia, Tecnología e Innovación de Asturias, Spain; the Ministry of Science, Technology and Research, Sri Lanka; the Swiss Funding Agencies (ETH Board, ETH Zurich, PSI, SNF, UniZH, Canton Zurich, and SER); the Ministry of Science and Technology, Taipei; the Ministry of Higher Education, Science, Research and Innovation, and the National Science and Technology Development Agency of Thailand; the Scientific and Technical Research Council of Turkey, and Turkish Energy, Nuclear and Mineral Research Agency; the National Academy of Sciences of Ukraine; the Science and Technology Facilities Council, U.K.; the US Department of Energy, and the US National Science Foundation.

Individuals have received support from the Marie-Curie program and the European Research Council and Horizon 2020 Grant, contract Nos. 675440, 724704, 752730, 758316, 765710, 824093, and COST Action CA16108 (European Union) the Leventis Foundation; the Alfred P. Sloan Foundation; the Alexander von Humboldt Foundation; the Belgian Federal Science Policy Office; the Fonds pour la Formation à la Recherche dans l’Industrie et dans l’Agriculture (FRIA-Belgium); the Agentschap voor Innovatie door Wetenschap en Technologie (IWT-Belgium); the F.R.S.-FNRS and FWO (Belgium) under the “Excellence of Science — EOS” — be.h project n. 30820817; the Beijing Municipal Science & Technology Commission, No. Z191100007219010; the Ministry of Education, Youth and Sports (MEYS) of the Czech Republic; the Shota Rustaveli National Science Foundation, grant FR-22-985 (Georgia); the Hungarian Academy of Sciences, the New National Excellence Program - ÚNKP, the NKFIH research grants K 124845, K 124850, K 128713, K 128786, K 129058, K 131991, K 133046, K 138136, K 143460, K 143477, 2020-2.2.1-ED-2021-00181, and TKP2021-NKTA-64 (Hungary); the Council of Scientific and Industrial Research, India; the Latvian Council of Science; the Ministry of Education and Science, project no. 2022/WK/14, and the National Science Center, contracts Opus 2021/41/B/ST2/01369 and 2021/43/B/ST2/01552 (Poland); the Fundação para a Ciência e a Tecnologia, grant FCT CEECIND/01334/2018; the National Priorities Research Program by Qatar National Research Fund; the Programa Estatal de Fomento de la Investigación Científica y Técnica de Excelencia María de Maeztu, grant MDM-2017-0765 and projects PID2020-113705RB, PID2020-113304RB, PID2020-116262RB and PID2020-113341RB-I00, and Programa Severo Ochoa del Principado de Asturias (Spain); the Chulalongkorn Academic into Its 2nd Century Project Advancement Project, and the National Science, Research and Innovation Fund via the Program Management Unit for Human Resources & Institutional Development, Research and Innovation, grant B05F650021 (Thailand); the Kavli Foundation; the Nvidia Corporation; the SuperMicro Corporation; the Welch Foundation, contract C-1845; and the Weston Havens Foundation (U.S.A.).

References

- [1] CMS collaboration, *The CMS Experiment at the CERN LHC*, 2008 *JINST* **3** S08004.
- [2] CMS collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[arXiv:1207.7235](#)].
- [3] CMS collaboration, *Observation of a New Boson with Mass Near 125 GeV in pp Collisions at $\sqrt{s} = 7$ and 8 TeV*, *JHEP* **06** (2013) 081 [[arXiv:1303.4571](#)].
- [4] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](#)].
- [5] CMS collaboration, *Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV*, 2020 *JINST* **15** P10017 [[arXiv:2006.10165](#)].
- [6] CMS collaboration, *The CMS trigger system*, 2017 *JINST* **12** P01020 [[arXiv:1609.02366](#)].
- [7] CMS collaboration, *Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, 2015 *JINST* **10** P06005 [[arXiv:1502.02701](#)].
- [8] CMS collaboration, *Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV*, 2018 *JINST* **13** P06015 [[arXiv:1804.04528](#)].
- [9] CMS collaboration, *Performance of Photon Reconstruction and Identification with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, 2015 *JINST* **10** P08010 [[arXiv:1502.02702](#)].
- [10] CMS collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, 2014 *JINST* **9** P10009 [[arXiv:1405.6569](#)].
- [11] CMS collaboration, *Cutaway diagrams of CMS detector*, CMS Collection [CMS-OUTREACH-2019-001](#) (2019).
- [12] CMS and TOTEM collaborations, *Proton reconstruction with the CMS-TOTEM Precision Proton Spectrometer*, 2023 *JINST* **18** P09009 [[arXiv:2210.05854](#)].
- [13] M. Hoch, *CMS detector and cavern (open view)*, CMS Collection [CMS-PHO-GEN-2008-019](#) CERN, Geneva, Switzerland (2008).
- [14] S. Yammine, G. Le Godec, and H. Thiesen, *Study report for the new free wheeling thyristors system for the upgrade of the CMS solenoid power converter*, Tech. Rep. EDMS Document 1845659 v3.2, CERN, Geneva (2023).
- [15] CMS collaboration, *CMS Technical Design Report for the Pixel Detector Upgrade*, Tech. Rep. [CERN-LHCC-2012-016](#) CERN, Geneva (2012).
- [16] CMS TRACKER GROUP collaboration, *The CMS Phase-I Pixel Detector Upgrade*, 2021 *JINST* **16** P02027 [[arXiv:2012.14304](#)].
- [17] ROSE collaboration, *Radiation hard silicon detectors developments by the RD48 (ROSE) Collaboration*, *Nucl. Instrum. Meth. A* **466** (2001) 308.
- [18] T. Rohe et al., *Fluence dependence of charge collection of irradiated pixel sensors*, *Nucl. Instrum. Meth. A* **552** (2005) 232 [[physics/0411214](#)].
- [19] H.C. Kästli, *Frontend electronics development for the CMS pixel detector upgrade*, *Nucl. Instrum. Meth. A* **731** (2013) 88.
- [20] R. Horisberger, *Readout architectures for pixel detectors*, *Nucl. Instrum. Meth. A* **465** (2000) 148.

- [21] CMS TRACKER GROUP collaboration, *The DAQ and control system for the CMS Phase-1 pixel detector upgrade*, [2019 JINST 14 P10017](#).
- [22] A. Marchioro and P. Moreira, *QPLL: a Quartz Crystal Based PLL for Jitter Filtering Applications in LHC*, in the proceedings of the 9th Workshop on Electronics for LHC Experiments, Amsterdam, The Netherlands, 29 September–3 October 2003, pp. 280–284 [[DOI:10.5170/CERN-2003-006.280](#)].
- [23] C. Nägeli, *Analysis of the rare decay $B_s^0 \rightarrow \mu^+ \mu^-$ using the compact muon solenoid experiment at CERN's Large Hadron Collider*, Ph.D. thesis, ETH, Zürich, Switzerland (2013).
- [24] M. Pesaresi et al., *The FC7 AMC for generic DAQ & control applications in CMS*, [2015 JINST 10 C03036](#).
- [25] E. Hazen et al., *The AMC13XG: a new generation clock/timing/DAQ module for CMS MicroTCA*, [2013 JINST 8 C12036](#).
- [26] S. Michelis et al., *DC-DC converters in 0.35 μm CMOS technology*, [2012 JINST 7 C01072](#).
- [27] L. Feld et al., *The DC-DC conversion power system of the CMS Phase-1 pixel upgrade*, [2015 JINST 10 C01052](#).
- [28] L. Feld et al., *Experience from design, prototyping and production of a DC-DC conversion powering scheme for the CMS Phase-1 Pixel Upgrade*, [2016 JINST 11 C02033](#).
- [29] C. Paillard, A. Marchioro and C. Ljuslin, *The CCU25: a network oriented communication and control unit integrated circuit in a 0.25 μm CMOS technology*, in the proceedings of the 8th Workshop on Electronics for LHC Experiments, Colmar, France, 9–13 September 2002, pp. 174–178 [[DOI:10.5170/CERN-2002-003.174](#)].
- [30] F. Faccio, S. Michelis, and G. Ripamonti, *Summary of measurements on FEAST2 modules to understand the failures observed in the CMS pixel system*, Tech. Rep., CERN, Geneva (2018), <https://espace.cern.ch/project-DCDC-new/Shared%20Documents/SummaryMeasurements18.pdf>.
- [31] P. Tropea et al., *Advancements and plans for the LHC upgrade detector thermal management with CO_2 evaporative cooling*, *Nucl. Instrum. Meth. A* **936** (2019) 644.
- [32] CMS collaboration, *The Phase-2 Upgrade of the CMS Tracker*, Tech. Rep. [CERN-LHCC-2017-009](#), CERN, Geneva (2017) [[DOI:10.17181/CERN.QZ28.FLHW](#)].
- [33] M.J. French et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, *Nucl. Instrum. Meth. A* **466** (2001) 359.
- [34] CMS collaboration, *Operation and performance of the CMS silicon strip tracker with proton-proton collisions at the CERN LHC*, to be submitted to JINST (2023).
- [35] CMS collaboration, *Silicon Strip Tracker Performance results 2018*, Tech. Rep. [CMS-DP-2018-052](#) CERN, Geneva (2018).
- [36] C. Barth, *Performance of the CMS Tracker under Irradiation*, Ph.D. thesis, KIT, Karlsruhe, Germany (2013).
- [37] K. Deiters et al., *Properties of the most recent avalanche photodiodes for the CMS electromagnetic calorimeter*, *Nucl. Instrum. Meth. A* **442** (2000) 193.
- [38] F. Addesa and F. Cavallari, *Performance prospects for the CMS electromagnetic calorimeter barrel avalanche photodiodes for LHC phase I and phase II: Radiation hardness and longevity*, *Nucl. Instrum. Meth. A* **787** (2015) 114.
- [39] CMS ECAL collaboration, *Radiation hardness qualification of PbWO_4 scintillation crystals for the CMS Electromagnetic Calorimeter*, [2010 JINST 5 P03010](#) [[arXiv:0912.4300](#)].

- [40] R. Benetta et al., *The CMS ECAL readout architecture and the clock and control system*, in the proceedings of the 11th International Conference on Calorimetry in High-Energy Physics, Perugia, Italy, 29 March–2 April 2004, p. 162–172 [[DOI:10.1142/9789812701978_0021](#)].
- [41] CMS collaboration, *Reconstruction of signal amplitudes in the CMS electromagnetic calorimeter in the presence of overlapping proton-proton interactions*, 2020 *JINST* **15** P10002 [[arXiv:2006.14359](#)].
- [42] CMS ECAL collaboration, *Reconstruction of the signal amplitude of the CMS electromagnetic calorimeter*, *Eur. Phys. J. C* **46S1** (2006) 23.
- [43] CMS collaboration, *Time Reconstruction and Performance of the CMS Electromagnetic Calorimeter*, 2010 *JINST* **5** T03011 [[arXiv:0911.4044](#)].
- [44] P. Paganini, *CMS electromagnetic trigger commissioning and first operation experiences*, *J. Phys. Conf. Ser.* **160** (2009) 012062.
- [45] M. Hansen, *The new readout architecture for the CMS ECAL*, in the proceedings of the 9th Workshop on Electronics for LHC Experiments, Amsterdam, Netherlands, 29 September–3 October 2003, pp. 78–82 [[DOI:10.5170/CERN-2003-006.78](#)].
- [46] D.A. Petyt, *Anomalous APD signals in the CMS Electromagnetic Calorimeter*, *Nucl. Instrum. Meth. A* **695** (2012) 293.
- [47] CMS collaboration, *Energy Calibration and Resolution of the CMS Electromagnetic Calorimeter in pp Collisions at $\sqrt{s} = 7$ TeV*, 2013 *JINST* **8** P09009 [[arXiv:1306.2016](#)].
- [48] F. Thiant et al., *New development in the CMS ECAL Level-1 trigger system to meet the challenges of LHC Run 2*, *PoS TWEPP2018* (2019) 052.
- [49] CMS collaboration, *Search for long-lived particles using delayed photons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **100** (2019) 112003 [[arXiv:1909.06166](#)].
- [50] CMS collaboration, *The hadron calorimeter*, Tech. Rep. [CERN-LHCC-97-031](#), CERN, Geneva (1997).
- [51] CMS HCAL collaboration, *Design, performance, and calibration of CMS hadron-barrel calorimeter wedges*, *Eur. Phys. J. C* **55** (2008) 159.
- [52] CMS HCAL collaboration, *Design, performance and calibration of the CMS forward calorimeter wedges*, *Eur. Phys. J. C* **53** (2008) 139.
- [53] CMS HCAL collaboration, *Design, performance, and calibration of the CMS Hadron-outer calorimeter*, *Eur. Phys. J. C* **57** (2008) 653.
- [54] CMS collaboration, *Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV*, 2020 *JINST* **15** P05002 [[arXiv:1910.00079](#)].
- [55] CMS collaboration, *Calibration of the CMS hadron calorimeters using proton-proton collision data at $\sqrt{s} = 13$ TeV*, 2020 *JINST* **15** P05002 [[arXiv:1910.00079](#)].
- [56] T. Zimmerman and J.R. Hoff, *The Design of a charge integrating, modified floating point ADC chip*, *IEEE J. Solid State Circuits* **39** (2004) 895.
- [57] A. Baumbaugh et al., *QIE10: a new front-end custom integrated circuit for high-rate experiments*, 2014 *JINST* **9** C01062.
- [58] T. Roy et al., *QIE: Performance Studies of the Next Generation Charge Integrator*, 2015 *JINST* **10** C02009.

- [59] D. Hare et al., *First large volume characterization of the QIE10/11 custom front-end integrated circuits*, [2016 JINST 11 C02052](#).
- [60] CMS collaboration, *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*, Tech. Rep. [CERN-LHCC-2012-015](#) CERN, Geneva (2012).
- [61] CMS collaboration, *The Phase-2 Upgrade of the CMS Barrel Calorimeters*, Tech. Rep. [CERN-LHCC-2017-011](#), CERN, Geneva (2017).
- [62] CMS collaboration, *Identification and Filtering of Uncharacteristic Noise in the CMS Hadron Calorimeter*, [2010 JINST 5 T03014](#) [[arXiv:0911.4881](#)].
- [63] R.A. Shukla et al., *Microscopic characterisation of photodetectors used in the hadron calorimeter of the Compact Muon Solenoid experiment*, [Rev. Sci. Instrum. 90](#) (2019) 023303.
- [64] CMS collaboration, *Measurements with silicon photomultipliers of dose-rate effects in the radiation damage of plastic scintillator tiles in the CMS hadron endcap calorimeter*, [2020 JINST 15 P06009](#) [[arXiv:2001.06553](#)].
- [65] P. Cushman, A. Heering and A. Ronzhin, *Custom HPD readout for the CMS HCAL*, [Nucl. Instrum. Meth. A 442](#) (2000) 289.
- [66] A. Heering et al., *Parameters of the preproduction series SiPMs for the CMS HCAL phase I upgrade*, [Nucl. Instrum. Meth. A 824](#) (2016) 115.
- [67] Y. Musienko et al., *Radiation damage studies of silicon photomultipliers for the CMS HCAL phase I upgrade*, [Nucl. Instrum. Meth. A 787](#) (2015) 319.
- [68] F. Vasey et al., *The Versatile Link common project: feasibility report*, [2012 JINST 7 C01075](#).
- [69] G. Cummings, *CMS HCAL VTRx-induced communication loss and mitigation*, [2022 JINST 17 C05020](#).
- [70] P. Moreira et al., *The GBT-SerDes ASIC prototype*, [2010 JINST 5 C11022](#).
- [71] J. Gutleber and L. Orsini, *Software Architecture for Processing Clusters Based on I₂O*, [Clust. Comput. 5](#) (2002) 55.
- [72] C. Ghabrous Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, [2015 JINST 10 C02019](#) [[CMS-CR-2014-334](#)].
- [73] V. Brigljevic et al., *Run control and monitor system for the CMS experiment*, [eConf C0303241](#) (2003) THGT002 [[cs/0306110](#)].
- [74] D. Charousset, R. Hiesgen and T.C. Schmidt, *Revisiting Actor Programming in C++*, [Comput. Lang. Syst. Struct. 45](#) (2016) 105 [[arXiv:1505.07368](#)].
- [75] JSON-RPC Working Group, *JSON-RPC 2.0 specification*, 2013, <https://www.jsonrpc.org/specification>.
- [76] I. Fette and A. Melnikov, *The WebSocket protocol*, RFC Proposed Standard 6455, 2011 [[DOI:10.17487/RFC6455](#)].
- [77] B. Bilki, *Review of Scalar Meson Production at $\sqrt{s} = 7$ TeV in CMS, U(1)' Gauge Extensions of the MSSM and Calorimetry for Future Colliders*, Ph.D. thesis, University of Iowa, Iowa City, IA, U.S.A. (2011), [CERN-THESIS-2011-229](#).
- [78] CMS HCAL collaboration, *Study of Various Photomultiplier Tubes with Muon Beams And Cerenkov Light Produced in Electron Showers*, [2010 JINST 5 P06002](#).

- [79] CMS ECAL and HCAL collaborations, *The CMS barrel calorimeter response to particle beams from 2 to 350-GeV/c*, *Eur. Phys. J. C* **60** (2009) 359.
- [80] CMS collaboration, *Performance of CMS Hadron Calorimeter Timing and Synchronization using Test Beam, Cosmic Ray, and LHC Beam Data*, *2010 JINST* **5** T03013 [[arXiv:0911.4877](#)].
- [81] J. Lawhorn, *New method of out-of-time energy subtraction for the CMS hadronic calorimeter*, *J. Phys. Conf. Ser.* **1162** (2019) 012036.
- [82] CMS collaboration, *Performance of the local reconstruction algorithms for the CMS hadron calorimeter with Run 2 data*, [arXiv:2306.10355](#).
- [83] CMS collaboration, *Noise in Phase 1 HF detector in 2017*, Tech. Rep. [CMS-DP-2017-034](#) CERN, Geneva (2017).
- [84] CMS collaboration, *Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $s=13$ TeV*, *Phys. Rev. D* **105** (2022) 092007 [[arXiv:2201.11585](#)].
- [85] CMS collaboration, *The Performance of the CMS Muon Detector in Proton-Proton Collisions at $\sqrt{s} = 7$ TeV at the LHC*, *2013 JINST* **8** P11002 [[arXiv:1306.6905](#)].
- [86] CMS collaboration, *Calibration of the CMS Drift Tube Chambers and Measurement of the Drift Velocity with Cosmic Rays*, *2010 JINST* **5** T03016 [[arXiv:0911.4895](#)].
- [87] CMS collaboration, *Technical proposal for the upgrade of the CMS detector through 2020*, Tech. Rep. [CERN-LHCC-2011-006](#) CERN, Geneva (2011).
- [88] CMS collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, Tech. Rep. [CERN-LHCC-2020-004](#), CERN, Geneva (2020).
- [89] Á. Navarro-Tobar and C. Fernández-Bedoya, *CMS DT upgrade: the sector collector relocation*, *2016 JINST* **11** C02046.
- [90] Á. Navarro-Tobar, C. Fernández-Bedoya and I. Redondo, *Low-cost, high-precision propagation delay measurement of 12-fibre MPO cables for the CMS DT electronics upgrade*, *2013 JINST* **8** C02001.
- [91] CMS collaboration, *CMS Technical Design Report for the Level-1 Trigger Upgrade*, Tech. Rep. [CERN-LHCC-2013-011](#) CERN, Geneva (2013).
- [92] A. Triossi et al., *A New Data Concentrator for the CMS Muon Barrel Track Finder*, *PoS TIPP2014* (2014) 412.
- [93] Xilinx Inc., *7 series FPGAs GTX/GTH transceivers*, User Guide UG476 v1.12.1, 2018, https://docs.xilinx.com/v/u/en-US/ug476_7Series_Transceivers.
- [94] P. Moreira et al., *A radiation tolerant gigabit serializer for LHC data transmission*, in the proceedings of the 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10–14 September 2001, pp. 145–149, <http://cds.cern.ch/record/588665> [[DOI:10.5170/CERN-2001-005.145](#)].
- [95] CMS collaboration, *Performance of the CMS TwinMux Algorithm in late 2016 pp collision runs*, Tech. Rep. [CMS-DP-2016-074](#) CERN, Geneva (2016).
- [96] Á. Navarro-Tobar et al., *Phase 1 upgrade of the CMS drift tubes read-out system*, *2017 JINST* **12** C03070.
- [97] CMS collaboration, *Efficiency of the CMS Drift Tubes at LHC in 2017*, Tech. Rep. [CMS-DP-2018-016](#) CERN, Geneva (2018).
- [98] CMS collaboration, *Performance of the CMS Drift Tubes at the end of LHC Run 2*, Tech. Rep. [CMS-DP-2019-008](#) CERN, Geneva (2019).

- [99] E. Conti and F. Gasparini, *Test of the wire ageing induced by radiation for the CMS barrel muon chambers*, *Nucl. Instrum. Meth. A* **465** (2001) 472.
- [100] CMS collaboration, *The Phase-2 Upgrade of the CMS Muon Detectors*, Tech. Rep. [CERN-LHCC-2017-012](#), CERN, Geneva (2017).
- [101] G. Altenhöfer et al., *The drift velocity monitoring system of the CMS barrel muon chambers*, *Nucl. Instrum. Meth. A* **888** (2018) 1.
- [102] CMS collaboration, *Background measurements in the CMS DT chambers during LHC Run 2*, Tech. Rep. [CMS-DP-2020-011](#) CERN, Geneva (2020).
- [103] CMS collaboration, *Local reconstruction algorithms in the cathode strip chambers of CMS*, *EPJ Web Conf.* **214** (2019) 02014.
- [104] J. Heilman, *Work on CMS muon detector (CSCs) during long shutdown 1 (LS1)*, CMS Collection [CMS-PHO-MUON-2015-001](#) (2014).
- [105] CMS collaboration, *Upgrade of the CSC Muon System for the CMS Detector at the HL-LHC*, *PoS EPS-HEP2021* (2022) 837.
- [106] CAEN S.p.A., *Mod. A1733-A1833-A1733B-A1833B HV boards*, Technical Information Manual rev. 10, 2013, <https://www.caen.it/download/>.
- [107] D. Pfeiffer et al., *The radiation field in the Gamma Irradiation Facility GIF++ at CERN*, *Nucl. Instrum. Meth. A* **866** (2017) 91 [[arXiv:1611.00299](#)].
- [108] B. Bylsma et al., *Radiation Testing of Electronics for the CMS Endcap Muon System*, *Nucl. Instrum. Meth. A* **698** (2013) 242 [[arXiv:1208.4051](#)].
- [109] F. Ferrarese, *Statistical analysis of total ionizing dose response in 25-nm NAND flash memory*, Ph.D. thesis, Università degli Studi di Padova, Padova, Italy (2014), <http://tesi.cab.unipd.it/45696/>.
- [110] K. Bunkowski, *Optimization, Synchronization, Calibration and Diagnostic of the RPC PAC Muon Trigger System for the CMS detector*, Ph.D. thesis, Warsaw University, Warsaw, Poland (2009).
- [111] E.P. Corrin, *Development of Digital Readout Electronics for the CMS Tracker*, Ph.D. thesis, Imperial College, London, U.K. (2002).
- [112] S. Colafranceschi et al., *Resistive plate chambers for 2013-2014 muon upgrade in CMS at LHC*, *2014 JINST* **9** C10033.
- [113] CMS collaboration, *CMS technical design report, volume II: Physics performance*, *J. Phys. G* **34** (2007) 995.
- [114] L.-B. Cheng, P.-C. Cao, J.-Z. Zhao and Z.-A. Liu, *Design of online control and monitoring software for the CPPF system in the CMS Level-1 trigger upgrade*, *Nucl. Sci. Tech.* **29** (2018) 166.
- [115] A. Samalan et al., *RPC system in the CMS Level-1 Muon Trigger*, *2020 JINST* **15** C10007.
- [116] M. Abbrescia et al., *Resistive plate chambers performances at cosmic rays fluxes*, *Nucl. Instrum. Meth. A* **359** (1995) 603.
- [117] S. Colafranceschi et al., *Performance of the Gas Gain Monitoring system of the CMS RPC muon detector and effective working point fine tuning*, *2012 JINST* **7** P12004 [[arXiv:1209.3893](#)].
- [118] J. Goh et al., *CMS RPC tracker muon reconstruction*, *2014 JINST* **9** C10027.
- [119] M.I. Pedraza-Morales, M.A. Shah and M. Shopova, *First results of CMS RPC performance at 13 TeV*, *2016 JINST* **11** C12003 [[arXiv:1605.09521](#)].

- [120] M.A. Shah et al., *Experiences from the RPC data taking during the CMS RUN-2*, 2020 *JINST* **15** C10027 [[arXiv:2005.12532](#)].
- [121] M.A. Shah et al., *The CMS RPC Detector Performance and Stability during LHC RUN-2*, 2019 *JINST* **14** C11012 [[arXiv:1808.10488](#)].
- [122] A. Gelmi, R. Guida and B. Mandelli, *Gas mixture quality studies for the CMS RPC detectors during LHC Run 2*, 2021 *JINST* **16** C04004.
- [123] M. Abbrescia et al., *HF production in CMS-Resistive Plate Chambers*, *Nucl. Phys. B Proc. Suppl.* **158** (2006) 30.
- [124] CMS collaboration, *The CMS muon project: Technical Design Report*, Tech. Rep. CERN-LHCC-97-032, CERN, Geneva (1997).
- [125] A. Gelmi et al., *Longevity studies on the CMS-RPC system*, 2019 *JINST* **14** C05012.
- [126] R. Guida, *GIF++: A new CERN Irradiation Facility to test large-area detectors for the HL-LHC program*, *PoS ICHEP2016* (2016) 260.
- [127] R. Aly et al., *Aging Study on Resistive Plate Chambers of the CMS Muon Detector for HL-LHC*, 2020 *JINST* **15** C11002 [[arXiv:2005.11397](#)].
- [128] S. Costantini et al., *Radiation background with the CMS RPCs at the LHC*, 2015 *JINST* **10** C05031 [[arXiv:1406.2859](#)].
- [129] CMS MUON collaboration, *Latest results of Longevity studies on the present CMS RPC system for HL-LHC phase*, *Nucl. Instrum. Meth. A* **1055** (2023) 168452.
- [130] G. Carboni et al., *An extensive aging study of bakelite resistive plate chambers*, *Nucl. Instrum. Meth. A* **518** (2004) 82.
- [131] F. Thyssen, *Performance of the resistive plate chambers in the CMS experiment*, 2012 *JINST* **7** C01104.
- [132] G. Pugliese et al., *Aging studies for Resistive Plate Chambers of the CMS muon trigger detector*, *Nucl. Instrum. Meth. A* **515** (2003) 342.
- [133] R. Guida and B. Mandelli, *R&D strategies for optimizing greenhouse gases usage in the LHC particle detection systems*, *Nucl. Instrum. Meth. A* **958** (2020) 162135.
- [134] R. Guida, B. Mandelli and G. Rigoletti, *Studies on alternative eco-friendly gas mixtures and development of gas recuperation plant for RPC detectors*, *Nucl. Instrum. Meth. A* **1039** (2022) 167045.
- [135] F. Sauli, *GEM: A new concept for electron amplification in gas detectors*, *Nucl. Instrum. Meth. A* **386** (1997) 531.
- [136] CMS MUON collaboration, *Layout and Assembly Technique of the GEM Chambers for the Upgrade of the CMS First Muon Endcap Station*, *Nucl. Instrum. Meth. A* **918** (2019) 67 [[arXiv:1812.00411](#)].
- [137] C. Calabria, *Large-size triple GEM detectors for the CMS forward muon upgrade*, *Nucl. Part. Phys. Proc.* **273-275** (2016) 1042.
- [138] CMS collaboration, *Benchmarking LHC background particle simulation with the CMS triple-GEM detector*, 2021 *JINST* **16** P12026 [[arXiv:2107.03621](#)].
- [139] P. Aspell et al., *VFAT3: A Trigger and Tracking Front-end ASIC for the Binary Readout of Gaseous and Silicon Sensors*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference*, Sydney, NSW, Australia, 10–17 November 2018, [[DOI:10.1109/NSSMIC.2018.8824655](#)].

- [140] P. Aspell et al., *Development of a GEM Electronic Board (GEB) for triple-GEM detectors*, [2014 JINST 9 C12030](#).
- [141] D. Abbaneo, *CMS muon system phase 2 upgrade with triple-GEM detectors*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference*, San Diego, CA, U.S.A., 31 October 2015–7 November 2015, [[DOI:10.1109/NSSMIC.2015.7581803](#)].
- [142] T. Lenzi, *A micro-TCA based data acquisition system for the Triple-GEM detectors for the upgrade of the CMS forward muon spectrometer*, [2017 JINST 12 C01058](#).
- [143] A. Svetek et al., *The Calorimeter Trigger Processor Card: the next generation of high speed algorithmic data processing at CMS*, [2016 JINST 11 C02011](#).
- [144] R.K. Mommsen et al., *The CMS event-builder system for LHC run 3 (2021-23)*, [EPJ Web Conf. 214 \(2019\) 01006](#).
- [145] P. Aspell et al., *VFAT2: A front-end “system on chip” providing fast trigger information and digitized data storage for the charge sensitive readout of multi-channel silicon and gas particle detectors.*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference and 16th International Workshop on Room-Temperature Semiconductor X-Ray and Gamma-Ray Detectors*, Dresden, Germany, 19–25 October 2008, pp. 1489–1494 [[DOI:10.1109/NSSMIC.2008.4774696](#)].
- [146] D. Abbaneo et al., *Design of a constant fraction discriminator for the VFAT3 front-end ASIC of the CMS GEM detector*, [2016 JINST 11 C01023](#).
- [147] P. Moreira et al., *The GBT Project*, in the proceedings of the *Topical Workshop on Electronics for Particle Physics*, Paris, France, 21–25 September 2009, pp. 342–346 [[DOI:10.5170/CERN-2009-006.342](#)].
- [148] P. Alfke, *Xilinx Virtex-6 and Spartan-6 FPGA families*, in the proceedings of the *IEEE Hot Chips 21 Symposium*, Stanford CA, U.S.A., 23–25 August 2009, pp. 1–20.
- [149] D. Abbaneo et al., *The Triple-GEM Project for the Phase 2 Upgrade of the CMS Muon System*, [PoS TIPP2014 \(2014\) 065](#).
- [150] K. Ecklund et al., *Upgrade of the cathode strip chamber level 1 trigger optical links at CMS*, [2012 JINST 7 C11011](#).
- [151] D. Acosta, *Boosted Decision Trees in the Level-1 Muon Endcap Trigger at CMS*, [J. Phys. Conf. Ser. 1085 \(2018\) 042042](#).
- [152] P. Golonka et al., *FwWebViewPlus: integration of web technologies into WinCC OA based Human-Machine Interfaces at CERN*, [J. Phys. Conf. Ser. 513 \(2014\) 012009](#).
- [153] M. Abbas et al., *Detector Control System for the GE1/1 slice test*, [2020 JINST 15 P05023](#).
- [154] R. Venditti, *Production, Quality Control and Performance of GE1/1 Detectors for the CMS Upgrade*, [J. Phys. Conf. Ser. 1498 \(2020\) 012055](#).
- [155] CMS MUON collaboration, *Performance of prototype GE1/1 chambers for the CMS muon spectrometer upgrade*, [Nucl. Instrum. Meth. A 972 \(2020\) 164104](#) [[arXiv:1903.02186](#)].
- [156] CMS and TOTEM collaborations, *CMS-TOTEM Precision Proton Spectrometer*, Tech. Rep. [CERN-LHCC-2014-021](#) CERN, Geneva (2014).
- [157] TOTEM collaboration, *Total cross-section, elastic scattering and diffraction dissociation at the Large Hadron Collider at CERN: TOTEM Technical Design Report*, Tech. Rep. [CERN-LHCC-2004-002](#), CERN, Geneva (2004).
- [158] M.G. Albrow, *The CMS-TOTEM Precision Proton Spectrometer: CT-PPS*, [PoS DIS2015 \(2015\) 064](#).

- [159] V. Vacek, G.D. Hallewell, S. Ilie and S. Lindsay, *Perfluorocarbons and their use in Cooling Systems for Semiconductor Particle Detectors*, *Fluid Phase Equilib.* **174** (2000) 191.
- [160] G. Ruggiero et al., *Characteristics of edgeless silicon detectors for the Roman Pots of the TOTEM experiment at the LHC*, *Nucl. Instrum. Meth. A* **604** (2009) 242.
- [161] TOTEM collaboration, *Performance of the TOTEM Detectors at the LHC*, *Int. J. Mod. Phys. A* **28** (2013) 1330046 [[arXiv:1310.2908](#)].
- [162] F. Ravera, *The CT-PPS tracking system with 3D pixel detectors*, *2016 JINST* **11** C11027.
- [163] F. Ravera, *3D silicon pixel detectors for the CT-PPS tracking system*, Ph.D. thesis, Università degli Studi di Torino, Torino, Italy (2017), [CERN-THESIS-2017-473](#).
- [164] D.M.S. Sultan et al., *First Production of New Thin 3D Sensors for HL-LHC at FBK*, *2017 JINST* **12** C01022 [[arXiv:1612.00638](#)].
- [165] G. Pellegrini et al., *3D double sided detector fabrication at IMB-CNM*, *Nucl. Instrum. Meth. A* **699** (2013) 27.
- [166] G.-F. Dalla Betta et al., *Small pitch 3D devices*, *PoS Vertex2016* (2017) 028.
- [167] D. Hits and A. Starodumov, *The CMS Pixel Readout Chip for the Phase 1 Upgrade*, *2015 JINST* **10** C05029.
- [168] J. Hoss et al., *Radiation tolerance of the readout chip for the Phase I upgrade of the CMS pixel detector*, *2016 JINST* **11** C01003.
- [169] TOTEM collaboration, *Diamond Detectors for the TOTEM Timing Upgrade*, *2017 JINST* **12** P03007 [[arXiv:1701.05227](#)].
- [170] E. Bossini, *The CMS Precision Proton Spectrometer timing system: performance in Run 2, future upgrades and sensor radiation hardness studies*, *2020 JINST* **15** C05054 [[arXiv:2004.11068](#)].
- [171] E. Bossini and N. Minafra, *Diamond Detectors for Timing Measurements in High Energy Physics*, *Front. in Phys.* **8** (2020) 248 [[arXiv:2020.00248](#)].
- [172] CMS collaboration, *Time resolution of the diamond sensors used in the Precision Proton Spectrometer*, Tech. Rep. [CMS-DP-2019-034](#) CERN, Geneva (2019).
- [173] E. Bossini, D.M. Figueiredo, L.M.C. Forthomme and F.I. Garcia Fuentes, *Test beam results of irradiated single-crystal CVD diamond detectors at DESY-II*, Tech. Rep. [CMS-NOTE-2020-007](#), CERN, Geneva (2020).
- [174] F. Anghinolfi et al., *NINO: An ultra-fast and low-power front-end amplifier/discriminator ASIC designed for the multigap resistive plate chamber*, *Nucl. Instrum. Meth. A* **533** (2004) 183.
- [175] J. Christiansen, *HPTDC: High performance time to digital converter. version 2.2 for HPTDC version 1.3*, Tech. Rep., <https://cds.cern.ch/record/1067476/>, CERN, Geneva (2000).
- [176] C. Royon, *SAMPIC: a readout chip for fast timing detectors in particle physics and medical imaging*, *J. Phys. Conf. Ser.* **620** (2015) 012008 [[arXiv:1503.04625](#)].
- [177] TOTEM collaboration, *Timing Measurements in the Vertical Roman Pots of the TOTEM Experiment*, Tech. Rep. [CERN-LHCC-2014-020](#) CERN, Geneva (2014).
- [178] M. Bousonville and J. Rausch, *Universal picosecond timing system for the Facility for Antiproton and Ion Research*, *Phys. Rev. ST Accel. Beams* **12** (2009) 042801.
- [179] P. Moritz and B. Zipfel, *Recent Progress on the Technical Realization of the Bunch Phase Timing System BuTiS*, *Conf. Proc. C* **110904** (2011) 418.

- [180] M. Quinto, F. Cafagna, A. Fiergolski and E. Radicioni, *Upgrade of the TOTEM DAQ using the Scalable Readout System (SRS)*, [2013 JINST 8 C11006](#).
- [181] CMS collaboration, *Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS*, [Eur. Phys. J. C 81 \(2021\) 800](#) [[arXiv:2104.01927](#)].
- [182] CMS collaboration, *CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV*, Tech. Rep. [CMS-PAS-LUM-17-004](#), CERN, Geneva (2018).
- [183] CMS collaboration, *CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV*, Tech. Rep. [CMS-PAS-LUM-18-002](#), CERN, Geneva (2019).
- [184] CMS collaboration, *CMS luminosity measurement using nucleus-nucleus collisions at $\sqrt{s_{NN}} = 5.02$ TeV in 2018*, Tech. Rep. [CMS-PAS-LUM-18-001](#), CERN, Geneva (2022).
- [185] CMS collaboration, *CMS Luminosity Calibration for the pp Reference Run at $\sqrt{s} = 5.02$ TeV*, Tech. Rep. [CMS-PAS-LUM-16-001](#), CERN, Geneva (2016).
- [186] CMS collaboration, *Luminosity measurement in proton-proton collisions at 5.02 TeV in 2017 at CMS*, Tech. Rep. [CMS-PAS-LUM-19-001](#), CERN, Geneva (2021).
- [187] CMS collaboration, *CMS luminosity measurement using 2016 proton-nucleus collisions at nucleon-nucleon center-of-mass energy of 8.16 TeV*, Tech. Rep. [CMS-PAS-LUM-17-002](#), CERN, Geneva (2018).
- [188] A. Kornmayer, *The CMS pixel luminosity telescope*, [Nucl. Instrum. Meth. A 824 \(2016\) 304](#).
- [189] K. Rose, *The new pixel luminosity telescope of CMS at the LHC*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference*, San Diego CA, U.S.A., 31 October–7 November 2015 [[DOI:10.1109/NSSMIC.2015.7581875](#)].
- [190] P. Lujan, *Performance of the Pixel Luminosity Telescope for Luminosity Measurement at CMS during Run2*, [PoS EPS-HEP2017 \(2017\) 504](#).
- [191] CMS BRIL collaboration, *The Pixel Luminosity Telescope: a detector for luminosity measurement at CMS using silicon pixel sensors*, [Eur. Phys. J. C 83 \(2023\) 673](#) [[arXiv:2206.08870](#)].
- [192] N. Karunaratna, *Run 3 luminosity measurements with the Pixel Luminosity Telescope*, [PoS ICHEP2022 \(2022\) 936](#).
- [193] CMS collaboration, *Pixel luminosity telescope (PLT) calibration and first measurements*, CMS Detector Performance Note [CMS-DP-2015-025](#) (2015).
- [194] G. Bolla et al., *Sensor development for the CMS pixel detector*, [Nucl. Instrum. Meth. A 485 \(2002\) 89](#).
- [195] Y. Allkofer et al., *Design and performance of the silicon sensors for the CMS barrel pixel detector*, [Nucl. Instrum. Meth. A 584 \(2008\) 25](#) [[physics/0702092](#)].
- [196] H.C. Kastli et al., *Design and performance of the CMS pixel detector readout chip*, [Nucl. Instrum. Meth. A 565 \(2006\) 188](#) [[physics/0511166](#)].
- [197] M. Barbero, *Development of a radiation-hard pixel read out chip with trigger capability*, Ph.D. thesis, Universität Basel, Basel, Switzerland (2003).
- [198] E. Bartz, *The 0.25um Token Bit Manager Chip for the CMS Pixel Readout*, in the proceedings of the *11th Workshop on Electronics for LHC and Future Experiments*, Heidelberg, Germany, 12–16 September 2005, pp. 153–157, <http://cds.cern.ch/record/920426> [[DOI:10.5170/CERN-2005-011.153](#)].

- [199] M. Pernicka et al., *The CMS Pixel FED*, in the proceedings of the *Topical Workshop on Electronics for Particle Physics*, Prague, Czech Republic, 3–7 September 2007, pp. 487–491, <http://cds.cern.ch/record/1091743> [DOI: 10.5170/CERN-2007-007.487].
- [200] E. van der Bij, R.A. McLaren, O. Boyle and G. Rubin, *S-LINK, a data link interface specification for the LHC era*, *IEEE Trans. Nucl. Sci.* **44** (1997) 398.
- [201] A.A. Zagozdinska et al., *New Fast Beam Conditions Monitoring (BCM1F) system for CMS*, *2016 JINST* **11** C01088.
- [202] M. Guthoff, *The new Fast Beam Condition Monitor using poly-crystalline diamond sensors for luminosity measurement at CMS*, *Nucl. Instrum. Meth. A* **936** (2019) 717.
- [203] J. Wańczyk, *Upgraded CMS Fast Beam Condition Monitor for LHC Run 3 Online Luminosity and Beam Induced Background Measurements*, *JACoW IBIC2022* (2022) 540.
- [204] A.G. Delannoy Sotomayor and J. Wańczyk, *CMS BRIL BCM1F c-shape*, CMS Collection *CMS-PHO-GEN-2022-001* (2022).
- [205] D. Przyborowski, J. Kaplon and P. Rymaszewski, *Design and Performance of the BCM1F Front End ASIC for the Beam Condition Monitoring System at the CMS Experiment*, *IEEE Trans. Nucl. Sci.* **63** (2016) 2300.
- [206] M. Friedl, *Analog optohybrids, CMS tracker TOB/TEC*, Tech. Rep. *CMS-TK-CS-0002*, CERN, Geneva (2004).
- [207] A. Dabrowski, *A real-time histogramming unit for luminosity measurements of each bunch crossing at CMS*, talk given at the *Topical Workshop on Electronics for Particle Physics*, Perugia, Italy, 23–27 September 2013 [DESY-2013-00940].
- [208] A.A. Zagozdinska et al., *The fast beam condition monitor BCM1F backend electronics upgraded MicroTCA based architecture*, *Proc. SPIE* **9290** (2014) 92902L.
- [209] A.A. Zagozdinska, *The CMS fast beams condition monitor back-end electronics based on MicroTCA technology*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference*, San Diego CA, U.S.A., 31 October–7 November 2015 [DOI: 10.1109/NSSMIC.2015.7581987].
- [210] A.J. Rüde, *New Peak Finding Algorithm for the BCM1F Detector of the CMS Experiment at CERN*, M.Sc. thesis, Ernst-Abbe-Hochschule, Jena, Germany (2018), *CERN-THESIS-2018-021*.
- [211] CMS collaboration, *BRIL luminosity performance plots: Cross-detector stability in early Run 3 data*, Tech. Rep. *CMS-DP-2022-038* CERN, Geneva (2022).
- [212] CMS collaboration, *The Phase-2 Upgrade of the CMS Beam Radiation, Instrumentation, and Luminosity Detectors: Conceptual Design*, Tech. Rep. *CMS-NOTE-2019-008*, CERN, Geneva (2020).
- [213] A. Triossi, *The CMS Barrel Muon trigger upgrade*, *2017 JINST* **12** C01095.
- [214] J. Salfeld-Nebgen and D. Marlow, *Data-Driven Precision Luminosity Measurements with Z Bosons at the LHC and HL-LHC*, *2018 JINST* **13** P12016 [arXiv:1806.02184].
- [215] CMS collaboration, *Luminosity determination using Z boson production at the CMS experiment*, arXiv:2309.01008.
- [216] CMS collaboration, *Luminosity monitoring with Z counting in early 2022 data*, Tech. Rep. *CMS-DP-2023-003* CERN, Geneva (2023).
- [217] CMS collaboration, *First measurement of the top quark pair production cross section in proton-proton collisions at $\sqrt{s} = 13.6$ TeV*, *JHEP* **08** (2023) 204 [arXiv:2303.10680].

- [218] CMS collaboration, *CMS $Z(\mu\mu)$ yields for comparisons with ATLAS*, Tech. Rep. [CMS-DP-2012-014](#) CERN, Geneva (2012).
- [219] C. Schwick and B. Petersen, *LPC's view on Run 2*, in the proceedings of the 9th *Evian Workshop on LHC beam operation*, Evian Les Bains, France, 30 January–1 February 2019, pp. 27–33, <http://cds.cern.ch/record/2750274>.
- [220] S. Orfanelli et al., *A novel Beam Halo Monitor for the CMS experiment at the LHC*, [2015 JINST 10 P11011](#).
- [221] N. Tosi et al., *Electronics and Calibration system for the CMS Beam Halo Monitor*, [PoS TIPP2014 \(2014\) 411](#).
- [222] N. Tosi, *The new Beam Halo Monitor for the CMS experiment at the LHC*, Ph.D. thesis, Università di Bologna, Bologna, Italy (2015), [CERN-THESIS-2015-283] [[DOI:10.6092/unibo/amsdottorato/6767](#)].
- [223] S. Müller, *The Beam Condition Monitor 2 and the Radiation Environment of the CMS Detector at the LHC*, Ph.D. thesis, KIT, Karlsruhe, Germany (2011) [CERN-THESIS-2011-085] [[DOI:10.5445/IR/1000022456](#)].
- [224] M. Guthoff, *Radiation Damage to the diamond-based Beam Condition Monitor of the CMS Detector at the LHC*, Ph.D. thesis, KIT, Karlsruhe, Germany (2014) [CERN-THESIS-2014-216] [[DOI:10.5445/IR/1000044731](#)].
- [225] F.R. Kassel, *The Rate Dependent Radiation Induced Signal Degradation of Diamond Detectors*, Ph.D. thesis, KIT, Karlsruhe, Germany (2017) [CERN-THESIS-2017-071] [[DOI:10.5445/IR/1000071068](#)].
- [226] B. Dehning et al., *LHC beam loss monitor system design*, [AIP Conf. Proc. 648 \(2002\) 229](#).
- [227] J. Emery et al., *Functional and Linearity test system for the LHC Beam Loss Monitoring data acquisition card*, in the proceedings of the 12th *Workshop on Electronics for LHC and Future Experiments*, Valencia, Spain, 25–29 September 2006, pp. 447–451, <http://cds.cern.ch/record/1057135> [[DOI:10.5170/CERN-2007-001.447](#)].
- [228] C. Zamantzas, *The Real-Time Data Analysis and Decision System for Particle Flux Detection in the LHC Accelerator at CERN*, Ph.D. thesis, Brunel University, London, U.K. (2006), <http://cds.cern.ch/record/976628> CERN-THESIS-2006-037.
- [229] B. Todd, A. Dinius, C. Martin and B. Puccio, *User interface to the beam interlock system*, Tech. Rep. [EDMS Document 636589 v1.5](#) CERN, Geneva (2011).
- [230] V.N. Kurlov, *Sapphire: Properties, growth, and applications*, in *Reference module: Material science and materials engineering*, Elsevier, Amsterdam, Netherlands (2016), ISBN 978-0128035818 [[DOI:10.1016/B978-0-12-803581-8.03681-X](#)].
- [231] O. Karacheban et al., *Investigation of a direction sensitive sapphire detector stack at the 5 GeV electron beam at DESY-II*, [2015 JINST 10 P08008](#) [[arXiv:1504.04023](#)].
- [232] CMS collaboration, *FLUKA Run 2 simulation benchmark with beam loss monitors in the CMS forward region*, Tech. Rep. [CMS-DP-2021-008](#) CERN, Geneva (2021).
- [233] A.M. Gribushin et al., *A neutron field monitoring system for collider experiments*, [Instrum. Exp. Tech. 60 \(2017\) 167](#).
- [234] G. Segura Millan, D. Perrin, and L. Scibile, , in the proceedings of the 10th *International Conference*

on Accelerator and Large Experimental Physics Control Systems, Geneva, Switzerland, 10–15 October 2005 [*Conf. Proc. C* **051010** (2005) TH3B.1-3O].

- [235] A. Ledoul et al., *CERN Supervision, Control and Data Acquisition System for Radiation and Environmental Protection*, in the proceedings of the 12th International Workshop on Personal Computers and Particle Accelerator Controls, Hsinchu, Taiwan, 16–19 October 2018, pp. 248–252 [[DOI:10.18429/JACoW-PCaPAC2018-FRCC3](#)].
- [236] G. Spiezia et al., *The LHC Radiation Monitoring System — RadMon*, *PoS* **RD11** (2011) 024.
- [237] C. Martinella, *High energy hadrons fluence measurements in the LHC during 2015, 2016 and 2017 proton physics operations*, Tech. Rep. [CERN-ACC-NOTE-2018-088](#), CERN, Geneva (2018).
- [238] T.T. Böhlen et al., *The FLUKA Code: Developments and Challenges for High Energy and Medical Applications*, *Nucl. Data Sheets* **120** (2014) 211.
- [239] J. Gutleber, S. Murray and L. Orsini, *Towards a homogeneous architecture for high-energy physics data acquisition systems*, *Comput. Phys. Commun.* **153** (2003) 155.
- [240] B. Copy, E. Mandilara, I. Prieto Barreiro and F. Varela, *Monitoring of CERN’s Data Interchange Protocol (DIP) system*, in the proceedings of the 16th International Conference on Accelerator and Large Experimental Physics Control Systems, Barcelona, Spain, 8–13 October 2017, pp. 1797–1800 [[DOI:10.18429/JACoW-ICALEPCS2017-THPHA162](#)].
- [241] Y. Fain and A. Moiseev, *Angular 2 development with TypeScript*, Manning Publications, Shelter Island, NY, U.S.A. (2016), ISBN 978-1617293122.
- [242] J.M. Andre et al., *A Scalable Monitoring for the CMS Filter Farm Based on Elasticsearch*, *J. Phys. Conf. Ser.* **664** (2015) 082036.
- [243] G. Bauer et al., *The CMS Event Builder and Storage System*, *J. Phys. Conf. Ser.* **219** (2010) 022038.
- [244] G. Bauer et al., *The new CMS DAQ system for LHC operation after 2014 (DAQ2)*, *J. Phys. Conf. Ser.* **513** (2014) 012014.
- [245] T. Bawej et al., *The New CMS DAQ System for Run-2 of the LHC*, *IEEE Trans. Nucl. Sci.* **62** (2015) 1099.
- [246] G. Bauer et al., *The Terabit/s Super-Fragment Builder and Trigger Throttling System for the Compact Muon Solenoid Experiment at CERN*, *IEEE Trans. Nucl. Sci.* **55** (2008) 190.
- [247] G. Bauer et al., *10 Gbps TCP/IP streams from the FPGA for the CMS DAQ eventbuilder network*, *2013 JINST* **8** C12039.
- [248] G. Bauer et al., *10 Gbps TCP/IP streams from the FPGA for High Energy Physics*, *J. Phys. Conf. Ser.* **513** (2014) 012042.
- [249] D. Gigi et al., *The FEROL40, a microTCA card interfacing custom point-to-point links and standard TCP/IP*, *PoS TWEPP-17* (2017) 075.
- [250] G. Bauer et al., *Upgrade of the CMS Event Builder*, *J. Phys. Conf. Ser.* **396** (2012) 012039.
- [251] J.-M. André et al., *Performance of the new DAQ system of the CMS experiment for run-2*, in the proceedings of the 20th IEEE-NPSS Real Time Conference, Padua, Italy, 5–10 June 2016 [[DOI:10.1109/RTC.2016.7543164](#)].
- [252] J.M. Andre et al., *Performance of the CMS Event Builder*, *J. Phys. Conf. Ser.* **898** (2017) 032020.

- [253] Juniper Networks, Inc., *QFX10000 modular Ethernet switches*, Datasheet 1000529-022-EN, 2021, <https://www.juniper.net/us/en/products/switches/qfx-series/qfx10000-modular-ethernet-switches-datasheet.html>.
- [254] IEEE Standards Association, *IEEE 802.3-2012, Standard for Ethernet*, 2012, <https://standards.ieee.org/ieee/802.3/5084/>.
- [255] *tmpfs*, Software available in the Linux kernel, 2021, <https://man7.org/linux/man-pages/man5/tmpfs.5.html>.
- [256] C.D. Jones et al., *Using the CMS Threaded Framework In A Production Environment*, *J. Phys. Conf. Ser.* **664** (2015) 072026.
- [257] J.M. Andre et al., *File-Based Data Flow in the CMS Filter Farm*, *J. Phys. Conf. Ser.* **664** (2015) 082033.
- [258] *inotify*, software available in the Linux kernel, <https://man7.org/linux/man-pages/man7/inotify.7.html>.
- [259] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81.
- [260] P. Deutsch and J.-L. Gailly, *ZLIB compressed data format specification version 3.3*, RFC Informational 1950, 1996 [DOI: 10.17487/RFC1950].
- [261] *XZ Utils* <https://github.com/tukaani-project/xz>. More information at <https://tukaani.org/xz/>
- [262] Y. Collet and M. Kucherauw, *Zstandard compression and the 'application/zstd' media type*, RFC Informational 8878, 2021 [DOI: 10.17487/RFC8878].
- [263] *Elasticsearch*, <https://github.com/elastic/elasticsearch>. More information at <https://www.elastic.co/elasticsearch/>.
- [264] T. Bray, *The JavaScript object notation (JSON) data interchange format*, RFC Proposed Standard 7159, 2014 [DOI: 10.17487/RFC7159].
- [265] M. Michelotto et al., *A comparison of HEP code with SPEC benchmarks on multi-core worker nodes*, *J. Phys. Conf. Ser.* **219** (2010) 052009.
- [266] Nvidia Corporation, *NVIDIA T4 tensor core GPU*, Datasheet, 2019, <https://www.nvidia.com/en-us/data-center/tesla-t4/>.
- [267] A.J. Peters, E.A. Sindrilaru and G. Adde, *EOS as the present and future solution for data storage at CERN*, *J. Phys. Conf. Ser.* **664** (2015) 042042.
- [268] J.M. Andre et al., *Online Data Handling and Storage at the CMS Experiment*, *J. Phys. Conf. Ser.* **664** (2015) 082009.
- [269] DataDirect Networks, *EXAScaler product family*, DDN Data Sheet, 2023, <https://www.ddn.com/exascaler>.
- [270] J. Hegeman et al., *The CMS Timing and Control Distribution System*, in the proceedings of the *IEEE Nuclear Science Symposium and Medical Imaging Conference*, San Diego CA, U.S.A., 31 October–7 November 2015 [DOI: 10.1109/NSSMIC.2015.7581984].
- [271] CentOS Project, *CentOS-7 (2009)*, Release Notes, 2020, <https://wiki.centos.org/action/show/Manuals/ReleaseNotes/CentOS7.2009>.
- [272] Red Hat, Inc., *Red Hat Enterprise Linux 8.8*, Release Notes, 2023, https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8.

- [273] R. García Leiva et al., *Quattor: Tools and techniques for the configuration, installation and management of large-scale grid computing fabrics*, *J. Grid Comput.* **2** (2004) 313.
- [274] *Puppet*, <https://github.com/puppetlabs/puppet>. More information at <https://www.puppet.com/>.
- [275] *oVirt*, <https://github.com/ovirt/>. More information at <https://www.ovirt.org/>.
- [276] J.-M. Andre et al., *Experience with dynamic resource provisioning of the CMS online cluster using a cloud overlay*, *EPJ Web Conf.* **214** (2019) 07017.
- [277] G. Bauer et al., *First operational experience with a high-energy physics run control system based on Web technologies*, *IEEE Trans. Nucl. Sci.* **59** (2012) 1597.
- [278] G. Bauer et al., *A comprehensive zero-copy architecture for high performance distributed data acquisition over advanced network technologies for the CMS experiment*, in the proceedings of the 18th IEEE-NPSS Real-Time Conference, Berkeley CA, U.S.A., 11–15 June 2012 [DOI:10.1109/RTC.2012.6418171].
- [279] T. Bawej et al., *Achieving High Performance With TCP Over 40 GbE on NUMA Architectures for CMS Data Acquisition*, *IEEE Trans. Nucl. Sci.* **62** (2015) 1091.
- [280] G. Bauer et al., *Distributed Error and Alarm Processing in the CMS Data Acquisition System*, *J. Phys. Conf. Ser.* **396** (2012) 012038.
- [281] G. Bauer et al., *Automating the CMS DAQ*, *J. Phys. Conf. Ser.* **513** (2014) 012031.
- [282] J.M. Andre et al., *New Operator Assistance Features in the CMS Run Control System*, *J. Phys. Conf. Ser.* **898** (2017) 032028.
- [283] J.-M. Andre et al., *DAQExpert — An expert system to increase CMS data-taking efficiency*, *J. Phys. Conf. Ser.* **1085** (2018) 032021.
- [284] J.-M. Andre et al., *Operational experience with the new CMS DAQ-Expert*, *EPJ Web Conf.* **214** (2019) 01015.
- [285] G. Badaro et al., *DAQExpert the service to increase CMS data-taking efficiency*, *EPJ Web Conf.* **245** (2020) 01028.
- [286] J.A. Lopez-Perez et al., *The Web Based Monitoring project at the CMS experiment*, *J. Phys. Conf. Ser.* **898** (2017) 092040.
- [287] J. Duarte et al., *Fast inference of deep neural networks in FPGAs for particle physics*, 2018 JINST **13** P07027 [arXiv:1804.06913].
- [288] R. Bainbridge, *Recording and reconstructing 10 billion unbiased b hadron decays in CMS*, *EPJ Web Conf.* **245** (2020) 01025.
- [289] F.M.A. Erich, C. Amrit and M. Daneva, *A qualitative study of DevOps usage in practice*, *J. Softw. Evol. Process* **29** (2017) e1885.
- [290] T. Dingsøyr, S. Nerur, V. Balijepally, and N. B. Moe, *A decade of agile methodologies: Towards explaining agile software development*, *J. Syst. Softw.* **85** (2012) 1213.
- [291] *Prometheus*, <https://github.com/prometheus/prometheus>. More information at <https://prometheus.io/>.
- [292] S.H. Sunil Kumar and C. Saravanan, *A comprehensive study on data visualization tool: Grafana*, *JETIR* **8** (2021) f908.
- [293] A.X. Ming Chang et al., *Deep neural networks compiler for a trace-based accelerator*, *J. Syst. Archit.* **102** (2020) 101659.

- [294] Xilinx Inc., *VCU128 evaluation board*, User Guide UG1302 v1.2, 2022, <https://docs.xilinx.com/r/en-US/ug1302-vcu128-eval-bd>.
- [295] CMS collaboration, *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*, Tech. Rep. [CERN-LHCC-2021-007](#), CERN, Geneva (2021).
- [296] D. Golubovic et al., *40 MHz Scouting with Deep Learning in CMS*, in the proceedings of the 6th *International Workshop Connecting the Dots*, Princeton NJ, U.S.A., 20–30 April 2020 <https://zenodo.org/record/5040675> [DOI:10.5281/ZENODO.5040675].
- [297] CMS collaboration, *40 MHz Scouting with Deep Learning in CMS*, Tech. Rep. [CMS-DP-2022-066](#) CERN, Geneva (2022).
- [298] R. Ardino et al., *A 40 MHz Level-1 trigger scouting system for the CMS Phase-2 upgrade*, *Nucl. Instrum. Meth. A* **1047** (2023) 167805.
- [299] A. Bocci et al., *Heterogeneous Reconstruction of Tracks and Primary Vertices With the CMS Pixel Tracker*, *Front. Big Data* **3** (2020) 601728 [[arXiv:2008.13461](#)].
- [300] CMS collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, *2017 JINST* **12** P10003 [[arXiv:1706.04965](#)].
- [301] CMS collaboration, *Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *2021 JINST* **16** P07001 [[arXiv:2102.04790](#)].
- [302] W. Adam, R. Fruhwirth, A. Strandlie and T. Todorov, *Reconstruction of electrons with the Gaussian sum filter in the CMS tracker at LHC*, *J. Phys. G* **31** (2005) N9 [[physics/0306087](#)].
- [303] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)].
- [304] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)].
- [305] CMS collaboration, *Performance of tau-lepton reconstruction and identification in CMS*, *2012 JINST* **7** P01001 [[arXiv:1109.6034](#)].
- [306] CMS collaboration, *Identification of hadronic tau lepton decays using a deep neural network*, *2022 JINST* **17** P07023 [[arXiv:2201.08458](#)].
- [307] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)].
- [308] CMS collaboration, *Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV*, *2018 JINST* **13** P05011 [[arXiv:1712.07158](#)].
- [309] E. Bols et al., *Jet Flavour Classification Using DeepJet*, *2020 JINST* **15** P12012 [[arXiv:2008.10519](#)].
- [310] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)].
- [311] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)].
- [312] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in the proceedings of the *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)*, Hamburg, Germany, 27–30 April 1998, p. 270–279 [[hep-ph/9907280](#)].

- [313] CMS collaboration, *Search for Narrow Resonances and Quantum Black Holes in Inclusive and b -Tagged Dijet Mass Spectra from pp Collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **01** (2013) 013 [[arXiv:1210.2387](#)].
- [314] CMS collaboration, *Search for narrow resonances in dijet final states at $\sqrt{s} = 8$ TeV with the novel CMS technique of data scouting*, *Phys. Rev. Lett.* **117** (2016) 031802 [[arXiv:1604.08907](#)].
- [315] CMS collaboration, *Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. Lett.* **124** (2020) 131802 [[arXiv:1912.04776](#)].
- [316] CMS collaboration, *Search for pair-produced three-jet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **99** (2019) 012010 [[arXiv:1810.10092](#)].
- [317] GEANT4 collaboration, *GEANT4 — a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.
- [318] D.J. Lange, M. Hildreth, V.N. Ivantchenko and I. Osborne, *Upgrades for the CMS simulation*, *J. Phys. Conf. Ser.* **608** (2015) 012056.
- [319] S. Sekmen, *Recent Developments in CMS Fast Simulation*, *PoS ICHEP2016* (2016) 181 [[arXiv:1701.03850](#)].
- [320] K. Pedro, *Current and future performance of the CMS simulation*, *EPJ Web Conf.* **214** (2019) 02036.
- [321] C. Caputo, *Enabling continuous speedup of CMS Event Reconstruction through continuous benchmarking*, in the proceedings of the 21th International Workshop on Advanced Computing and Analysis Techniques in Physics Research, Bari, Italy, 24–28 October 2022 [[CMS-CR-2023-038](#)].
- [322] D. Piparo, *Automated quality monitoring and validation of the CMS reconstruction software*, *J. Phys. Conf. Ser.* **368** (2012) 012008.
- [323] M.N. Rodozov and D. Lange, *Modernizing the CMS software stack*, <https://zenodo.org/record/3599193> [[DOI:10.5281/ZENODO.3599193](#)].
- [324] J. Blomer, C. Aguado Sanchez, P. Buncic and A. Harutyunyan, *Distributing LHC application software and conditions databases using the CernVM file system*, *J. Phys. Conf. Ser.* **331** (2011) 042003.
- [325] J. Blomer et al., *The CernVM File System: v2.7.5*, <https://zenodo.org/record/4114078> [[DOI:10.5281/ZENODO.4114078](#)].
- [326] C.D. Jones and E. Sexton-Kennedy, *Stitched Together: Transitioning CMS to a Hierarchical Threaded Framework*, *J. Phys. Conf. Ser.* **513** (2014) 022034.
- [327] *oneAPI Threading Building Blocks (oneTBB)*, <https://github.com/oneapi-src/oneTBB>. More information at <https://oneapi-src.github.io/oneTBB/>.
- [328] I. Bird et al., *LHC computing Grid: Technical Design Report. Version 1.06* (20 Jun 2005), Tech. Rep. [CERN-LHCC-2005-024](#), CERN, Geneva (2005).
- [329] C.D. Jones, *CMS event processing multi-core efficiency status*, *J. Phys. Conf. Ser.* **898** (2017) 042008.
- [330] A. Bocci et al., *Bringing heterogeneity to the CMS software framework*, *EPJ Web Conf.* **245** (2020) 05009 [[arXiv:2004.04334](#)].
- [331] E. Zenker et al., *Alpaka — An Abstraction Library for Parallel Kernel Acceleration*, in the proceedings of the *IEEE International Parallel and Distributed Processing Symposium Workshops*, Chicago IL, U.S.A., 23–27 May 2016 [[arXiv:1602.08477](#)] [[DOI:10.1109/IPDPSW.2016.50](#)].
- [332] H.C. Edwards, C.R. Trott and D. Sunderland, *Kokkos: Enabling manycore performance portability through polymorphic memory access patterns*, *J. Parallel Distributed Comput.* **74** (2014) 3202.

- [333] M.J. Kortelainen et al., *Porting CMS Heterogeneous Pixel Reconstruction to Kokkos*, *EPJ Web Conf.* **251** (2021) 03034 [[arXiv:2104.06573](#)].
- [334] A. Bocci et al., *Performance portability for the CMS Reconstruction with Alpaka*, *J. Phys. Conf. Ser.* **2438** (2023) 012058.
- [335] J. Duarte et al., *FPGA-accelerated machine learning inference as a service for particle physics computing*, *Comput. Softw. Big Sci.* **3** (2019) 13 [[arXiv:1904.08986](#)].
- [336] D.S. Rankin et al., *FPGAs-as-a-Service Toolkit (FaaSST)*, in the proceedings of the *IEEE/ACM International Workshop on Heterogeneous High-Performance Reconfigurable Computing*, Atlanta GA, U.S.A., 13 November 2020 [[arXiv:2010.08556](#)] [[DOI:10.1109/H2RC51942.2020.00010](#)].
- [337] J. Krupa et al., *GPU coprocessors as a service for deep learning inference in high energy physics*, *Mach. Learn. Sci. Tech.* **2** (2021) 035005 [[arXiv:2007.10359](#)].
- [338] Nvidia Corporation, *Triton inference server*, Release Notes RN-08995-001 v23.04, 2023, <https://docs.nvidia.com/deeplearning/triton-inference-server/pdf/Triton-Inference-Server-Release-Notes.pdf>.
- [339] M. Case, A.T.M. Aerts, A. Muhammad and M. Liendl, *CMS Detector Description: New Developments*, in the proceedings of the *14th International Conference on Computing in High-Energy and Nuclear Physics*, Interlaken, Switzerland, 27 September–1 October 2004, <http://cds.cern.ch/record/865635> [[DOI:10.5170/CERN-2005-002.498](#)].
- [340] M. Frank, F. Gaede, C. Grefe and P. Mato, *DD4hep: A Detector Description Toolkit for High Energy Physics Experiments*, *J. Phys. Conf. Ser.* **513** (2014) 022010.
- [341] G. Petrucciani, A. Rizzi and C. Vuosalo, *Mini-AOD: A New Analysis Data Format for CMS*, *J. Phys. Conf. Ser.* **664** (2015) 7 [[arXiv:1702.04685](#)].
- [342] A. Rizzi, G. Petrucciani and M. Peruzzi, *A further reduction in CMS event data for analysis: the NANO AOD format*, *EPJ Web Conf.* **214** (2019) 06021.
- [343] M. Hildreth, V.N. Ivanchenko and D.J. Lange, *Upgrades for the CMS simulation*, *J. Phys. Conf. Ser.* **898** (2017) 042040.
- [344] M. Aderholz et al., *Models of Networked Analysis at Regional Centres for LHC Experiments (MONARC), Phase 2 Report*, 24th March 2000, Tech Rep. [CERN-LCB-2000-001](#), CERN, Geneva (2000).
- [345] D. Futyan, *Commissioning the CMS alignment and calibration framework*, *J. Phys. Conf. Ser.* **219** (2010) 032041.
- [346] K. Bloom et al., *Any data, any time, anywhere: Global data access for science*, in the proceedings of the *IEEE/ACM 2nd International Symposium on Big Data Computing*, Limassol, Cyprus, 7–10 December 2015 [[DOI:10.1109/BDC.2015.33](#)].
- [347] L. Bauerdick et al., *Using Xrootd to Federate Regional Storage*, *J. Phys. Conf. Ser.* **396** (2012) 042009.
- [348] L.A.T. Bauerdick et al., *XRootd, disk-based, caching proxy for optimization of data access, data placement and data replication*, *J. Phys. Conf. Ser.* **513** (2014) 042044.
- [349] G.M. Kurtzer, V. Sochat and M.W. Bauer, *Singularity: Scientific containers for mobility of compute*, *PLOS One* **12** (2017) e0177459.
- [350] T. Boccali et al., *Dynamic Distribution of High-Rate Data Processing from CERN to Remote HPC Data Centers*, *Comput. Softw. Big Sci.* **5** (2021) 7.
- [351] T. Boccali et al., *Extension of the INFN Tier-1 on a HPC system*, *EPJ Web Conf.* **245** (2020) 09009 [[arXiv:2006.14603](#)].

- [352] C. Acosta-Silva et al., *Exploitation of network-segregated CPU resources in CMS*, *EPJ Web Conf.* **251** (2021) 02020.
- [353] M. Fischer et al., *Effective Dynamic Integration and Utilization of Heterogenous Compute Resources*, *EPJ Web Conf.* **245** (2020) 07038.
- [354] T. Boccali et al., *Enabling CMS Experiment to the utilization of multiple hardware architectures: a Power9 Testbed at CINECA*, *J. Phys. Conf. Ser.* **2438** (2023) 012031.
- [355] D. Hufnagel et al., *HPC resource integration into CMS Computing via HEPCloud*, *EPJ Web Conf.* **214** (2019) 03031.
- [356] M.C. Davis et al., *CERN Tape Archive - from development to production deployment*, *EPJ Web Conf.* **214** (2019) 04015.
- [357] J. Rehn et al., “PhEDEx high-throughput data transfer management system”, in the proceedings of the 15th International Conference on Computing in High Energy and Nuclear Physics, Mumbai, India, 13–17 February 2006, https://indico.cern.ch/event/408139/contributions/979918/attachments/815907/1118010/CHEP06_PhEDEx_overview.pdf
- [358] Y. Iiyama et al., *Dynamo: Handling Scientific Data Across Sites and Storage Media*, *Comput. Softw. Big Sci.* **5** (2021) 11 [arXiv:2003.11409].
- [359] I. Béjar Alonso et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, Tech. Rep. CERN-2020-010 CERN, Geneva (2020) [DOI: 10.23731/CYRM-2020-0010].
- [360] E. Karavakis et al., *FTS improvements for LHC Run-3 and beyond*, *EPJ Web Conf.* **245** (2020) 04016.
- [361] M. Barisits et al., *Rucio - Scientific data management*, *Comput. Softw. Big Sci.* **3** (2019) 11 [arXiv:1902.09857].
- [362] *Helm: the package manager for Kubernetes*, <https://github.com/helm/helm>. More information at <https://helm.sh/>.
- [363] *Kubernetes (K8s)*, <https://github.com/kubernetes/kubernetes>. More information at <https://kubernetes.io/>.
- [364] *Docker*, <https://github.com/docker>. More information at <https://www.docker.com/>.
- [365] I. Foster, C. Kesselman, G. Tsudik and S. Tuecke, “A security architecture for computational grids”, in the proceedings of the 5th ACM conference on Computer and Communications Security, San Francisco CA, U.S.A., 2–5 November 1998, [DOI: 10.1145/288090.288111].
- [366] R. Butler et al., *A national-scale authentication infrastructure*, *Computer* **33** (2000) 60.
- [367] L. Dusseault, *HTTP extensions for web distributed authoring and versioning (WebDAV)*, RFC Proposed Standard 4918, 2007 [DOI: 10.17487/RFC4918].
- [368] X. Espinal et al., *The Quest to solve the HL-LHC data access puzzle*, *EPJ Web Conf.* **245** (2020) 04027.
- [369] I. Sfiligoi et al., *The pilot way to Grid resources using glideinWMS*, in the proceedings of the WRI World Congress on Computer Science and Information Engineering Los Angeles CA, U.S.A., 31 March–2 April 2009, pp. 428–432 [DOI: 10.1109/CSIE.2009.950].
- [370] D. Thain, T. Tannenbaum, and M. Livny, *Distributed computing in practice: the Condor experience*, *Concurr. Comput.* **17** (2005) 323.
- [371] J. Balcas et al., *Using the GlideinWMS System as a Common Resource Provisioning Layer in CMS*, *J. Phys. Conf. Ser.* **664** (2015) 062031.

- [372] A. Pérez-Calero Yzquierdo et al., *Evolution of the CMS Global Submission Infrastructure for the HL-LHC Era*, *EPJ Web Conf.* **245** (2020) 03016.
- [373] D. Spiga et al., *Exploiting private and commercial clouds to generate on-demand CMS computing facilities with DODAS*, *EPJ Web Conf.* **214** (2019) 07027.
- [374] D.P. Anderson, *BOINC: A platform for volunteer computing*, *J. Grid Comput.* **18** (2020) 99.
- [375] D. Smith et al., *Sharing server nodes for storage and compute*, *EPJ Web Conf.* **214** (2019) 08025.
- [376] D. Giordano et al., *Cern-It Evaluation Of Microsoft Azure Cloud Iaas*, <https://zenodo.org/record/48495> [DOI: 10.5281/ZENODO.48495].
- [377] C. Cordeiro et al., *CERN computing in commercial clouds*, *J. Phys. Conf. Ser.* **898** (2017) 082030.
- [378] A. Perez-Calero Yzquierdo et al., *CMS Readiness for Multi-Core Workload Scheduling*, *J. Phys. Conf. Ser.* **898** (2017) 052030.
- [379] B.P. Bockelman et al., *Improving the Scheduling Efficiency of a Global Multi-Core HTCondor Pool in CMS*, *EPJ Web Conf.* **214** (2019) 03056.
- [380] A. Pérez-Calero Yzquierdo et al., *Reaching new peaks for the future of the CMS HTCondor Global Pool*, *EPJ Web Conf.* **251** (2021) 02055.
- [381] P. Couvares et al., *Workflow management in Condor*, in *Workflows for e-Science*, Springer Nature, London, U.K., 2007. ISBN 978-1846285196 [DOI: 10.1007/978-1-84628-757-2_22].
- [382] J. Balcas et al., *CMS Connect*, *J. Phys. Conf. Ser.* **898** (2017) 082032.
- [383] V. Kuznetsov, D. Evans and S. Metson, *The CMS data aggregation system*, in the proceedings of the 10th International Conference on Computational Science, Amsterdam, Netherlands, 31 May–2 June 2010 [DOI: 10.1016/j.procs.2010.04.172].
- [384] M. Imran et al., *Migration of CMSWEB cluster at CERN to Kubernetes: a comprehensive study*, *Clust. Comput.* **24** (2021) 3085.
- [385] WLCG Authorization WG, *WLCG Token Transition Timeline*, <https://zenodo.org/record/7014668> [DOI: 10.5281/ZENODO.7014668].
- [386] C. Ariza-Porras, V. Kuznetsov and F. Legger, *The CMS monitoring infrastructure and applications*, *Comput. Softw. Big Sci.* **5** (2021) 5 [arXiv:2007.03630].
- [387] A. Aimar et al., *Unified Monitoring Architecture for IT and grid services*, *J. Phys. Conf. Ser.* **898** (2017) 092033.

The CMS collaboration

Yerevan Physics Institute, Yerevan, Armenia

A. Hayrapetyan, A. Tumasyan ¹

Institut für Hochenergiephysik, Vienna, Austria

W. Adam , J.W. Andrejkovic, B. Arnold, H. Bergauer, T. Bergauer , S. Chatterjee , K. Damanakis ,
M. Dragicevic , A. Escalante Del Valle , P.S. Hussain , M. Jeitler ², N. Krammer , D. Liko ,
I. Mikulec , J. Schieck ², R. Schöfbeck , D. Schwarz , M. Sonawane , S. Templ ,
W. Waltenberger , C.-E. Wulz ²

Universiteit Antwerpen, Antwerpen, Belgium

M.R. Darwish ³, T. Janssen , P. Van Mechelen 

Vrije Universiteit Brussel, Brussel, Belgium

E.S. Bols , J. D'Hondt , S. Dansana , A. De Moor , M. Delcourt , H. El Faham , S. Lowette ,
I. Makarenko , A. Morton , D. Müller , A.R. Sahasransu , S. Tavernier , M. Tytgat ⁴,
S. Van Putte , D. Vannerom 

Université Libre de Bruxelles, Bruxelles, Belgium

B. Clerbaux , G. De Lentdecker , L. Favart , D. Hohov , J. Jaramillo , A. Khalilzadeh, K. Lee ,
M. Mahdavihorrani , A. Malara , S. Paredes , L. Pétré , N. Postiau, L. Thomas ,
M. Vanden Bemden , C. Vander Velde , P. Vanlaer 

Ghent University, Ghent, Belgium

M. De Coen , D. Dobur , J. Knolle , L. Lambrecht , G. Mestdach, C. Rendón, A. Samalan,
K. Skovpen , N. Van Den Bossche , B. Vermassen, L. Wezenbeek 

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

A. Benecke , G. Bruno , C. Caputo , C. Delaere , I.S. Donertas , A. Giammanco , K. Jaffel ,
Sa. Jain , V. Lemaitre, J. Lidrych , P. Mastrapasqua , K. Mondal , T.T. Tran , S. Wertz 

Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, Brazil

G.A. Alves , E. Coelho , C. Hensel , T. Menezes De Oliveira, A. Moraes , P. Rebello Teles ,
M. Soeiro

Universidade do Estado do Rio de Janeiro, Rio de Janeiro, Brazil

W.L. Aldá Júnior , M. Alves Gallo Pereira , M. Barroso Ferreira Filho , H. Brandao Malbouisson ,
W. Carvalho , J. Chinellato⁵, E.M. Da Costa , G.G. Da Silveira ⁶, D. De Jesus Damiao ,
S. Fonseca De Souza , J. Martins ⁷, C. Mora Herrera , K. Mota Amarilo , L. Mundim ,
H. Nogima , A. Santoro , S.M. Silva Do Amaral , A. Sznajder , M. Thiel , A. Vilela Pereira 

Universidade Estadual Paulista, Universidade Federal do ABC, São Paulo, Brazil

C.A. Bernardes ⁶, L. Calligaris , T.R. Fernandez Perez Tomei , E.M. Gregores , P.G. Mercadante ,
S.F. Novaes , B. Orzari , Sandra S. Padula 

Institute for Nuclear Research and Nuclear Energy, Bulgarian Academy of Sciences, Sofia, Bulgaria

A. Aleksandrov , G. Antchev , R. Hadjiiska , P. Iaydjiev , M. Misheva , M. Shopova ,
G. Sultanov 

University of Sofia, Sofia, Bulgaria

A. Dimitrov , T. Ivanov , L. Litov , B. Pavlov , P. Petkov , A. Petrov , E. Shumka 

Instituto De Alta Investigación, Universidad de Tarapacá, Casilla 7 D, Arica, Chile

S. Keshri , S. Thakur 

Beihang University, Beijing, China

T. Cheng , Q. Guo, T. Javaid , M. Mittal , L. Yuan 

Department of Physics, Tsinghua University, Beijing, China

G. Bauer⁸, Z. Hu , K. Yi ^{8,9}

Institute of High Energy Physics, Beijing, China

G.M. Chen ¹⁰, H.S. Chen ¹⁰, M. Chen ¹⁰, F. Iemmi , C.H. Jiang, A. Kapoor , H. Liao ,
Z.-A. Liu ¹¹, F. Monti , R. Sharma , J.N. Song¹¹, J. Tao , J. Wang , H. Zhang 

State Key Laboratory of Nuclear Physics and Technology, Peking University, Beijing, China

A. Agapitos , Y. Ban , A. Levin , C. Li , Q. Li , X. Lyu, Y. Mao, S.J. Qian , X. Sun , D. Wang ,
H. Yang, C. Zhou 

Sun Yat-Sen University, Guangzhou, China

Z. You 

University of Science and Technology of China, Hefei, China

N. Lu 

Institute of Modern Physics and Key Laboratory of Nuclear Physics and Ion-beam Application (MOE) - Fudan University, Shanghai, China

X. Gao ¹², D. Leggat, H. Okawa , Y. Zhang 

Zhejiang University, Hangzhou, Zhejiang, China

Z. Lin , C. Lu , M. Xiao 

Universidad de Los Andes, Bogota, Colombia

C. Avila , D.A. Barbosa Trujillo, A. Cabrera , C. Florez , J. Fraga , J.A. Reyes Vega

Universidad de Antioquia, Medellin, Colombia

J. Mejia Guisao , F. Ramirez , M. Rodriguez , J.D. Ruiz Alvarez 

University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, Split, Croatia

D. Giljanovic , N. Godinovic , D. Lelas , A. Sculac 

University of Split, Faculty of Science, Split, CroatiaM. Kovac , T. Sculac **Institute Rudjer Boskovic, Zagreb, Croatia**P. Bargassa , V. Brigljevic , B.K. Chitroda , D. Ferencek , S. Mishra , A. Starodumov ¹³, T. Susa **University of Cyprus, Nicosia, Cyprus**A. Attikis , K. Christoforou , S. Konstantinou , J. Mousa , C. Nicolaou, F. Ptochos , P.A. Razis , H. Rykaczewski, H. Saka , A. Stepennov **Charles University, Prague, Czech Republic**M. Finger , M. Finger Jr. , A. Kveton **Escuela Politecnica Nacional, Quito, Ecuador**E. Ayala **Universidad San Francisco de Quito, Quito, Ecuador**A. Cardenas , E. Carrera Jarrin , D. Cazar Ramírez , E.F. Mendez Garces, X. Riofrio **Academy of Scientific Research and Technology of the Arab Republic of Egypt, Egyptian Network of High Energy Physics, Cairo, Egypt**Y. Assran^{14,15}, S. Elgammal¹⁵**Center for High Energy Physics (CHEP-FU), Fayoum University, El-Fayoum, Egypt**M. Abdullah Al-Mashad , M.A. Mahmoud **National Institute of Chemical Physics and Biophysics, Tallinn, Estonia**R.K. Dewanjee ¹⁶, K. Ehataht , M. Kadastik, T. Lange , S. Nandan , C. Nielsen , J. Pata , M. Raidal , L. Tani , C. Veelken **Department of Physics, University of Helsinki, Helsinki, Finland**H. Kirschenmann , K. Osterberg , M. Voutilainen **Helsinki Institute of Physics, Helsinki, Finland**S. Bharthuar , E. Brücken , F. Garcia , J. Havukainen , K.T.S. Kallonen , M.S. Kim , R. Kinnunen, P.P. Koponen, T. Lampén , K. Lassila-Perini , S. Lehti , T. Lindén , M. Lotti, L. Martikainen , M. Myllymäki , M.m. Rantanen , H. Siikonen , E. Tuominen , J. Tuominiemi , R. Turpeinen**Lappeenranta-Lahti University of Technology, Lappeenranta, Finland**P. Luukka , H. Petrow , T. Tuuva[†]**IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France**M. Besancon , F. Couderc , M. Dejardin , D. Denegri, J.L. Faure, F. Ferri , S. Ganjour , P. Gras , G. Hamel de Monchenault , V. Lohezic , J. Malcles , J. Rander, A. Rosowsky , M.Ö. Sahin , A. Savoy-Navarro ¹⁷, P. Simkina , M. Titov 

Laboratoire Leprince-Ringuet, CNRS/IN2P3, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

C. Baldenegro Barrera , F. Beaudette , A. Buchot Perraguin , P. Busson , A. Cappati , C. Charlot , A. Chiron, F. Damas , O. Davignon , G. Falmagne , B.A. Fontana Santos Alves , S. Ghosh , A. Gilbert , R. Granier de Cassagnac , A. Hakimi , B. Harikrishnan , L. Kalipoliti , G. Liu , J. Motta , M. Nguyen , C. Ochando , L. Portales , T. Romanteau, R. Salerno , U. Sarkar , J.B. Sauvan , Y. Sirois , A. Tarabini , E. Vernazza , A. Zabi , A. Zghiche 

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

J.-L. Agram ¹⁸, J. Andrea , D. Apparu , D. Bloch , J.-M. Brom , E.C. Chabert , L. Charles, C. Collard , S. Falke , U. Goerlach , C. Grimault, L. Gross, R. Haeberle , A.-C. Le Bihan , M.A. Sessini , P. Van Hove 

Institut de Physique des 2 Infinis de Lyon (IP2I), Villeurbanne, France

G. Baulieu , S. Beauceron , B. Blancon , G. Boudoul , N. Chanon , J. Choi , D. Contardo , P. Depasse , C. Dozen ¹⁹, H. El Mamouni, J. Fay , S. Gascon , M. Gouzevitch , C. Greenberg, G. Grenier , B. Ille , I.B. Laktineh, M. Lethuillier , N. Lumb, L. Mirabito, S. Perries, D. Pugner , F. Schirra, M. Vander Donckt , P. Verdier , J. Xiao 

Georgian Technical University, Tbilisi, Georgia

G. Adamov, I. Bagaturia ²⁰, D. Chokheli , O. Kemularia, A. Khvedelidze ¹³, D. Lomidze , I. Lomidze , A. Melkadze, T. Toriashvili ²¹, Z. Tsamalaide ¹³

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

C. Autermann , V. Botta , L. Feld , W. Karpinski, M.K. Kiesel, K. Klein , M. Lipinski , D. Meuser , A. Pauls , G. Pierschel, M.P. Rauch, N. Röwert , C. Schomakers, J. Schulz, M. Teroerde , M. Wlochal

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

S. Diekmann , A. Dodonova , N. Eich , D. Eliseev , F. Engelke , M. Erdmann , P. Fackeldey , B. Fischer , T. Hebbeker , K. Hoepfner , F. Ivone , A. Jung , M.y. Lee , L. Mastrolorenzo, M. Merschmeyer , A. Meyer , S. Mukherjee , D. Noll , A. Novak , F. Nowotny, A. Pozdnyakov , Y. Rath, W. Redjeb , F. Rehm, H. Reithler , V. Sarkisovi , A. Schmidt , S.C. Schuler, A. Sharma , A. Stein , F. Torres Da Silva De Araujo ²², L. Vigilante, S. Wiedenbeck , S. Zaleski

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

C. Dziwok , G. Flügge , W. Haj Ahmad ²³, T. Kress , A. Nowack , O. Pooth , A. Stahl , T. Ziemons , A. Zott 

Deutsches Elektronen-Synchrotron, Hamburg, Germany

H. Aarup Petersen , M. Aldaya Martin , J. Alimena , S. Amoroso, Y. An , S. Baxter , M. Bayatmakou , H. Becerril Gonzalez , O. Behnke , A. Belvedere , S. Bhattacharya , F. Blekman ²⁴, K. Borras ²⁵, D. Brunner , A. Campbell , A. Cardini , C. Cheng, F. Colombina , S. Consuegra Rodríguez , G. Correia Silva , M. De Silva , G. Eckerlin, D. Eckstein , L.I. Estevez Banos , O. Filatov , E. Gallo ²⁴, A. Geiser , A. Giralaldi , G. Greau, V. Guglielmi , M. Guthoff , A. Hinzmann , A. Jafari ²⁶, L. Jeppe , N.Z. Jomhari , H. Jung , B. Kaeck 

M. Kasemann^{[ID](#)}, H. Kaveh^{[ID](#)}, C. Kleinwort^{[ID](#)}, R. Kogler^{[ID](#)}, M. Komm^{[ID](#)}, D. Krücker^{[ID](#)}, W. Lange,
D. Leyva Pernia^{[ID](#)}, K. Lipka^{[ID](#)}²⁷, W. Lohmann^{[ID](#)}²⁸, R. Mankel^{[ID](#)}, I.-A. Melzer-Pellmann^{[ID](#)},
M. Mendizabal Morentin^{[ID](#)}, J. Metwally, A.B. Meyer^{[ID](#)}, G. Milella^{[ID](#)}, M. Mormile^{[ID](#)}, A. Mussgiller^{[ID](#)},
A. Nürnberg^{[ID](#)}, Y. Otari, D. Pérez Adán^{[ID](#)}, E. Ranken^{[ID](#)}, A. Raspereza^{[ID](#)}, B. Ribeiro Lopes^{[ID](#)}, J. Rübenach,
A. Saggio^{[ID](#)}, M. Scham^{[ID](#)}^{29,25}, V. Scheurer, S. Schnake^{[ID](#)}²⁵, P. Schütze^{[ID](#)}, C. Schwanenberger^{[ID](#)}²⁴,
M. Shchedrolosiev^{[ID](#)}, R.E. Sosa Ricardo^{[ID](#)}, L.P. Sreelatha Pramod^{[ID](#)}, D. Stafford, F. Vazzoler^{[ID](#)},
A. Ventura Barroso^{[ID](#)}, R. Walsh^{[ID](#)}, Q. Wang^{[ID](#)}, Y. Wen^{[ID](#)}, K. Wichmann, L. Wiens^{[ID](#)}²⁵, C. Wissing^{[ID](#)},
S. Wuchterl^{[ID](#)}, Y. Yang^{[ID](#)}, A. Zimmermann Castro Santos^{[ID](#)}

University of Hamburg, Hamburg, Germany

A. Albrecht^{[ID](#)}, S. Albrecht^{[ID](#)}, M. Antonello^{[ID](#)}, S. Bein^{[ID](#)}, L. Benato^{[ID](#)}, M. Bonanomi^{[ID](#)}, P. Connor^{[ID](#)},
M. Eich, K. El Morabit^{[ID](#)}, Y. Fischer^{[ID](#)}, A. Fröhlich, C. Garbers^{[ID](#)}, E. Garutti^{[ID](#)}, A. Grohsjean^{[ID](#)},
M. Hajheidari, J. Haller^{[ID](#)}, H.R. Jabusch^{[ID](#)}, G. Kasieczka^{[ID](#)}, P. Keicher, R. Klanner^{[ID](#)}, W. Korcaric^{[ID](#)},
T. Kramer^{[ID](#)}, V. Kutzner^{[ID](#)}, F. Labe^{[ID](#)}, J. Lange^{[ID](#)}, A. Lobanov^{[ID](#)}, C. Matthies^{[ID](#)}, A. Mehta^{[ID](#)},
L. Moureaux^{[ID](#)}, M. Mrowietz, A. Nigamova^{[ID](#)}, Y. Nissan, A. Paasch^{[ID](#)}, K.J. Pena Rodriguez^{[ID](#)},
T. Quadfasel^{[ID](#)}, B. Raciti^{[ID](#)}, M. Rieger^{[ID](#)}, D. Savoie^{[ID](#)}, J. Schindler^{[ID](#)}, P. Schleper^{[ID](#)}, M. Schröder^{[ID](#)},
J. Schwandt^{[ID](#)}, M. Sommerhalder^{[ID](#)}, H. Stadie^{[ID](#)}, G. Steinbrück^{[ID](#)}, A. Tews, M. Wolf^{[ID](#)}

Karlsruher Institut fuer Technologie, Karlsruhe, Germany

S. Brommer^{[ID](#)}, M. Burkart, E. Butz^{[ID](#)}, T. Chwalek^{[ID](#)}, A. Dierlamm^{[ID](#)}, A. Droll, N. Faltermann^{[ID](#)},
M. Giffels^{[ID](#)}, A. Gottmann^{[ID](#)}, F. Hartmann^{[ID](#)}³⁰, M. Horzela^{[ID](#)}, U. Husemann^{[ID](#)}, M. Klute^{[ID](#)},
R. Koppenhöfer^{[ID](#)}, M. Link, A. Lintuluoto^{[ID](#)}, S. Maier^{[ID](#)}, S. Mitra^{[ID](#)}, Th. Müller^{[ID](#)}, M. Neukum, M. Oh^{[ID](#)},
G. Quast^{[ID](#)}, K. Rabbertz^{[ID](#)}, I. Shvetsov^{[ID](#)}, H.J. Simonis^{[ID](#)}, N. Trevisani^{[ID](#)}, R. Ulrich^{[ID](#)}, J. van der Linden^{[ID](#)},
R.F. Von Cube^{[ID](#)}, M. Wassmer^{[ID](#)}, S. Wieland^{[ID](#)}, F. Wittig, R. Wolf^{[ID](#)}, S. Wunsch, X. Zuo^{[ID](#)}

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

G. Anagnostou, P. Assiouras^{[ID](#)}, G. Daskalakis^{[ID](#)}, A. Kyriakis, A. Papadopoulos^{[ID](#)}³⁰, A. Stakia^{[ID](#)}

National and Kapodistrian University of Athens, Athens, Greece

D. Karasavvas, P. Kontaxakis^{[ID](#)}, G. Melachroinos, A. Panagiotou, I. Papavergou^{[ID](#)}, I. Paraskevas^{[ID](#)},
N. Saoulidou^{[ID](#)}, K. Theofilatos^{[ID](#)}, E. Tziaferi^{[ID](#)}, K. Vellidis^{[ID](#)}, I. Zisopoulos^{[ID](#)}

National Technical University of Athens, Athens, Greece

G. Bakas^{[ID](#)}, T. Chatzistavrou, G. Karapostoli^{[ID](#)}, K. Kousouris^{[ID](#)}, I. Papakrivopoulos^{[ID](#)}, E. Siamarkou,
G. Tsiopolitis, A. Zacharopoulou

University of Ioánnina, Ioánnina, Greece

K. Adamidis, I. Bestintzanos, I. Evangelou^{[ID](#)}, C. Foudas, P. Gianneios^{[ID](#)}, C. Kamtsikis, P. Katsoulis,
P. Kokkas^{[ID](#)}, P.G. Kosmoglou Kioseoglou^{[ID](#)}, N. Manthos^{[ID](#)}, I. Papadopoulos^{[ID](#)}, J. Strologas^{[ID](#)}

MTA-ELTE Lendület CMS Particle and Nuclear Physics Group, Eötvös Loránd University, Budapest, Hungary

M. Csanád^{[ID](#)}, K. Farkas^{[ID](#)}, M.M.A. Gadallah^{[ID](#)}³¹, Á. Kadlecik^{[ID](#)}, P. Major^{[ID](#)}, K. Mandal^{[ID](#)}, G. Pásztor^{[ID](#)},
A.J. Rádl^{[ID](#)}³², O. Surányi^{[ID](#)}, G.I. Veres^{[ID](#)}

Wigner Research Centre for Physics, Budapest, Hungary

M. Bartók ³³, C. Hajdu , D. Horvath ^{34,35}, F. Sikler , V. Veszpremi 

Institute of Nuclear Research ATOMKI, Debrecen, Hungary

G. Bencze, S. Czellar, J. Karancsi ³³, J. Molnar, Z. Szillasi

Institute of Physics, University of Debrecen, Debrecen, Hungary

P. Raics, B. Ujvari ³⁶, G. Zilizi 

Karoly Robert Campus, MATE Institute of Technology, Gyongyos, Hungary

T. Csorgo ³², F. Nemes ³², T. Novak 

Panjab University, Chandigarh, India

J. Babbar , S. Bansal , S.B. Beri, V. Bhatnagar , G. Chaudhary , S. Chauhan , N. Dhingra ³⁷,
R. Gupta, A. Kaur , A. Kaur , H. Kaur , M. Kaur , S. Kumar , P. Kumari , M. Meena ,
K. Sandeep , T. Sheokand, J.B. Singh ³⁸, A. Singla 

University of Delhi, Delhi, India

A. Ahmed , A. Bhardwaj , A. Chhetri , B.C. Choudhary , A. Kumar , M. Naimuddin ,
K. Ranjan , S. Saumya 

Saha Institute of Nuclear Physics, HBNI, Kolkata, India

S. Baradia , S. Barman ³⁹, S. Bhattacharya , D. Bhowmik, S. Dutta , S. Dutta, B. Gomber ⁴⁰,
P. Palit , G. Saha , B. Sahu ⁴⁰, S. Sarkar

Indian Institute of Technology Madras, Madras, India

P.K. Behera , S.C. Behera , S. Chatterjee , P. Jana , P. Kalbhor , J.R. Komaragiri ⁴¹, D. Kumar ⁴¹,
M. Mohammad Mobassir Ameen , L. Panwar ⁴¹, R. Pradhan , P.R. Pujahari , N.R. Saha ,
A. Sharma , A.K. Sikdar , S. Verma 

Tata Institute of Fundamental Research-A, Mumbai, India

T. Aziz, I. Das , S. Dugad, M. Kumar , G.B. Mohanty , P. Suryadevara

Tata Institute of Fundamental Research-B, Mumbai, India

A. Bala , S. Banerjee , R.M. Chatterjee⁴², M. Guchait , S. Karmakar , S. Kumar , G. Majumder ,
K. Mazumdar , S. Mukherjee , A. Thachayath 

National Institute of Science Education and Research, An OCC of Homi Bhabha National Institute, Bhubaneswar, Odisha, India

S. Bahinipati ⁴³, A.K. Das, C. Kar , D. Maity ⁴⁴, P. Mal , T. Mishra ,
V.K. Muralleedharan Nair Bindhu ⁴⁴, K. Naskar ⁴⁴, A. Nayak ⁴⁴, P. Sadangi, P. Saha , S.K. Swain ,
S. Varghese ⁴⁴, D. Vats ⁴⁴

Indian Institute of Science Education and Research (IISER), Pune, India

A. Alpana , S. Dube , B. Kansal , A. Laha , S. Pandey , A. Rastogi , S. Sharma 

Isfahan University of Technology, Isfahan, IranH. Bakhshiansohi ^{45,46}, E. Khazaie ⁴⁶, M. Zeinali ⁴⁷**Institute for Research in Fundamental Sciences (IPM), Tehran, Iran**S. Chenarani ⁴⁸, S.M. Etesami , M. Khakzad , M. Mohammadi Najafabadi **University College Dublin, Dublin, Ireland**M. Felcini , M. Grunewald **INFN Sezione di Bari^a, Università di Bari^b, Politecnico di Bari^c, Bari, Italy**M. Abbrescia ^{a,b}, R. Aly ^{a,c,49}, A. Colaleo ^a, D. Creanza ^{a,c}, B. D' Anzi ^{a,b}, N. De Filippis ^{a,c}, M. De Palma ^{a,b}, G. De Robertis ^a, A. Di Florio ^{a,c}, W. Elmetenawee ^{a,b}, L. Fiore ^a, M. Franco ^a, G. Iaselli ^{a,c}, F. Licciulli ^a, F. Loddo ^a, G. Maggi ^{a,c}, M. Maggi ^a, I. Margjeka ^{a,b}, V. Mastrapasqua ^{a,b}, S. My ^{a,b}, S. Nuzzo ^{a,b}, A. Pellicchia ^{a,b}, A. Pompili ^{a,b}, G. Pugliese ^{a,c}, R. Radogna ^a, G. Ramirez-Sanchez ^{a,c}, D. Ramos ^a, A. Ranieri ^a, L. Silvestris ^a, F.M. Simone ^{a,b}, Ü. Sözbilir ^a, A. Stamerra ^a, R. Venditti ^a, P. Verwilligen ^a, A. Zaza ^{a,b}**INFN Sezione di Bologna^a, Università di Bologna^b, Bologna, Italy**G. Abbiendi ^a, C. Baldanza ^a, C. Battilana ^{a,b}, D. Bonacorsi ^{a,b}, L. Borgonovi ^a, V.D. Cafaro ^a, P. Capiluppi ^{a,b}, A. Castro ^{a,b}, F.R. Cavallo ^a, A. Crupano ^a, M. Cuffiani ^{a,b}, G.M. Dallavalle ^a, T. Diotallevi ^{a,b}, F. Fabbri ^a, A. Fanfani ^{a,b}, D. Fasanella ^{a,b}, P. Giacomelli ^a, L. Giommi ^{a,b}, V. Giordano ^a, C. Grandi ^a, C. Guandalini ^a, L. Guiducci ^{a,b}, S. Lo Meo ^{a,50}, L. Lunerti ^a, S. Marcellini ^a, G. Masetti ^a, F.L. Navarria ^{a,b}, A. Perrotta ^a, F. Primavera ^{a,b}, A.M. Rossi ^{a,b}, T. Rovelli ^{a,b}, G.P. Siroli ^{a,b}**INFN Sezione di Catania^a, Università di Catania^b, Catania, Italy**S. Costa ^{a,b,51}, A. Di Mattia ^a, R. Potenza ^{a,b}, A. Tricomi ^{a,b,51}, C. Tuve ^{a,b}**INFN Sezione di Firenze^a, Università di Firenze^b, Firenze, Italy**G. Barbagli ^a, G. Bardelli ^{a,b}, B. Camaiani ^{a,b}, A. Cassese ^a, R. Ceccarelli ^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. Lizzo ^{a,b}, M. Meschini ^a, S. Paoletti ^a, A. Papanastassiou ^{a,b}, G. Sguazzoni ^a, L. Viliani ^a**INFN Laboratori Nazionali di Frascati, Frascati, Italy**L. Benussi ^a, S. Bianco ^a, M. Caponero ⁵², S. Meola ⁵³, D. Piccolo ^a, G. Saviano ⁵⁴**INFN Sezione di Genova^a, Università di Genova^b, Genova, Italy**S. Cerchi ^a, P. Chatagnon ^a, F. Ferro ^a, R. Puppo ^a, E. Robutti ^a, C. Rossi ^a, S. Tosi ^{a,b}, A. Trovato ^a**INFN Sezione di Milano-Bicocca^a, Università di Milano-Bicocca^b, Milano, Italy**A. Benaglia ^a, G. Boldrini ^a, F. Brivio ^{a,b}, F. Cetorelli ^{a,b}, F. De Guio ^{a,b}, M.E. Dinardo ^{a,b}, P. Dini ^a, S. Gennai ^a, A. Ghezzi ^{a,b}, P. Govoni ^{a,b}, L. Guzzi ^{a,b}, M.T. Lucchini ^{a,b}, M. Malberti ^a, S. Malvezzi ^a, A. Massironi ^a, D. Menasce ^a, L. Moroni ^a, M. Paganoni ^{a,b}, D. Pedrini ^a, B.S. Pinolini ^a, S. Ragazzi ^{a,b}, N. Redaelli ^a, T. Tabarelli de Fatis ^{a,b}, D. Zuolo ^{a,b}

INFN Sezione di Napoli^a, Università di Napoli ‘Federico II’^b, Napoli, Italy; Università della Basilicata^c, Potenza, Italy; Università G. Marconi^d, Roma, Italy

S. Buontempo ^a, A. Cagnotta ^{a,b}, F. Carnevali^{a,b}, N. Cavallo ^{a,c}, A. De Iorio ^{a,b}, F. Fabozzi ^{a,c}, A.O.M. Iorio ^{a,b}, L. Lista ^{a,b,55}, P. Paolucci ^{a,30}, B. Rossi ^a, C. Sciacca ^{a,b}

INFN Sezione di Padova^a, Università di Padova^b, Padova, Italy; Università di Trento^c, Trento, Italy

R. Ardino ^a, P. Azzi ^a, N. Bacchetta ^{a,56}, L. Barcellan^a, M. Bellato ^a, M. Benettoni ^a, A. Bergnoli ^a, L. Berti^a, M. Biasotto ^{a,57}, D. Bisello ^{a,b}, P. Bortignon ^a, A. Bragagnolo ^{a,b}, R. Carlin ^{a,b}, L. Castellani^a, P. Checchia ^a, L. Ciano^a, A. Colombo^a, D. Corti^a, A. Crescente^a, T. Dorigo ^a, S. Fantinel ^a, F. Fanzago ^a, F. Gasparini ^{a,b}, U. Gasparini ^{a,b}, F. Gonella ^a, A. Gozzelino ^a, A. Griggio^a, G. Grosso^a, M. Gulmini ^{a,57}, R. Isocrate^a, L. Layer^{a,58}, E. Lusiani ^a, M. Margoni ^{a,b}, G. Maron ^{a,57}, A.T. Meneguzzo ^{a,b}, M. Michelotto ^a, M. Migliorini ^{a,b}, L. Modenese^a, F. Montecassiano ^a, M. Negrello^a, M. Passaseo ^a, J. Pazzini ^{a,b}, L. Ramina^a, M. Rampazzo^a, M. Rebeschini^a, P. Ronchese ^{a,b}, R. Rossin ^{a,b}, M. Sgaravatto ^a, F. Simonetto ^{a,b}, G. Strong ^a, R. Temporin^a, M. Tessaro^a, M. Toffano^a, N. Toniolo ^a, M. Tosi ^{a,b}, A. Triossi ^{a,b}, S. Ventura ^a, H. Yarar^{a,b}, M. Zanetti ^{a,b}, P.G. Zatti^a, P. Zotto ^{a,b}, A. Zucchetta ^{a,b}, G. Zumerle ^{a,b}

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

S. Abu Zeid ^{a,59}, C. Aimè ^{a,b}, A. Braghieri ^a, S. Calzaferri ^{a,b}, D. Fiorina ^{a,b}, P. Montagna ^{a,b}, V. Re ^a, C. Riccardi ^{a,b}, P. Salvini ^a, I. Vai ^{a,b}, P. Vitulo ^{a,b}

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

S. Ajmal ^{a,b}, P. Asenov ^{a,60}, G. Baldinelli ^a, F. Bianchi ^a, G.M. Bilei ^a, S. Bizzaglia^a, B. Checcucci ^a, D. Ciangottini ^{a,b}, L. Fanò ^{a,b}, L. Farnesini^a, M. Ionica ^a, M. Magherini ^{a,b}, G. Mantovani^{a,b}, V. Mariani ^{a,b}, M. Menichelli ^a, A. Morozzi ^{a,b}, F. Moscatelli ^{a,60}, D. Passeri ^{a,b}, A. Piccinelli ^{a,b}, P. Placidi ^{a,b}, M. Presilla ^{a,b}, A. Rossi ^{a,b}, A. Santocchia ^{a,b}, D. Spiga ^a, T. Tedeschi ^{a,b}, C. Turrioni ^a

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

P. Azzurri ^a, G. Bagliesi ^a, A. Basti ^{a,b}, R. Bhattacharya ^a, L. Bianchini ^{a,b}, M. Bitossi ^a, T. Boccali ^a, L. Borrello ^a, F. Bosi^a, E. Bossini ^a, D. Bruschini ^{a,c}, R. Castaldi ^a, M.A. Ciocci ^{a,b}, M. Cipriani ^{a,b}, V. D’Amante ^{a,d}, R. Dell’Orso ^a, S. Donato ^a, F. Fiori ^a, A. Giassi ^a, F. Ligabue ^{a,c}, G. Magazzu ^a, M. Massa ^a, D. Matos Figueiredo ^a, A. Messineo ^{a,b}, A. Moggi ^a, M. Musich ^{a,b}, F. Palla ^a, F. Palmonari^a, S. Parolia ^a, F. Raffaelli ^a, A. Rizzi ^{a,b}, G. Rolandi ^{a,c}, S. Roy Chowdhury ^a, T. Sarkar ^a, A. Scribano ^a, P. Spagnolo ^a, R. Tenchini ^{a,b}, G. Tonelli ^{a,b}, N. Turini ^{a,d}, A. Venturi ^a, P.G. Verdini ^a

INFN Sezione di Roma^a, Sapienza Università di Roma^b, Roma, Italy

P. Barria ^a, C. Basile ^{a,b}, R. BIANCO^a, M. Campana ^{a,b}, F. Cavallari ^a, L. Cunqueiro Mendez ^{a,b}, I. Dafinei ^a, D. Del Re ^{a,b}, E. Di Marco ^a, M. Diemoz ^a, F. Errico ^{a,b}, F. Failla^a, E. Longo ^{a,b}, P. Meridiani ^a, J. Mijuskovic ^{a,b}, C.A. Nicolau^a, M. Nuccetelli^a, G. Organtini ^{a,b}, F. Pandolfi ^a, R. Paramatti ^{a,b}, V. Pettinacci ^a, C. Quaranta ^{a,b}, S. Rahatlou ^{a,b}, C. Rovelli ^a, F. Santanastasio ^{a,b}, L. Soffi ^a, R. Tramontano ^{a,b}

INFN Sezione di Torino^a, Università di Torino^b, Torino, Italy; Università del Piemonte Orientale^c, Novara, Italy

N. Amapane , R. Arcidiacono , S. Argiro , M. Arneodo , N. Bartosik , R. Bellan , A. Bellora , C. Biino , N. Cartiglia , S. Coli , M. Costa , R. Covarelli , D. Dattola , N. Demaria , L. Finco , M. Grippo , B. Kiani , F. Legger , F. Luongo , C. Mariotti , S. Maselli , A. Mecca , E. Migliore , M. Monteno , R. Mulargia , M.M. Obertino , G. Ortona , L. Pacher , N. Pastrone , M. Pelliccioni , M. Ruspa , F. Siviero , V. Sola , A. Solano , C. Tarricone , M. Tornago , D. Trocino , G. Umoret , A. Vagnerini , E. Vlasov

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

S. Belforte , V. Candelise , M. Casarsa , F. Cossutti , K. De Leo , G. Della Ricca 

Kyungpook National University, Daegu, Korea

S. Dogra , C. Huh , B. Kim , D.H. Kim , J. Kim , J. Lee , S.W. Lee , C.S. Moon , Y.D. Oh , S.I. Pak , M.S. Ryu , S. Sekmen , Y.C. Yang 

Chonnam National University, Institute for Universe and Elementary Particles, Kwangju, Korea

G. Bak , P. Gwak , H. Kim , D.H. Moon 

Hanyang University, Seoul, Korea

E. Asilar , D. Kim , T.J. Kim , J.A. Merlin , J. Park 

Korea University, Seoul, Korea

S. Choi , S. Han , B. Hong , K. Lee , K.S. Lee , J. Park , S.K. Park , J. Yoo 

Kyung Hee University, Department of Physics, Seoul, Korea

J. Goh 

Sejong University, Seoul, Korea

H. S. Kim , Y. Kim , S. Lee 

Seoul National University, Seoul, Korea

J. Almond , J.H. Bhyun , J. Choi , S. Jeon , W. Jun , J. Kim , J.S. Kim , S. Ko , H. Kwon , H. Lee , J. Lee , S. Lee , B.H. Oh , S.B. Oh , H. Seo , U.K. Yang , I. Yoon

University of Seoul, Seoul, Korea

W. Jang , D.Y. Kang , Y. Kang , S. Kim , B. Ko , J.S.H. Lee , Y. Lee , I.C. Park , Y. Roh , I.J. Watson , S. Yang 

Yonsei University, Department of Physics, Seoul, Korea

S. Ha , H.D. Yoo 

Sungkyunkwan University, Suwon, Korea

M. Choi , M.R. Kim , H. Lee , Y. Lee , I. Yu 

**College of Engineering and Technology, American University of the Middle East (AUM),
Dasman, Kuwait**

T. Beyrouthy, Y. Maghrbi 

Riga Technical University, Riga, Latvia

K. Dreimanis , A. Gaile , G. Pikurs, A. Potrebko , M. Seidel , V. Veckalns ⁶¹

University of Latvia (LU), Riga, Latvia

N.R. Strautnieks 

Vilnius University, Vilnius, Lithuania

M. Ambrozas , A. Juodagalvis , A. Rinkevicius , G. Tamulaitis 

National Centre for Particle Physics, Universiti Malaya, Kuala Lumpur, Malaysia

N. Bin Norjoharuddeen , I. Yusuff ⁶², Z. Zolkapli

Universidad de Sonora (UNISON), Hermosillo, Mexico

J.F. Benitez , A. Castaneda Hernandez , H.A. Encinas Acosta, L.G. Gallegos Maríñez, M. León Coello ,
J.A. Murillo Quijada , A. Sehrawat , L. Valencia Palomo 

Centro de Investigacion y de Estudios Avanzados del IPN, Mexico City, Mexico

G. Ayala , H. Castilla-Valdez , E. De La Cruz-Burelo , I. Heredia-De La Cruz ⁶³,
R. Lopez-Fernandez , C.A. Mondragon Herrera, D.A. Perez Navarro , A. Sánchez Hernández 

Universidad Iberoamericana, Mexico City, Mexico

C. Oropeza Barrera , M. Ramírez García 

Benemerita Universidad Autonoma de Puebla, Puebla, Mexico

I. Bautista , I. Pedraza , H.A. Salazar Ibarguen , C. Uribe Estrada 

University of Montenegro, Podgorica, Montenegro

I. Bubanja, N. Raicevic 

University of Canterbury, Christchurch, New Zealand

P.H. Butler 

National Centre for Physics, Quaid-I-Azam University, Islamabad, Pakistan

A. Ahmad , M.I. Asghar, A. Awais , M.I.M. Awan, H.R. Hoorani , W.A. Khan 

**AGH University of Krakow, Faculty of Computer Science, Electronics and
Telecommunications, Krakow, Poland**

V. Avati, L. Grzanka , M. Malawski 

National Centre for Nuclear Research, Swierk, Poland

H. Bialkowska , M. Bluj , B. Boimska , M. Górski , M. Kazana , M. Szleper , P. Zalewski 

Institute of Experimental Physics, Faculty of Physics, University of Warsaw, Warsaw, Poland

K. Bunkowski , K. Doroba , A. Kalinowski , K. Kierzkowski , M. Konecki , J. Krolikowski ,
W. Oklinski

Warsaw University of Technology, Warsaw, Poland

K. Pozniak , W. Zabolotny 

Laboratório de Instrumentação e Física Experimental de Partículas, Lisboa, Portugal

M. Araujo , D. Bastos , C. Beirão Da Cruz E Silva , A. Boletti , M. Bozzo , G. Da Molin ,
P. Faccioli , M. Gallinaro , J. Hollar , N. Leonardo , T. Niknejad , M. Pisano , J. Rasteiro Da Silva,
J. Seixas , J. Varela , J.W. Wulff

Faculty of Physics, University of Belgrade, Belgrade, Serbia

P. Adzic , P. Milenovic 

VINCA Institute of Nuclear Sciences, University of Belgrade, Belgrade, Serbia

M. Dordevic , J. Milosevic , V. Rekovic

Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

M. Aguilar-Benitez, J. Alcaraz Maestre , M. Barrio Luna, Cristina F. Bedoya , M. Cepeda ,
M. Cerrada , N. Colino , J. Cuchillo Ortega, B. De La Cruz , C.I. De Lara Rodríguez,
A. Delgado Peris , D. Fernández Del Val , J.P. Fernández Ramos , J. Flix , M.C. Fouz ,
O. Gonzalez Lopez , S. Goy Lopez , J.M. Hernandez , M.I. Josa , J. León Holgado ,
O. Manzanilla , D. Moran , C. M. Morcillo Perez , Á. Navarro Tobar , R. Paz Herrera,
C. Perez Dengra , A. Pérez-Calero Yzquierdo , J. Puerta Pelayo , I. Redondo ,
D.D. Redondo Ferrero , L. Romero, S. Sánchez Navas , J. Sastre , L. Urda Gómez ,
J. Vazquez Escobar , C. Willmott

Universidad Autónoma de Madrid, Madrid, Spain

J.F. de Trocóniz 

Universidad de Oviedo, Instituto Universitario de Ciencias y Tecnologías Espaciales de Asturias (ICTEA), Oviedo, Spain

B. Alvarez Gonzalez , J. Cuevas , J. Fernandez Menendez , S. Folgueras , I. Gonzalez Caballero ,
J.R. González Fernández , E. Palencia Cortezon , C. Ramón Álvarez , V. Rodríguez Bouza ,
A. Soto Rodríguez , A. Trapote , C. Vico Villalba , P. Vischia 

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

S. Bhowmik , S. Blanco Fernández , J.A. Brochero Cifuentes , I.J. Cabrillo , A. Calderon ,
J. Duarte Campderros , M. Fernandez , C. Fernandez Madrazo , G. Gomez , C. Lasasa García ,
C. Martinez Rivero , P. Martinez Ruiz del Arbol , F. Matorras , P. Matorras Cuevas ,
E. Navarrete Ramos , J. Piedra Gomez , C. Prieels, L. Scodellaro , I. Vila , J.M. Vizan Garcia 

University of Colombo, Colombo, Sri Lanka

M.K. Jayananda , B. Kailasapathy ⁶⁴, D.U.J. Sonnadara , D.D.C. Wickramarathna 

University of Ruhuna, Department of Physics, Matara, Sri Lanka

W.G.D. Dharmaratna ⁶⁵, K. Liyanage , N. Perera , N. Wickramage 

CERN, European Organization for Nuclear Research, Geneva, Switzerland

D. Abbaneo , C. Amendola , V. Amoiridis , E. Auffray , G. Auzinger , J. Baechler , D. Barney ,
A. Bermúdez Martínez , M. Bianco , B. Bilin , A.A. Bin Anuar , A. Bocci , E. Brondolin ,
P. Brummer , C. Caillol , T. Camporesi , G. Cerminara , N. Chernyavskaya , B. Curé ,
D. d’Enterria , A. Dabrowski , A. David , A. De Roeck , M.M. Defranchis , M. Deile ,
C. Deldicque , M. Dobson , A. Dvorak , F. Fallavollita⁶⁶, L. Forthomme , G. Franzoni , W. Funk ,
S. Giani , D. Gigi , K. Gill , F. Glege , L. Gouskos , N. Gutic , M. Haranko , J. Hegeman ,
V. Innocente , G. Izquierdo Moreno , T. James , P. Janot , J. Kieseler , C. Kishimoto Bisbe ,
N. Kratochwil , S. Laurila , P. Lecoq , E. Leutgeb , C. Lourenço , B. Maier , L. Malgeri ,
M. Mannelli , A.C. Marini , F. Meijers , S. Mersi , E. Meschi , V. Milosevic , R.K. Mommsen ,
F. Moortgat , M. Mulders , S. Orfanelli , L. Orsini , F. Pantaleo , K. Peron , M. Peruzzi , A. Petrilli ,
G. Petrucciani , A. Pfeiffer , M. Pierini , D. Piparo , A. Poupakis , H. Qu , D. Rabady , A. Racz ,
G. Reales Gutiérrez , M. Rovere , H. Sakulin , S. Scarfi , C. Schwick , M. Selvaggi , A. Sharma ,
K. Shchelina , P. Silva , P. Sphicas⁶⁷, A.G. Stahl Leiton , A. Steen , S. Summers , D. Treille ,
P. Tropea , A. Tsirou , C. Vázquez , D. Walter , J. Wanczyk⁶⁸, K.A. Wozniak , P. Zehetner ,
P. Zejdl , W.D. Zeuner 

Paul Scherrer Institut, Villigen, Switzerland

T. Bevilacqua ⁶⁹, L. Caminada ⁶⁹, A. Ebrahimi , W. Erdmann , R. Horisberger , Q. Ingram ,
H.C. Kaestli , D. Kotlinski , C. Lange , B. Meier , M. Missiroli⁶⁹, L. Noehte⁶⁹, T. Rohe , S. Streuli 

ETH Zurich - Institute for Particle Physics and Astrophysics (IPA), Zurich, Switzerland

T.K. Aarrestad ⁶⁸, K. Androsov ⁶⁸, M. Backhaus , G. Bonomelli , S. Burkhalter ⁷⁰, A. Calandri ,
C. Cazzaniga , K. Datta , A. De Cosa , G. Dissertori , M. Dittmar , M. Donegà , F. Eble , T. Gadek ,
M. Galli , K. Gedia , F. Glessgen , C. Grab , N. Härringer , T.G. Harte , D. Hits , W. Lustermann ,
A.-M. Lyon , R.A. Manzoni , M. Marchegiani , L. Marchese , C. Martin Perez , A. Mascellani⁶⁸,
F. Nessi-Tedaldi , F. Pauss , V. Perovic , S. Pigazzini , M.G. Ratti , C. Reissel , T. Reitenspiess ,
B. Ristic , F. Riti , D. Ruini , R. Seidita , K. Stachon , J. Steggemann⁶⁸, D. Valsecchi , R. Wallny 

Universität Zürich, Zurich, Switzerland

C. Amsler ⁷¹, P. Bärtschi , K. Boesiger , C. Botta , D. Brzzechko , M.F. Canelli , K. Cormier ,
A. De Wit , R. Del Burgo , M. Gienal , J.K. Heikkilä , D. Hernandez Garland , M. Huwiler , W. Jin ,
A. Jofrehei , B. Kilminster , S. Leontsinis , S.P. Liechti , A. Macchiolo , R. Maier , P. Meiring ,
V.M. Mikuni , U. Molinatti , B. Neuenschwander , I. Neutelings , G. Rauco , A. Reimers , P. Robmann ,
D. Salerno , S. Sanchez Cruz , K. Schweiger , M. Senger , Y. Takahashi , S.A. Wiederkehr , D. Wolf 

National Central University, Chung-Li, Taiwan

C. Adloff⁷², C.M. Kuo , W. Lin , P.K. Rout , P.C. Tiwari⁴¹, S.S. Yu 

National Taiwan University (NTU), Taipei, Taiwan

L. Ceard , Y. Chao , K.F. Chen , P.s. Chen, Z.g. Chen, W.-S. Hou , T.h. Hsu, Y.w. Kao, R. Khurana, G. Kole , Y.y. Li , R.-S. Lu , E. Paganis , A. Psallidas, X.f. Su, J. Thomas-Wilsker , H.y. Wu, E. Yazgan 

Chulalongkorn University, Faculty of Science, Department of Physics, Bangkok, Thailand

C. Asawatangtrakuldee , N. Srimanobhas , V. Wachirapusanand 

Çukurova University, Physics Department, Science and Art Faculty, Adana, Turkey

D. Agyel , F. Boran , Z.S. Demiroglu , F. Dolek , I. Dumanoglu ⁷³, E. Eskut , Y. Guler ⁷⁴, E. Gurpınar Guler ⁷⁴, C. Isik , O. Kara, A. Kayis Topaksu , U. Kiminsu , G. Onengut , K. Ozdemir ⁷⁵, A. Polatoz , B. Tali ⁷⁶, U.G. Tok , S. Turkcapar , E. Uslan , I.S. Zorbakir 

Middle East Technical University, Physics Department, Ankara, Turkey

K. Ocalan ⁷⁷, M. Yalvac ⁷⁸

Bogazici University, Istanbul, Turkey

B. Akgun , I.O. Atakisi , E. Gülmez , M. Kaya ⁷⁹, O. Kaya ⁸⁰, S. Tekten ⁸¹, E.A. Yetkin ⁸²

Istanbul Technical University, Istanbul, Turkey

A. Cakir , K. Cankocak ⁷³, Y. Komurcu , S. Sen ⁸³

Istanbul University, Istanbul, Turkey

O. Aydılek , S. Cerci ⁷⁶, V. Epshteyn , B. Hacisahinoglu , I. Hos ⁸⁴, B. Isildak ⁸⁵, B. Kaynak , S. Ozkorucuklu , H. Sert , C. Simsek , D. Sunar Cerci ⁷⁶, C. Zorbilmez 

Yildiz Technical University, Istanbul, Turkey

T. Yetkin 

Institute for Scintillation Materials of National Academy of Science of Ukraine, Kharkiv, Ukraine

A. Boyaryntsev , B. Grynyov 

National Science Centre, Kharkiv Institute of Physics and Technology, Kharkiv, Ukraine

L. Levchuk 

University of Bristol, Bristol, United Kingdom

D. Anthony , J.J. Brooke , A. Bundock , F. Bury , E. Clement , D. Cussans , H. Flacher , M. Glowacki, J. Goldstein , H.F. Heath , L. Kreczko , B. Krikler , S. Paramesvaran , S. Seif El Nasr-Storey, V.J. Smith , N. Stylianou ⁸⁶, K. Walkingshaw Pass, R. White 

Rutherford Appleton Laboratory, Didcot, United Kingdom

A.H. Ball, K.W. Bell , A. Belyaev ⁸⁷, C. Brew , R.M. Brown , D.J.A. Cockerill , C. Cooke , K.V. Ellis, K. Harder , S. Harper , M.-L. Holmberg ⁸⁸, Sh. Jain ⁸⁹, J. Linacre , K. Manolopoulos, D.M. Newbold , E. Olaiya, D. Petyt , T. Reis , G. Salvi , T. Schuh, C.H. Shepherd-Themistocleous , I.R. Tomalin , T. Williams 

Imperial College, London, United Kingdom

R. Bainbridge , P. Bloch , C.E. Brown , O. Buchmuller, V. Cacchio, C.A. Carrillo Montoya , G.S. Chahal ⁹⁰, D. Colling , J.S. Dancu, P. Dauncey , G. Davies , J. Davies, M. Della Negra , S. Fayer, G. Fedi , G. Hall , M.H. Hassanshahi , A. Howard, G. Iles , M. Knight , J. Langford , L. Lyons , A.-M. Magnan , S. Malik, A. Martelli , M. Mieskolainen , J. Nash ⁹¹, M. Pesaresi, B.C. Radburn-Smith , A. Richards, A. Rose , C. Seez , R. Shukla , A. Tapper , K. Uchida , G.P. Uttley , L.H. Vage, T. Virdee ³⁰, M. Vojinovic , N. Wardle , D. Winterbottom 

Brunel University, Uxbridge, United Kingdom

K. Coldham, J.E. Cole , A. Khan, P. Kyberd , I.D. Reid 

Baylor University, Waco, Texas, U.S.A.

S. Abdullin , A. Brinkerhoff , B. Caraway , J. Dittmann , K. Hatakeyama , J. Hiltbrand , A.R. Kanuganti , B. McMaster , M. Saunders , S. Sawant , C. Sutantawibul , M. Toms , J. Wilson 

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

R. Bartek , A. Dominguez , C. Huerta Escamilla, A.E. Simsek , R. Uniyal , A.M. Vargas Hernandez 

The University of Alabama, Tuscaloosa, Alabama, U.S.A.

R. Chudasama , S.I. Cooper , S.V. Gleyzer , C.U. Perez ⁹², E. Usai , C. West 

Boston University, Boston, Massachusetts, U.S.A.

A. Akpinar , A. Albert , D. Arcaro , C. Cosby , Z. Demiragli , C. Erice , E. Fontanesi , D. Gastler , J. Rohlf , K. Salyer , D. Sperka , D. Spitzbart , I. Suarez , A. Tsatsos , S. Yuan 

Brown University, Providence, Rhode Island, U.S.A.

G. Benelli , X. Coubez²⁵, D. Cutts , M. Hadley , U. Heintz , J.M. Hogan ⁹³, T. Kwon , G. Landsberg , K.T. Lau , D. Li , J. Luo , S. Mondal , M. Narain [†], N. Pervan , S. Sagir ⁹⁴, F. Simpson , W.Y. Wong, X. Yan , D. Yu , W. Zhang

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

S. Abbott , J. Bonilla , C. Brainerd , R. Breedon , M. Calderon De La Barca Sanchez , M. Chertok , M. Citron , J. Conway , P.T. Cox , R. Erbacher , G. Haza , F. Jensen , O. Kukral , G. Mocellin , M. Mulhearn , D. Pellett , B. Regnery , W. Wei , Y. Yao , F. Zhang 

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

M. Bachtis , R. Cousins , A. Datta , J. Hauser , M. Ignatenko , M.A. Iqbal , T. Lam , E. Manca , W.A. Nash , D. Saltzberg , B. Stone , V. Valuev 

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

R. Clare , M. Gordon, G. Hanson , W. Si , S. Wimpenny [†]

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

J.G. Branson , S. Cittolin , S. Cooperstein , D. Diaz , J. Duarte , R. Gerosa , L. Giannini , J. Guiang , R. Kansal , V. Krutelyov , R. Lee , J. Letts , M. Masciovecchio , F. Mokhtar 

S. Morovic , A. Petrucci , M. Pieri , M. Quinnan , B.V. Sathia Narayanan , V. Sharma ,
M. Tadel , E. Vourliotis , F. Würthwein , Y. Xiang , A. Yagil 

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

L. Brennan , C. Campagnari , G. Collura , A. Dorsett , J. Incandela , M. Kilpatrick , J. Kim ,
A.J. Li , P. Masterson , H. Mei , M. Oshiro , J. Richman , U. Sarica , R. Schmitz , F. Setti ,
J. Sheplock , D. Stuart , S. Wang 

California Institute of Technology, Pasadena, California, U.S.A.

J. Balcas, A. Bornheim , O. Cerri, A. Latorre, J.M. Lawhorn , J. Mao , H.B. Newman ,
T. Q. Nguyen , M. Spiropulu , J.R. Vlimant , C. Wang , S. Xie , R.Y. Zhu 

Carnegie Mellon University, Pittsburgh, Pennsylvania, U.S.A.

J. Alison , S. An , M.B. Andrews , P. Bryant , V. Dutta , T. Ferguson , A. Harilal , C. Liu ,
T. Mudholkar , S. Murthy , M. Paulini , A. Roberts , A. Sanchez , W. Terrill 

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

J.P. Cumalat , W.T. Ford , A. Hassani , G. Karathanasis , E. MacDonald, N. Manganeli ,
F. Marini , A. Perloff , C. Savard , N. Schonbeck , K. Stenson , K.A. Ulmer , S.R. Wagner ,
N. Zipper 

Cornell University, Ithaca, New York, U.S.A.

J. Alexander , S. Bright-Thonney , X. Chen , D.J. Cranshaw , J. Fan , X. Fan , D. Gadkari ,
S. Hogan , J. Monroy , J.R. Patterson , J. Reichert , M. Reid , A. Ryd , J. Thom , P. Wittich ,
R. Zou 

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

J. Adelman-McCarthy , M. Albrow , M. Alyari , O. Amram , G. Apollinari , A. Apresyan ,
L.A.T. Bauerdick , D. Berry , J. Berryhill , P.C. Bhat , K. Burkett , J.N. Butler , A. Canepa ,
G.B. Cerati , H.W.K. Cheung , F. Chlebana , G. Cummings , W. Dagenhart , J. Dickinson ,
I. Dutta , V.D. Elvira , Y. Feng , J. Freeman , A. Gandrakota , P. Gartung , Z. Gecse , L. Gray ,
D. Green, S. Grünendahl , D. Guerrero , O. Gutsche , R.M. Harris , R. Heller , T.C. Herwig ,
J. Hirschauer , L. Horyn , D. Hufnagel , B. Jayatilaka , S. Jindariani , M. Johnson , C.D. Jones,
U. Joshi , T. Klijnsma , B. Klima , M.J. Kortelainen , K.H.M. Kwok , S. Lammel , D. Lincoln ,
R. Lipton , T. Liu , C. Madrid , K. Maeshima , C. Mantilla , D. Mason , P. McBride ,
P. Merkel , S. Mrenna , S. Nahn , J. Ngadiuba , D. Noonan , V. Papadimitriou , N. Pastika ,
K. Pedro , C. Pena ⁹⁵, F. Ravera , A. Reinsvold Hall ⁹⁶, L. Ristori , E. Sexton-Kennedy ,
N. Smith , A. Soha , L. Spiegel , S. Stoynev , L. Taylor , S. Tkaczyk , N.V. Tran , L. Uplegger ,
E.W. Vaandering , I. Zoi 

University of Florida, Gainesville, Florida, U.S.A.

C. Aruta , P. Avery , D. Bourilkov , L. Cadamuro , P. Chang , V. Cherepanov , R.D. Field,
E. Koenig , M. Kolosova , J. Konigsberg , A. Korytov , K.H. Lo, K. Matchev , N. Menendez 

G. Mitselmakher , A. Muthirakalayil Madhu , N. Rawal , D. Rosenzweig , S. Rosenzweig , K. Shi , J. Wang 

Florida State University, Tallahassee, Florida, U.S.A.

T. Adams , A. Al Kadhimi , A. Askew , N. Bower , R. Habibullah , V. Hagopian , R. Hashmi , R.S. Kim , S. Kim , T. Kolberg , G. Martinez , H. Prosper , P.R. Prova , O. Viazlo , M. Wulansatiti , R. Yohay , J. Zhang

Florida Institute of Technology, Melbourne, Florida, U.S.A.

B. Alsufyani , M.M. Baarmand , S. Butalla , T. Elkafrawy ⁵⁹, M. Hohlmann , R. Kumar Verma , M. Rahmani , F. Yumiceva 

University of Illinois Chicago, Chicago, U.S.A., Chicago, U.S.A.

M.R. Adams , C. Bennett , R. Cavanaugh , S. Dittmer , R. Escobar Franco , A. Evdokimov , O. Evdokimov , C.E. Gerber , D.J. Hofman , J.h. Lee , D. S. Lemos , A.H. Merrit , C. Mills , S. Nanda , G. Oh , B. Ozek , D. Pilipovic , T. Roy , S. Rudrabhatla , M.B. Tonjes , N. Varelas , X. Wang , Z. Ye , J. Yoo 

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

M. Alhusseini , B. Bilki ⁹⁷, D. Blend , P. Debbins , K. Dilsiz ⁹⁸, L. Emediato , G. Karaman , O.K. Köseyan , J.-P. Merlo , A. Mestvirishvili ⁹⁹, M.J. Miller , J. Nachtman , O. Neogi , H. Ogul ¹⁰⁰, Y. Onel , A. Penzo , I. Schmidt , C. Snyder , D. Southwick , E. Tiras ¹⁰¹, J. Wetzel 

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

B. Blumenfeld , L. Corcodilos , J. Davis , A.V. Gritsan , L. Kang , S. Kyriacou , P. Maksimovic , M. Roguljic , J. Roskes , S. Sekhar , M. Swartz , T.Á. Vámi 

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

A. Abreu , L.F. Alcerro Alcerro , J. Anguiano , P. Baringer , A. Bean , Z. Flowers , D. Grove , J. King , G. Krintiras , M. Lazarovits , C. Le Mahieu , C. Lindsey , J. Marquez , N. Minafra , M. Murray , M. Nickel , M. Pitt , S. Popescu ¹⁰², C. Rogan , C. Royon , R. Salvatico , S. Sanders , C. Smith , Q. Wang , G. Wilson 

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

B. Allmond , A. Ivanov , K. Kaadze , A. Kalogeropoulos , D. Kim , Y. Maravin , K. Nam , J. Natoli , D. Roy , G. Sorrentino 

Lawrence Livermore National Laboratory, Livermore, California, U.S.A.

F. Rebassoo , D. Wright 

University of Maryland, College Park, Maryland, U.S.A.

E. Adams , A. Baden , O. Baron , A. Belloni , A. Bethani , Y.m. Chen , S.C. Eno , N.J. Hadley , S. Jabeen , R.G. Kellogg , T. Koeth , Y. Lai , S. Lascio , A.C. Mignerey , S. Nabili , C. Palmer , C. Papageorgakis , M.M. Paranjpe , L. Wang , K. Wong 

Massachusetts Institute of Technology, Cambridge, Massachusetts, U.S.A.

J. Bendavid , W. Busza , I.A. Cali , Y. Chen , M. D'Alfonso , G.L. Darlea , J. Eysermans ,
C. Freer , G. Gomez-Ceballos , M. Goncharov, P. Harris, D. Hoang, D. Kovalskyi , J. Krupa ,
L. Lavezzo , Y.-J. Lee , K. Long , C. Mironov , C. Paus , D. Rankin , C. Roland , G. Roland ,
S. Rothman , Z. Shi , G.S.F. Stephans , J. Wang, Z. Wang , B. Wyslouch , T. J. Yang 

University of Minnesota, Minneapolis, Minnesota, U.S.A.

B. Crossman , E. Frahm, B.M. Joshi , C. Kapsiak , M. Krohn , D. Mahon , J. Mans ,
M. Revering , R. Rusack , R. Saradhy , N. Schroeder , N. Strobbe , M.A. Wadud 

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

L.M. Cremaldi 

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

K. Bloom , M. Bryson, D.R. Claes , C. Fangmeier , F. Golf , C. Joo , I. Kravchenko , I. Reed ,
J.E. Siado , G.R. Snow[†], W. Tabb , A. Wightman , F. Yan , A.G. Zecchinelli 

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

G. Agarwal , H. Bandyopadhyay , L. Hay , I. Iashvili , A. Kharchilava , C. McLean , M. Morris ,
D. Nguyen , J. Pekkanen , S. Rappoccio , H. Rejeb Sfar, A. Williams 

Northeastern University, Boston, Massachusetts, U.S.A.

G. Alverson , E. Barberis , J. Dervan, Y. Haddad , Y. Han , A. Krishna , J. Li , M. Lu ,
G. Madigan , B. Marzocchi , D.M. Morse , V. Nguyen , T. Orimoto , A. Parker , L. Skinnari ,
A. Tishelman-Charny , B. Wang , D. Wood 

Northwestern University, Evanston, Illinois, U.S.A.

S. Bhattacharya , J. Bueghly, Z. Chen , K.A. Hahn , Y. Liu , Y. Miao , D.G. Monk ,
M.H. Schmitt , A. Taliercio , M. Velasco

University of Notre Dame, Notre Dame, Indiana, U.S.A.

R. Band , R. Bucci, S. Castells , M. Cremonesi, A. Das , R. Goldouzian , M. Hildreth , K.W. Ho ,
K. Hurtado Anampa , C. Jessop , K. Lannon , J. Lawrence , N. Loukas , L. Lutton , J. Mariano,
N. Marinelli, I. Mcalister, T. McCauley , C. Mcgrady , K. Mohrman , C. Moore , Y. Musienko ¹³,
H. Nelson , M. Osherson , R. Ruchti , A. Townsend , M. Wayne , H. Yockey, M. Zarucki ,
L. Zygala 

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

A. Basnet , B. Bylsma, M. Carrigan , L.S. Durkin , C. Hill , M. Joyce , A. Lesauvage ,
M. Nunez Ornelas , K. Wei, B.L. Winer , B. R. Yates 

Princeton University, Princeton, New Jersey, U.S.A.

F.M. Addesa , H. Bouchamaoui , P. Das , G. Dezoort , P. Elmer , A. Frankenthal , B. Greenberg ,
N. Haubrich , S. Higginbotham , G. Kopp , S. Kwan , D. Lange , A. Loeliger , D. Marlow ,
I. Ojalvo , J. Olsen , D. Stickland , C. Tully 

University of Puerto Rico, Mayaguez, Puerto Rico, U.S.A.S. Malik **Purdue University, West Lafayette, Indiana, U.S.A.**

A.S. Bakshi , V.E. Barnes , S. Chandra , R. Chawla , S. Das , A. Gu , L. Gutay, M. Jones ,
 A.W. Jung , D. Kondratyev , A.M. Koshy, M. Liu , G. Negro , N. Neumeister , G. Paspalaki ,
 S. Piperov , A. Purohit , J.F. Schulte , M. Stojanovic ¹⁷, J. Thieman , A. K. Virdi , F. Wang ,
 A. Wildridge , W. Xie 

Purdue University Northwest, Hammond, Indiana, U.S.A.J. Dolen , N. Parashar , A. Pathak **Rice University, Houston, Texas, U.S.A.**

D. Acosta , K. Banicz, A. Baty , U. Behrens, T. Carnahan , S. Dildick , K.M. Ecklund ,
 P.J. Fernández Manteca , S. Freed, P. Gardner, F.J.M. Geurts , P. Kelling, R. Krawczyk , A. Kumar ,
 W. Li , J.H. Liu, M. Matveev, O. Miguel Colin , T. Nussbaum, B.P. Padley , R. Redjimi, J. Rotter ,
 E. Yigitbasi , Y. Zhang 

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

A. Bodek , P. de Barbaro , R. Demina , J.L. Dulemba , C. Fallon, A. Garcia-Bellido ,
 O. Hindrichs , A. Khukhunaishvili , P. Parygin , E. Popova , R. Taus , G.P. Van Onsem 

The Rockefeller University, New York, New York, U.S.A.K. Goulianos **Rutgers, The State University of New Jersey, Piscataway, New Jersey, U.S.A.**

B. Chiarito, J.P. Chou , Y. Gershtein , E. Halkiadakis , A. Hart , M. Heindl , D. Jaroslawski ,
 O. Karacheban ²⁸, I. Laflotte , A. Lath , R. Montalvo, K. Nash, H. Routray , S. Salur , S. Schnetzer,
 S. Somalwar , R. Stone , S.A. Thayil , S. Thomas, J. Vora , H. Wang 

University of Tennessee, Knoxville, Tennessee, U.S.A.

H. Acharya, D. Ally , A.G. Delannoy , S. Fiorendi , T. Holmes , N. Karunaratna , L. Lee ,
 E. Nibigira , S. Spanier 

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

D. Aebi , M. Ahmad , T. Akhter , O. Bouhali ¹⁰³, M. Dalchenko , R. Eusebi , J. Gilmore ,
 T. Huang , T. Kamon ¹⁰⁴, H. Kim , S. Luo , S. Malhotra, R. Mueller , D. Overton , D. Rathjens ,
 A. Safonov 

Texas Tech University, Lubbock, Texas, U.S.A.

N. Akchurin , J. Damgov , V. Hegde , A. Hussain , Y. Kazhykarim, K. Lamichhane , S.W. Lee ,
 A. Mankel , T. Mengke, S. Muthumuni , T. Peltola , I. Volobouev , A. Whitbeck 

Vanderbilt University, Nashville, Tennessee, U.S.A.

E. Appelt , S. Greene, A. Gurrola , W. Johns , R. Kunnawalkam Elayavalli , A. Melo , F. Romeo ,
 P. Sheldon , S. Tuo , J. Velkovska , J. Viinikainen 

University of Virginia, Charlottesville, Virginia, U.S.A.

B. Cardwell , B. Cox , S. Goadhouse , J. Hakala , R. Hirosky , A. Ledovskoy , A. Li , C. Neu ,
C.E. Perez Lara 

Wayne State University, Detroit, Michigan, U.S.A.

P.E. Karchin 

University of Wisconsin - Madison, Madison, Wisconsin, U.S.A.

A. Aravind , S. Banerjee , K. Black , T. Bose , S. Dasu , I. De Bruyn , P. Everaerts , C. Galloni,
H. He , M. Herndon , A. Herve , C.K. Koraka , A. Lanaro , R. Loveless ,
J. Madhusudanan Sreekala , A. Mallampalli , A. Mohammadi , S. Mondal , G. Parida , D. Pinna,
A. Savin , V. Shang , V. Sharma , W.H. Smith , D. Teague , H.F. Tsoi , W. Vetens , A. Warden 

Authors affiliated with an institute or an international laboratory covered by a cooperation agreement with CERN

S. Afanasiev , V. Andreev , Yu. Andreev , T. Aushev , M. Azarkin , A. Babaev , A. Belyaev ,
V. Blinov¹⁰⁵, E. Boos , V. Borshch , D. Budkouski , M. Chadeeva ¹⁰⁵, V. Chekhovsky , A. Dermenev ,
T. Dimova ¹⁰⁵, D. Druzhkin ¹⁰⁶, M. Dubinin ⁹⁵, L. Dudko , A. Ershov , G. Gavrillov ,
V. Gavrillov , S. Gninenko , V. Golovtsov , N. Golubev , I. Golutvin , I. Gorbunov , A. Gribushin ,
Y. Ivanov , V. Kachanov , L. Kardapoltsev ¹⁰⁵, V. Karjavine , A. Karneyeu , L. Khein , V. Kim ¹⁰⁵,
M. Kirakosyan , D. Kirpichnikov , M. Kirsanov , V. Klyukhin , O. Kodolova ¹⁰⁷, D. Konstantinov ,
V. Korenkov , A. Kozyrev ¹⁰⁵, N. Krasnikov , A. Lanev , P. Levchenko ¹⁰⁸, A. Litomin , O. Lukina ,
N. Lychkovskaya , A. Malakhov , V. Matveev ¹⁰⁵, V. Murzin , A. Nikitenko ^{109,107}, S. Obraztsov ,
V. Oreshkin , A. Oskin , V. Palichik , V. Perelygin , S. Petrushanko , S. Polikarpov ¹⁰⁵, V. Popov,
O. Radchenko ¹⁰⁵, V. Rusinov , M. Savina , V. Savrin , D. Selivanova , V. Shalaev , S. Shmatov ,
S. Shulha ¹⁰⁵, Y. Skovpen ¹⁰⁵, S. Slabospitskii , V. Smirnov , D. Sosnov , V. Sulimov ,
E. Tcherniaev , A. Terkulov , O. Teryaev , I. Tlisova , A. Toropin , L. Uvarov , A. Uzunian ,
A. Vorobyev[†], N. Voytishin , B.S. Yuldashev¹¹⁰, A. Zarubin , I. Zhizhin , A. Zhokin 

[†] Deceased

¹ Also at Yerevan State University, Yerevan, Armenia

² Also at TU Wien, Vienna, Austria

³ Also at Institute of Basic and Applied Sciences, Faculty of Engineering, Arab Academy for Science, Technology and Maritime Transport, Alexandria, Egypt

⁴ Also at Ghent University, Ghent, Belgium

⁵ Also at Universidade Estadual de Campinas, Campinas, Brazil

⁶ Also at Federal University of Rio Grande do Sul, Porto Alegre, Brazil

⁷ Also at UFMS, Nova Andradina, Brazil

⁸ Also at Nanjing Normal University, Nanjing, China

⁹ Now at The University of Iowa, Iowa City, Iowa, U.S.A.

¹⁰ Also at University of Chinese Academy of Sciences, Beijing, China

¹¹ Also at University of Chinese Academy of Sciences, Beijing, China

¹² Also at Université Libre de Bruxelles, Bruxelles, Belgium

¹³ Also at an institute or an international laboratory covered by a cooperation agreement with CERN

¹⁴ Also at Suez University, Suez, Egypt

¹⁵ Now at British University in Egypt, Cairo, Egypt

¹⁶ Also at Birla Institute of Technology, Mesra, Mesra, India

- ¹⁷ Also at *Purdue University, West Lafayette, Indiana, U.S.A.*
- ¹⁸ Also at *Université de Haute Alsace, Mulhouse, France*
- ¹⁹ Also at *Department of Physics, Tsinghua University, Beijing, China*
- ²⁰ Also at *Ilia State University, Tbilisi, Georgia*
- ²¹ Also at *Tbilisi State University, Tbilisi, Georgia*
- ²² Also at *The University of the State of Amazonas, Manaus, Brazil*
- ²³ Also at *Erzincan Binali Yildirim University, Erzincan, Turkey*
- ²⁴ Also at *University of Hamburg, Hamburg, Germany*
- ²⁵ Also at *RWTH Aachen University, III. Physikalisches Institut A, Aachen, Germany*
- ²⁶ Also at *Isfahan University of Technology, Isfahan, Iran*
- ²⁷ Also at *Bergische University Wuppertal (BUW), Wuppertal, Germany*
- ²⁸ Also at *Brandenburg University of Technology, Cottbus, Germany*
- ²⁹ Also at *Forschungszentrum Jülich, Juelich, Germany*
- ³⁰ Also at *CERN, European Organization for Nuclear Research, Geneva, Switzerland*
- ³¹ Also at *Physics Department, Faculty of Science, Assiut University, Assiut, Egypt*
- ³² Also at *Wigner Research Centre for Physics, Budapest, Hungary*
- ³³ Also at *Institute of Physics, University of Debrecen, Debrecen, Hungary*
- ³⁴ Also at *Institute of Nuclear Research ATOMKI, Debrecen, Hungary*
- ³⁵ Now at *Universitatea Babes-Bolyai - Facultatea de Fizica, Cluj-Napoca, Romania*
- ³⁶ Also at *Faculty of Informatics, University of Debrecen, Debrecen, Hungary*
- ³⁷ Also at *Punjab Agricultural University, Ludhiana, India*
- ³⁸ Also at *UPES - University of Petroleum and Energy Studies, Dehradun, India*
- ³⁹ Also at *University of Visva-Bharati, Santiniketan, India*
- ⁴⁰ Also at *University of Hyderabad, Hyderabad, India*
- ⁴¹ Also at *Indian Institute of Science (IISc), Bangalore, India*
- ⁴² Also at *University of Minnesota, Minneapolis, Minnesota, U.S.A.*
- ⁴³ Also at *IIT Bhubaneswar, Bhubaneswar, India*
- ⁴⁴ Also at *Institute of Physics, Bhubaneswar, India*
- ⁴⁵ Also at *Deutsches Elektronen-Synchrotron, Hamburg, Germany*
- ⁴⁶ Now at *Department of Physics, Isfahan University of Technology, Isfahan, Iran*
- ⁴⁷ Also at *Sharif University of Technology, Tehran, Iran*
- ⁴⁸ Also at *Department of Physics, University of Science and Technology of Mazandaran, Behshahr, Iran*
- ⁴⁹ Also at *Helwan University, Cairo, Egypt*
- ⁵⁰ Also at *Italian National Agency for New Technologies, Energy and Sustainable Economic Development, Bologna, Italy*
- ⁵¹ Also at *Centro Siciliano di Fisica Nucleare e di Struttura Della Materia, Catania, Italy*
- ⁵² Also at *ENEA - Casaccia Research Center, S. Maria di Galeria, Italy*
- ⁵³ Also at *Università degli Studi Guglielmo Marconi, Roma, Italy*
- ⁵⁴ Also at *Facoltà Ingegneria, Università di Roma, Roma, Italy*
- ⁵⁵ Also at *Scuola Superiore Meridionale, Università di Napoli 'Federico II', Napoli, Italy*
- ⁵⁶ Also at *Fermi National Accelerator Laboratory, Batavia, Illinois, U.S.A.*
- ⁵⁷ Also at *Laboratori Nazionali di Legnaro dell'INFN, Legnaro, Italy*
- ⁵⁸ Also at *Università di Napoli 'Federico II', Napoli, Italy*
- ⁵⁹ Also at *Ain Shams University, Cairo, Egypt*
- ⁶⁰ Also at *Consiglio Nazionale delle Ricerche - Istituto Officina dei Materiali, Perugia, Italy*
- ⁶¹ Also at *Riga Technical University, Riga, Latvia*
- ⁶² Also at *Department of Applied Physics, Faculty of Science and Technology, Universiti Kebangsaan Malaysia, Bangi, Malaysia*
- ⁶³ Also at *Consejo Nacional de Ciencia y Tecnología, Mexico City, Mexico*
- ⁶⁴ Also at *Trincomalee Campus, Eastern University, Sri Lanka, Nilaveli, Sri Lanka*
- ⁶⁵ Also at *Saegis Campus, Nugegoda, Sri Lanka*
- ⁶⁶ Also at *INFN Sezione di Pavia, Università di Pavia, Pavia, Italy*

- ⁶⁷ Also at National and Kapodistrian University of Athens, Athens, Greece
- ⁶⁸ Also at Ecole Polytechnique Fédérale Lausanne, Lausanne, Switzerland
- ⁶⁹ Also at Universität Zürich, Zurich, Switzerland
- ⁷⁰ Also at Paul Scherrer Institut, Villigen, Switzerland
- ⁷¹ Also at Stefan Meyer Institute for Subatomic Physics, Vienna, Austria
- ⁷² Also at Laboratoire d'Annecy-le-Vieux de Physique des Particules, IN2P3-CNRS, Annecy-le-Vieux, France
- ⁷³ Also at Near East University, Research Center of Experimental Health Science, Mersin, Turkey
- ⁷⁴ Also at Konya Technical University, Konya, Turkey
- ⁷⁵ Also at Izmir Bakircay University, Izmir, Turkey
- ⁷⁶ Also at Adiyaman University, Adiyaman, Turkey
- ⁷⁷ Also at Necmettin Erbakan University, Konya, Turkey
- ⁷⁸ Also at Bozok Universitetesi Rektörlüğü, Yozgat, Turkey
- ⁷⁹ Also at Marmara University, Istanbul, Turkey
- ⁸⁰ Also at Milli Savunma University, Istanbul, Turkey
- ⁸¹ Also at Kafkas University, Kars, Turkey
- ⁸² Also at Istanbul Bilgi University, Istanbul, Turkey
- ⁸³ Also at Hacettepe University, Ankara, Turkey
- ⁸⁴ Also at Istanbul University - Cerrahpasa, Faculty of Engineering, Istanbul, Turkey
- ⁸⁵ Also at Ozyegin University, Istanbul, Turkey
- ⁸⁶ Also at Vrije Universiteit Brussel, Brussel, Belgium
- ⁸⁷ Also at School of Physics and Astronomy, University of Southampton, Southampton, United Kingdom
- ⁸⁸ Also at University of Bristol, Bristol, United Kingdom
- ⁸⁹ Now at Tata Institute of Fundamental Research-B, Mumbai, India
- ⁹⁰ Also at IPPP Durham University, Durham, United Kingdom
- ⁹¹ Also at Monash University, Faculty of Science, Clayton, Australia
- ⁹² Also at Università di Torino, Torino, Italy
- ⁹³ Also at Bethel University, St. Paul, Minnesota, U.S.A.
- ⁹⁴ Also at Karamanoğlu Mehmetbey University, Karaman, Turkey
- ⁹⁵ Also at California Institute of Technology, Pasadena, California, U.S.A.
- ⁹⁶ Also at United States Naval Academy, Annapolis, Maryland, U.S.A.
- ⁹⁷ Also at Beykent University, Istanbul, Turkey
- ⁹⁸ Also at Bingol University, Bingol, Turkey
- ⁹⁹ Also at Georgian Technical University, Tbilisi, Georgia
- ¹⁰⁰ Also at Sinop University, Sinop, Turkey
- ¹⁰¹ Also at Erciyes University, Kayseri, Turkey
- ¹⁰² Also at Horia Hulubei National Institute of Physics and Nuclear Engineering (IFIN-HH), Bucharest, Romania
- ¹⁰³ Also at Texas A&M University at Qatar, Doha, Qatar
- ¹⁰⁴ Also at Kyungpook National University, Daegu, Korea
- ¹⁰⁵ Also at another institute or international laboratory covered by a cooperation agreement with CERN
- ¹⁰⁶ Also at Universiteit Antwerpen, Antwerpen, Belgium
- ¹⁰⁷ Also at Yerevan Physics Institute, Yerevan, Armenia
- ¹⁰⁸ Also at Northeastern University, Boston, Massachusetts, U.S.A.
- ¹⁰⁹ Also at Imperial College, London, United Kingdom
- ¹¹⁰ Also at Institute of Nuclear Physics of the Uzbekistan Academy of Sciences, Tashkent, Uzbekistan