

The CMS experiment at the LHC—design and initial performance

Author's contribution to this subject

The author actively participated in various phases of design, hardware development and construction of the CMS experiment. The author took a key role in the RPC trigger system design, being for a long time responsible for system simulation, worked on system construction and, as a coordinator, took major responsibilities in commissioning, operation and maintenance of the full system.

The author has also participated in beam tests and commissioning of the Pixel Detector and became the main author of the reconstruction code that converts raw data detector format to digitized data format suitable for further reconstruction. During LHC startup and soon after (until 2010) the author was officially responsible for this code.

The author has also developed the RPC data conversion code and maintained the DAQ module for the RPC data readout. The commissioning of the RPC detector included participation in several preparatory runs, before and after LHC startup, but also a series of runs dedicated for RPC sub-detector understanding and initial improvements.

The author was also actively participating in CMS detector data taking and quality monitoring of the trigger data. During the preparation and regular data-taking period the author was a RPC trigger shifter and on-call expert as well as a Level-1 trigger shifter. In both cases the author was responsible for sub-detector part and data quality monitoring. During selected runs the author was also the Level-1 Detector-On-Call expert responsible for operation of the entire Level-1 trigger.

The CMS results described in this Chapter were presented by the author on behalf of the CMS collaboration at international conferences as invited talks [16, 17]. These conference contributions are revised and updated in this Chapter.

1.1. Introduction

The Large Hadron Collider (LHC¹) [22] is an accelerator built at the CERN laboratory (European Organization for Nuclear Research) in Geneva, Switzerland. The four main experiments at the LHC provide frontier physics results. In its main mode the machine is designed to collide protons. The machine center-of-mass energy was $\sqrt{s} = 7$ TeV in 2010 and 2011, and $\sqrt{s} = 8$ TeV until the end of 2012, when proton-proton (pp) runs in the period of LHC Run-1 ended. It is expected that at the beginning of 2015 the LHC will resume its operation with the energy of $\sqrt{s} = 13$ TeV, which is close to the design value of $\sqrt{s} = 14$ TeV.

The Compact Muon Solenoid (CMS) experiment [1] had been successfully operated during LHC Run-1. The delivered statistics of about 30 fb^{-1} (see Figure² 1.1), allowed physicists to perform a variety of measurements, searches, and first of all, to discover, together with ATLAS collaboration [23], the Higgs particle [24–26]—fundamental to understand Spontaneous Symmetry Breaking Mechanism

¹The glossary of acronyms is given in Appendix D. Moreover some of the acronyms are re-introduced in each chapter for easier reading.

²All the presented figures show the results of the CMS collaboration. Whenever appropriate a more detailed reference is given.

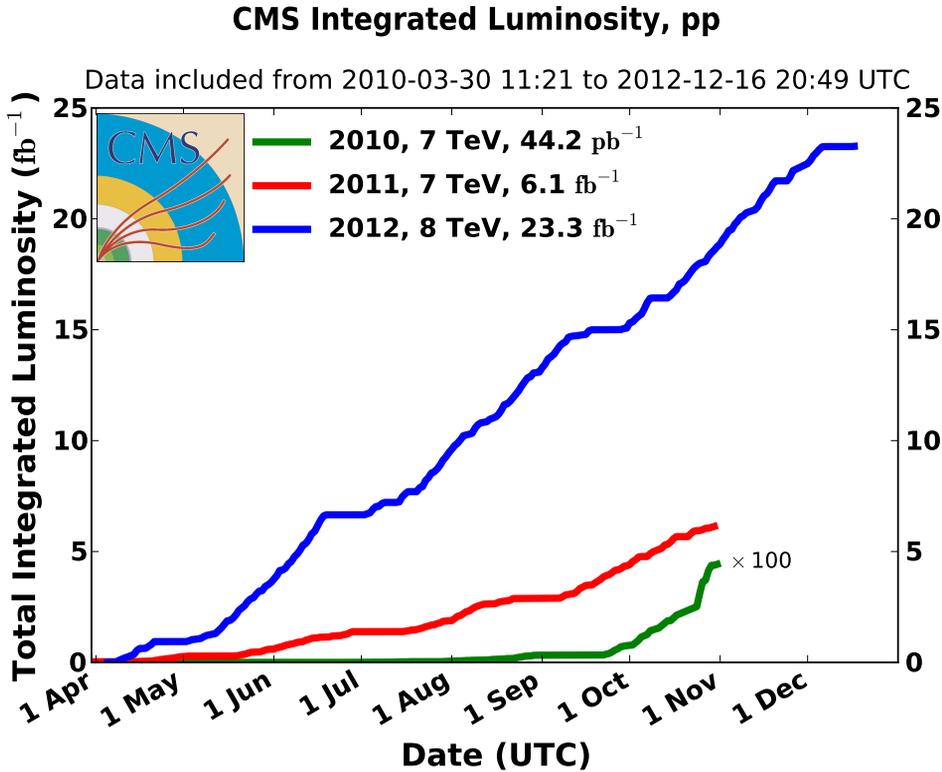


Figure 1.1. Integrated luminosity delivered to CMS in: 2010 (green, multiplied by a factor of 100), 2011 (red) and 2012 (green). The maximum reached instantaneous luminosity: $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, $4 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, $7.7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ respectively. Source: [27].

and the Standard Model. In this chapter the design of the CMS experiment, with an emphasis on the trigger component, is described. The initial performance and commissioning results are presented as well.

1.2. CMS detector

CMS is a general purpose experiment for physics discoveries at the highest luminosities provided by the LHC. Its main component is a large solenoid (6 m diameter and 13 m long). It generates strong 3.8 T magnetic field in the inner part of the CMS detector and about 1.8 T inside the iron return yoke surrounding the solenoid. CMS is traditionally divided into the barrel part (with subdetectors aligned roughly parallel to the beam pipe) and two endcaps. Next to the beam-beam interaction region the tracker system is located. It consists of the silicon Pixel- and Strip Detectors. The CMS tracker provides excellent reconstruction of charged particle tracks and primary and secondary vertices. The tracker is surrounded by the electromagnetic calorimeter (ECAL). It is a homogeneous calorimeter made of lead-tungstate crystals. The energy measurement is supplemented with the sampling brass-scintillator hadron calorimeter (HCAL). The above subdetectors are positioned in the inner part of the CMS detector, inside the solenoid. In the outer part, outside the coil, the muon system is placed. It is dedicated to the muon reconstruction and identification. The muon system is based on gaseous detectors: Drift Tubes (DT) in the barrel, Cathode Strip Chambers (CSC) in the endcaps and Resistive Plate Chambers (RPC) in both barrel and endcaps. The pseudorapidity¹ coverage of CMS

¹Pseudorapidity $\eta = -\ln[\tan(\theta/2)]$, where θ is the polar angle measured from the beam-line (z -axis in CMS). In CMS the azimuth angle φ is measured in the standard way, anticlockwise relative to the z axis, in the transverse $x - y$ plane.

depends on a given subsystem. The tracking detectors (muon system, tracker) provide reconstruction up to $|\eta| \approx 2.4\text{--}2.5$. The calorimeter coverage is larger for the purpose of hermeticity and extends up to $|\eta| \approx 3$ in the case of ECAL and up to $|\eta| \approx 5$ for HCAL.

CMS has a two-step triggering system to reduce the designed 40 MHz LHC event input rate down to the rate suitable for storage and offline analyses.

1.3. CMS trigger and data acquisition system

The high energy and high luminosity provided by the LHC machine result in a large number of pp collisions and produced particles. Due to technical restrictions not all of them can be stored and analyzed offline.

The CMS design principles [9] include efficient lepton and photon selection and measurements. These particles are of key importance for discovery physics. Another important objects for the CMS physics program are hadronic jets. The event topology, including correlations, invariant mass constraints and isolation are other quantities which are important to find the interesting events, and reduce the background trigger rate. Thus, the event selection system has to be flexible, based on configurable algorithms in order to preserve signal signatures, efficiently reduce the background, and adopt to running conditions.

The CMS collaboration has designed the Trigger and Data Acquisition (Tridas) system to handle these requirements. The target output data volume rate, suitable for permanent storage and for further, offline analysis, is designed to be kept at the level of $\mathcal{O}(100\text{ MB/s})$,

The design parameters of the CMS trigger assume that each LHC beam crossing at high luminosity produces about 20 pp collisions resulting in the event data volume of approximately 1 MB of zero-suppressed data. Thus, the CMS trigger must be able to suppress the initial LHC rate down to the data storage rate of $\mathcal{O}(10^2\text{ Hz})$. The event selection at CMS is done in two triggering steps only: in the Level-1 trigger (Level-1) and in the High-Level trigger (HLT).

The Level-1 reduces the designed 40 MHz beam collision rate¹ to less than 100 kHz. It is based on custom, partially programmable hardware devices (dedicated ASICs, or FPGAs where appropriate). It analyses coarsely segmented data from the calorimeter and muon systems only. During the Level-1 trigger processing the full granularity data are stored in detector front-end pipelines.

The HLT is implemented in an expandable computer farm. The executed algorithms use detector data at full granularity, including information from the tracker. The algorithms executed at HLT are similar to those used in the offline analysis but are optimized for fast, online processing.

The 100 kHz event rate of assumed 1 MB event data volume is a design constraint for the HLT and the CMS Data Acquisition (DAQ). A schematic view of the central part of the CMS DAQ is shown in Figure 1.2. Data from the detectors are stored in modules of the detector front-ends until the Level-1 accept signal is issued by the final step of the Level-1—Global Trigger. Then data are read-out from Front-End Drivers (FED) and buffered in Readout Units (RU). At this stage data from each event are spread out over several separated units. The Builder Network (Readout Builder Network) is responsible for collecting data belonging to one event and providing this information to the Filter Systems. The Builder Network itself is a large switching facility with 100 GB/s throughput bandwidth, imposed by the Level-1 rate and event size. The Builder Network at the time of system design was a major technological challenge. After the Builder Network, in the Filter Systems each event is buffered in one Builder Unit (BU) to provide the data from a complete event to the Filter Unit (implemented in the same machine). Each Filter Unit (FU) contains a set of commercial CPU processor and the Filter Units form the Filter Farm. Each event is assigned to a single processor

¹The Level-1 works with a frequency of 40 MHz irrespectively of the bunch structure used by LHC.

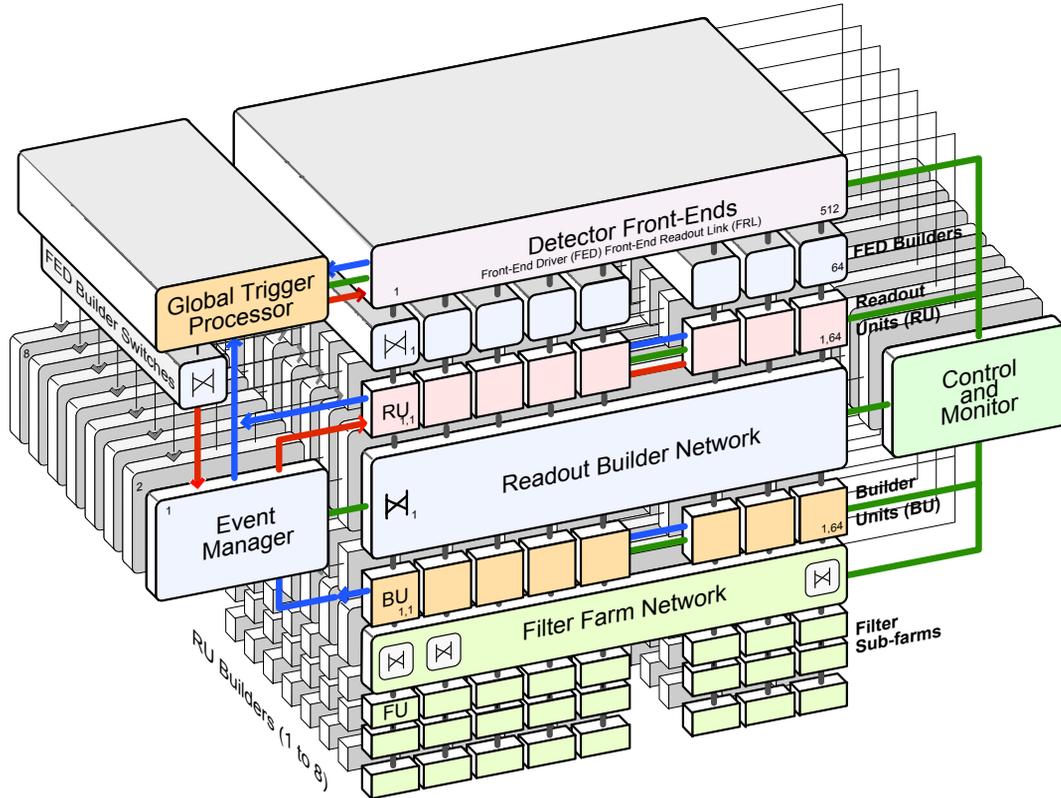


Figure 1.2. Modular design of the central part of the CMS DAQ. The data read-out from detector front-ends is initiated by Global Trigger signal. The data are collected, distributed by the readout network, which provides a full event to the Filter Farm, consisting of the Filter Units. At Filter Units the HLT algorithms are executed. The design allows to expand the system or to temporarily disable its part. Source: CMS.

where the HLT algorithms are executed and the actual event selection takes place. The HLT reduces the Level-1 output rate below the level of the maximal available event storage and offline event-processing rate. The final treatment of the events selected at the HLT is addressed to Computing Services that forward the data to mass storage and perform monitoring tasks. The DAQ system is completed with the Event Manager, responsible for the data flow control, and the Control System which takes care of configuration, DAQ monitoring and various control tasks. The events accepted by the HLT are temporary stored on local disks and further send to CERN Data Center for permanent storage and for offline analysis. The CMS data analysis is performed using the Worldwide LHC Computing Grid (WLCG).

The CMS trigger system does not have any intermediate step between the Level-1 and the HLT. The traditional "Level-2" trigger is often based on dedicated hardware and uses limited data granularity. The CMS approach to skip this intermediate filter and to execute HLT algorithms immediately after Level-1 was a challenge at a time, but it provided many advantages. This design maximally benefits from the computer technology and its developments over time; it has the maximal flexibility without design and architectural limitations of dedicated hardware solutions; there is no limitation on data accessing type and granularity; there is maximal freedom in the choice of the selection algorithms. On the other hand processing huge amounts of data makes the system challenging and, therefore, the HLT algorithms have to be very efficient.

The available storage and offline computing resources allowed CMS to increase the rate of events accepted by the HLT. Moreover, the storage capabilities surpass the processing resources. Hence in CMS, the recorded data is divided into those dedicated for the main CMS program and those for

later processing for additional analyses. The typical event rates were 400 Hz and 600 Hz for the two cases above, respectively.

1.4. CMS operation

The LHC has started its operation in 2008, but it was closed shortly after its opening due to helium leakage incident, on account of splice-bar bus overheating. After re-opening of the LHC in 2009, CMS had started to collect the actual pp collisions data. The initial center-of-mass energy delivered by the LHC of $\sqrt{s} = 0.9$ TeV was followed by 2.36 TeV at the end of 2009. In March 2010 the LHC energy was increased up to half of its designed energy, i.e. to $\sqrt{s} = 7$ TeV and later, in 2012, up to $\sqrt{s} = 8$ TeV.

Although the instantaneous LHC luminosity at the startup was small, during a few months of operation in 2010 it was increased by 5 orders of magnitude and reached a value of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. The instantaneous luminosity was further increased to the maximal peak values of $4 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in 2011 and $7.7 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in 2012. The luminosity changes as a function of time are presented in Figure 1.3. CMS is facing the changing conditions given by the LHC. The CMS trigger is adapting to them by continuous adjustment of executed trigger algorithms (*trigger menus*, see Sections 2.1.3 and 3.2). The luminosity increase affects not only the trigger rates but also the number of overlapping events in one bunch crossing. At peak luminosity, the event pile-up increased from 3.5 in 2010 up to 34.5 in 2012.

The bunch structure was changing during Run-1. Initially there were single colliding bunches, but for most of operation in 2011 and 2012 the LHC bunch spacing was 50 ns (twice the designed values), corresponding to real collision rate of 20 MHz. The lower than designed collision frequencies were compensated by larger number of protons per bunch and better beam focusing. The average number of pile-up events in 2012 reached 21.5—a value close to the expected pile-up for LHC design parameters of energy and luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with 25 ns bunch spacing.

In Figure 1.4 delivered, recorded and validated integrated luminosities are shown. The large pile-up and luminosity changes have not compromised the performance of the CMS detector. The efficiency of data taking, given by the ratio of recorded and delivered luminosity varies in the range

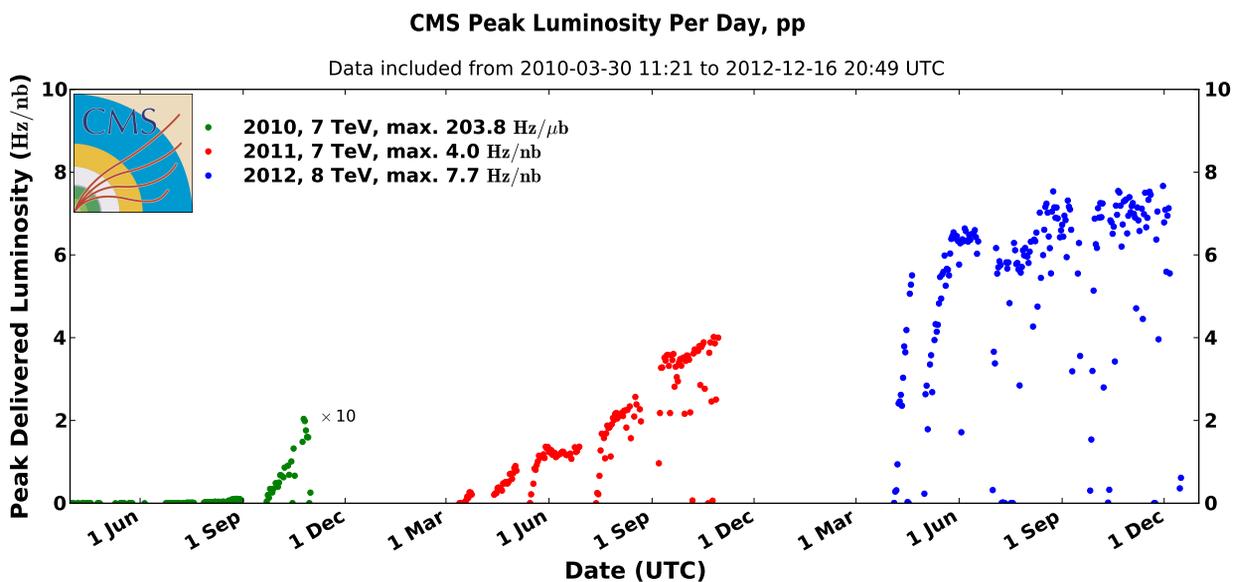


Figure 1.3. Evolution of peak instantaneous luminosity in CMS in: 2010 (green, multiplied by a factor of 10), 2011 (red) and 2012 (blue). The maximal values in each year are also shown. Source: [27].

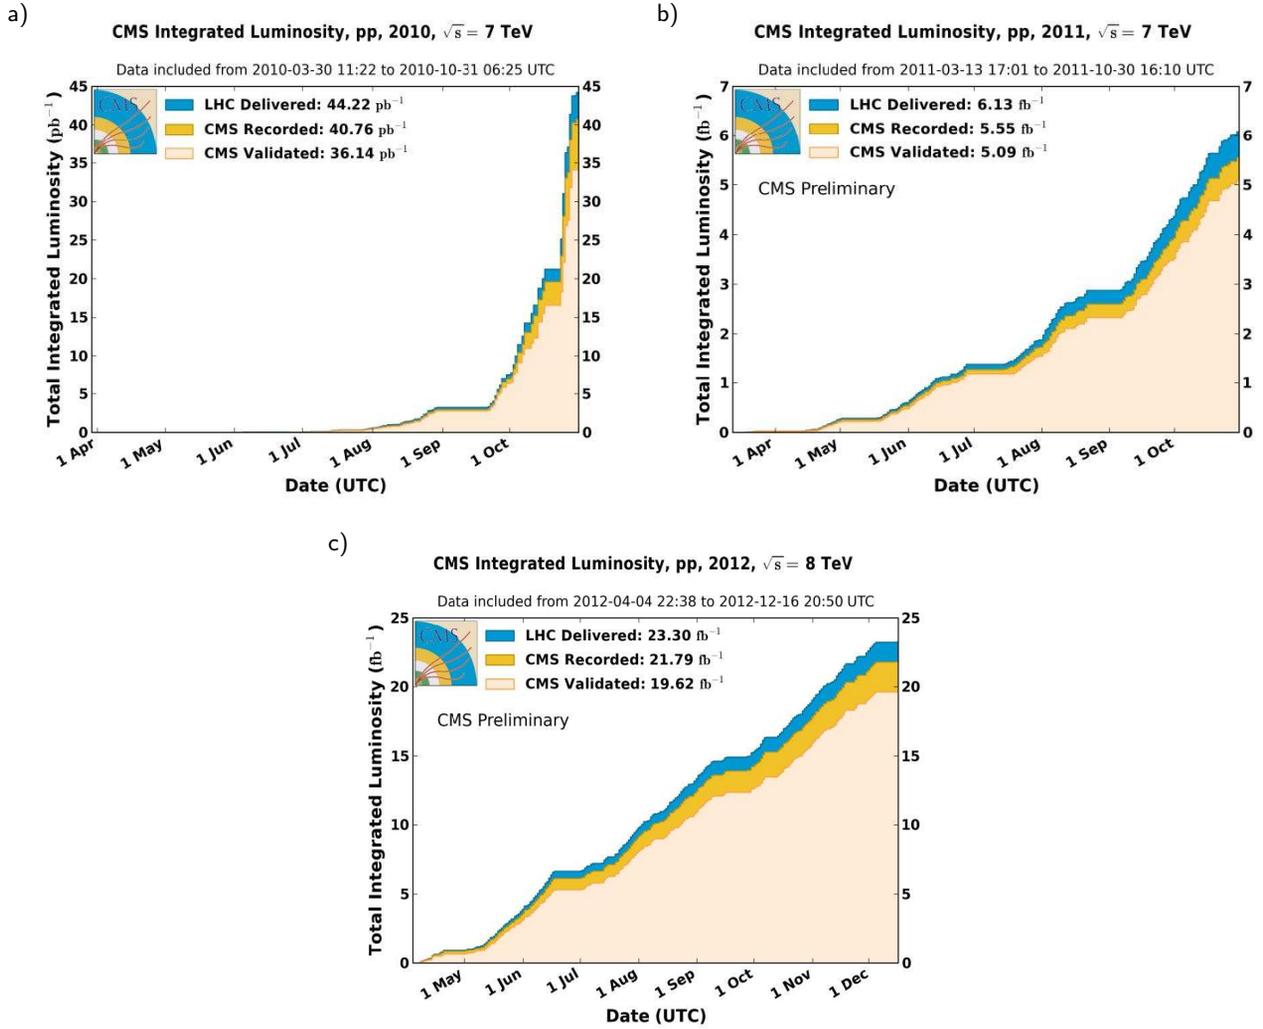


Figure 1.4. Delivered, recorded and validated integrated luminosity in 2010 (a), 2011 (b) and 2012 (c). The full statistics of 5.1 fb⁻¹ data at $\sqrt{s} = 7$ TeV and 19.7 fb⁻¹ at 8 TeV is used in the Higgs analysis. Source: CMS.

of 90–93%. The efficiency losses are attributed to: start-stop run procedure, down-times due to detector configuration and limited bandwidth. A significant improvement of the recording efficiency in 2012 was achieved by automatisation of run recovery procedures. The recorded CMS data underwent further quality validation procedures. Only validated data were used in the physics analyses. Approximately 90% of data was validated positively. This fraction was rather constant with time. Among the reasons of qualifying the data as bad, there were: detector configuration issues, noisy or inefficient cells in subdetectors and problems with data processing at computer farm. CMS also exhibited very good performance in terms of the number of active channels. Within the main subdetectors it varied from 96.3% (Pixels) up to 99.9% (HCAL) at the end of LHC Run-1 (Figure 1.5). These fractions of active channels were not visibly degraded since the beginning of CMS operation. In case of the muon system the number of active channels has decreased by about half percent starting from initial values of 98%, 98.5%, 99.5% for the CSC, RPC and DT sub-systems, respectively.

At the end of 2010 the LHC has entered into a heavy-ion program, colliding lead-lead beams at nucleon-nucleon center-of-mass energy of $\sqrt{s_{NN}} = 2.76$ TeV and collecting 8.3 μb^{-1} of data. The heavy-ion program was continued in the same configuration in the end of 2011, with 158 μb^{-1} of collected data. The CMS heavy-ion program was complemented with proton-lead collisions at the energy of 5.02 TeV. LHC Run-1 heavy-ion program ended in February 2013 with 30 nb⁻¹ of collected data.

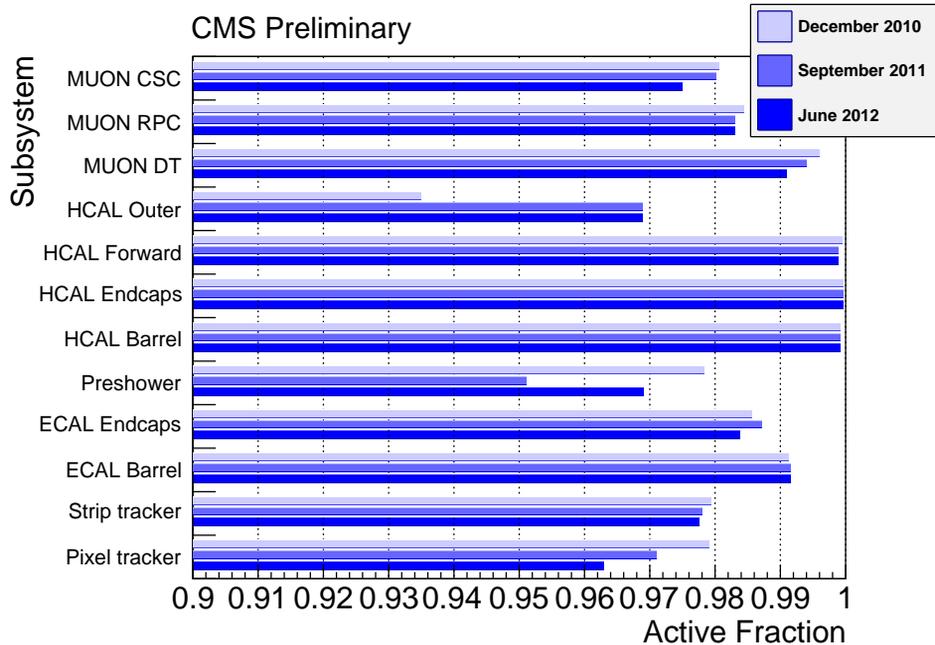


Figure 1.5. Summary of fraction of active channels in CMS subdetectors. Source [28].

1.5. Commissioning and initial detector performance results

The CMS detector was initially commissioned with test beams data and cosmic runs (see [29] and references therein) before the LHC startup. Since 2008 the baseline construction of the CMS detector has been finished. While waiting for the LHC beams, CMS was focusing on integration runs with cosmic muons. Among others this allowed us to test muon reconstruction and inter-detector alignment [30]. Since the structure of the CMS detector is not rigid, the mechanical deformations due to magnetic field, gravity (weight), temperature and humidity have to be taken into account. Thus CMS has developed alignment mechanisms using dedicated optical system and tracking. An example of such an integration exercise is shown in Figure 1.6. An impact of the alignment corrections on the muon reconstruction is clearly visible. The alignment precision was estimated to be about $500 \mu\text{m}$ for DT and $300\text{--}600 \mu\text{m}$ for CSC detectors [31].

Commissioning was continued with early LHC data. At this step the key aspects were: calibration and alignment of subdetectors, validation of reconstruction algorithms, comparison of detector response (reconstructed physics objects) with simulation predictions, validation and tuning of trigger algorithms and menus.

The reconstruction of charged particles is one of key aspects to understand an event content. CMS has demonstrated an excellent performance and a good understanding of tracking capabilities [33, 34]. Tracker operation conditions were validated. Just after the LHC startup timing readout windows were optimized. The early commissioning also included measurements of the Lorentz angle, energy loss (dE/dx) and subdetector efficiencies. The validation of tracking reconstruction algorithms included studies of: resolution and efficiency of track and primary-vertices finding, multiple interaction extraction, determination of beam-line position and width.

A good illustration of the overall tracking performance are searches for well-known resonances, for example in $\Xi^\pm \rightarrow \Lambda^0/\bar{\Lambda}^0\pi^\pm$ decays. This early analysis involves reconstruction of a secondary vertex (decay of Λ). It is formed by two opposite charge tracks (assumed to be pion and (anti)proton). Their transverse impact parameters should be not compatible with the beam spot. These tracks should form together a well-separated secondary vertex with a correct invariant mass. In addition,

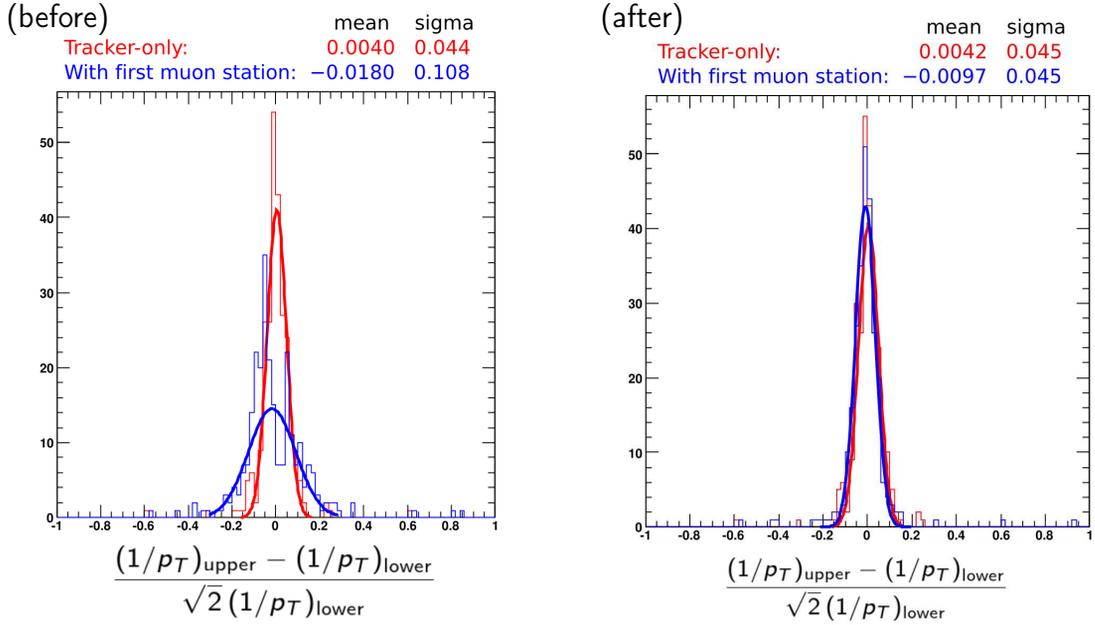


Figure 1.6. An example of muon system alignment exercise before the LHC startup. Events with muons reconstructed independently in the upper and lower part of detector are selected. The distribution of $1/p_T$ difference between measurements is plotted. Muons are reconstructed using tracker-only data and tracker data plus first muon station measurements. The distributions are shown before and after alignment (left and right plots respectively). The obtained alignment precisions is equivalent to integrated LHC luminosity of 10 pb^{-1} . Only tracks with $p_T > 200 \text{ GeV}/c$ enter. Source: [32].

since Ξ^\pm is a long lived baryon, there should be one more charged particle (pion) not compatible with beam spot. This pion should form a common vertex with Λ . In Figure 1.7a one can see invariant mass histogram of Ξ candidates from early LHC collisions at $\sqrt{s} = 0.9$ and 2.36 TeV .

The excellent muon reconstruction is one of the key design points of the CMS experiment, and an important observable supporting Higgs searches. The muon reconstruction, described in Chapter 3, includes measurements in the tracker and muon system. The period of an early commissioning with proton beams was dominated by validating detector performance [35] of resolution, local reconstruction inside muon stations, alignment, calibration and data synchronization. It has

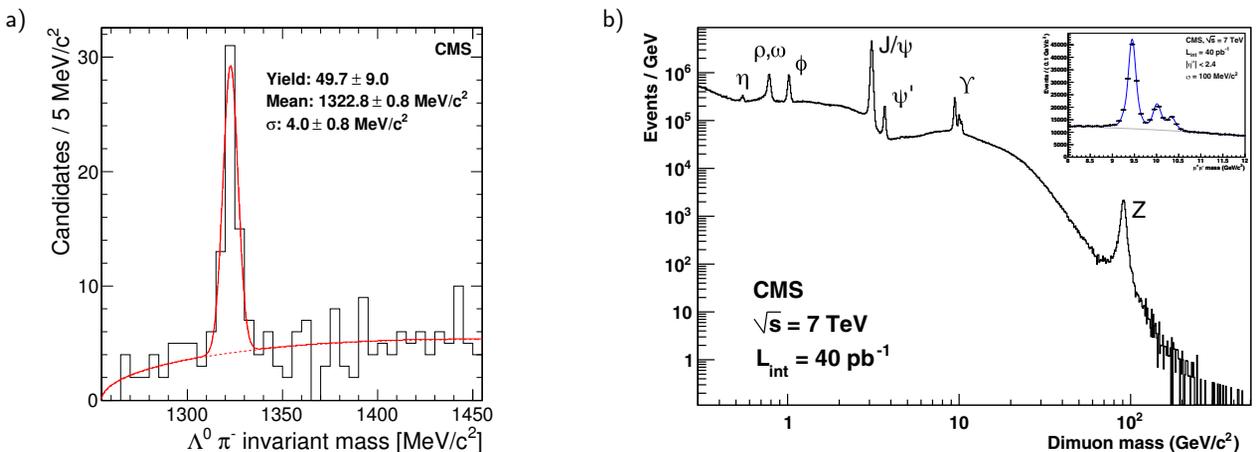


Figure 1.7. (a) Illustration of low mass resonances searches, 2009 data: invariant mass distribution of $\Lambda^0 \pi^-$ (and $\bar{\Lambda}^0 \pi^+$) with a peak from $\Xi \rightarrow \Lambda \pi$ decays. (b) Muon reconstruction and identification, 2010 data: invariant mass distributions of $\mu^+ \mu^-$ with $\eta, (\rho, \omega), \phi, J/\psi, \psi', \Upsilon(1, 2, 3S)$ and Z mass peaks visible. Source: [31, 33].

been followed by studies of muon identification and reconstruction [31]. This includes analysis of reconstruction algorithms by comparison with generator expectations, validation of muon isolation algorithms, analysis of muon deposit in calorimeters, cosmic backgrounds, hadron decays in-flight, punch through probability and muon trigger performance. An illustration of very good performance of the muon system, from the first days of LHC, can be a di-muon invariant mass distribution as shown in Figure 1.7b. One can note a clearly visible fine structure of Υ family. The relative resolution of muon transverse momenta for $p_T < 100$ GeV/c is 1.3–2% in the barrel and increases up to 6% in endcaps [31]. The muon momentum resolution, in this p_T range, is determined by the tracker, which in the central region has a resolution in p_T of about 1.5% for $1 < p_T < 10$ GeV/c tracks, and about 2.8% for $p_T = 100$ GeV/c. The resolutions in transverse and longitudinal impact parameter for high- p_T muons are $10 \mu\text{m}$ and $30 \mu\text{m}$, respectively [36]. The muonic decays of J/ψ , Υ and Z are used to validate the muon momentum scale and resolution.

The precise measurement of electromagnetic cascades is another vital aspect for the Higgs boson searches as well as many exotic channels. Thus, the early 2010 data were used to finalize

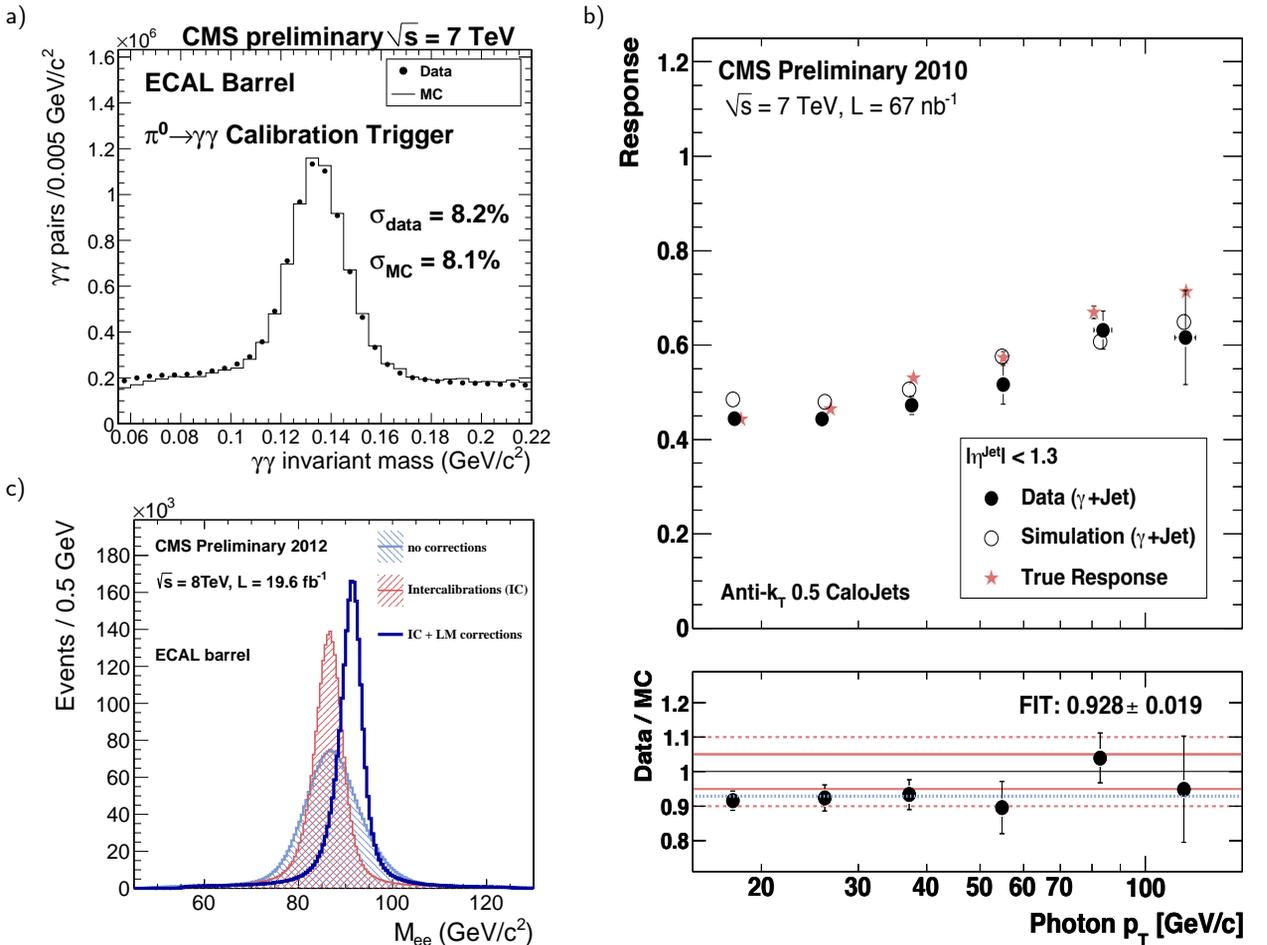


Figure 1.8. (a) Result of initial calibration of electromagnetic calorimeter: Invariant mass distribution for photon pairs (barrel only). The plot was obtained after analyzing only 18.7 nb^{-1} of data. The distribution expected from Monte Carlo generator, corresponding to the same number of events is also shown. (b) Initial calibration of hadronic calorimeter. For a photon plus jet sample the relative response $\langle p_T^{\text{jet}}/p_T^\gamma \rangle$ is shown as a function of p_T^γ . The data agrees well with simulation prediction. In addition simulation truth response (p_T ratio of simulated response and particle level jet) is also indicated. (c) The 2012 ECAL calibration monitoring with $Z \rightarrow e^+e^-$ decays: the non-calibrated data are compared to those with inter-calibration method applied and to fully calibrated by using intercalibration and laser monitoring data. Source: [40, 42, 43].

calorimeter commissioning [37]. This includes validation of crystal transparency, thermal stability and timing alignment. Interestingly, CMS observed unphysical high deposits in single crystals. They are understood to be caused by direct ionization of the avalanche photodiode by highly ionizing particles resulting from LHC collisions. Algorithms have been developed to flag these signals based on topological and timing characteristics and reject them. The commissioning of electromagnetic calorimeter with 2010 data has been completed by analyses of reconstruction performance [38, 39], including efficiency measurement from data and calibration [40] (see also [41]). The CMS electromagnetic calorimeter has been pre-calibrated with laboratory measurements, test beams, cosmic rays and early LHC data. The final calibration is made in situ using LHC collision data. The strategy to calibrate electromagnetic calorimeter since 2010 includes φ -symmetry intercalibration and π^0/η calibration methods. The first one is exploring the φ symmetry of the detector around the beam axis and can be done with minimum bias events. It allows to intercalibrate crystals at the same pseudorapidity regions. The second method uses the photon pairs from decays of π^0 and η particles. It extends crystal intercalibration to different values of pseudorapidities and allows to investigate calorimeter energy scale. Both methods can be combined. The invariant mass of photons from π^0 decays is a quick illustration of calorimeter performance. The one obtained with only 18 nb^{-1} of collected data is shown in Figure 1.8a, where a good agreement with Monte Carlo generator prediction is visible. Another method of calibration uses decays of Z and W to electrons with electron measurement in the tracker to tune the energy scale. The $Z \rightarrow e^+e^-$ and $J/\psi \rightarrow e^+e^-$ decays can be used to monitor and correct the absolute electromagnetic energy scale (see Figure 1.8c).

Jets are another vital observables for CMS. They are among main tools to verify predictions of the Standard Model in the LHC energy regime. Moreover they are possible signatures of many New Physics processes. For jet reconstruction CMS has adopted the anti- k_t clustering algorithm [44]. An important part of CMS commissioning is the study of jet energy response and resolution [45]. Since the energy measured in the detector differs from the jet energy, a factorized procedure for the jet energy calibration was developed. There are three types of corrections applied. The energy offset correction is supposed to remove contributions from calorimeter electronic noise and pile-up. The relative correction compensates non-uniform pseudorapidity response of the calorimeter. The

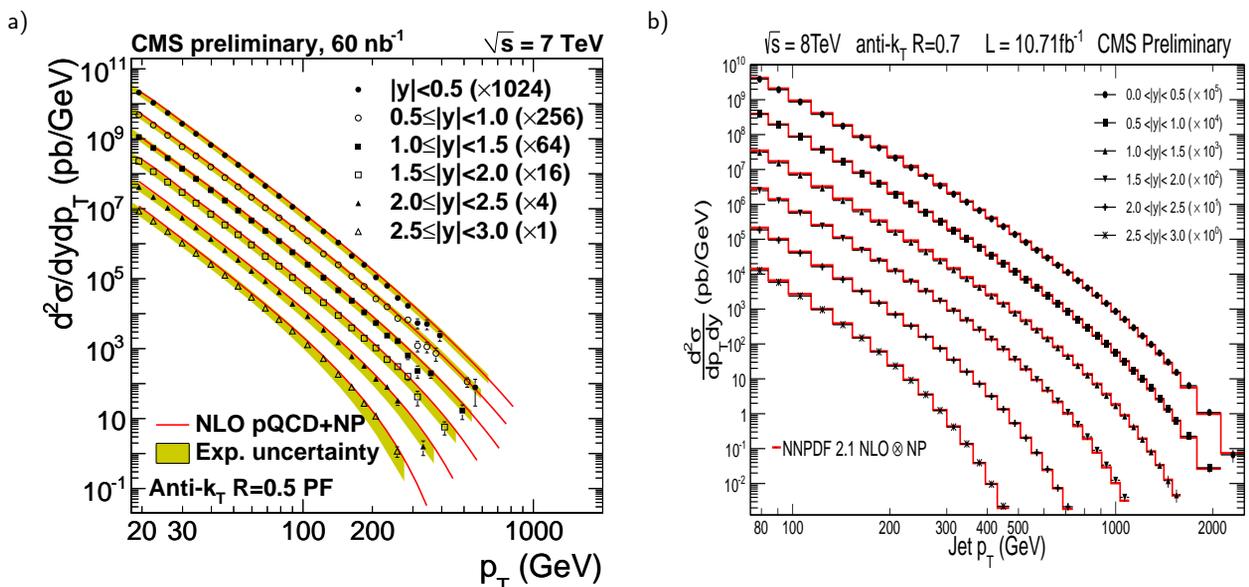


Figure 1.9. (a) Measurement of jet spectra in the early analysis (60 nb^{-1} , $\sqrt{s} = 7 \text{ TeV}$) compared to theoretical predictions. The spectra are scaled with a factor indicated in the legend. (b) The 2012 ($\sqrt{s} = 8 \text{ TeV}$) analysis is shown for the comparison. Source: [47, 48].

absolute correction removes variation in jet response as a function of jet transverse momentum p_T . In order to determine jet energy corrections CMS is using Monte Carlo information and physics processes for validation and in-situ calibration (resulting in small additional corrections). The di-jet p_T balance is used for validation of relative corrected jet energy response, while γ/Z plus jet balance method provides measurement of the absolute energy scale. An initial CMS result, illustrating not only quality of preparation of CMS for data taking but also the high quality of CMS simulation, is presented in Figure 1.8b. The relative response agrees well with expectations and justifies usage of 10% of jet energy uncertainties for early physics publications.

Calorimeter-only jet measurements can be improved with particle-flow method [46]. It attempts to identify and individually reconstruct all particles produced in collisions using information from all CMS subdetectors. This information used at the initial level of jet clusterization, allows for more precise jet reconstruction. The measured CMS jet spectrum is shown in Figure 1.9. The power of the particle-flow method is well visible in Figure 1.9a, obtained soon after LHC startup [47] (see also [49]). The jet energies extend to regimes not accessible with calorimeter-only jets, being still in very good agreement with theoretical predictions. The particle-flow method is used in most of CMS physics analyses.