Università degli Studi di Padova

Dipartimento di Fisica e Astronomia "Galileo Galilei"
Corso di Laurea Triennale in Fisica

Study of the performance of the CMS Pixel Luminosity Telescope

Laureando:
Gianmarco Pompeo
Matricola 1096791

Relatore:
Prof. Marco Zanetti

Anno Accademico 2016-2017

Contents

Introduction

The Large Hadron Collider (LHC) operating at the European Organization for Nuclear Research (CERN) is the largest and most powerful accelerator ever built and its record-setting features have opened the way to groundbreaking discoveries in subnuclear physics, thanks to its four iconic experiments, ALICE, ATLAS, CMS and LHCb.
The CMS Beam Radiation Instrumentation and Luminosity (BRIL) group is devoted to the measurement of luminosity, beam conditions and radiation fields at the CMS experiment. The project is engaged in operating and developing new detectors, in view of the high luminosity upgrade at the LHC.
Luminosity stands alongside energy as one of the two most important performance parameters of an accelerator, so it should come as no surprise that a big experiment like CMS has at least four detectors in charge of delivering this kind of measurement. Luminosity provides with an estimate of the number of particles that can collide inside an accelerator in a given amount of time and therefore, while not being exactly a collision rate, it can easily give an understanding of how many events are taking place in a particle collider. It is defined as the number of such events detected per unit of time and it is also connected to the cross-section σ associated to that particular process. Both the instantaneous and integrated values of luminosity are helpful to describe how efficiently an accelerator is performing: as a general rule, accelerators tend to maximize their integrated luminosity as this means that more data is available for offline analysis.

In this thesis I will firstly outline the experimental framework, describing the main characteristics of the LHC accelerator and CMS, the experiment I worked at, with its complex features. Afterwards, I will focus on measurements of luminosity being conducted at CMS: I will both present the main theoretical aspects of such topic, the procedure of its calibration and I will then move on to describe the features of the PLT detector and the way it works in measuring luminosity. Finally, I will delve into the main scope of the analysis that has been done to investigate the performance of this detector: an investigation on the possible causes for a drop in the per-fill visible cross-section measured is presented together with another analysis on an unforeseen increase in the PLT baseline corrections.

Chapter 1

The experimental framework

1.1 The Large Hadron Collider

Located in the outskirts of Geneva, Switzerland, straddling the border with France, the Large Hadron Collider is the largest and most powerful accelerator ever built. Projects for its construction were made in as far back as 1984; the accelerator started up for the first time on 10 September 2008 and over the years it has been shut down also for consecutive months to keep improving its performance.

The LHC consists of a 27 km ring that is located underground at a mean depth of 100 meters and it basically consists of a series of accelerating structures to boost the energy of the particles and a huge number of superconductive magnets, namely of three different kinds. There are 1232 dipole magnets which are 15 meters in length and weigh as much as 30 tons: they are used to bend the beam in the circumference. Moreover, there are 392 quadrupole magnets of about half the size that are required to focus the beam and a number of higher-order magnets that come into play right prior to the collisions and which squeeze the particles even closer together in order to maximize the chances of such collisions taking place.
The electromagnets are built from coils made of a copper-clad niobium-titanium (NbTi) alloy which operates in a superconductive state, that is conducting electricity without loss of current or energy. In order to keep this regime, as it is well-known, the magnets have to be kept at the extremely low temperature of 1.9 K (-271.3 °C). This is made possible thanks to a distribution system that pumps superfluid helium-4 in a closed circuit all around the accelerator, as it can be seen from the transversal section in Fig. (1.2).

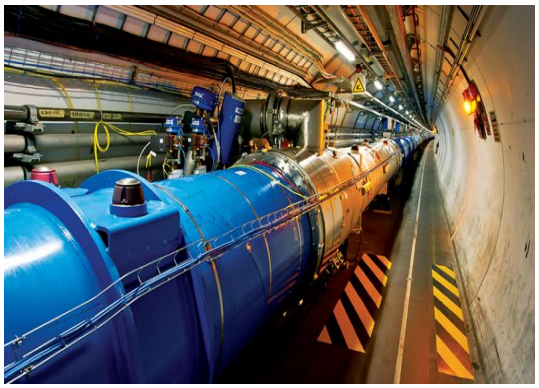

Figure 1.1: A portion of the LHC in its underground tunnel.

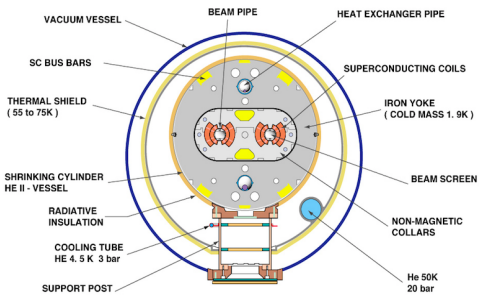

Figure 1.2: A transverse section of the LHC: the cryogenic system and beam pipes are shown (adapted from [1]).

Such low temperatures make it possible for the coils to withstand currents of about 11850 A, which are needed to reach magnetic fields as intense as 8.33 T.
In the accelerator, two high-energy particle beams travel in opposite directions at very close to the speed of light c inside two separate beam pipes. As it can be seen from the above-mentioned picture, these are two tubes 6.3 cm each in diameter running parallel to each other. It is crucial that these pipes are free from any possible obstacle that could disrupt the collisions between particle. Since also air molecules could be an impediment, an impressive array of vacuum technologies is in place, which makes the LHC the biggest operational vacuum system in the world. As a matter of fact, the LHC has three separate vacuum systems: one for the beam pipes, which serves the purpose to free the beams from collateral collisions with foreign particles; the other two act as thermal insulators both for the cryogenically cooled magnets and for the helium distribution line. The most severe vacuum conditions are used for the beam pipes, kept at the so called *ultra-high vacuum*: the pressure in there is of about 10^{-10} bar, thinner than interstellar void.

The particles are made to circulate in the LHC vacuum pipes not in a continuous flow but separated into bunches. However, the Large Hadron Collider is just the last step of this acceleration process. As it can be seen from Fig (1.3) at CERN there are a series of linear and circular accelerator. In the usual pp collisions, protons are taken from a bottle containing hydrogen by stripping the electrons from the atoms; after going through LINAC2 (a linear accelerator) protons are injected in different accelerators at increasingly high energy, following the path shown in figure.
Inside the Super Proton Synchrotron (SPS), the beams are accelerated up to 450 GeV and they are then injected into the LHC moving in both a clockwise and an anticlockwise direction.

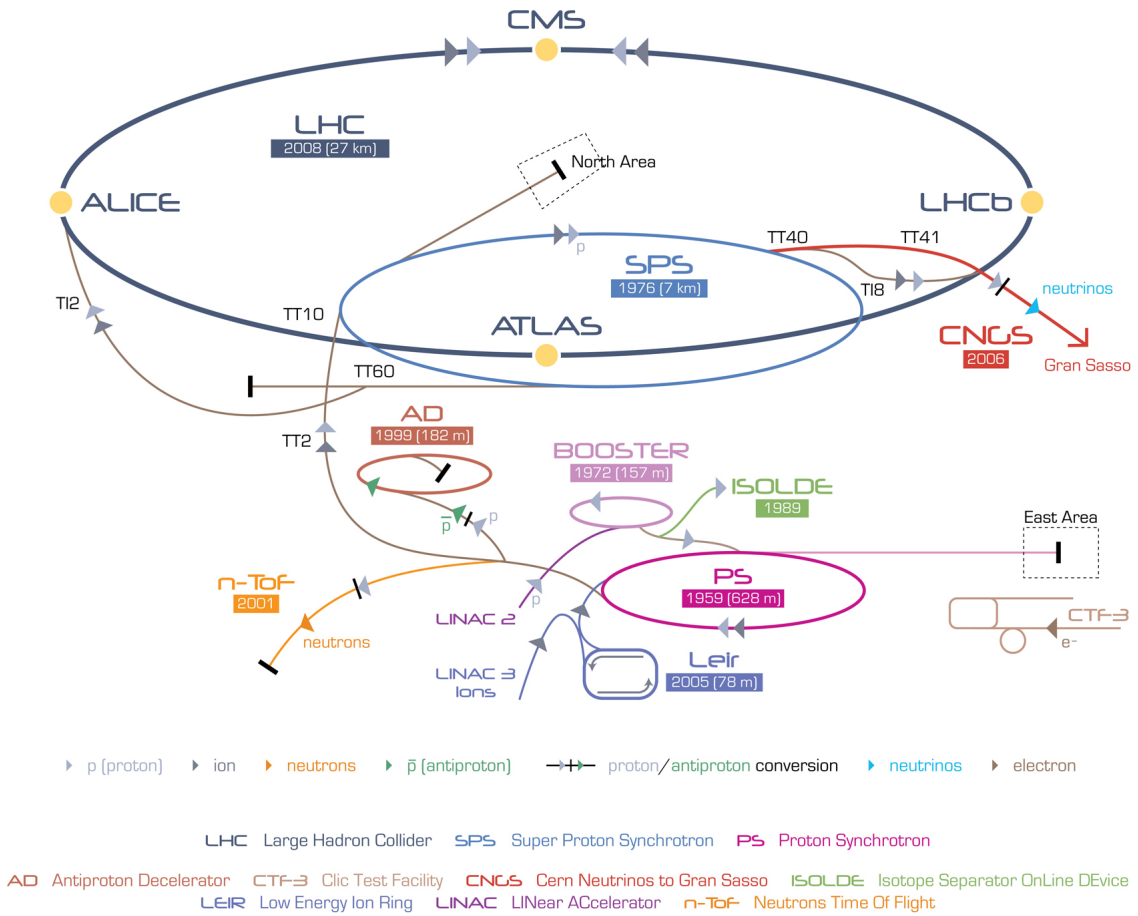

Figure 1.3: Infographic of the CERN complex, with the names, inauguration year and size of the accelerators and the position of the four main experiments around the LHC; the light grey arrows show the path followed by protons, starting from LINAC2.

In here, after about 20 minutes, they are accelerated at an energy of $E = 6.5$ TeV; given that they move in opposite directions, when they are forced to collide, the energy is actually double this value. This value is kept constant and uniform thanks to 16 radiofrequency cavities, dedicated chambers housed in four cylindrical refrigerators called *cryomodules* and containing an electromagnetic field. This field oscillates at the given frequency of 400 MHz and so protons with an excess or a deficiency of energy compared to the ideal value will either decelerate or accelerate slightly to match it; the particle beam is therefore arranged into discrete packets called *bunches*.
Inside the LHC the beam can circulate for several hours under normal operating conditions, until it degrades and it has to be dumped using a kicker magnet which pushes the beams off the accelerator onto a graphite block. During the fills, the beams keep rotating in the LHC and are brought into collisions inside the four large detectors located around the accelerator.

1.2 The Compact Muon Solenoid Experiment

The Compact Muon Solenoid (CMS) is an omni-purpose detector located in the area of Cessy, France, to first order diametrically opposite to the region where the bunches are first injected in the Large Hadron Collider. The construction of the site where it is located, the LHC P5, started in 1998 and it was carried on for six years, but the first beam only went through it in September 2008 and the first collisions were seen about a year later.

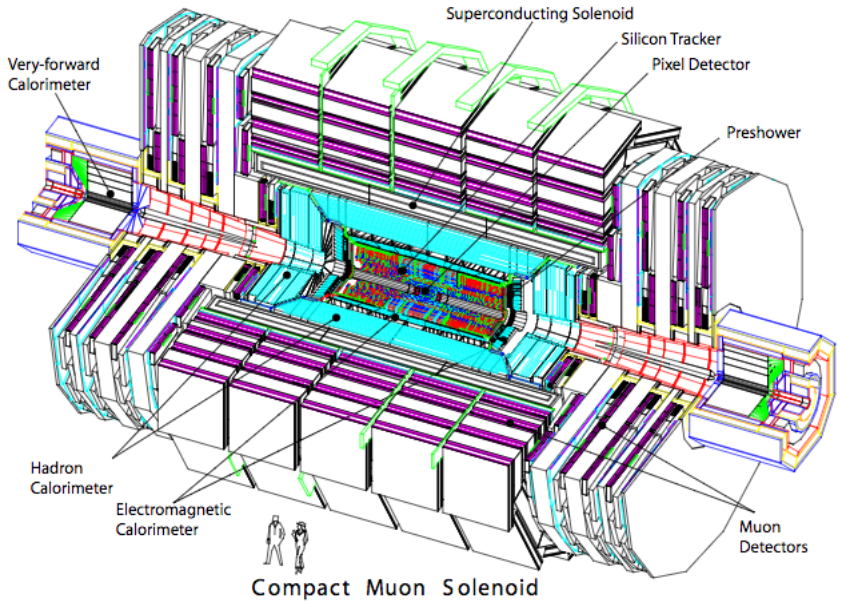

Figure 1.4: A perspective view of the CMS detector, where the two human figures can be used for scaling purposes (adapted from [1]).

The detector is 21.6 m long and 15 m in diameter, so its height is comparable to the one of a 5-story building. The whole structure weighs about 14,000 tons even if its volume is somehow very contained: to put things in perspective, the Eiffel Tower in Paris is twice as light as the CMS experiment while having 400 times its volume!
At the core of the machine the *interaction point* (IP) can be found: this is where the proton-proton collisions occur between the two counter-rotating beams in the LHC.

In order to increase the chances of producing and observing a rare particle, a large amount of collisions is required: the amount of raw data produced by each crossing is of about 1 MB and, given that the crossing rate is $\nu_{bx} = 40$ MHz, this adds up to the unmanageable amount of 1 TB of data per second.
A trigger system is therefore in place with the purpose of reducing this amount to about 1000 collisions per second, which can feasibly be stored and further processed later on.

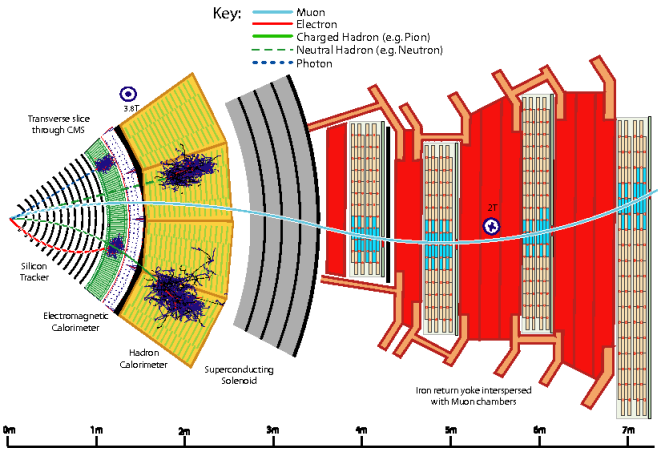

Figure 1.5: The figure shows the path and behaviour of different particles through the CMS detector (adapted from [2]).

Moving radially outwards from the IP, as it can be seen from Fig. (1.4), the Pixel Detector is found, which is useful for vertices reconstruction. Then, there are the tracker strips, made of silicon - a semiconductor - and coupled with different inclinations with respect to each other. Their goal is to measure the curvature of the collision products which come from the IP moving away from it in different directions.
A very important part of the CMS detector are its two calorimeters. The first one is the Electromagnetic Calorimeter (ECAL), which has the main purpose of stopping electrons and photons by creating showers in a transparent medium so that the least amount of energy is lost in the process thanks to its high efficiency. The second one is the Hadronic Calorimeter (HCAL), which is used to stop hadronic particles by the same principle. However, since these particles interact through collisions with the nuclei of the material, a stratified calorimeter is needed, which alternates a transparent layer (where efficiency is high and photons can be detected) with a heavy and opaque one (where hadrons can produce a shower and be slowed down, but in which portions of the energy are lost).
The actual detector is built around a huge solenoid magnet in the form of a cylindrical coil of superconducting cable. This can exert a magnetic field up to 4 T, which is confined by a steel yoke that forms the bulk of the detector mass.
The muon chamber has been positioned in the outermost layer of the machine, a natural design considering that muons are particles with an enormous penetration power and therefore among the hardest to detect. From the machine point of view, these chambers are found both on the barrel of the detector in the form of Drift Tubes (DT) or Resistive Plate Chambers (RPC), and in its endcaps, where again RPCs are implemented together with Cathode Strip Chambers (CSC). However, if the region is to be divided according to the event point of reference, there is also an overlap area, as collision products could potentially travel through both regions.

As already mentioned, CMS, just like ATLAS, is an *omni-purpose detector*: compared to its cousin it certainly has different technical solutions and a different detector magnet system, but the goals are almost the same and both are needed as a corroboration of each other's findings.
The main general goal of the CMS experiment is to explore subnuclear physics at the TeV scale, which has only become possible with such a big and powerful accelerator as LHC. In particular, the properties of the recently-found Higgs boson are being analyzed and studied through the observed (and not observed) decays of such particle. Heavy ions are also studied when, on occasion, special fills with either PbPb or, in the 2017 data taking period, XeXe collisions take place in the collider instead of the usual proton physics.
CMS is also experimentally studying the Physics BSM (beyond Standard Model), the ensemble of all the theoretical developments that have been done to better explain the deficiencies of the Standard Model of particles. Topics like these include the search for supersymmetry - also known as *SuSy* -, extra dimensions, dark matter and, perhaps more concretely from an experimental standpoint, the search for events with large amounts of missing transverse energy, which would imply the presence of particles capable of passing through the detector without leaving any trace, which the SM does not predict except for neutrinos. It is again crucial to stress the rarity of such events and, consequently, the high importance of collecting a very wide statistics while, at the same time, having in place an effective selection to discard those data deemed not worthy of further analysis.

1.3 CMS-LHC interface

In such a complex and sophisticated experimental environment, where hundreds of people are involved so that every part of the data taking process can be perfected as it should be, there needs to be a high level of reciprocal feedback between CMS, the experiment in which the collisions take place and where the data are collected and analyzed, and the LHC, the machine thanks to which the beams are able to circulate in the first place and that has control over the features of the beams themselves.
Specifically in matter of luminosity, CMS keeps track of both delivered and recorded luminosity. Delivered luminosity refers to the luminosity delivered to CMS by the LHC; ideally, the amount of luminosity recorded should be the same as the amount delivered, but in some cases the CMS detector is unable to take data, either because its data acquisition chain is busy or because one or more of its detector subsystems are temporarily unavailable. The recorded luminosity includes only the luminosity actually logged by CMS.
The LHC accelerator team, on its side, uses this feedback to maximize the luminosity being delivered to CMS and to optimize the overall performance of the machine. This information is also useful to adjust beam properties to accommodate request from experiments should they be willing to perform specific kinds of measurements or in need to test an apparatus.

Chapter 2

Luminosity measurements at CMS

2.1 Theoretical aspects of luminosity

BEAM energy and luminosity are the two most important parameters to describe the performance of particle accelerators.
High energy allows to probe smaller and smaller scales opening the way to the exploration of new fundamental phenomena; increasing the collision rate allows to increase the chances to produce and detect rare events.
Luminosity is the collider parameter that defines its capability of producing events in a given amount of time. If N is the total number of events, then we can define their rate as

$$R \coloneqq \dot{N} = \frac{dN}{dt} \tag{2.1.1}$$

and it holds true that

$$R = \sigma \mathscr{L} \tag{2.1.2}$$

so we have

$$\mathscr{L} = \frac{1}{\sigma}\frac{dN}{dt} \tag{2.1.3}$$

where $\mathscr{L}$ is the instantaneous luminosity and σ is the cross-section for a given physics process. As a reference, the total inelastic proton-proton cross-section at 13 TeV has a value of about 80 mb. The unit of measurement for instantaneous luminosity is $\mathrm{cm}^{-2}\mathrm{s}^{-1}$, which means it gives an understanding of the number of potential collisions that can take place per unit of area and time.
In a cyclical collider such as LHC, the average instantaneous luminosity of one pair of colliding bunches can be expressed through the general formula

$$\mathscr{L}_b = N_1 N_2 f_{orb} \sqrt{(\vec{v}_1 - \vec{v}_2)^2 - \frac{(\vec{v}_1 \times \vec{v}_2)^2}{c^2}} \int \varrho_1(x,y,z,t)\varrho_2(x,y,z,t)\,dx\,dy\,dz\,dt \tag{2.1.4}$$

where $f_{orb} = 11246$ Hz is the orbit frequency for the bunches inside the LHC, N_1 and N_2 are the numbers of protons in the two bunches - also known as *beam currents* -, ϱ_1 and ϱ_2 the particle densities for the two beams (normalized so that their individual integrals over all space are unitary) and $\vec{v}_1$ and $\vec{v}_2$ their respective velocities, grouped in what is commonly known as the kinematic factor K.
Now, we consider the case of a collision happening at $t_0 = 0$ and perfectly head-on ($\vec{v}_1 = -\vec{v}_2$), with bunches traveling almost at the speed of light so that $K = 2$; we also assume that all densities are uncorrelated in all planes and that the shape of the beams is perfectly gaussian. Keeping in mind that

$$\int_{-\infty}^{+\infty} e^{at^2}\, dt = \sqrt{\frac{\pi}{a}}, \qquad a \in \mathbb{R}^{\times} \tag{2.1.5}$$

then Eq (2.1.4) can be re-written in a more straight-forward way[1] as

$$\mathscr{L}_b = \frac{N_1 N_2 f_{orb}}{2\pi \Sigma_x \Sigma_y} \tag{2.1.6}$$

where Σ_x and Σ_y are the beam width parameters along the x and y axes (we use the CMS coordinate system).
This formula shows that luminosity may also be derived from LHC machine parameters alone, together with properties of the beams: the revolution frequency in a collider is accurately known and the number of particles is continuously measured with beam current transformers which should reach an accuracy of less than 1% for LHC nominal beam parameters.
Let now N be the number of protons in the two beams, which we now suppose having the same intensity; then, if n is the number of bunches the beams are divided into, the total instantaneous luminosity is

$$\mathscr{L} = n\mathscr{L}_b = \frac{nN^2 f_{orb}}{2\pi \Sigma_x \Sigma_y} \tag{2.1.7}$$

where we see that luminosity goes linearly with the number of bunches and quadratically with the number of particles in each bunch.
If we now consider Eq (2.1.2) and we contextualize it in the specific case of a luminometer, this will become

$$\tilde{R} = \mu f_{orb} = \sigma_{vis} \mathscr{L} \tag{2.1.8}$$

where μ is the average number of whatever observable the luminometer is based on (for example, hits, tracks, ...) and σ_{vis} is the fraction of the total inelastic cross-section of CMS (or similar) which is visible to the detector - hence the name *visible cross-section*. We indicate the rate with $\tilde{R}$ to distinguish it from the nominal rate from the previous equation.
The visible cross-section incorporates also a contribution of acceptance and emittance from the detector that is performing the measurements. This is why this quantity has to be as

[1] A full proof of this formula can be found in [3].

constant as possible even over a long period of time, and its calibration is of fundamental importance: we will delve into more details on the importance of the visible cross-section and how it is calibrated so that instantaneous luminosity measurements can concretely be performed in the next section.

2.2 Calibration of luminosity

2.2.1 Van der Meer scans

The Van der Meer scanning technique, pioneered by its namesake at the CERN Intersecting Storage Ring accelerator in the 1960s, involves scanning the LHC beams through one another to determine the size of the beams at their point of collision. The aim is to determine at the same time in a dedicated experimental setup both the event rate R_0 and the absolute luminosity $\mathscr{L}_0$ on the basis of the beam and optics parameters, so that

$$\frac{R_0}{\mathscr{L}_0} = \sigma_{vis} \tag{2.2.1}$$

The quantity σ_{vis} (the visible cross-section) is then used as a constant factor to obtain the instantaneous luminosity during the standard physics operations, so that

$$\mathscr{L}(t) = \frac{\tilde{R}(t)}{\sigma_{vis}} \tag{2.2.2}$$

which therefore means that if we want to compute the rate R of a given physics process, this will be given by

$$R(t) = \sigma \left(\frac{\tilde{R}(t)}{\sigma_{vis}} \right) \tag{2.2.3}$$

so by a physical component (the cross-section of the process) and by a term which is solely dependent on the detector.
The whole Van der Meer scanning process relies on the concept of Eq. (2.1.7). Such a scan requires particular operating conditions to be verified: there has to be low pile-up, to lessen side effects and have a lower background; the bunches have to be isolated, that is, with a big separation one from the other; the crossing angle is set to zero so that easier formulae can be used; finally, special optics have to be implemented to optimize measurements.
The two colliding beams are then separated by 6σ and scanned one across another in steps of 0.5σ, both in the x axis (*crossing scan*) and the y axis (*separation scan*) with the goal to determine their overlap region A_{eff}; the scan is usually repeated two or three times over the dedicated fill, sometimes moving both beams in tandem across one another, other times leaving either one stationary (*beam imaging*, which is not an actual VdM scan and it relies on integrating the beam that is moving over the separation in order to find the density of the other one and viceversa). The whole purpose of this scan procedure in the end is to determine σ_{vis} through Eq (2.1.8), a constant that is then to be used as a calibration for the whole year of data taking that follows.

The visible cross-section is therefore a quantity that must not have a dependency on time. Sources of uncertainty have to be carefully taken into account to provide with an estimate that is as accurate as possible, given its crucial importance; we can classify such sources in two categories[2].
The first set is associated with the detector operation. The main contribution in this sense is the cross-detector stability, which is due to the discrepancies between the luminometers and accounts for a 1.5% uncertainty; then, each luminometer suffers from nonlinear effects especially at high pileup or high luminosity (0.6%) and it unavoidably changes its performance over the course of the year due to long term effects (0.5%).
The other main sources of uncertainties are the ones related to the calibration procedure in and of itself. In this instance, the main contribution comes from the correlations between x and y. The basic VdM formula assumes that the beam is factorizable independently in these two directions, which is obviously not entirely true; to correct for this effect, the beam imaging scans are used, as reconstructed verteces positions can be employed to probe the density function of the stationary beam, and more sophisticated fits are applied (for example, a *super gaussian* which includes a negative term). This effect yields a contribution of 0.9%, but it is expected to decrease to as low as 0.5% with the newest analysis performed. Another major contribution comes from the length scale: the nominal amount that the beam is moved by in a scan may not be equal to the actual amount due to the degradation of the magnets at CMS. This contribution weighs about 0.8% and it is not directly related to the scanning technique itself, but it can be taken into account for by performing a *length scale scan* in which the beams are separated by 1σ and then stepped together at intervals of 1σ in order to measure the resulting beamspot position and fit. Other minor causes for uncertainty on this side can be the so called *dynamic* β^* due to the defocusing of the beams (0.5%), the beam-beam deflection that takes place since both the beams are positive in charge and therefore tend to repel each other (0.4%), measurements in the beam currents (0.3%), also entirely not related to luminosity.
All these effects amount for a total uncertainty of about 2.5% in the luminosity measurements coming just from the technique which allows its calibration [4].

[2]While obvious, it may be worth pointing out that the values for uncertainties that follow are related to the 2016 data taking, as a full analysis on the 2017 collisions is yet to be perfected and a constant effort to reduce these uncertainties (where possible) is in place.

2.2.2 Emittance scans

Emittance scans have the same underlying concept and aim of the VdM scans, but they are performed under normal beam conditions and they only last for a matter of minutes, not hours; again, the beams are scanned one across the other in optimized conditions. Emittance scans usually occur at the beginning and at the end of each LHC fill.

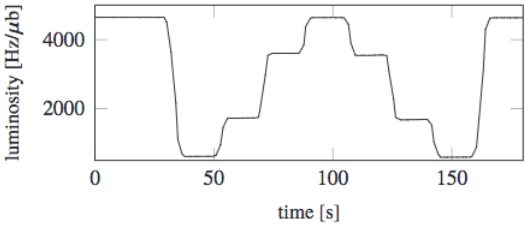

Figure 2.1: Luminosity evolution during a 7-point emittance scan.

This serves the purpose of being able to have daily estimations of the visible cross-section, to study its trend over different fills and over different conditions (i.e.: crossing angle) in the same fill and, most importantly, to take actions if an unexpected behaviour is seen from any luminometer in this quantity, which is supposed to remain as constant as possible over a very long period of time.

2.3 Luminosity at LHC

As mentioned, a higher luminosity means a higher number of events taking place which, in turn, means more data to be analyzed. To increase this value on an operational standpoint, there are basically two possibilities: more protons can be injected in the collider, as more particles obviously mean a higher chance of collisions, or the bunches can be squeezed as much as possible in order to increase the density of the particles in them. At the full 2017 regime, the filling scheme in the LHC was made such that 2556 bunches of protons could circulate in each ring inside the machine; each bunch is 25 ns apart from the other, which corresponds roughly to 7 meters, and it consists of some 100 billion protons. This adds up to an instantaneous luminosity that reached peaks as high as $2.05 \cdot 10^{34}$ cm^{-2}s^{-1} at stable beams (as of 2 Nov. 2017).
Given that, in non-SI units,

$$1 \ \frac{\text{Hz}}{\mu\text{b}} \sim 10^{30} \ \text{cm}^{-2}\text{s}^{-1}$$

it is more common to use the SBIL (Single Bunch Instantaneous Luminosity) to measure this quantity for individual bunches circulating in the machine. On average, values of SBIL range from 4 to 6 $\frac{\text{Hz}}{\mu\text{b}}$.
Another important quantity to be defined is the integrated luminosity $\mathscr{L}_{int}$. If we integrate Eq. (2.1.2) over time, we get the mathematical relationship

$$N = \sigma \mathscr{L}_{int} \tag{2.3.1}$$

This means that the integrated luminosity gives a measure of the total number of collisions that have happened over the integrated time interval, which can naturally be converted into an estimate of the amount of data made available by the accelerator. As it can be seen from Fig. (2.2), over the years the LHC has stepped up its integrated luminosity by a lot, which translates into a considerably higher amount of data.

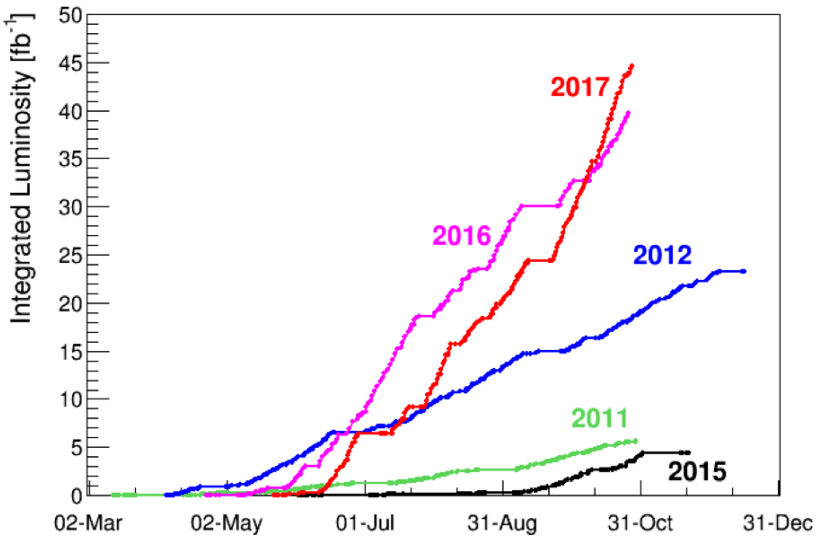

Figure 2.2: Multi-annual overview of the integrated delivered luminosity evolution. The 2017 line shows not only that the 45 fb^{-1} target was reached and surpassed but also the fastest accumulation of all years.

It can be easily seen how LHC has increased the amount of integrated luminosity in its developments. The 2017 integrated luminosity for the 13-TeV *pp* collisions has surpassed the one in 2015 and 2016 combined long before the year was over, with the CMS experiment falling just short of the 50 fb^{-1} milestone, which ATLAS managed to reach. To put things in perspective, 1 fb^{-1} of integrated luminosity roughly corresponds to $8 \cdot 10^{13} = 80$ million million collisions!

In luminosity measurements, ad-hoc units of time are used. It has already been mentioned that each bunch is separated from the other with a 25 ns time interval. This means that the time to complete one *orbit* is

$$1 \text{ orbit} = 3564 \cdot 25 \text{ ns} = 89 \ \mu\text{s} \tag{2.3.2}$$

where 3564 is the total number of possible configurations, both with beam and empty. We define a *lumi nibble* (NB) as the time it takes to complete 4096 orbits, so

$$1 \text{ NB} = 4096 \text{ orbits} = 0.36 \text{ s} \tag{2.3.3}$$

This corresponds to the readout frequency of the PLT Front End Device (FED), but since the readout of the BRILDAQ (the data acquisition for all luminosity-related data) is four times this quantity, the *nibble-4* (NB4) is more commonly used, such that

$$1 \text{ NB4} = 4 \text{ NB} = 1.4 \text{ s} \tag{2.3.4}$$

Finally, a *lumi section* is defined as

$$1 \text{ LS} = 64 \text{ NB} = 23 \text{ s} \tag{2.3.5}$$

and it corresponds to the granularity (that is, the scale or level of detail in a set of data) of the luminosity database.

2.4 Luminometers at CMS

The importance of luminosity as a fundamental parameter for an accelerator has already been widely discussed. It should therefore come as no surprise that the CMS experiment has a dedicated subgroup for luminosity analysis and for the maintenance of the detectors in charge of delivering such measurements, called *luminometers*.
We can distinguish two kinds of analysis detector: existing analysis detectors, which are very well understood yet in a way limited by the centralized DAQ (Data Aquisition) and dedicated analysis detectors, which are specifically optimized for luminosity measures and have a dedicated readout, while at the same time being, however, limited in size.
Let us now focus on the so called *zero-counting method* which is in place to give an estimate of μ, the number of measured tracks or hits. The probability $\mathcal{P}$ of a detector measuring k hits follows the Poisson distribution

$$\mathcal{P}(k) = \frac{\mu^k}{k!} e^{-\mu} \tag{2.4.1}$$

so, if we indicate with f_0 the fraction of *zeroes*, scopes with no triple coincidence, then the Poisson distribution will give us

$$f_0 = \mathcal{P}(0) = e^{-\mu} \tag{2.4.2}$$

which implies

$$\mu = -\ln f_0 \tag{2.4.3}$$

and the luminosity will be proportional to μ, with the calibration constant to be determined using the Van der Meer scan technique as explained in subsection 2.2.1.
The reason why zeroes are used instead of hits greatly depends on the detector. For the PLT this is mostly due to technical limitations of readouts, which are not capable of reading any more than 2 hits (so measuring triple coincidences would be impossible), and

to their recovery time which could prevent a good reading of a track that happens to be in the same place as the immediately previous one. The drawback is that, when the rate is really high, f_0 tends towards zero and it can be seen that Eq (2.4.3) diverges to $+\infty$.
The zero-counting method is still susceptible of wrong estimations. Possible causes for undercount are the dynamic inefficiencies that could take place, as an example, if there are two consecutive hits and the detector reads poorly the latter, or also inefficiencies in reconstruction. On the other hand, causes that can lead to overcount are *afterglow*, which happens if a hit is so intense that it can have effects on the reading of the following ones; any kind of machine-induced background, such as *beam halo*; last but not least, time dependant effects, which usually are the most difficult to take into account since they are mostly related to the materials the detectors are made of and they take place over a considerably longer timespan.
The analysis of the luminosity in and of itself happens in two distinctive ways. It can be an *online* analysis, done in real time, on a bunch by bunch basis and granting a high precision in short time; or it could be *offline*, with raw data that have already been collected, in which case the precision is usually much higher but the analysis is done per fill over an extended period of time. We will now outline the luminometers and the analysis tools available at the CMS experiment, distinguishing them between online and offline.

- Online analysis
  - **Pixel Luminosity Telescope** (PLT), which is made of silicon sensors on 3 planes and it is based on the concept of triple coincidence (it will be discussed in detail in Chapter 3);
  - **Hadronic Forward** (HF) calorimeters, which collect data at a very high rate with a 1% statistical accuracy thanks to special electronics that read signals as they come off the detector;
  - **Beam Control Monitor - Fast** (BCM1F), which consists of 4 half-rings installed on both sides of the IP and implementing both monocrystaline and polycrystaline diamond sensors with fast readout.
- Offline analysis
  - **Drift Tubes** (DT), which uses the rates of muon trigger primitives in the drift tubes;
  - **Pixel Cluster Counting** (PCC), based upon the Silicon Pixel detector where luminosity is evaluated from the cluster rates occurring on average in a zero-bias event (i.e.: an event triggered by requiring only that two bunches cross at the CMS IP).

In the following graph we can see a plot of the instantaneous luminosity per LS as measured in real time by the three online luminometers mentioned above; in the caption a brief description of the most notable common features of the plot can be found.

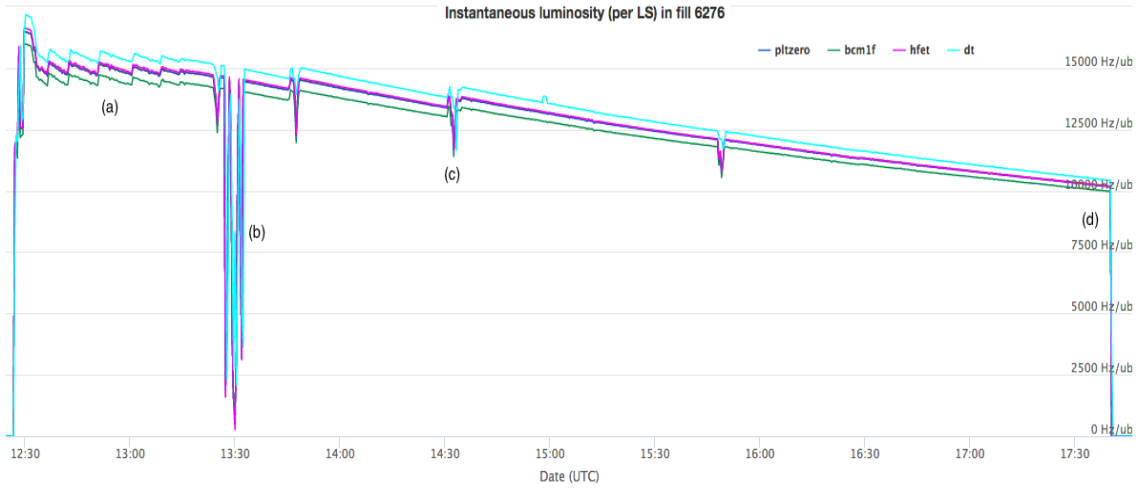

Figure 2.3: Plot of the instantaneous luminosity per LS as measured by the three CMS online detectors (the DTs are also included). In the plot we can recognize: (a) the *lumi leveling*, when the beams are displaced with respect to each other in order to adjust the luminosity values at the desired target; (b) an *emittance scan*, which is performed at the beginning and (not in this case) at the end of the fill in order to do a routinely calibration of absolute luminosity (see subsection 2.2.2); wiggles with an upward step like the one in (c), shared among all luminometers, could indicate a change in the crossing angle[3]; (d) beam dump.

[3]The crossing angle is necessary in order to prevent encounters in the region where the two beams share the same vacuum chamber and also to minimize the effects of long-distance interaction as much as possible. Since a large crossing angle decreases the luminosity, as it decreases the overlap area of the bunches, reducing the crossing angle over the course of a fill allows the recovery of some of the total potential luminosity (an estimated 5%) that gets lost because the beams do not collide head-on.

Chapter 3

The CMS Pixel Luminosity Detector

3.1 General overview

BEFORE the end of Long Shutdown 1 (January 2015), the first major maintenance and upgrade period for the LHC and its experiments, the Pixel Luminosity Telescope (PLT) was installed in the CMS experiment as part of the Run 2 upgrades for the CMS BRIL project.
Positioned close to the beam pipe, directly behind the Forward Pixel detector, its purpose is to measure the instantaneous luminosity at the highest energies and highest collision rates foreseen at the LHC using silicon pixel sensors. The PLT is the only subdetector in CMS whose sole function is to measure the delivered instantaneous luminosity, which is capable of doing with excellent statistical precision and in real time.

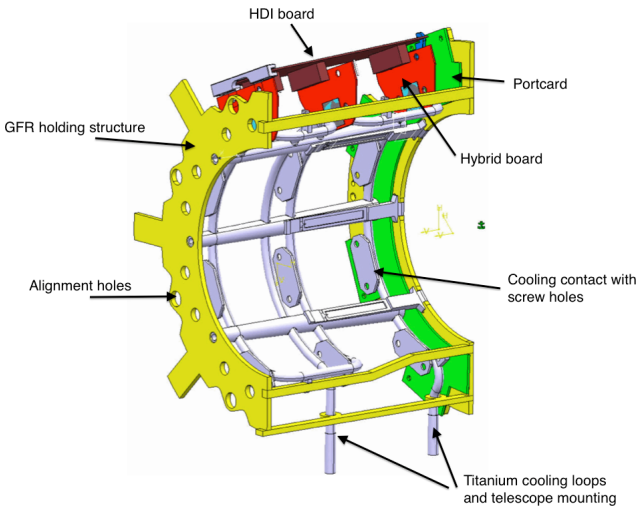

Figure 3.1: An outline of one quarter of the PLT detector with its components pointed out by the arrows (adapted from [7]).

The high precision is essential since the luminosity is a key quantity for many physics measurements at the LHC, where data collected under different beam conditions have to be combined; the latter feature can be explained by the fact that luminosity measurements need to be directly available to the LHC operators in order to optimize the beam conditions for the experiment (as mentioned in Section 1.3).

3.2 Description of the detector

The Pixel Luminosity Telescope is located 1.75 m from the interaction point at both ends of CMS, covering a pseudorapidity[4] of $|\eta| \approx 4.2$. The detector is separated into quarters (also called *cassettes*), where each quarter is a separate structure that houses 4 beam telescopes placed in half a circle around the beam pipe. There are therefore a total of 16 telescopes, 8 at each end, installed outside of the CMS pixel endcaps and referred to as *channels* with a numbering from 0 to 16. It should be noted right away that the two telescopes corresponding to numbers 0 and 4 stopped responding during the 2017 data taking and were dropped from luminosity calibrations.
Each telescope is composed of 3 individual sensor planes. They are mounted in the xy plane, with the imaginary line connecting the center of each sensor parallel to the z axis, for a total length of about 7.5 cm - always with respect to the CMS coordinate system. Each sensor has 80 rows and 52 columns of pixels and it is 150 μm wide $\times$ 100 μm high, for a total active area of 8×8 mm and with a depletion depth of 285 μm. However, in order to decrease the contribution from combinatory effects from fake triple coincidences, the active area of each sensor is decreased by masking out the outer pixels. The center plane in each telescope was reduced to an active area of 4.2×4.0 mm (28 columns $\times$ 40 rows), while the outer two planes to 5.4×5.2 mm (36 columns $\times$ 52 rows); in particular, the active area of the outer sensor planes was displaced such that they are aligned towards the IP and the alignment is adjusted to optimize the rate while at the same time reducing background tracks (*accidentals*) [6, 7].

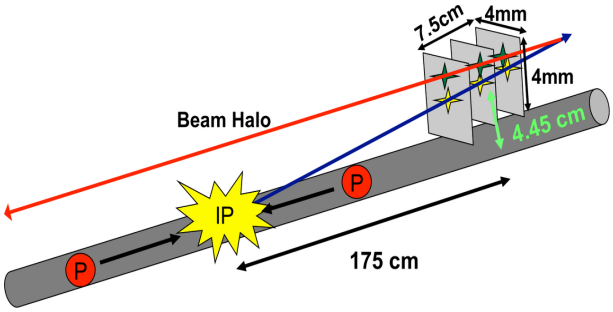

Figure 3.2: The drawing simplifies how a *pp* collision is detected by the PLT sensor; distances are included for completeness (adapted from [7]).

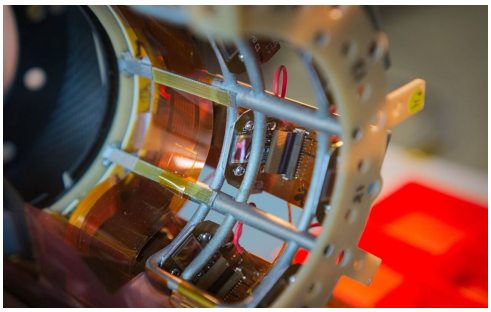

Figure 3.3: A picture of a quadrant of the PLT: the Ti cooling structure and the sensors can be clearly seen.

[4] In experimental particle physics, the *pseudorapidity* η is a spatial coordinate which describes the angle of a particle relative to the beam axis and it is defined as

$$\eta \equiv -\ln\left(\tan\frac{\theta}{2}\right) \tag{3.2.1}$$

where θ is the angle between the particle 3-momentum $\vec{p}$ and the positive direction of the beam axis. As it can be seen from Eq. (3.2.1), η decreases the more θ increases and it is 0 for a particle which is traveling orthogonally with respect to the beam axis. Particles with high pseudorapidity values, on the other hand, have a really small θ and thus are generally lost, escaping through the space in the detector along with the beam.

3.3 The readout process

The PLT uses much of the same technology as the CMS pixel detector, including the sensors and readout chips, but it also takes advantage of a special fast-or readout mode in the readout chips, not used in the pixel detector. The following schematic summarizes the paths taken by the different signals and it should always be kept as a reference throughout this Section.

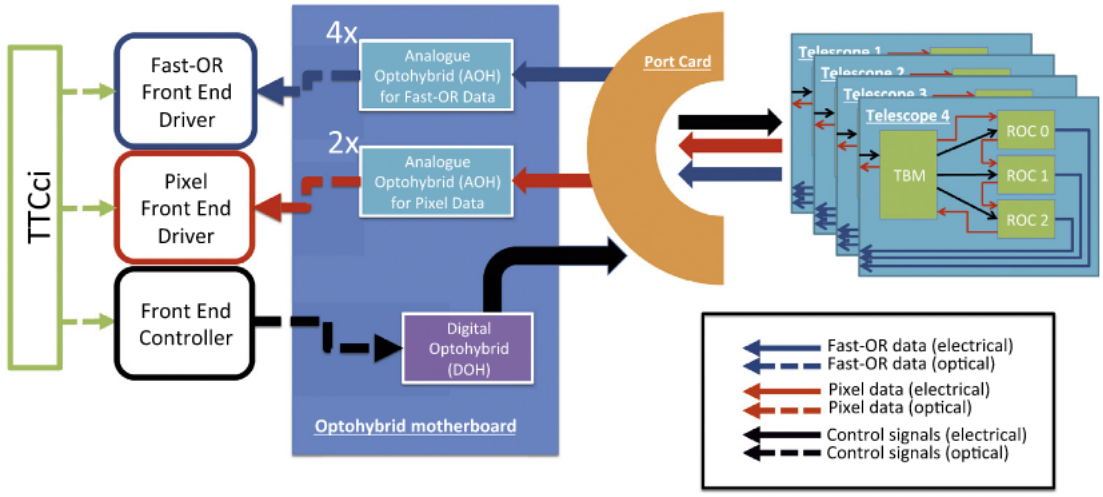

Figure 3.4: A schematic of the control and readout logic of a single PLT quarter with its four telescopes (adapted from [8]).

The PLT sensors are read out by a PSI46v2 readout chip (ROC), which provides data in two different formats:

1. **Pixel data** (offline measurements): if a pixel is hit by a charge which is higher than a programmed threshold, a *pixel hit* is created, which carries information about the pulse height, the pixel(s) address(es) and a timestamp, all of which are then conveniently saved in a buffer. The pixel data are usable to reconstruct tracks, which is helpful to distinguish particles coming from the IP from the beam halo, a powerful way to determine the efficiencies of the pixels and other systematic corrections. Because the full pixel data cannot be read out at the full BX rate, a trigger must be put in place at a lower rate: to avoid systematic uncertainties, a purely random trigger was used which selected indiscriminately any bunch crossing with a specified rate of 2 kHz.

2. **Fast-OR data** (online measurements): the ROC records hit position and amount of charge deposited on a sensor with a time resolution of 25 ns (40 MHz readout rate); such signal is suppressed if the charge is lower than a certain threshold, which can be set globally and adjusted for each individual pixel. In this case the ROC sets

a fast-OR signal - a differential analog signal generated by each ROC - with a pulse height proportional to the number of double columns hit in a specific bunch crossing. The fast-OR FED then collects these signals; since according to the zero-counting method the only information needed is whether the number of hits is zero or nonzero, the FED works in *digital mode* and it does not have to distinguish how many pixels were actually hit. Then, the driver counts the triple coincidences between the three planes of a given telescope and this can be translated into a luminosity value for every bunch crossing.

The three ROCs from each telescope are connected to an HDI (High Density Interconnect) card; this contains a Token Bit Manager (TBM) chip which handles the readout of the series of ROCs. Four telescopes, one per quadrant, are connected to a port card that manages the control signals for each cassette. This port card is itself connected to the Opto-hybrid Motherboard (OMB), whose function is to convert the electric signals into optical signals which can then be carried by fibers into the service caverns of CMS, where the data are collected and stored.
The front-end electronics consist of a Front-End Controller (FEC) card which issues commands to the ROCs, and the three already mentioned FEDs, among which one is in charge of reading out the pixel data from the ROCs, decoding them and writing them in the appropriate location, while the other two (one at each end) histogram events with triple coincidences per each BX - that is, when three planes in a telescope register a hit.

3.4 The cooling system

The silicon sensors of the PLT need to be cooled mainly for three reasons: to create a stable operating temperature for the detector; to reduce the leakage current, which is strongly temperature-dependent for silicon detectors and whose degrading effects can be reduced greatly even with non-extreme freezing conditions; finally, to prevent a thermal runaway or overheating of the sensors themselves after irradiation.
The PLT cooling and mechanical structure was manufactured in a novel production process called Selective Laser Melting (SLM), which uses a high power density laser to melt and fuse metallic powders together. This process allows to create metal structures with complex geometries, thin walls and hidden voids or channels granting great precision; it is also perfectly suited for small and highly customizable production batches. The cooling structure was 3D printed as a single piece of titanium alloy and built as a meandering tube of 2.8 mm in diameter. The coolant liquid (C_6F_{14}) comes directly from the CMS Tracker cooling plant at a temperature of -15 °C; not only does it reach the sensors and the readout chips, but also the Opto-hybrid motherboard electronics.
A temperature monitoring system is in place with three different sensors for each quadrant of the PLT, in order to prevent overheating and, if necessary, to study trends during a fill or throughout a more extended period of time, which is usually when radiation damage can be picked up on [8].

Chapter 4

Analysis of the PLT performance

4.1 Drop in σ_{vis}

4.1.1 Introduction

As already mentioned in subsection 2.2.2, emittance scans are used on a regular basis to check the trend in σ_{vis} to make sure that there is not a significant deviation from its regularity. In fact, this could be a proof of some efficiency problem in the detector, or in the fitting procedure or in some other aspect, and it is therefore a fundamental feedback for the BRIL group.
On 11 September, an ongoing downward trend was spotted in the recent fills that had been taking place at LHC, as it can be clearly seen from Fig. (4.1), which shows the average trend of this quantity for the PLT detector as taken from the BRIL centralized Webmonitor tool. The visible cross-section decreases by slightly more than 1% seemingly around fill 6171 and it does so progressively over about ten fills.

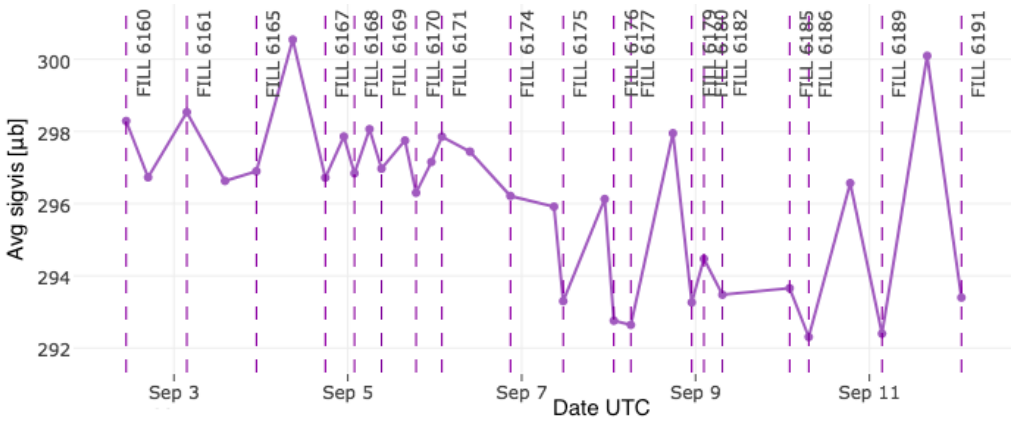

Figure 4.1: The plot shows the trend of the PLT average σ_{vis} in about 30 fills. It is clear how, starting from fill 6171, a downward trend starts to appear.

Such a trend being a possible indication of a loss of efficiency in PLT, perhaps due to radiation damage in some detector components, or a symptom of a fitting problem, a per-channel analysis was performed to give a broader insight on the issue.

4.1.2 Per-channel σ_{vis} trend

First off, the values of σ_{vis} were checked as calculated by each telescope in the PLT. This can already give some important pieces of information about the possible cause of an unexpected trend: if such trend is equally visible to all the channels, then it is less likely to be due to radiation damage, which would reasonably affect the telescopes in different ways according to their position with respect to the irradiating beam.
The plots that follow show the individual values for the visible cross-section in each channel over the same timespan as Fig. (4.1). The channels are grouped to optimize the scale of the plots so that also smaller changes can be spotted more easily; for the same reason some of them had to be plotted individually. It should be noted that channel 10 is not present because of a software glitch that gave unreasonable values of σ_{vis} in a couple of the fills of interest, values which highly misrepresented the plots; also, it is reminded that channels 0 and 4 are the non-working ones: channel 0 is absent from the Webmonitor, while channel 4 is still plotted for completeness even if, not being part of the luminosity offline calibrations, its contribution is of little relevance anyway.

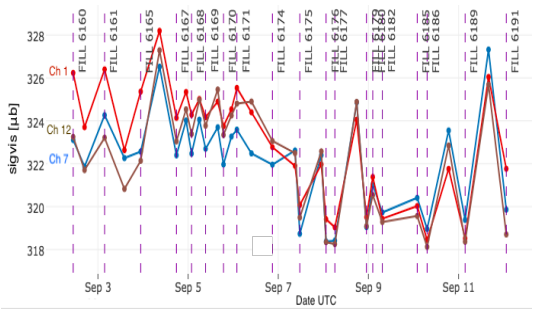

(a) Channels 1-12-7

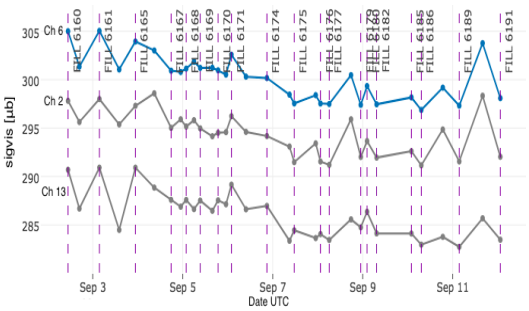

(b) Channels 6-2-13

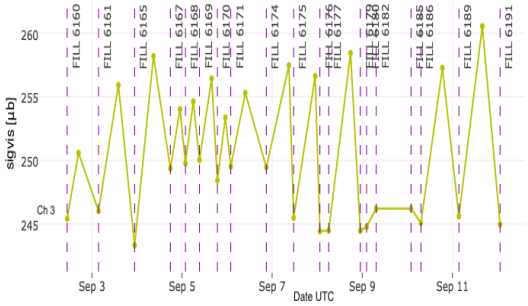

(c) Channel 3

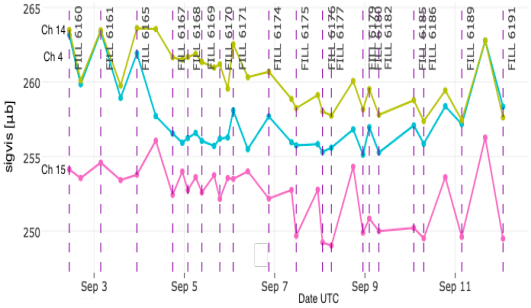

(d) Channels 14-4-15

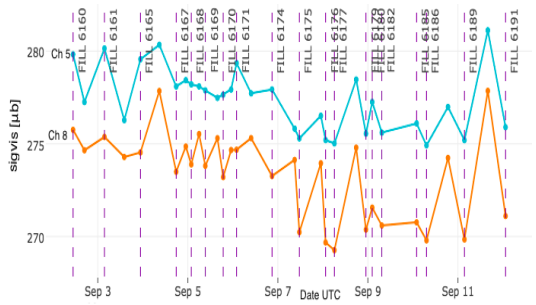

(e) Channels 5-8

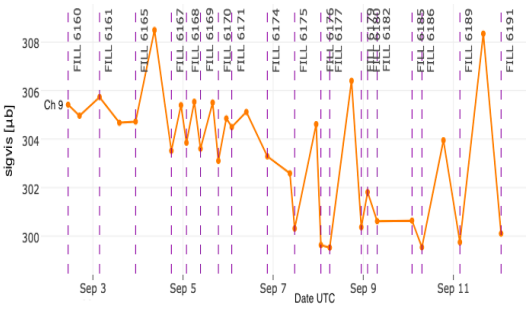

(f) Channel 9

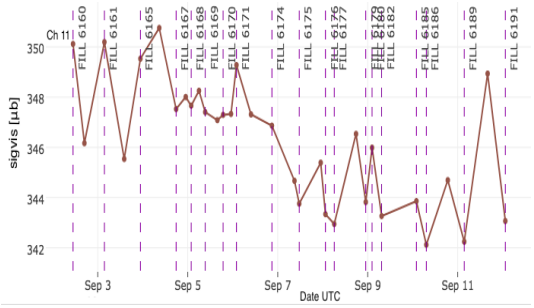

(g) Channel 11

Figure 4.2: Per-channel σ_{vis} plots over the period of interest for the analysis; the channel number is indicated from top to bottom and can also be found on the left-hand side of the plots.

From these plots, it is clear that the values of the visible cross-sections seem to be dropping in pretty much all the channels, by a commensurable amount and always starting at around fill 6171. Some channels do appear to have a drop earlier than that, at fill 6165, but it is commonly followed by a series of fills where the visible cross-sections showed great stability and it was therefore deemed not significant.
Channel 3 in Fig. (4.2-(c)) represents the only exception with its erratic behaviour, but this telescope was already known to be unreliable during the 2017 data taking and analyses to decide its exclusion from the overall luminosity calibration are still ongoing.
The uniformity of the trend can be better appreciated in the plot found on the following page, where all the σ_{vis}, both the individual ones and the average, are normalized with their values at fill 6160, the first one in the period considered. Given that it is the furthest in time from the drop and that nothing worrisome was spotted earlier than that, it is a safe assumption to use it as a reference and, by normalizing the other fills to it, to have a finer scaling so as to better estimate the entity of the drop itself.

In the plots are excluded: channels 0 and 4 because non-working; channel 10 because of the mentioned spikes; channel 3 because its behaviour would have unnecessarily altered the scale, which is wanted to be as precise as possible not to miss possible details that can slip an approximate analysis.

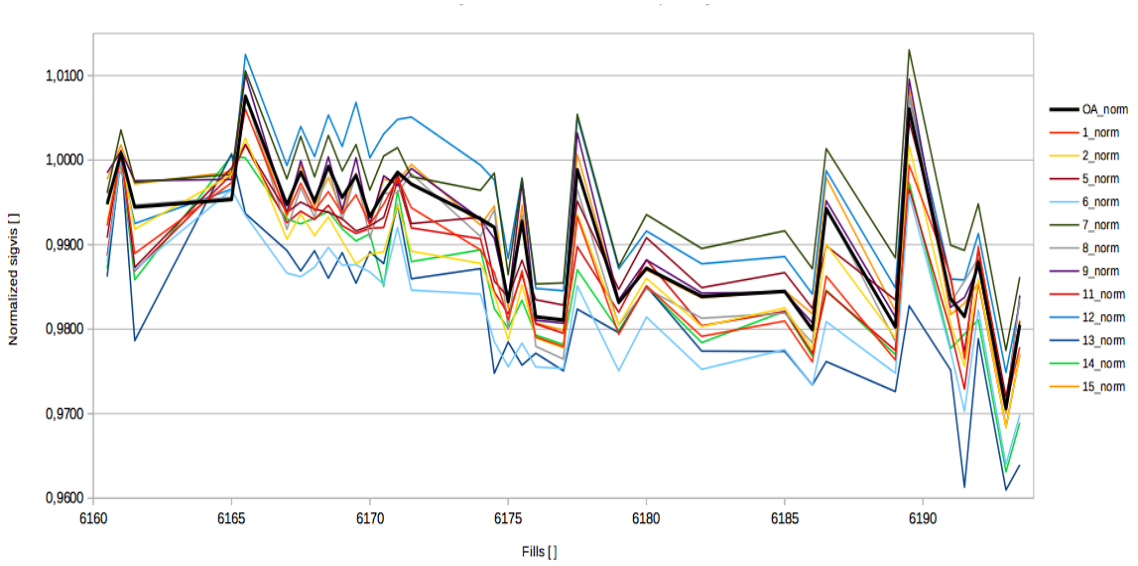

Figure 4.3: Values of the per channel and average visible cross-section normalized with the respective values measured at fill 6160.

It is also useful to check the deviation that each channel shows (if any) with respect to the average: again, this will tell wether there are channels that are affecting the average more substantially than others. To do so, values of the per-channel visible cross-sections were normalized this time with the average σ_{vis} plotted in Fig. (4.1) on a per-fill basis; again, channels 0, 4 and 10 were excluded, but this time also channel 3 can be seen in the first plot, as a complementary proof of its lack of stability.
The plotted ratios for the channels which are operating under normal conditions all appear constant at least on the per-cent scale. There is no significative deviation between these, which means there is no particular channel which is behaving as an outlier and skewing the overall calibration constant.

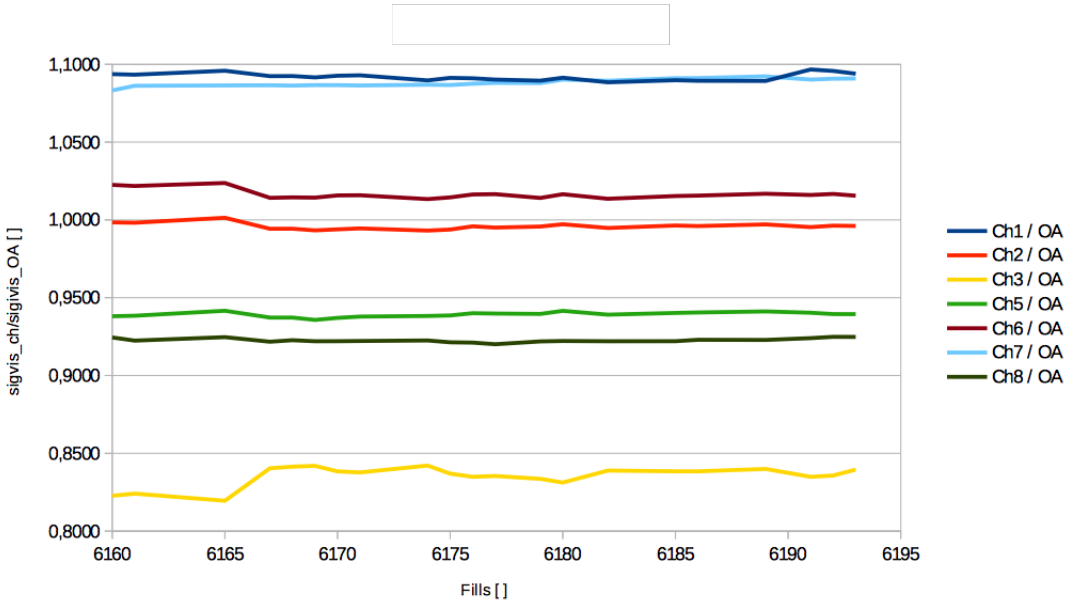

(a) Channels 1-2-3-5-6-7-8

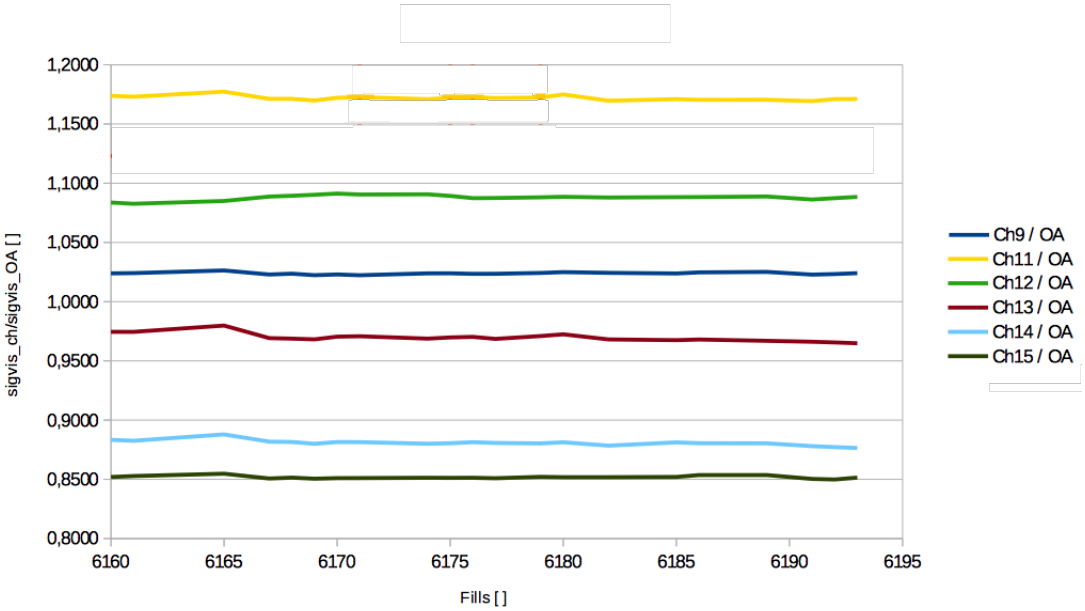

(b) Channels 9-11-12-13-14-15

Figure 4.4: Plots of the per-channel σ_{vis} normalized with respect to the average for each fill.

4.1.3 Per-channel efficiencies trend

Efficiencies in PLT are measured for each telescope, where the telescope efficiency is taken as the product of the individual efficiencies of the three planes. To measure the efficiency of a given plane, two-hit tracks are tagged using the two remaining planes; the number of tagged two-hit tracks is labeled $N_{fiducial}$. Then, for each tagged track, an automated script checks whether there is a hit in the probed plane (the one for which the efficiency measurement has to be made) consistent with the track. The number of tracks which satisfy this condition is labeled as $N^{hit}_{fiducial}$ and the efficiency is taken to be

$$\varepsilon = \frac{N^{hit}_{fiducial}}{N_{fiducial}} \tag{4.1.1}$$

It should be noted that the efficiency online-plotting tool is still at a preliminary stage of its development, which justifies the numerous glitches found in the plots as well as the lack of the last four channels, corresponding to one of the PLT quadrant which had not been implemented at that time of the data taking year.
Anyways, as far as it can be seen from the portions of the plots where physically-sensible data are collected, the efficiencies show no change whatsoever in their behaviour over the same timespan as the visible cross-section. This is another strong element which tends to rule out a damage due to beam radiation, which would obviously have a noticeable impact on at least some of the channels of the detector.

4.1.4 Conclusion

From this analysis, it was shown that the drop in σ_{vis} which took place after fill 6171 is likely due to some fitting issue rather than to an actual loss in the PLT efficiency. This conclusion can be drawn mostly for two reasons: first off, all the channels except one were shown to have undergone the same trend over the same period of time, while a radiation damage would have been expected to cause a nonlinear and inhomogeneous behaviour among the channels; secondly, the efficiencies, even if plotted with a preliminary tool, showed no appreciable variation in the critical fills.
For completeness sake, it is noted that further analysis, conducted with more data and with different and way more sophisticated kinds of measurements, did show that the PLT had indeed been suffering by a slight degradation due to radiation damage over the year. No hardware intervention has been planned so far, but studies are still ongoing.

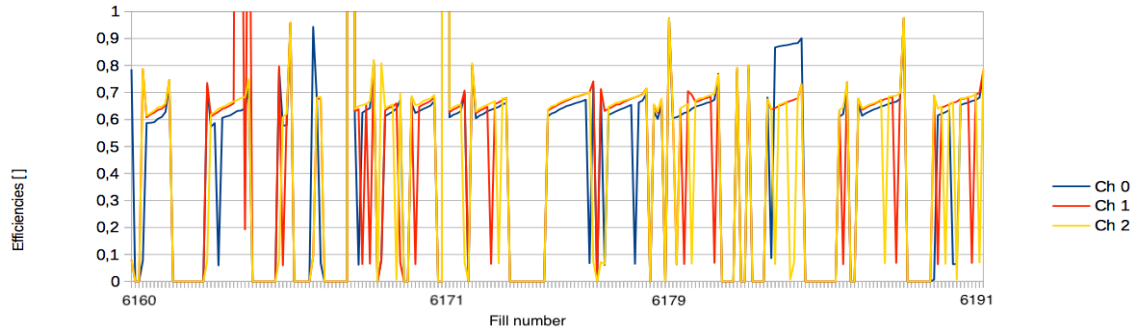

(a) Channels 0-1-2

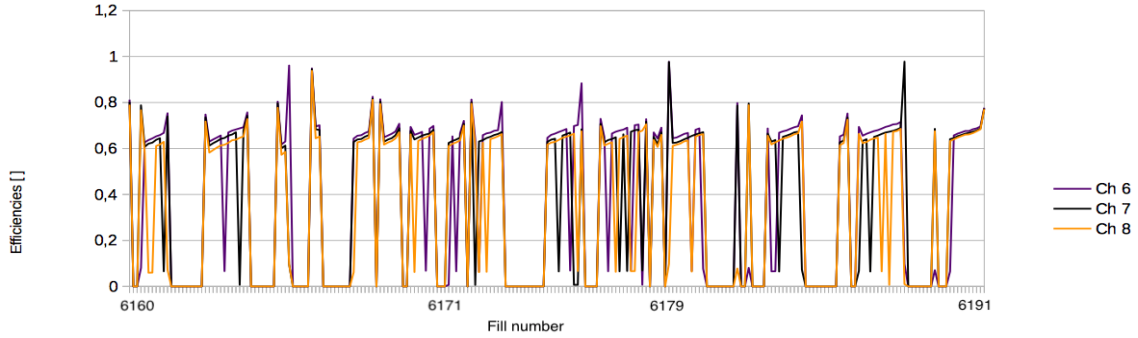

(c) Channel 6-7-8

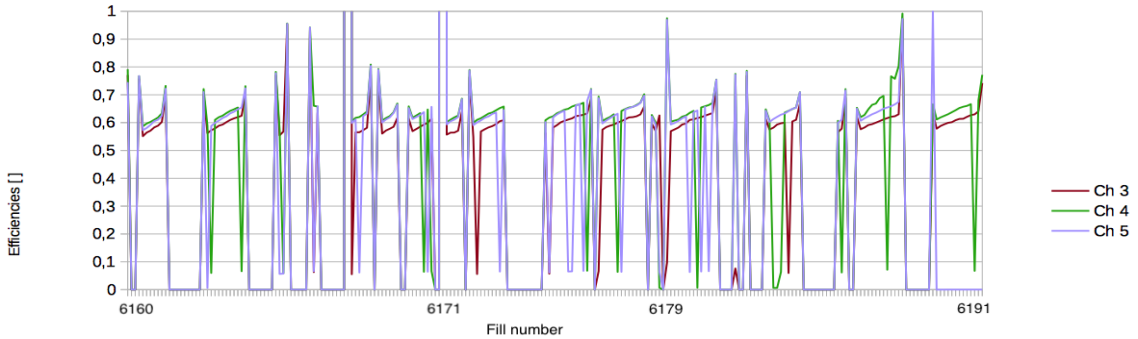

(b) Channels 3-4-5

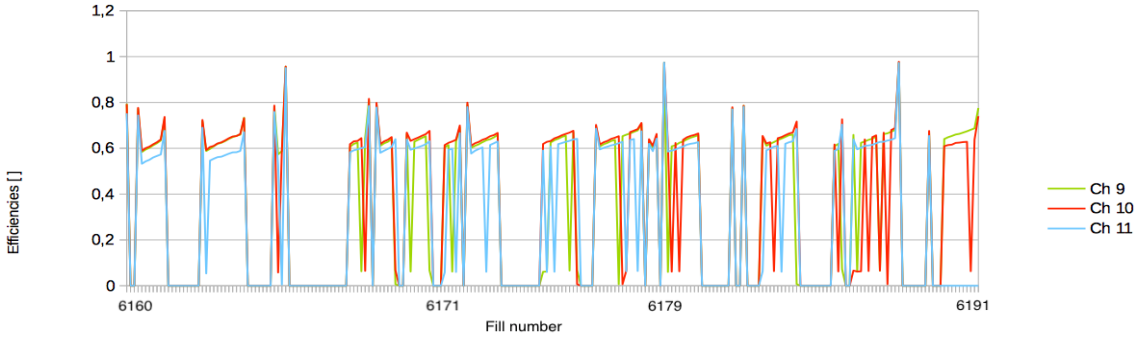

(d) Channels 9-10-11

Figure 4.5: Efficiency plots for the available channels from the PLTSlink machine.

4.2 Corrections in the PLT baseline

4.2.1 Introduction

An important part of the calibration of the PLT detector is to make sure that all pixel addresses and the header and trailer information from the TBM are decoded correctly and that every hit is associated to the correct ROC.
In the Front End Driver the incoming optical signal from the Opto-hybrid motherboard is sampled at a constant rate. When no information is transmitted the baseline value is supposed to be relayed by the analog opto-hybrid (AOH), which converts the electronic signals from the TBM/ROC into the optical signals sent from the detector. However, due to temperature changes in the AOH or to fluctuations in the cooling liquid temperature, after turning on the detector the baseline can drift since the laser output drivers are temperature-dependent. This baseline drift is constantly monitored and adjusted for during the PLT operation.
The fast-OR baseline correction works by defining a target baseline level, usually 175 analog-digital converter (ADC) counts and a bunch crossing (BX) number selected to be in the abort gap, where there should be no beam present and so no signal. The FED then looks at the signal level coming from the AOH during the BX (actually an average of the four BXes starting with the target one) which should always be at the baseline level, and compares it to the target level. If the measured level is different, then it will adjust it by ± 1 until they agree. Before actually perfecting the adjustment, some number of measurements above or below have to be made. The Pixel baseline correction is even simpler, since there the TBM output always includes some baseline level before the header and after the trailer, so the FED can just use those to obtain the measured baseline level. During stable beams, the work-loop automatically takes care of such corrections unless they go over a certain threshold (150 ADC counts for Pixel, 175 for FastOr) [9].

4.2.2 Baseline corrections plots

The PLT automated work-loop saves the corrections that are applied to the baseline in dedicated files. Through a script, these corrections were collected, divided into channels and then plotted, distinguishing them between the Pixel correction and the ones that come from the three FED channels in each telescope.
The plots show an unexpected behaviour. While the functional form of the corrections has the usual shape, their values are much higher than average (usually about 10 units for Pixel and slightly more for the FEDs), so high in fact that during those fills a manual intervention was required and alarms were thrown by the work-loop.
The vertical spikes are another unforeseen feature: those can be found, for examples in Fig. (4.6-(b)-(f)–(j)-(m)-(n) and others), both in the Pixel and in the FEDs. The expected trend, in fact, is a smooth one due to conservative behaviour before applying corrections, since the work-loop only makes the change effective after about 400 measurements. Also, these spikes are always in correspondence of an increase in such corrections and never viceversa.

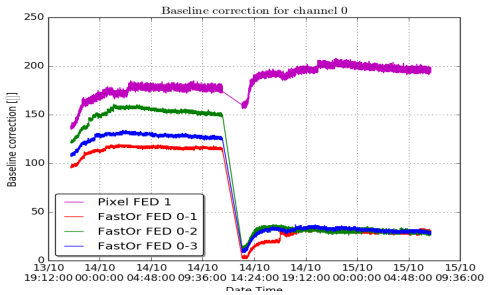

(a) Channel 0

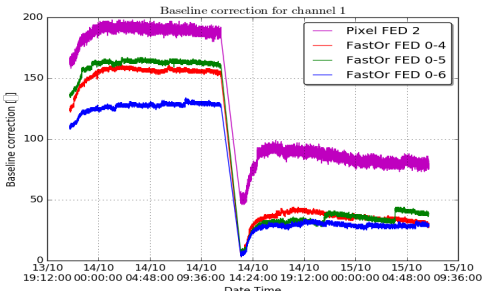

(b) Channel 1

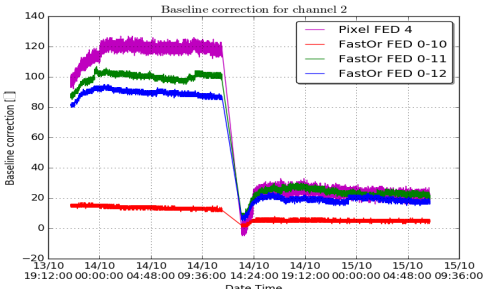

(c) Channel 2

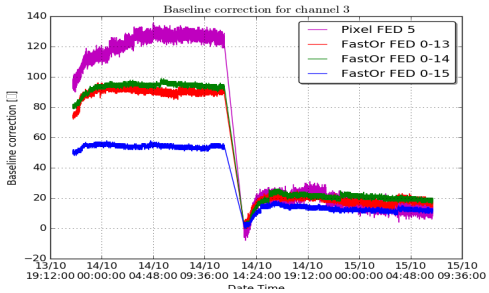

(d) Channel 3

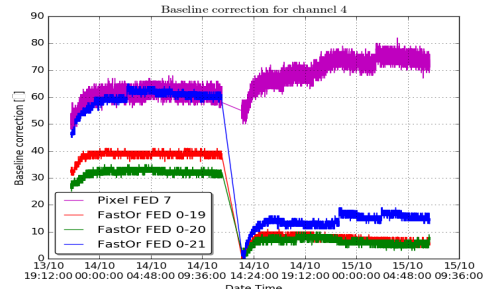

(e) Channel 4

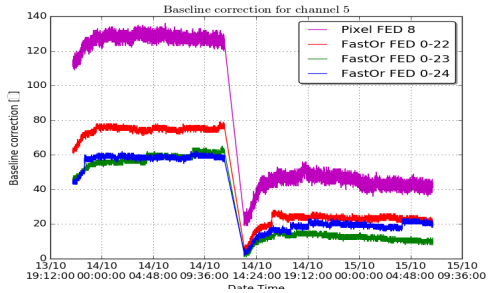

(f) Channel 5

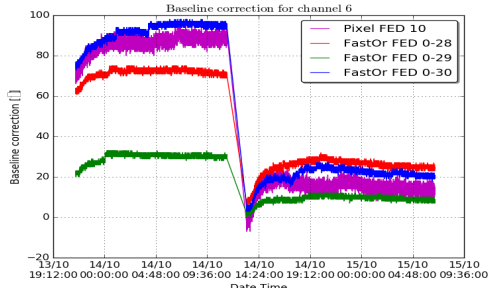

(g) Channel 6

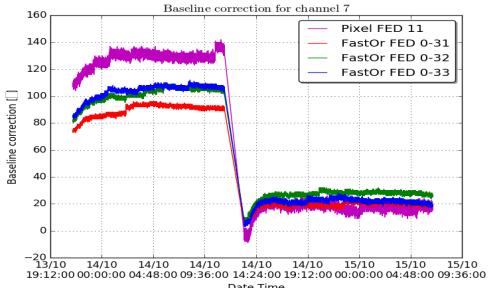

(h) Channel 7

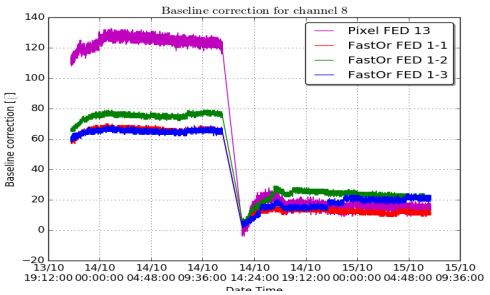

(i) Channel 8

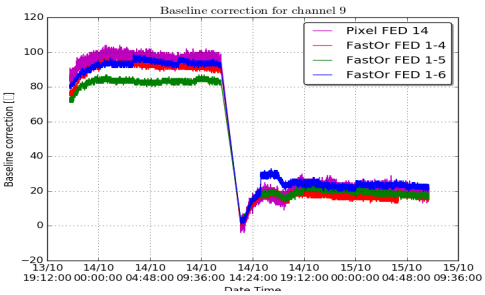

(j) Channel 9

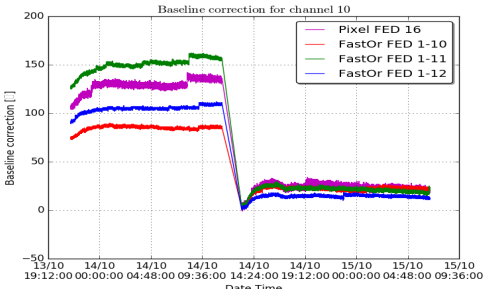

(k) Channel 10

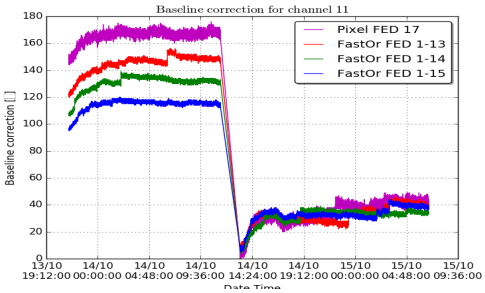

(l) Channel 11

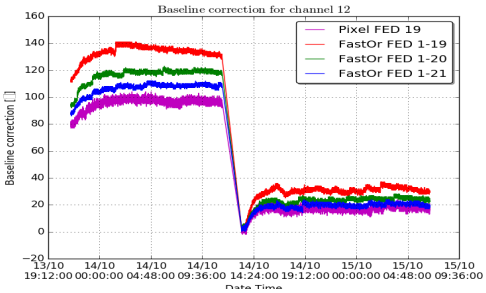

(m) Channel 12

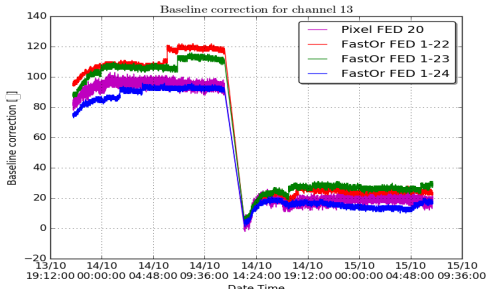

(n) Channel 13

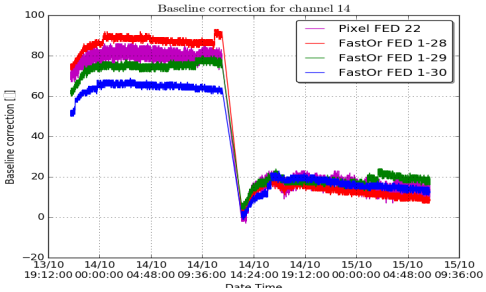

(o) Channel 14

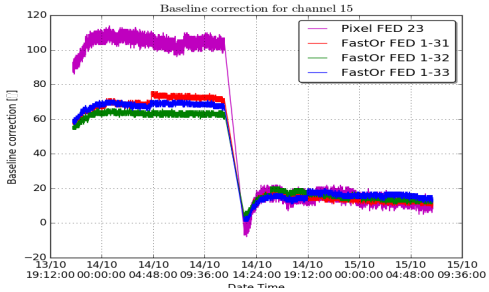

(p) Channel 15

Figure 4.6: Plots with the per-channel baseline corrections for Pixel and FEDs in fills 6297-6298; the straight lines represent the interfill, when the work-loop is in quiet mode and it therefore applies no corrections.

4.2.3 Temperature trends

It was decided to check the temperature trend of the detector during those fills to see if any relevant correlation could be found. The PLT has available three sensors in charge of measuring the temperature of its components: one on the bulkhead, one on the body of the detector itself and the last one in proximity of the OMB. It is this one that is expected to deliver the most important pieces of information.
The following plots show the values of temperature for the four PLT quadrants, referred to with their usual convention (the z coordinate is again with respect to the CMS coordinate system). Actually, the plots only show if a variation in this quantity has occurred, so the points represent a change taking place.
Utnsurprisingly, he detector and bulkhead sensors do not show any particular trend. The OMB, however, shows that only a partial cool-down took place in between the fills, as it can be seen from the fact that the temperature never reaches the values that it had at the very beginning of the first fill considered, 6297.

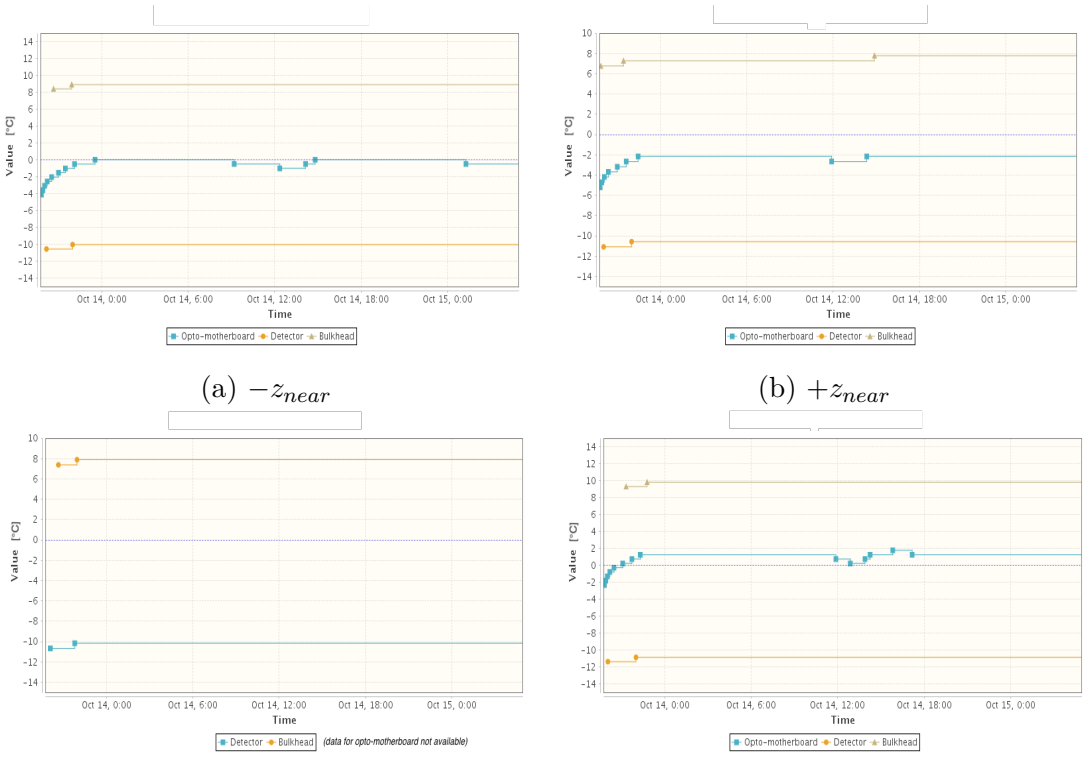

(a) $-z_{near}$ (b) $+z_{near}$

(c) $-z_{far}$ (d) $-z_{far}$

Figure 4.7: Plots with the temperature trends for the three components of the PLT detector; the OMB temperature in $-z_{far}$ is absent because of a broken sensor.

4.2.4 Summary

To summarize it all, over the two fills something unexpected happened in the baseline corrections, which showed higher values especially in fill 6297 and a somewhat surprising behaviour. This is probably related to the temperature trends only to a limited extent: while a partial cool-down could be a symptom of an issue, it hardly explains the frequently repeated spikes found in the plots all around the channels; this feature has not been fully understood yet.
On a bigger picture, however, it has to be said that the baseline correction mechanism proved to be extremely effective even under conditions that pushed it to the limit of its intended working regime.

Bibliography

[1] The CMS Collaboration, *The CMS experiment at the CERN LHC*, CERN (2008)

[2] The CMS Collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, CERN (2017)

[3] W. Herr, B Muratori, *Concept of luminosity*, CERN (2006)

[4] P. Lujan, *Outlook for the luminosity uncertainty*, Top Workshop, CERN (2017)

[5] The CMS Collaboration, *CMS Luminosity Based on Pixel Cluster Counting - Summer 2013 Update*, CERN (2013)

[6] P. Lujan, G. Riley, K. Rose, S. Spanier, K. Thapa, A. Kornmayer, *PLT Luminosity Measurement, Corrections, and Systematic Uncertainty for the 2015 Run*, CERN (2015)

[7] A. G. Dellanoy, *Detector performance and upgrade plans of the Pixel Luminosity Telescope for online per-bunch luminosity measurement at CMS*, CERN Detector Seminar, CERN (2017)

[8] A. Kornmayer, *The CMS Pixel Luminosity Telescope*, Nucl. Instrum. and Meth. A (2015)

[9] A. Kormayer, *Studies on the response behaviour of pixel detector prototypes at high collision rates for the CMS experiment*, CERN Thesis, CERN (2015)

When life asks to move forward after reaching the end of a certain path, one should always take some time to have a look at what's left behind.
This is the right time and place for me to do so.

First of all, the hugest possible thank you goes to mamma *e* papá*: without them I would be nothing and this should not only be intended in the most literal way possible. They have helped me in so many ways throughout this journey and no acknowledgement is ever going to be enough. Ah, obviously also to my brother, just for sticking around from time to time: lately it has been somewhat of a rarity, so I guess every moment should be treasured. And how could I forget* Snow*, capable of looking at the world with different eyes from anyone of us, yet with an unconditional love for all.*

A special and loving thank you goes to Nonna Uovo*, one of my warmest fan, who always managed to infuse optimism and willpower in me with her worthy-of-envy energy.*

A big thank you to Olio Annacquato*, a surprising mixture right from the very name, made of three people who have given me incredible joy and laughter in the good days as much as relief and help in the bad ones. Words can hardly describe the bond that has been growing and growing with these guys and I truly hope that it will be preserved just as strong in the years to come.*

On to another fantastic group of people, the Ammiratori di via Marzolo*, my Physics companions thanks to whom I found pleasure in going to classes with a smile every day. In this instance I need to call out in particular (in no specific order)* Mirko*, my loyal and awesome lab pal with all his crazy adventures,* Davide*, whose skills in science are only surpassed by his kind heart and his availability to people,* Giorgio*, a spectacular and solar guy who has an unmistakeable... style, and* Gallo *and* Azzo*, a great couple that is clearly set to be a great addition to an already roaring group.*

Padova would not have been quite as beautiful and pleasurable without Paolo *and* Marco*, my one-of-a-kind roommates, fun and special people that shared with me the most private moments with a great time for all.*

A more serious yet just as deep thank you goes rightfully to Marco Zanetti*, the coolest thesis supervisor known to mankind, with whom I shared a very good time and who gave me the opportunity to live an outstanding experience here at CERN. And this experience would not have been quite as rich without* Paul Lujan*, a never-ending fountain of knowledge: without him, most of this work would not have been possible and his kindness and availability made the process all the more enjoyable.*

Last, but absolutely not least, a big thank you to Luca*,* Francesca*,* Tyler *and* Francesco*: when I first got here at CERN, I was hoping to find good roommates, but they were actually much closer to a second family for me. They were all super caring, sweet and very cool and their contribution to my good time here was absolutely huge.*

If you want to go fast,
go alone.
If you want to go far,
go together.