LHC Luminosity Performance

Inauguraldissertation
der Philosophisch-naturwissenschaftlichen Fakultät
der Universität Bern

vorgelegt von

Michael Hostettler

von Bern

Leiter der Arbeit:

Prof. Dr. Antonio Ereditato

Albert Einstein Center for Fundamental Physics
Laboratorium für Hochenergiephysik

Dr. Kajetan Fuchsberger

Dr. Giulia Papotti

Europäische Organisation für Kernforschung (CERN)

LHC Luminosity Performance

Inauguraldissertation
der Philosophisch-naturwissenschaftlichen Fakultät
der Universität Bern

vorgelegt von

Michael Hostettler
von Bern

Leiter der Arbeit:

Prof. Dr. Antonio Ereditato

Albert Einstein Center for Fundamental Physics
Laboratorium für Hochenergiephysik

Dr. Kajetan Fuchsberger

Dr. Giulia Papotti

Europäische Organisation für Kernforschung (CERN)

Von der Philosophisch-naturwissenschaftlichen Fakultät angenommen.

Der Dekan

Bern, 21.06.2018 Prof. Dr. Gilberto Colangelo

Abstract

This thesis adresses several approaches with the common goal of assessing, understanding and improving the luminosity of the Large Hadron Collider (LHC). To better exploit existing margins for maximum luminosity while fulfilling the requirements of the LHC experiments, new techniques for luminosity levelling are studied and developed to an operational state, such as changing the crossing angle or β^* (beam size) at the interaction points with the beams in collisions. In 2017 LHC operation, the crossing angle reduction in collisions improved the integrated luminosity by $\sim 2\,\mathrm{fb}^{-1}$ ($\sim 4\,\%$ of the yearly production).

For additional diagnostics, a new method for measuring beam sizes and orbits for each circulating bunch using the luminosity measurement during beam separation scans is shown. The results of these Emittance Scans improved the understanding of the LHC luminosity reach and of the orbit offsets introduced by beam-beam long-range effects.

Acknowledgements

I would like to thank all my dear colleagues and friends, who always provided helpful advice, fruitful discussions, and made CERN an amazing place to study and work at.

First of all, thanks to my supervisors at CERN, Dr. Giulia Papotti and Dr. Kajetan Fuchsberger for their excellent support. Giulia, who I already had the pleasure to work with during a Technical Studentship before, provided invaluable feedback and advice during the first years of this work. After she left LHC operations to join the RF studies group, Kajetan kindly agreed to take over my supervision. He never was too busy to answer my questions, discuss results and give advice. This work would not have been possible without the two of you – thank you!

Second, this study would not have been possible without excellent collaboration of CERN and the Albert Einstin Center for Fundamental Physics, University of Bern. In particular, I would like to thank my supervisor and doctoral advisor, Prof. Dr. Antonio Ereditato. You made this collaboration possible, and you were always there for advice and directions.

As a member of the LHC operations section at CERN, I got the unique opportunity to study data from, and to apply my results directly to the world's largest collider. For this opportunity, and for many fruitful discussions, I would like to thank our section leader, Dr. Jorg Wenniger, as well as all the operators and engineers in charge of the LHC: Belen, David, Delphine, Georges, Georges-Henry, Giulia, Guy, Kajetan, Lasse, Laurette, Matteo, Markus, Michaela, Mirko, Theo, Reyes, Ron and Rossano.

While I started the work on the LHC Luminosity Server back in 2015, many of the new features were collaboratively developed as a SCRUM team in 2017. As the "product owner" of the software developed, I would like to thank all the team members for their ideas, discussions, contributions and dedication to the project. Thank you, Andrea, Delphine, Georges-Henry, Marek, Mathieu and Kajetan. The Luminosity Server would not be where it is today without this collaboration.

Most beam physics studies were done in collaboration with the Accelerator Beam Physics group. In particular, I much appreciate all the fruitful discussions and advice from the LHC beam-beam and luminosity studies working group. Thanks Dario, Fanouria, Gianni, Ilias, Nikos, Stefania and Yannis! Also, the discussions with Tobias, Tatiana, Werner and Xavier contributed a lot towards my understanding of LHC beam physics and beam-beam effects in particular.

Luminosity studies at a collider are tightly linked to the luminosity measurements and the performance of the experiments. A significant part of the data used in this thesis has been provided courtesy of the LHC experiment collaborations, ATLAS, ALICE, CMS, and LHCb. Personally, I would like to thank the members of experiment luminosity measurment and calibration studies groups; your data, support and expertise was crucial for developing and studying the Emittance Scans. Thanks Anne, Chris, David, Moritz, Olena, and Paul from CMS; and Alessandro, Alex, Nicoletta and Witold from ATLAS.

Last but not least, thanks to my parents, Katharina and Urs, and to my brother Nik, for always supporting me in what I did, for being who they are, and for all the great time we spent together.

List of Publications

Over the course of my doctoral studies, I have presented my work in numerous internal meetings, workshops and notes, many of which have found immediate application in day-to-day LHC operation. Furthermore, I have published my work at relevant conferences and in a peer-reviewed journal paper. A selection of these publications are listed below. For all of these publications, I am either the primary author or have made a significant contribution.

Papers, Books and Proceedings:

- M. Hostettler *et al.*, "Luminosity Scans for Beam Diagnostics", submitted to Phys. Rev. ST Accel. Beams, preprint `arXiv:1804.10099 [physics.acc-ph]`.
- G. Valentino *et al.*, "An Upgraded Luminosity Levelling Procedure for the ALICE experiment", submitted to Nucl. Instr. Meth. Phys. Res. A.

  In collaboration with the authors from the ALICE experiment, I commissioned this new separation levelling procedure on the LHC side.
- O. Brünig and M. Hostettler, "Accelerator Operations: Luminosity Optimization", Springer Materials, Particle Physics References.

  This chapter was originally published in *Landolt-Boernstein, Vol. 21, Accelerators and Colliders*. For the upcoming reprint, I have been invited to update it with the concept and results of luminosity levelling by separation, crossing angle, and β^*.
- M. Hostettler *et al.*, "Beam Size Estimation from Luminosity Scans at the LHC During 2015 Proton Physics Operation", International Particle Physics Conference 2016 (IPAC'16).
- M. Hostettler *et al.*, "Impact of the Crossing Angle on Luminosity Asymmetries at the LHC in 2016 Proton Physics Operation", International Particle Physics Conference 2017 (IPAC'17).
- M. Hostettler *et al.*, "Online luminosity control and steering at the LHC", International Conference on Accelerator and Large Experiment Control Systems, Barcelona 2017 (ICALEPCS'17).
- N. Karastathis *et al.*, "Crossing Angle Anti-Levelling at the LHC in 2017", International Particle Physics Conference 2018 (IPAC'18).

  I developed the method which allowed changing the crossing angle in collisions, and participated in the analysis.

Presentations and Internal Notes:

- M. Hostettler *et al.*, "Beam-Beam Orbit Effects in the LHC", presented at the ICFA workshop on Beam-Beam Effects in Circular Colliders 2018.
- M. Hostettler *et al.*, "β^* levelling using the LHC Lumi Server (MD 2427)", CERN Internal MD Note.
- M. Hostettler *et al.*, "Luminosity Evolution and Optimum Fill Length", presented in the LMC Machine Committee, 2016.
- M. Hostettler *et al.*, "Crossing Angle Anti-Levelling: Strategy and Status", presented in the LMC Machine Committee, 2017.
- M. Hostettler *et al.*, "Beta* Levelling Implementation Proposal", presented in the LHC Beam Operation Committee, 2018.
- M. Hostettler *et al.*, "Luminosity, Emittance Evolution & OP Scans", presented at the Evian LHC Beam Operation Workshop, 2015.
- M. Hostettler *et al.*, "How Well do we Know our Beams?", presented at the Evian LHC Beam Operation Workshop, 2016.
- M. Hostettler *et al.*, "Emittance Evolution in the LHC: From Injection to Collisions", presented at the Evian LHC Beam Operation Workshop, 2017.

Contents

1 Motivation

Luminosity and energy are the two key figures of merit of a particle collider. Exploring new physics with particle accelerators requires to observe rare events involving potentially heavy new particles. The center of mass collision energy limits the mass of the produced particles and their branching ratios, while the luminosity determines the collision rate, and hence, the available statistics after a given period of time.

With increasing luminosity, a limiting factor is the *pile-up*: the number of simultaneous events in the detectors of the experiments. At high pile-up, recording and reading out the detector data, as well as the reconstruction of single events from the recorded data, becomes increasingly challenging. For example, the maximum acceptable pile-up for the LHC high-luminosity experiments ATLAS and CMS is $\sim$60. Above, it becomes crucial to level the luminosity down, reducing the pile-up.

Furthermore, the luminosity is not constant over several hours with the beams in collisions. Since the beams degrade over time, the instantaneous luminosity decays. However, refilling a large collider with fresh beams from the injectors takes several hours. In order to optimize the integrated luminosity yield, it is hence important to understand the luminosity evolution and to predict e.g. the optimum fill length.

In this work, the luminosity performance of the LHC, currently the world's largest accelerator, will be studied, and limiting factors will be assessed. Based on this, possibilities for optimizing the integrated luminosity will be proposed. Also, a method for measuring the beam size and beam-beam orbit effects based on luminosity measurements using beam separation scans will be developed. These measurements will then be used to validate simulations from the TRAIN code.

Furthermore, new approaches to luminosity levelling will be studied and developed to an operational level, including "anti-levelling" techniques to better exploit existing machine margins in order to improve the integrated luminosity. These improvements will then be implemented in the LHC control system to be used in regular day-to-day LHC operation.

2 Introduction to Accelerator Physics

2.1 Dynamics of Charged Particles

To bend, focus and accelerate a beam of charged particles, particle accelerators rely on magnetic and electric fields. The driving force is the Lorentz force

$$\frac{\mathrm{d}\vec{p}}{\mathrm{d}t} = e\left(\vec{E} + \vec{v} \times \vec{B}\right) \tag{2.1}$$

for a particle with charge e traveling at velocity $\vec{v}$ in the presence of an external electric field $\vec{E}$ and an external magnetic field $\vec{B}$.

It is worth noting that only electric fields can provide acceleration in the direction of $\vec{v}$, whereas the magnetic fields deflect the particle trajectory without affecting the net kinetic energy. In principle, electric fields could also be used for deflection perpendicular to $\vec{v}$. However, for high energy particles, the magnetic force scales with the magnitude of $|\vec{v}|$, whereas the electric force remains constant. For instance, for $v \approx c$ in a high energy accelerator, the equivalent of a 1 T magnetic field would be 300 MV/m, which is technically not feasible. Hence, accelerators typically use electric fields for longitudinal acceleration and magnetic fields for transverse bending and focusing of particles.

Particle accelerators can be built in a variety of forms. Conceptually the simplest form is a linear accelerator (or "beam line"), where particles start at one point, travel once through the accelerator, and are ejected at the other end. As each particle passes the accelerating elements of such a machine only once, the maximum energy gain is limited. Therefore, for particle physics, linear machines are only used nowadays to pre-accelerate particles before they are injected into other accelerators. Also, beam transfer lines between different accelerators are basically linear machines without accelerating elements.

To accelerate particles to high energies beyond a few GeV, they need to pass the accelerating elements multiple times. This is commonly realized by circular accelerators. The following chapter will focus on high energy synchrotrons, i.e. machines with a circular, closed design orbit and particles circulating close to the speed of light ($v \approx c$).

To describe the dynamics of particles in a synchrotron, a local co-moving coordinate system is introduced for convenience (figure 2.1). Here, x and y denote the transverse dimensions, where x is radial ("horizontal") with positive values towards the outside of the ring, and y is perpendicular ("vertical") with positive values upwards. The longitudinal coordinate s describes the position along the reference trajectory in the accelerator.

In the following, we will use u to refer to an arbitrary transverse coordinate (x or y) in a generic way. x' and y' denote the horizontal or vertical slope, respectively, with respect to the longitudinal coordinate s:

$$u' = \frac{\mathrm{d}u}{\mathrm{d}s} \quad \text{for } u = x, y \tag{2.2}$$

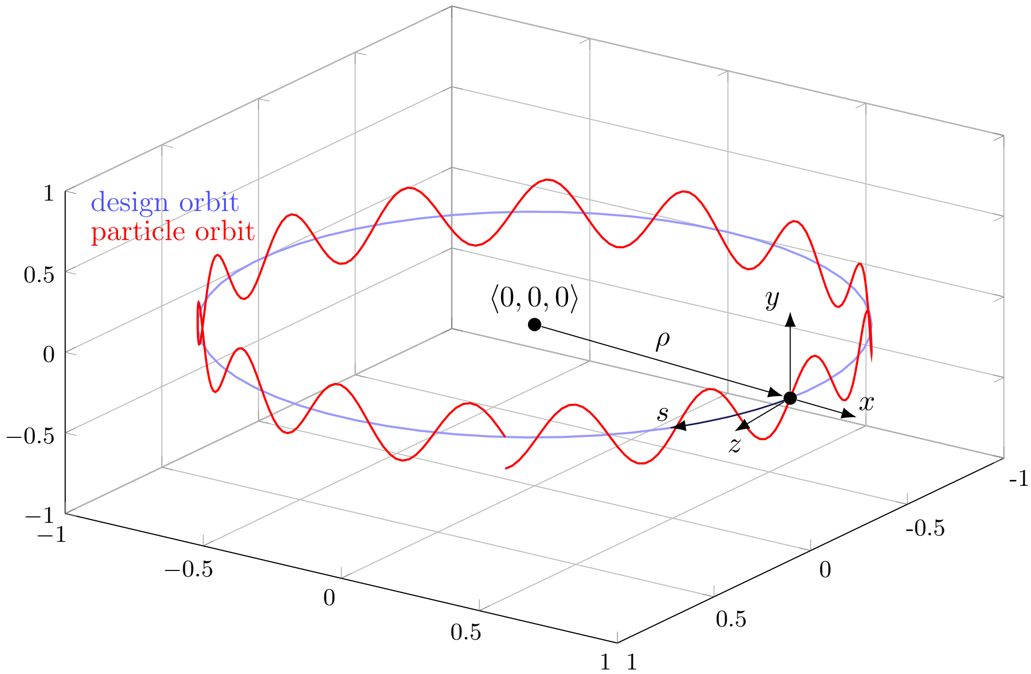

Figure 2.1: The local co-moving coordinate system for a circular accelerator.

2.2 Transverse Beam Dynamics

The motion in the transverse planes is dominated by magnetic fields. A magnetic field in the y plane ($\vec{B}_0(0,0,s) = \langle 0, B, 0\rangle$) and the centripetal and Lorentz forces in equilibrium yield a circular orbit in the x-z plane with the *bending radius* ρ satisfying

$$B\rho = \frac{p}{e} \tag{2.3}$$

where p is the particle momentum and e is the particle charge. $B\rho$ is referred to as the *beam rigidity*. Note that this condition of equilibrium implies that in order to keep ρ (and hence the circular orbit) constant, the bending field strength B needs to change with increasing momentum p as particles are accelerated.[1]

Arbitrary magnetic fields on particles close to this *design orbit* $\langle 0, 0, s\rangle$ can be described by a multipolar expansion [1, 2]:

$$B_y + iB_x = \frac{p}{e}\sum_n (ia_n + b_n)(x + iy)^{n-1} \tag{2.4}$$

where a_n are the *skew* and b_n are the *upright multipolar components*. The upright multipolar components are generated by magnets aligned to the axis, while the skew

[1] For a constant magnetic field B, the radius ρ would increase and particles would spiral outwards during acceleration. Such machines are referred to as cyclotrons and commonly used in a lower energy range, e.g. for medical applications.

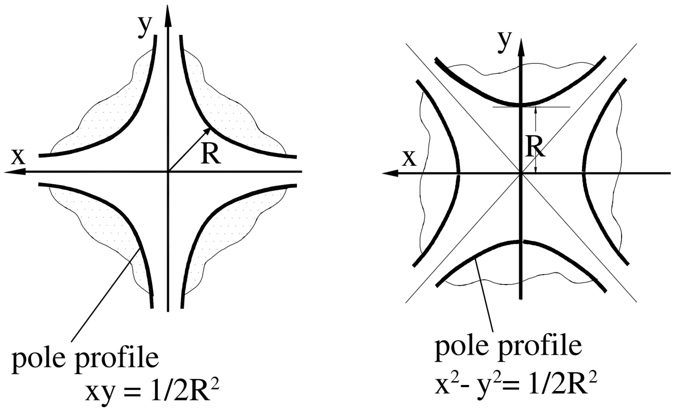

Figure 2.2: Pole profiles of upright and skew quadupole magnets [2]. The skew quadrupoles are rotated by $90°/2 = 45°$.

component	multipole	effect on the beam
$b_0 = \frac{1}{\rho}$	"upright" (H) dipole	bending ("main") dipoles in the horizontal plane, defining the design orbit as in (2.3); horizontal orbit correction
$b_1 = k$	upright quadrupole	focusing
a_0	"skew" (V) dipole	vertical orbit correction
a_1	skew quadrupole	linear coupling
b_2	upright sextupole	chromatic correction
b_3	upright octupole	landau damping through detuning

Table 2.1: Commonly used multipoles in accelerators.

multipoles are rotated by $90°/(n+1)$ (figure 2.2). The multipoles usually used in an accelerator are compiled in table 2.1.

The dominating forces in an accelerator are the upright dipolar (bending) and quadrupolar (focusing) fields. The higher order and skew multipoles are commonly used as correctors for higher order effects. Hence, the full multipolar expansion (2.4) can be treated in the *linear optics approximation*, keeping only the upright dipolar and quadrupolar terms. For such magnetic fields, the equations of motion solving (2.1) become

$$x''(s) + \left(\frac{1}{\rho^2(s)} - k(s)\right) x(s) = \frac{1}{\rho(s)} \frac{\Delta p}{p} \tag{2.5}$$

$$y''(s) + k(s) y(s) = 0 \tag{2.6}$$

where ρ is the dipolar bending radius, k is the quadrupole strength[2] and $\frac{\Delta p}{p}$ is the

[2]The sign of the quadrupole strength is arbitrary. By convention, $k < 0$ represents a horizontally focusing quadrupole.

deviation of a particle's momentum from the reference momentum as defined by (2.3). Due to the absence of skew and higher order multipoles, the equations for the two planes decouple and can be treated separately. In a circular accelerator of circumference L, both ρ and k are naturally periodic with periodicity L; then equations (2.5, 2.6) have the characteristics of *Hill's equations*. The homogeneous part can be solved applying Floquet's theorem [3]:

$$u_\beta(s) = a_u \sqrt{\beta_u(s)} \cos(\mu_u(s) + \mu_u(s_0)) \quad \text{for } u = x, y \tag{2.7}$$

with the parameters a_u and $\mu_u(s_0)$ being determined by boundary conditions[3]. $\beta_u(s)$ is called the *beta function* in plane u and $\mu_u(s)$ is the *phase advance* related to the beta function via

$$\mu_u(s) = \int_{s_0}^{s} \frac{1}{\beta_u(l)} \mathrm{d}l \tag{2.8}$$

where s_0 is an arbitrary starting point in the accelerator.

The inhomogeneous differential equation (2.5) for the horizontal[4] motion indicates that the dipolar fields affect the motion of particles with non-zero momentum offset $\frac{\Delta p}{p}$. This effect is know as *dispersion*. The solution is a combination of the homogeneous solution and a particular solution for the non-homogeneous case:

$$x(s) = x_\beta(s) + D_x(s)\frac{\Delta p}{p} \tag{2.9}$$

where $x_\beta(s)$ is the solution given by (2.7), and the *dispersion function* $D_x(s)$ satisfies

$$D_x(s) = \frac{\partial x(s)}{\partial \left(\frac{\Delta p}{p}\right)} \tag{2.10}$$

Betatron Tune and Stability

As indicated by the homogeneous solution (2.7), the trajectories of particles around the design orbit follow an oscillating behavior, the *betatron oscillations*. The number of oscillation periods over the circumference L of the accelerator are called the *betatron tune* Q_u:

$$Q_u = \frac{1}{2\pi}(\mu_u(L) - \mu_u(0)) = \frac{1}{2\pi} \int_0^L \frac{1}{\beta_u(s)} \mathrm{d}s \tag{2.11}$$

The horizontal and vertical tunes Q_x, Q_y are called the *working point* of an accelerator. To avoid instabilities driven by magnet imperfections, it is crucial to avoid tune resonances. For the n-th order resonance, the trajectory would be closed after n turns (nQ_u is integer),

[3] The notation and symbols used here are by convention. This form is convenient as the quantities a_u and $\beta_u(s)$ have physical properties which are explained in section 2.2.

[4] Note that in the presence of vertical ("skew") dipoles, the same effect occurs in the vertical plane. In real accelerators, the horizontal dispersion is usually significantly larger than the vertical one.

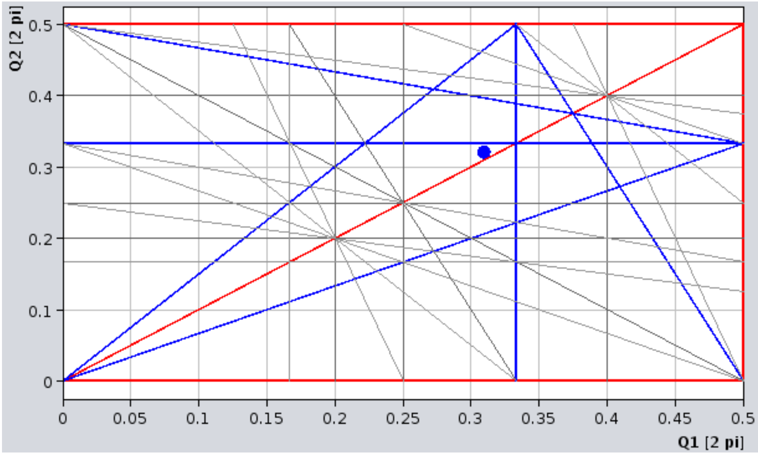

Figure 2.3: Screenshot of the LHC tune diagram display showing the fractional tune parts. The blue dot is the nominal working point ($Q_x = 62.31$, $Q_y = 60.32$). The lines indicate resonances to avoid.

and hence any imperfection would amplify after n turns. In the presence of skew fields and coupling between the horizontal and vertical motion, coupled resonances, e.g. where $nQ_x + mQ_y$ is integer, must also be avoided (figure 2.3).

It is worth noting that different particles on different trajectories along an accelerator encounter slightly different betatron tunes through *detuning* effects. For example, octupole magnets provide an amplitude dependent detuning. The total area covered in the tune diagram by all particles in the beam is referred to as the *tune footprint*. Having such a *tune spread* through detuning is favorable to mitigate multi-particle coherent motion and resulting beam instabilities. However, care must be taken that the expected tune footprint does not cross any resonance in the tune diagram, in which case the affected particles would get on unstable trajectories and would be lost from the beam.

Closed Orbit and Orbit Steering

Due to the periodicity of the magnetic fields in an accelerator, a stationary, periodic $\beta_u(s)$ exists so that the trajectory $u(s)$ (see equation 2.7) is closed after one revolution: the *closed orbit*. In a stationary, stable situation, the center of the beam is on a closed orbit which is invariant over time. Note that the closed orbit can be different from the design orbit.

The closed orbit in either plane can be *steered* by means of dipole corrector magnets. The change of closed orbit $\Delta u(s)$ due to an arbitrary dipole kick Δk at position s_k is [4]

$$\Delta u(s) = \frac{\sqrt{\beta(s)\beta(s_k)}\,\cos(|\mu(s) - \mu(s_k)| - \pi Q_u)}{2\sin(\pi Q_u)}\Delta k \tag{2.12}$$

Since Δu is linear in Δk, the combined effect of multiple dipole kicks is just the sum of the individual kick responses. This can be generalized into a matrix formalism. For an arbitrary vector of kicks θ_i at positions i, the resulting closed orbit effect Δu_j at position j is given by[5]

$$\Delta u_j = R_{ij}\theta_i \tag{2.13}$$

where R_{ij} is the *response matrix* given by (2.12):

$$R_{ij} = \frac{\sqrt{\beta_i\beta_j}\,\cos(|\mu_i - \mu_j| - \pi Q)}{2\sin(\pi Q)} \tag{2.14}$$

In particular, using 4 dipole correctors, it is possible to create a local closed *orbit bump* to control the offset and the angle of the beam at an arbitrary position enclosed by the 4 correctors, but leaving the closed orbit outside of the bump unaffected [4]. An example of such a bump is shown in figure 2.4.

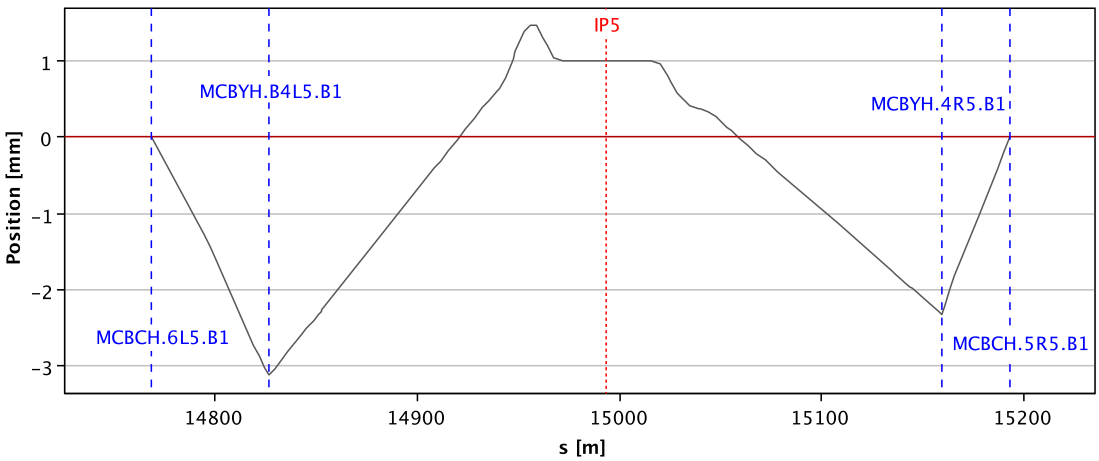

Figure 2.4: Example of a local closed orbit bump. The blue dashed lines indicate the positions of the used dipole correctors. The bump was made such that the orbit at "IP5" (red dotted line) is +1 mm with respect to the design orbit.

The Concept of Emittance

As given by (2.7), a particle in stable motion in an accelerator follows an ellipse in the phase space (u, u') of either transverse plane [2]

$$\gamma_u(s)u^2 + 2\alpha_u(s)uu' + \beta_u(s)u' = a^2 = \text{const.} \tag{2.15}$$

where a is the *Courant-Snyder invariant*, and α_u and γ_u are given by the β_u function as

$$\alpha_u(s) = -\frac{1}{2}\beta_u'(s) \tag{2.16}$$

$$\gamma_u(s) = \frac{1+\alpha_u^2(s)}{\beta_u(s)} \tag{2.17}$$

[5]For readability, the u index is not explicitly shown on θ_i and R_{ij}. Nevertheless this calculation is done per plane, and θ_i, R_{ij}, $\beta_{i/j}$, $\mu_{i/j}$ and Q are specific to the plane.

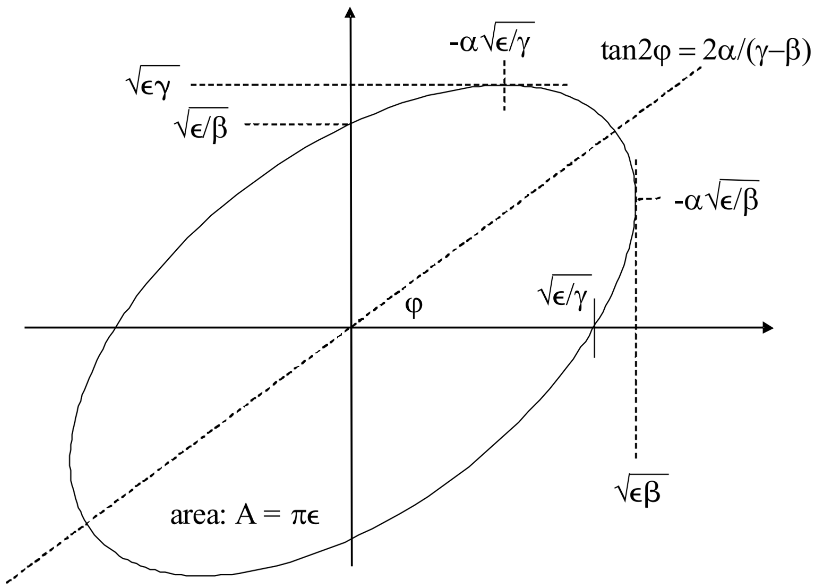

Figure 2.5: The Courant-Snyder ellipse in transverse phase space [2].

For a beam consisting of many particles, all particles follow similar ellipses, centered around the closed orbit. The *transverse emittance* ϵ_u is defined by the area the ellipse

$$\gamma_u(s)u^2 + 2\alpha_u(s)uu' + \beta_u(s)u' = \epsilon_u \tag{2.18}$$

covering $1\sigma = 68\%$ of all particles of the beam[6]. For a beam with an uncorrelated Gaussian transverse charge density

$$\rho(x,y) = \frac{N\,e}{2\pi\sigma_x\sigma_y}\exp\left(-\frac{x^2}{2\sigma_x^2}\right)\exp\left(-\frac{y^2}{2\sigma_y^2}\right) \tag{2.19}$$

the *beam size* σ_u in either plane is given by

$$\sigma_u = \sqrt{\beta_u\epsilon_u} \tag{2.20}$$

In an accelerator consisting only of linear optics elements (dipoles and quadrupoles), the emittance ε is conserved. It is worth noting that during acceleration, the transverse momentum is conserved while the longitudinal momentum is increased. This decreases the angles u', and hence the emittance. This *adiabatic damping* scales the emittance as

$$\varepsilon_u = \epsilon_u\,\gamma_{\text{rel}}\,\beta_{\text{rel}} \tag{2.21}$$

where γ_{rel} and β_{rel} are the relativistic factors, and the *normalized emittance* ε_u is conserved.

For convenience, in this thesis the normalized emittance ε_u will be used unless otherwise noted, and it will be simply called "the emittance". Its units are $[\text{mm}\cdot\text{mrad}] \mathrel{\widehat{=}} [\mu\text{m}]$.

[6] Other common definitions found in the literature include the "95% emittance" and the "100% emittance", with their ellipses covering the respective fraction of the beam phase space.

2.3 Lattice and Insertions

The *lattice* is the arrangement of magnets in an accelerator. In large synchrotrons, it is usually a periodic structure built of smaller *cells*. Since quadrupoles are always focusing in one plane but defocussing in the other, an arrangement has to be found which provides focusing in both planes. A common solution is the *FODO* cell, which consists of equidistant, upright horizontally focusing (QF) and horizontally defocussing (QD) quadrupoles of the same strength (figure 2.6). The space between the QF and the QD is taken by the main bending dipoles and corrector magnets.

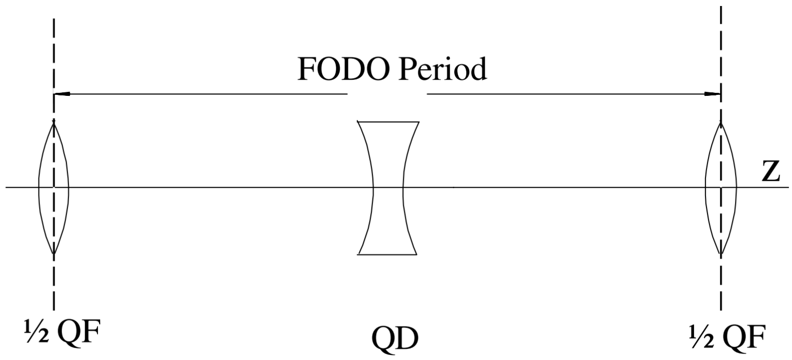

Figure 2.6: A basic FODO cell [2]. Here QF is a (horizontally) focusing quadrupole, and QD is a (horizontally) defocussing one.

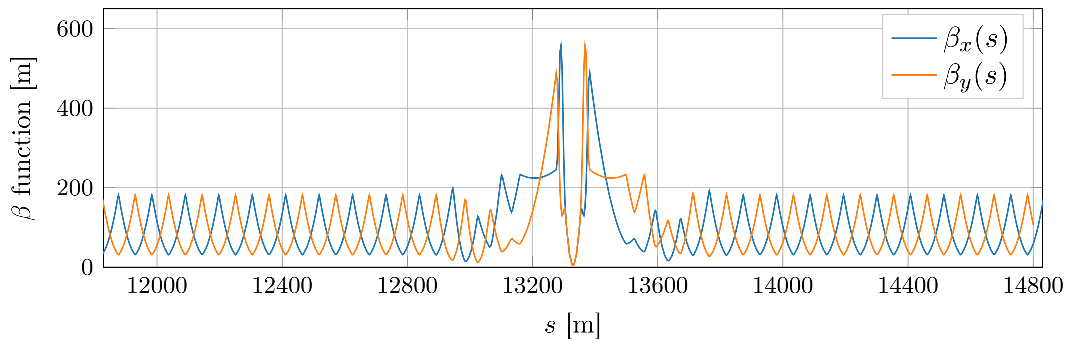

Figure 2.7: The β functions in a FODO lattice with a "mini-β" type insertion. The peaks of β_x are at the horizontally focusing quadrupoles (QF), while the peaks of β_y are at the horizontally defocussing (vertically focusing) quadrupoles (QD). The insertion only introduces local changes.

To inject or extract beam, to install accelerating cavities or experiment detectors, this periodic lattice can be interrupted by *insertions*. As long as the $\beta(s)$ functions (and their derivatives) are matched at the start and at the end of an insertion, the overall focusing properties of the periodic lattice remain unchanged. An example with the aim of minimizing the β function at a particular point (see chapter 3) is given in figure 2.7.

2.4 Longitudinal Beam Dynamics

As pointed out in section 2.1, electric fields are used to accelerate particles in an accelerator. Except for very low energy linear accelerators (e.g. particle sources), the accelerating electric fields are varying over time. Synchrotrons typically use RF cavities for acceleration. The energy gain ΔE of a particle with charge e passing through a cavity is [5]

$$\Delta E = eV_{\mathrm{rf}} \sin(\phi) \tag{2.22}$$

where V_{rf} is the effective RF voltage and ϕ is the phase of the particle with respect to the RF field. For $\phi > 0$, the particle is accelerated, while for $\phi < 0$ it is decelerated. The phase of a particle on the design orbit with zero momentum offset ($\frac{\Delta p}{p} = 0$) is called the *synchronous phase* ϕ_s. At this phase, the energy gain exactly matches the nominal acceleration (when accelerating) or compensates for the natural energy loss of the particle over one turn (at constant energy).

All particles in the beam oscillate around this synchronous phase in the longitudinal phase space. The regions around ϕ_s in which this motion is stable are called the *RF buckets*, limited by their *separatrices*. This is illustrated in figure 2.8.

In order to form stable, stationary buckets, the RF frequency must be an integer multiple of the machine revolution frequency f_{rev}:

$$f_{\mathrm{rf}} = h\, f_{\mathrm{rev}} \tag{2.23}$$

where h is the total number of RF buckets, the *harmonic number*. Each bucket can be filled by many particles (a *bunch*), forming a bunched beam.

The minimum *bunch spacing* is implied by the circumference L of the accelerator and the harmonic number h. The actual design bunch spacing can be a multiple n of this minimum spacing, in which case only every n-th RF bucket is filled.

For technical reasons, in particular to allow for the rise time of injection and extraction kicker magnets, some beam-free longitudinal gaps have to be foreseen. For highest energy

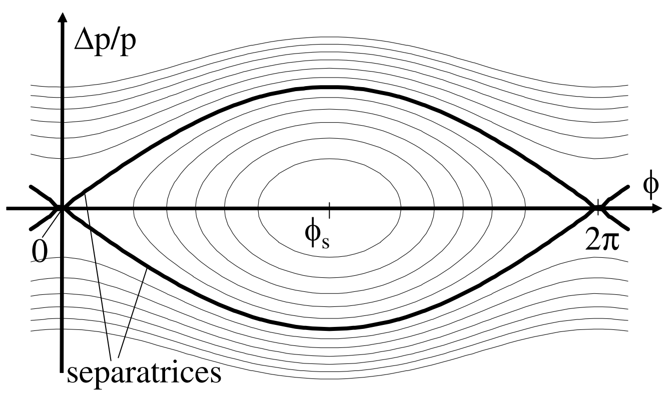

Figure 2.8: A (stationary) RF bucket around the synchronous phase ϕ_s in longitudinal phase space [2].

storage rings, where the beam is produced in a multi-stage process through an *injector chain* of lower energy accelerators, these gaps form a longitudinal macro structure. The resulting longitudinal pattern of bunches and empty RF buckets within a beam is called the *filling scheme*. An example is given in figure 2.9.

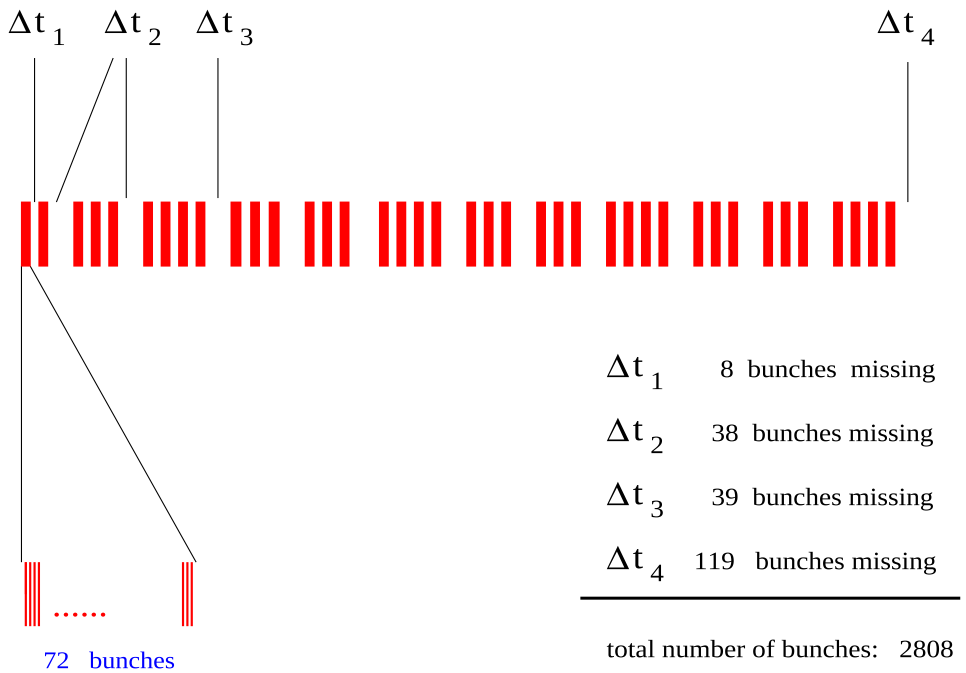

Figure 2.9: The nominal filling scheme of the LHC [6]. The gaps allow for injection and extraction kicker rise times throughout the injector chain and the LHC.

3 Colliders and Luminosity

3.1 Colliders

When two particles collide and scatter inelastically, they form new, potentially unknown, secondary particles. By recording and reconstructing such *events*, the predictions of theories in particle physics are probed. Searches for new physics beyond the Standard Model typically require observing rare events at the highest energy.

The center of mass energy E_{cm} (or $\sqrt{s}$) of two relativistic particles with four-momenta $\mathbf{p}$ and $\mathbf{q}$ colliding is given by

$$E_{\text{cm}} = \sqrt{s} = \sqrt{(p_\mu + q_\mu) \cdot (p^\mu + q^\mu)} \tag{3.1}$$

For a *fixed-target experiment*, where a particle with energy E ($p = \langle E, p_x, p_y, p_z \rangle$) collides with a particle at rest ($q = \langle m, 0, 0, 0 \rangle$), equation 3.1 yields

$$E_{\text{cm}} = \sqrt{s} = \sqrt{2\,E\,m + 2\,m^2} \tag{3.2}$$

while for a *colliding beam experiment*, where two particles of energy E collide head-on ($p = \langle E, p_x, p_y, p_z \rangle$, $q = \langle E, -p_x, -p_y, -p_z \rangle$), the center of mass energy is

$$E_{\text{cm}} = \sqrt{s} = 2\,E \tag{3.3}$$

Since for ultra-relativistic particles from accelerators the rest energy is negligible compared to the kinetic energy ($m \ll E$), the center-of-mass energy achievable from colliding beams is several orders of magnitude larger than the one of a fixed target experiment. An accelerator or storage ring built for colliding two beams is called a *collider*. Circular colliders can either have two beams of opposite charge (particle and anti-particle) circulating in opposite directions in a single beam pipe (e.g. LEP [7], SppS [8]) or use two separate rings which intersect at the particle interaction points (e.g. ISR [9], LHC [6]).

3.2 The Concept of Luminosity

The rate of events of a certain type p with a cross-section σ_p in a particle physics experiment is given by

$$\frac{\mathrm{d}N_p}{\mathrm{d}t} = \sigma_p\,\mathcal{L} \tag{3.4}$$

where the proportionality factor $\mathcal{L}$ is the *luminosity* [$\mathrm{m}^{-2}\mathrm{s}^{-1}$]. As observing more events per time increases the statistical significance of rare events, the luminosity is considered a key figure of merit for the performance of a collider.

For two colliding bunched beams, the luminosity per bunch crossing is given by the overlap integral [10]

$$\mathcal{L}_{bb} = K\, N_1\, N_2 \int_{-\infty}^{\infty} \mathrm{d}x \int_{-\infty}^{\infty} \mathrm{d}y \int_{-\infty}^{\infty} \mathrm{d}s \int_{-\infty}^{\infty} \mathrm{d}s_0\, \rho_1(x,y,s,-s_0)\, \rho_2(x,y,s,s_0) \tag{3.5}$$

where $N_{1/2}$ are the bunch intensities, $\rho_{1/2}$ are the normalized bunch densities, and K is the kinematic factor

$$K = \frac{1}{c}\sqrt{(\vec{v}_1 - \vec{v}_2)^2 - (\vec{v}_1 \times \vec{v}_2)^2/c^2} \tag{3.6}$$

Note that for ultra-relativistic beams colliding head-on, i.e. $\vec{v}_1 = -\vec{v}_2$ and $||\vec{v}_1|| = ||\vec{v}_2|| \approx c$, as in high-energy colliders, the kinematic factor becomes $K \approx 2$.

For bunches with an uncorrelated Gaussian charge distribution

$$\rho(x,y,s,s_0) = \frac{1}{\sqrt{2\pi}^3 \sigma_x \sigma_y \sigma_s} \exp\left(-\frac{x^2}{2\sigma_x^2}\right) \exp\left(-\frac{y^2}{2\sigma_y^2}\right) \exp\left(-\frac{(s-s_0)^2}{2\sigma_s^2}\right) \tag{3.7}$$

the analytical integration of (3.5) yields[1]

$$\mathcal{L}_{bb} = \frac{N_1 N_2}{2\pi\, \Sigma_x \Sigma_y} \tag{3.8}$$

where $\Sigma_{x/y}$ are the effective convoluted transverse beam sizes of the two colliding bunches, which for head-on colliding bunches are given by

$$\Sigma_u^2 = \sigma_{u,1}^2 + \sigma_{u,2}^2 = \frac{\beta^*}{\gamma_{\mathrm{rel}}\beta_{\mathrm{rel}}}(\varepsilon_{u,1} + \varepsilon_{u,2}) \tag{3.9}$$

where β^* is the β function at the collision point[2], $\varepsilon_{u,1/2}$ are the transverse (normalized) emittances, and β_{rel} and γ_{rel} are the relativistic factors. Note that the luminosity increases with a lower beam size at the collision point, and in particular with a lower β^*. Hence, colliders commonly use mini-β^* insertions to focus the beams towards the collision points, maximizing the luminosity.

The total luminosity is proportional to the rate of bunch crossings. In a collider with k colliding bunches per beam and revolution frequency f_{rev}, it is given by

$$\mathcal{L} = \mathcal{L}_{bb}\, k\, f_{\mathrm{rev}} = \frac{k N_1 N_2 f_{\mathrm{rev}}}{2\pi\, \Sigma_x \Sigma_y} \tag{3.10}$$

The average number of interactions per bunch crossing is given by the *pile-up*

$$\mu = \mathcal{L}_{bb}\sigma_{\mathrm{tot}} \tag{3.11}$$

where σ_{tot} is the total cross-section of the collided particles. The maximum pile-up the detectors of the experiments can handle sets an upper limit to the luminosity reach of a collider.

[1] Here and in the following, the approximation $K \approx 2$ is used. Note that even under the presence of a small crossing angle $\alpha \ll 1$ (section 3.2), this approximation holds, as $K \approx 2\cos^2(\alpha/2) \approx 2$.

[2] This assumes that the β function at the collision point is equal for the two beams and in both planes, i.e. $\beta^*_{x,1} = \beta^*_{x,2} = \beta^*_{y,1} = \beta^*_{y,2}$. The calculation for non-equal β^* is analogous.

If the two beams are colliding at an angle or with an offset, the overlap integral (3.5) is reduced. The resulting reduction in luminosity can be factored in via *reduction factors*.

In colliders with many bunches, it is necessary to collide the beams at a *crossing angle* in order to avoid unwanted collisions outside of the nominal collision points (figure 3.1). The transverse plane in which the crossing angle is applied (x or y) is referred to as the *crossing plane*, while the perpendicular plane is called the *separation plane*.

For two Gaussian bunches, the effective convoluted beam size in the crossing plane becomes [10, 11]

$$\Sigma_{\text{xing}}^2 = (\sigma_{\text{xing},1}^2 + \sigma_{\text{xing},2}^2)\,\cos^2(\alpha) + (\sigma_{s,1}^2 + \sigma_{s,2}^2)\,\sin^2(\alpha) \tag{3.12}$$

assuming that the crossing angle is small ($2\alpha \ll 1$) and the bunch length σ_s largely exceeds the transverse beam size ($\sigma_{\text{xing}} \ll \sigma_s$). The effective convoluted beam size in the separation plane remains unchanged.

The ratio of the luminosity under a crossing angle to the luminosity without it is called the *geometric factor*

$$\mathcal{F} = \frac{\mathcal{L}_{\text{xing}}}{\mathcal{L}_0} = \frac{\sqrt{\sigma_{\text{xing},1}^2 + \sigma_{\text{xing},2}^2}}{\Sigma_{\text{xing}}} \approx \left(1 + \alpha^2 \frac{\sigma_{s,1}^2 + \sigma_{s,2}^2}{\sigma_{\text{xing},1}^2 + \sigma_{\text{xing},2}^2}\right)^{-0.5} \tag{3.13}$$

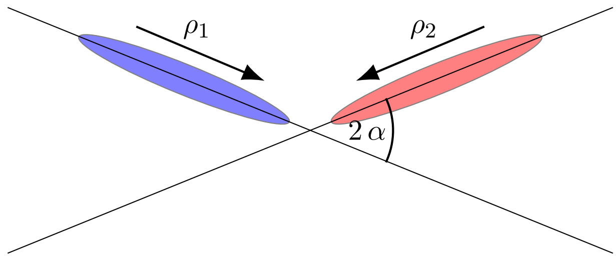

Figure 3.1: Two bunches crossing at an angle of $2\,\alpha$ (longitudinal and transverse planes not to the scale).

A transverse beam separation can either be introduced parasitically by orbit imperfections or intentionally for luminosity control or beam separation scans. The reduction of luminosity under the presence of such a transverse separation d_u in the transverse plane u is quantified by the *separation factor* $\mathcal{S}_u$. For Gaussian bunches, it is given by [11]

$$\mathcal{S}_u = \exp\left(\frac{-d_u^2}{2\Sigma_u^2}\right) \tag{3.14}$$

per transverse plane. The total luminosity reduction factor is $\mathcal{S} = \mathcal{S}_x \mathcal{S}_y$. In order to maximize the luminosity, it is hence crucial to reduce any parasitic separations. This is commonly achieved by locally steering the closed orbit at the IP using 4-corrector closed orbit bumps until the peak in luminosity is found.

A third type of geometric luminosity reduction, the *hourglass effect*, occurs in colliders with very small β^* and long bunches (such that $\sigma_s \gtrapprox \beta^*$). In the inner region around the IP with no focusing elements, the β function at a distance s from the IP is [10]

$$\beta(s) = \beta^* \left(1 + \left(\frac{s}{\beta^*}\right)^2\right) \tag{3.15}$$

In colliders with $\sigma_z \ll \beta^*$, the change in transverse beam size due to the change of β over the length of the bunch is negligible. Otherwise, the hourglass effect couples the transverse and longitudinal planes of the overlap integral (3.5). In case of round Gaussian bunches, the *hourglass factor* is given by [10]

$$\mathcal{H} = \sqrt{\pi}\, h\, \exp(h^2)\, \mathrm{erfc}(h) \quad \text{with} \quad h = \beta^*/\sigma_s \tag{3.16}$$

where erfc is the complementary error function.

Luminosity Time Evolution

The luminosity is not constant over time. The collisions lead to a *burn-off* loss of particles

$$k\, \frac{\mathrm{d}N_{1,2}(t)}{\mathrm{d}t} = -\sigma_{\mathrm{inel}}\, n_{\mathrm{ip}}\, \mathcal{L}(t) \tag{3.17}$$

where σ_{inel} is the inelastic cross-section of the colliding particles, and n_{ip} is the number of collision points. If burn-off is the only relevant mechanism for the variation of the luminosity, its time behavior is described by [12]

$$\mathcal{L}(t) = \frac{\mathcal{L}_0}{(1 + t\, \sigma_{\mathrm{inel}}\, n_{\mathrm{ip}} \mathcal{L}_0/(N_0\, k))^2} \tag{3.18}$$

where $\mathcal{L}_0$ is the initial luminosity, k is the number of bunches and N_0 their initial intensity, which is assumed to be equal for the two beams.

Additional particle losses e.g. due to diffusive effects can further degrade the beams [12]. Also, the beam emittances ε evolve over time due to intra-beam scattering, synchrotron radiation [13] and noise effects [14].

The *integrated luminosity* over the time of operation T

$$\mathcal{L}_{\mathrm{int}} = \int_0^T \mathcal{L}(t)\, \mathrm{d}t \tag{3.19}$$

is the main performance parameter of a collider. The total number of events produced in a particular process with cross-section σ_p is $\mathcal{L}_{\mathrm{int}} \cdot \sigma_p$.

3.3 Beam-Beam Effects

When two bunches of charged particles collide or pass nearby, they are perturbed by each others electromagnetic fields. The resulting forces are highly non-linear and result in a wide spectrum of consequences for beam dynamics [15]. In the approximation of round Gaussian beams, the beam-beam kick is given by

$$\Delta r' = -\frac{2 N r_0}{\gamma_{\mathrm{rel}}} \frac{r}{r^2} \left[1 - \exp\left(-\frac{r^2}{2\,\sigma^2} \right) \right] \tag{3.20}$$

where r is the distance from the opposing beam, $r_0 = e^2/4\pi\epsilon_0 mc^2$ is the particle classical radius, N is the number of particles in the bunch, σ is the transverse beam size, and γ_{rel} is the relativistic factor. The shape of this function is shown in figure 3.2.

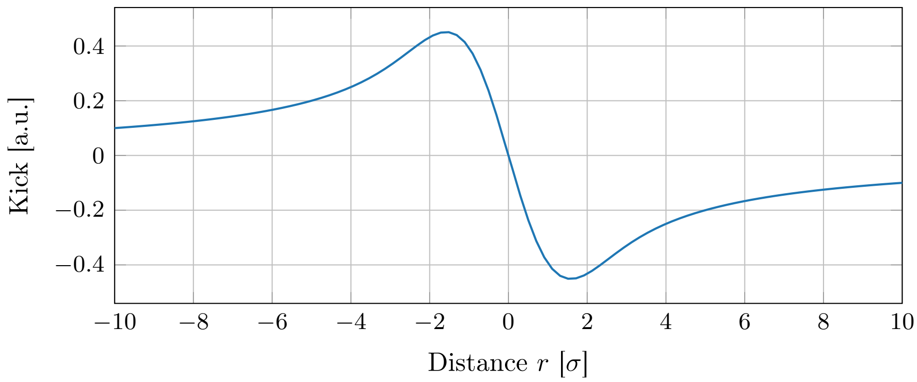

Figure 3.2: The beam-beam kick for round Gaussian Beams.

If two bunches collide centrally ($\langle r \rangle = 0$), the net dipolar kick is zero. Still, each particle exhibits a force which is approximately linear and acts like a quadrupole for small amplitudes (up to $\sim 1\sigma$). At higher amplitudes, it strongly diverges from this linear behavior. The linear part leads to a constant *beam-beam tune shift*, while the non-linear part causes an amplitude-dependent detuning comparable to that of octupole magnets.

When the two beams are separated by a crossing angle close to the nominal collision point as described in section 3.2, an incoming bunch in one beam interacts at a distance with the outgoing bunches of the opposite beam (figure 3.3). In these *long-range* beam-beam encounters, both bunches experience a dipolar *coherent beam-beam kick* which perturbs their closed orbit in addition to the detuning. The crossing angle needs to be large enough to keep the long-range encounters at an acceptable level by providing a minimum beam separation.

If the filling scheme has longitudinal gaps (see section 2.4), bunches next to these gaps miss one or more long-range encounters. As a result of the different beam-beam effects, these *PACMAN bunches* [15] have a different closed orbit, as well as different tunes with respect the “nominal” bunches. These differences can not be corrected by closed orbit steering or global tune corrections which would act on all bunches. An example for the nominal LHC filling scheme is given in figure 3.4.

In this work, a method for measuring the beam-beam closed orbit differences on a bunch-by-bunch level was developed and applied for studying these effects in the LHC for different optics and filling schemes. Results and details are given in chapter 5.

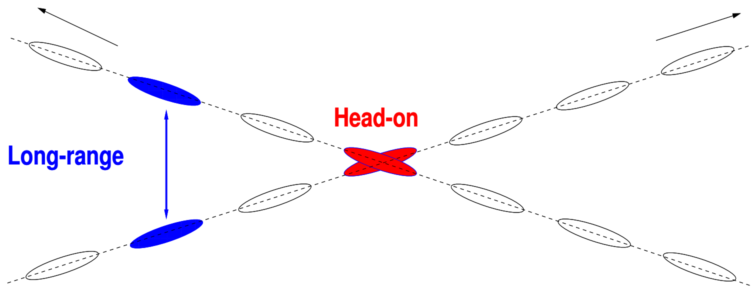

Figure 3.3: Head-on and long-range beam-beam encounters around a collision point [15].

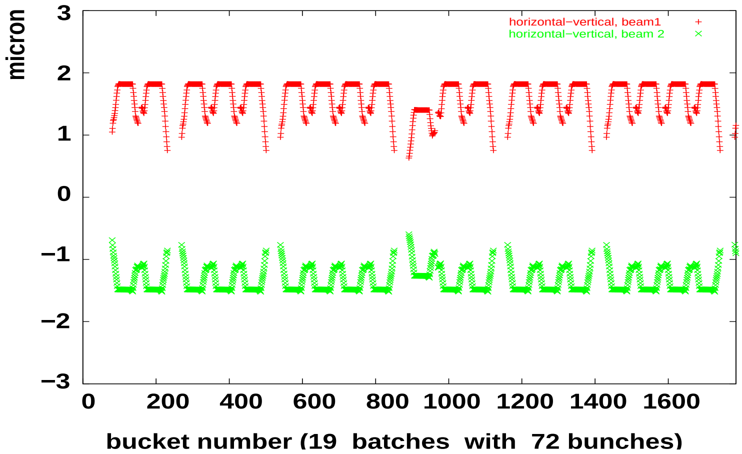

Figure 3.4: Closed orbit variations due to PACMAN effects in the nominal LHC filling scheme [6].

4 The Large Hadron Collider (LHC)

4.1 Injector Chain

CERN, the European Organization for Nuclear Research, hosts the world's largest accelerator complex (figure 4.1). Apart from some machines used for dedicated experiments, a series of smaller accelerators serve as *injector chain* to produce and pre-accelerate beams for the LHC.

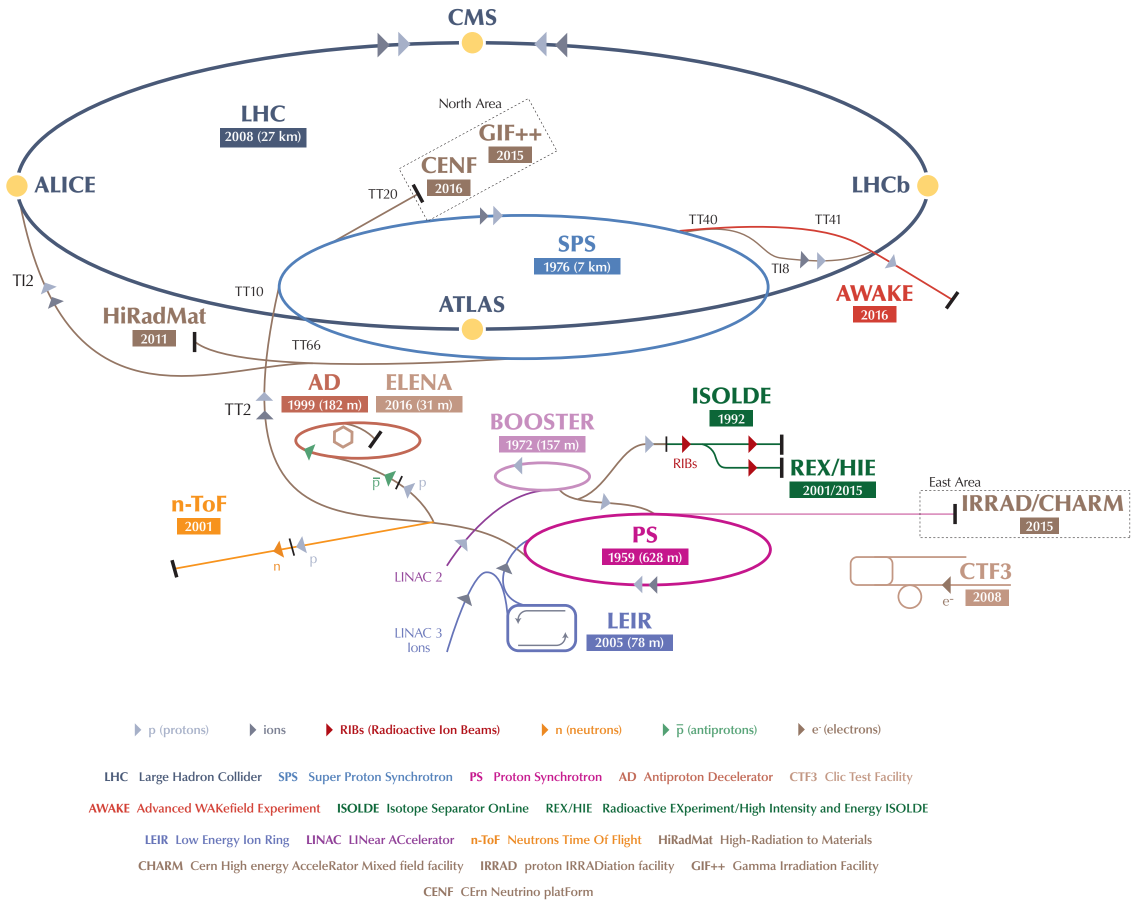

Figure 4.1: The CERN accelerator complex [16].

Protons are produced by ionization of hydrogen gas at the entry of LINAC2, a linear accelerator which pre-accelerates them to 50 MeV. In the PS Booster and PS rings, a bunch train structure compatible with the LHC RF system is established, and the beams are accelerated to 26 GeV. In the SPS, up to 4 of these PS trains are joined longitudinally to a SPS batch, which is then accelerated to 450 GeV and injected into one of the two LHC rings. Depending on the desired filling scheme, it takes up to 40 SPS batch injections per ring to fill the LHC.

4.2 Geometry and Lattice

The LHC is installed in the 27 km tunnel of the now-decommissioned Large Electron-Positron collider (LEP). It consists of 8 arcs with straight insertion regions (IRs) in between. The LHC stores two counter-rotating beams in separate vacuum chambers (*rings*); only next to the designated *Interaction Points* (IPs), the beams enter a common vacuum chamber to collide [6]. A schematic is shown in figure 4.2.

The arcs are built from FODO cells with 6 bending dipoles in between the focusing quadupoles of each cell (figure 4.3). Each cell is 106.9 m long, and each arc consists of 23 cells. The bending dipoles as well as the main quadupoles are superconducting.

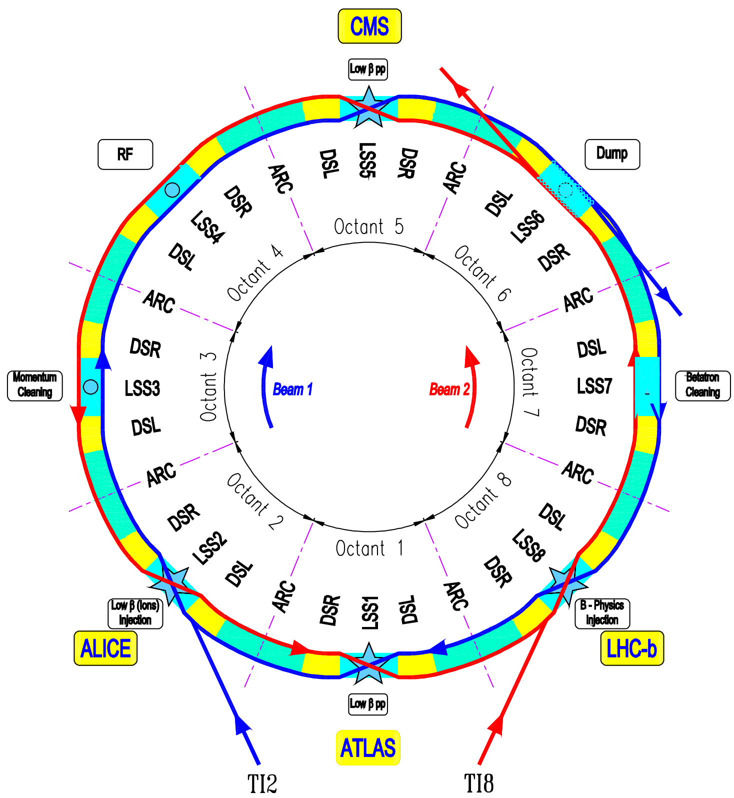

Figure 4.2: Schematic of the LHC rings [6]. Beam 1 (blue) rotates clockwise, Beam 2 (red) counter-clockwise.

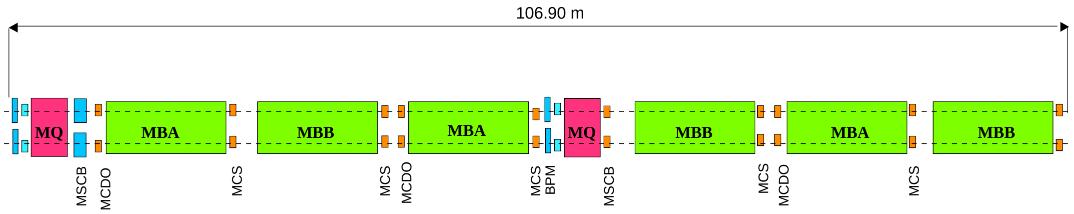

Figure 4.3: Schematic of an LHC cell [6]. "MBA/B" are the bending dipoles, "MQ" are the focusing quadrupoles (both superconducting). The smaller magnets in between are dipolar (orbit), quadrupolar (tune) and higher order correctors.

4.3 Interaction Regions and Experiments

The insertion regions are used for hosting both accelerator systems (e.g. RF, collimation, instrumentation [6]) as well as the Interaction Points with the four main LHC experiments: ATLAS [17] at IP1, CMS [18] at IP5, ALICE [19] at IP2 and LHCb [20] at IP8.

ATLAS and CMS are high-luminosity general purpose detectors, capable of recording data at a pile-up of up to $\mu \approx 60$ events per crossing. For most of the LHC running time, they are operated at the maximum luminosity possible, i.e. with the beams colliding fully head-on at a minimal β^* (e.g. 30 cm in 2017).

On the contrary, ALICE and LHCb need to limit the pile-up to $\mu \approx 0.1$ and $\mu \approx 1.5$, respectively. In present operation, this is achieved by keeping the β^* as high as 10 m in ALICE and 3 m in LHCb, and then using a beam separation to further decrease the luminosity.

A schematic of the insertion regions around IP1 and IP5 is shown in figure 4.4. The insertions are symmetric with respect to the IP. The innermost element is a low-β triplet focusing assembly (Q1, Q2, Q3) to minimize β^*. It is followed by a pair of separation-recombination dipoles (D1 and D2) which join the design orbits of the two rings. From the IP up to the D2, the beams travel in a common vacuum chamber.

After the D1, the separation needs to be ensured by a crossing angle. The crossing plane is vertical in IP1 (and IP2) and horizontal in IP5 (and IP8). The alternating crossing scheme was chosen to limit the PACMAN long-range beam-beam effects by self-compensation [21].

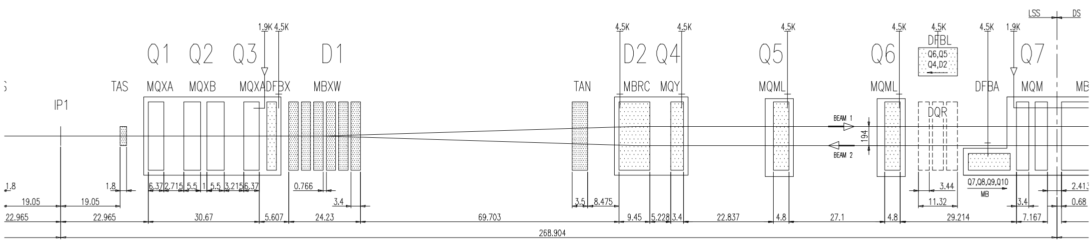

Figure 4.4: The insertion region to the right of IP1 [6]. Dimensions and distances (bottom) are given in meters. Cryogenic temperatures (top) are given in Kelvin.

4.4 Beam Instrumentation

The LHC is equipped with a wide range of instrumentation to measure the beam conditions and parameters. Instruments which measure global properties of the beams (e.g. intensity, emittance) are installed in IR4. Instrumentation for local measurements (e.g. beam position) is distributed over the rings. The instruments and measurements which are used later in this thesis are briefly described below.

Beam Intensity

The reference intensity measurement at the LHC is given by the *DC Beam Current Transformer (DCBCT)*. It provides a measurement of the total number of charges per beam (summed over all circulating bunches) with an uncertainty of ~0.5 % for high-intensity beams [22]. The bunch-by-bunch intensity sharing is measured by the *Fast Beam Current Transformer (FBCT)*. The total intensity measured by the FBCT is calibrated against the DCBCT measurement. Depending on how of the recalibration is done, the error on the intensity is typically $0.5 - 2\,\%$.

Additionally, the relative distribution of charges in both beams is also measured by the longitudinal profile monitors: the LHC Beam Quality Monitor (BQM) and the longitudinal synchrotron light monitor (BSRL). This data can be combined with the total intensity measurement from the DCBCT to derive bunch-by-bunch intensities.

Emittance

The *Wire Scanners* [23] are the reference devices to measure bunch-by-bunch beam sizes and derive transverse emittances for each beam and plane. The wire scanners can measure the emittance throughout the full LHC machine cycle including the energy ramp, provided that the total intensity in the machine is limited to ~240 bunches at 450 GeV and ~12 bunches at 6.5 TeV (as opposed to more than 2000 bunches in regular physics operation). The accuracy of the wire scanner measurement is limited by the accuracy of the scale of the wire position measurement. In 2015, measurements of the wire scanner length scale using closed orbit bumps and the beam position monitors (BPM) have shown discrepancies $< 5\%$ on the beam position which is within the uncertainty on the BPM length scale [24].

Furthermore, the *Synchrotron Radiation Telescopes (BSRT)* provide a continuous operational bunch-by-bunch measurement of transverse emittances by imaging the synchrotron radiation coming from a dedicated undulator at 450 GeV and from a bending dipole at energies > 2 TeV [25]. The BSRT is calibrated against the wire scanners during low-intensity calibration fills. The BSRT measurements are very precise when averaging over several acquisitions. The accuracy is limited to ~15 % by the accuracy of the calibration. In 2015-2017, 3 or 4 calibrations were done per year. In 2017, a long-term degradation effect was observed on the image intensifiers, which impacted the long-term stability of the instrument [26].

When the beams are in collisions, data from the LHC experiments provides additional information on the transverse emittance:

- Using the measured absolute luminosity and intensity measurements, equation 3.8 can be solved for the transverse emittances.
- Using the tracker data from the experimental detectors, the distribution of primary vertices – the *luminous region* – can be reconstructed.
- By scanning the beam separation and observing the change in luminosity, the emittances are found by inverting equation (3.14). This method is described in detail in chapter 5.

Bunch Length and Longitudinal Profiles

The bunch length is operationally measured by the LHC *Beam Quality Monitor (BQM)*. It uses a wall current monitor pick-up with an 8 GS/s ADC to acquire the longitudinal profiles. The full-width half-maximum of the waveform is then measured and converted to a bunch length assuming a Gaussian longitudinal distribution [27]. While this measurement is precise, the measured bunch length only accurately represents the actual r.m.s. bunch length if the longitudinal distribution is Gaussian.

For on-demand measurements and diagnostics, two 40 GS/s oscilloscopes connected to wall current monitor pickups are installed in SR4. The transfer functions of the pickups and cables were measured and are available for deconvolution [28]. The oscilloscopes can acquire a longitudinal profile of either beam over a full LHC turn.

Orbit and Beam Position

The closed orbit of the LHC beams is measured by more than 500 *Beam Position Monitors (BPMs)* per beam and plane distributed along the machine. The precision of the measurement is at the level of 20 μm. The accuracy is limited due to unknown position offsets. For this reason, the orbit is always corrected in relative (against a previously established reference) in LHC operation.

As of 2015, the BPMs at the focusing triplets around the IPs were equipped with a new front-end system, Diode ORbit and OScillation (DOROS) [29]. This improved the short-term precision of these BPMs to $< 1\,\mu$m. Over longer periods (e.g. hours), the precision can be compromised by electronic drifts.

4.5 Feedback Systems

The LHC employs a software feedback system to correct the tune (QFB) and orbit (OFB) in real time [30]. At a rate of 4 Hz (tune) or 12.5 Hz (orbit), the measurement from beam instrumentation is corrected against a pre-programmed reference. To ensure a smooth operation, it is crucial that the reference represents the actual state of the LHC; e.g. in case a closed orbit bump is intentionally introduced, it must also be added to the OFB reference orbit.

The orbit feedback uses a Singular Value Decomposition (SVD) to pseudo-invert the orbit response matrix (equation 2.14). The feedback loop is closed by a Proportional-Integral controller [30]. The *feedback gain* determines the damping time of the system, while the *number of eigenvalues* considered in the SVD determines the locality of the applied corrections. Both parameters need to be carefully tuned in order to avoid orbit excursions while not introducing unnecessary noise from errors in the BPM measurement.

4.6 Collimation

It is unavoidable that a small fraction of the particles in the beams is lost during LHC operation. To protect the superconducting magnets and other sensitive machine elements from the energy deposition of these particles, a 3-tier collimation hierarchy protects the machine aperture. A schematic is shown in figure 4.5.

The primary collimators (TCP), installed in IR 3 and IR 7 (the cleaning regions), are the devices closest to the beams. They intercept particles with large excitation amplitudes (the beam *halo*) both in the transverse and longitudinal directions.

The secondary collimators (TCS) intercept particle showers created by the interaction of protons from the beams with the primary collimators. Other families of collimator intercept mis-kicked beam from injection and extraction kickers (TDI, TCDQ) or debris from the collisions at the IPs (TCL).

The tertiary collimators (TCT) protect the superconducting focusing triplets in the IPs. As they are within the IP insertion regions, they need to follow any local optics changes at the IPs, in particular the β^* squeeze.

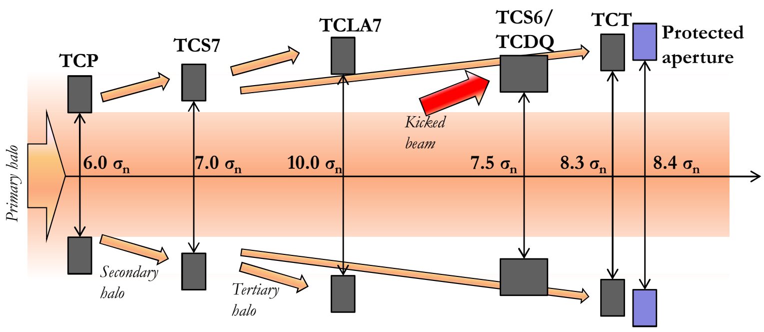

Figure 4.5: The LHC collimation hierarchy [31, edited].

4.7 Optics and Beam Parameters

The key LHC machine and beam parameters are given in table 4.1. For 2017 LHC operation, the maximum achieved values are given. A novel optic design, the "Achromatic Telescopic Squeeze" (ATS) [32], was introduced at the LHC in 2017 to allow for lower β^*; this also changed the integer tunes.

		Injection	Collisions
Geometry and Lattice			
Circumference L	[m]	26658.883	
Ring separation in the arcs	[mm]	194	
Energy (per beam)	[GeV/c^2]	450	6500 (7000)
Relativistic β_{rel}		$\sim$1	
Relativistic γ_{rel}		479	6927 (7461)
Revolution frequency f_{rev}	[Hz]	11245	
RF System			
RF frequency	[MHz]	400.8	
Harmonic number h		35640	
RF bucket length	ns	2.5	
Bunch spacing	ns	25	
Optics			
Horizontal tune Q_x		62.270 (64.28)	62.313 (64.31)
Vertical tune Q_y		60.295 (59.31)	60.317 (59.32)
β^* in IP1/IP5	[m]	11	0.3 (0.55)
β^* in IP2	[m]	10	10
β^* in IP8	[m]	10	3
half crossing angle α in IP1/IP5	[μrad]	170 (160)	150[1] (142.5)
Beam			
Transverse (normalized) emittance ε	[μm]	1.1 (3.75)	1.8 (3.75)
Bunch intensity N	[10^{11} p$^+$]	1.26 (1.15)	1.22 (1.15)
Number of bunches k		2556 (2880)	
Bunch length σ_s	cm	11.3	7.6
Luminosity $\mathcal{L}$ in IP1/IP5	10^{34} cm^{-2}s^{-1}	-	2.2[2] (1)
Pile-up μ in IP1/IP5		-	85[2] (20)

Table 4.1: Selected LHC Beam and Machine parameters achieved in 2017 operation. In case the design values [6] differ, they are given in parentheses.

[1]Reduced to 120 μrad progressively using crossing angle anti-levelling – see chapter 7.
[2]Due to the very high pile-up, the luminosity was levelled to $1.5 \cdot 10^{34}$ cm^{-2}s^{-1} (pile-up of $\sim$55) at the request of ATLAS and CMS in 2017. See chapter 7 for details.

4.8 Operational Cycle

The LHC is a machine with a long cycle time. As injection and collision energies differ and 2-in-1 magnets are used, first both rings need to be filled, then the energy is ramped up, and finally the beams are brought into collisions. The typical operational cycle of the LHC is illustrated in figure 4.6.

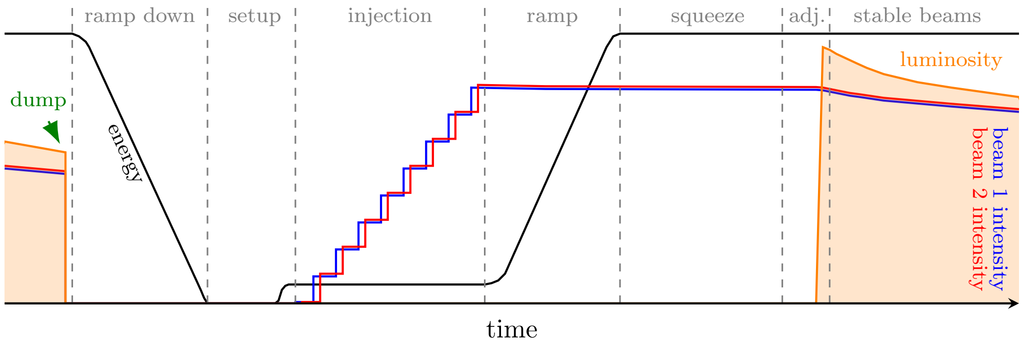

Figure 4.6: A typical operational LHC cycle, from dumping one LHC fill to the declaration of Stable Beams in the next one.

The cycle is organized in different phases of different activities, which are called *beam modes*. In the first phase, the *Setup*, the LHC is prepared to accept beam from the injector complex. In particular, the beam dump system is armed, all interlocks reset, and the main dipoles are ramped to the *Injection* configuration at 450 GeV. Beam is then produced by the injector complex and injected into both LHC rings. For machine protection reasons, first a single low-intensity bunch (probe beam) is injected into each rings. Thereafter, both Beam 1 and Beam 2 are filled with high-intensity SPS batches until the designated filling scheme is achieved.

Then, during the energy *Ramp*, the current in the main dipoles is gradually increased (at the same time, the beam draws energy from the RF system) until the top energy of 6.5 TeV is reached; this typically takes ∼20 minutes. Once at top energy, the optic is gradually changed from injection to collision settings[3]. This process is called the *Squeeze*, as the focusing triplets around the IPs are set for low-β^* operation. Also, the machine tunes are changed from injection to collision values.

Finally, in the *Adjust* beam mode, the closed orbit bumps which separated the two beams at the IPs earlier in the cycle are collapsed. Collisions are established in all IPs, and the beam positions at the IPs are optimized for luminosity. If requested by the experiments (in particular in IP2 and IP8), luminosity levelling is enabled.

Once the LHC is at the designated collision settings, the operators declare *Stable Beams*. This signals the experiments that conditions are stable and data taking can start. Only a limited set of commissioned and well understood machine manipulations (e.g. luminosity optimization, levelling) are allowed in Stable Beams. The beams typically stay

[3] As of 2016, also a first part of the squeeze is done in parallel to the ramp in order to save time [34].

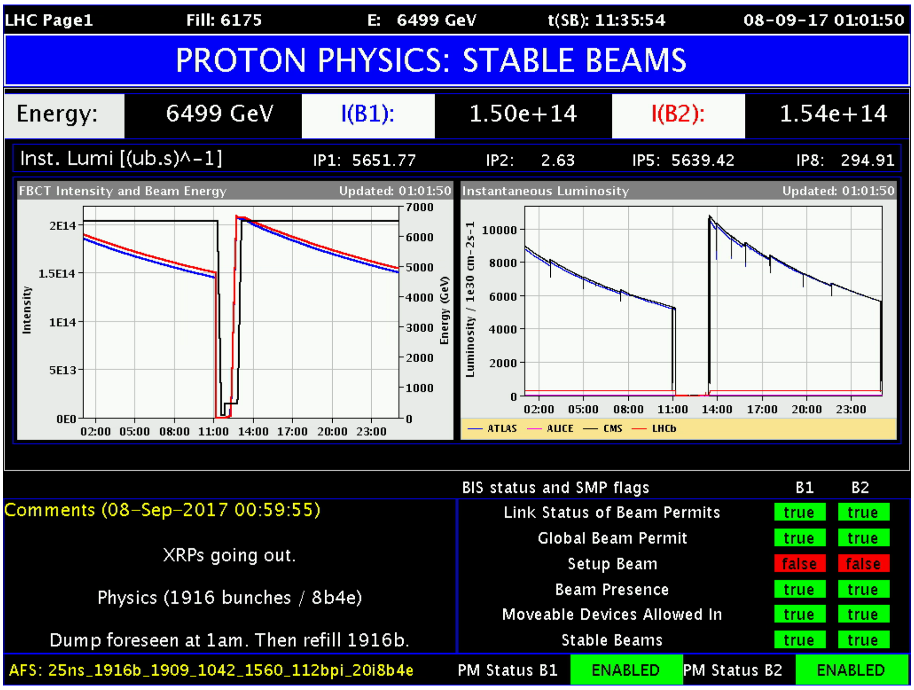

Figure 4.7: LHC Page 1 at September 8, 2017 [33]. LHC fill 6175 has been in Stable Beams for ~11.5 h. Part of the cycle of the previous fill 6174 is also visible.

in collisions in Stable Beams for 10 – 20 h. During this time, the beam intensities and the luminosity evolve as described in section 3.2.

Eventually, a *Beam Dump* can either be programmed by the operators at the designated end of a fill, or be triggered by a machine protection interlock in case of a technical problem. Once the beam permit is removed, the LHC beam dump system extracts the beam within 3 turns [6].

Without beam, the magnet currents are *Ramped Down* back to standby settings in order to accept new beam from the injector complex for the following cycle. To control magnetic hysteresis effects, this must be done by following a set of predefined functions, which take ~30 minutes to play. In case a circuit fails or is turned off without following the ramp down sequence, a shortened version of the full cycle, the *Precycle*, must be played in order to restore the magnetic hysteresis conditions before beam is injected again.

The time from the beam dump of one fill to the declaration of Stable Beams in the next one is called the *turnaround time*. A complete cycle is referred to as an *LHC fill* and is assigned an incremental *fill number*. The current fill number and beam mode are also displayed on the LHC public status webpage, “Page 1” [33]. An example is given in figure 4.7.

5 Emittance Scans

5.1 Motivation

As introduced in chapter 3, the intensity and the transverse emittances are the key parameters for the luminosity in any collider. While the intensity can be measured by the FBCT on a bunch-by-bunch level with less than 2 % uncertainty, measuring the transverse emittances at high energy remains a challenge due to the adiabatic shrinking and the high intensity of the LHC beams (see section 4.4). Before 2015, the BSRT was the only instrument at the LHC able to measure the transverse emittances of high-intensity beams in collisions. This led to major doubts about the accuracy and long-term stability of the BSRT calibration during the first LHC running period from 2010 to 2013, which could not be resolved due to the lack of a complementary measurement.

The following chapter describes a technique for measuring the convoluted transverse beam sizes $\Sigma_{x,y}$ at the Interaction Points by scanning the beam separation and observing the change of luminosity, given by the separation factor (equation 3.14).

The use of beam separation scans for this purpose was pioneered by Simon van der Meer [35] at the CERN Intersecting Storage Rings (ISR) in 1968. Such *van-der-Meer scans* are still used at the LHC for calibrating the luminosity monitors of the experiments in dedicated calibration fills, typically once or twice per year. In these calibration fills, the two LHC rings are filled with ~50 bunches, spaced widely to be able to remove the crossing angle and avoid any beam-beam long-range effect. The van-der-Meer scans are done with a high number of separation steps and long integration time, so a pair of one scan in the horizontal and vertical plane takes ~45 minutes.

In contrast, the *emittance scans* I developed in the LHC are done in regular physics conditions, with a full machine (typically more than 2000 bunches per ring) and under the presence of a crossing angle. The number of points and integration times are optimized to provide a reasonable trade-off between accuracy and time spent for the scan. A pair of horizontal and vertical emittance scans can be done in less than 5 minutes.

Emittance scans were first commissioned in 2015 [36] and have been done in most LHC fills since then [37, 38]. As they depend only on the linearity of the luminosity measurements and the transfer function of the steering magnets used, they are fully independent of any beam instrumentation used for emittance measurements. As such, they provide important complementary data for understanding the LHC emittance evolution and its impact on luminosity. Examples include quantifying the emittance optimizations in the LHC injector complex in 2015 or understanding the observed asymmetry in delivered luminosities between the two high-luminosity experiments ATLAS and CMS in 2016 (see chapter 6). Furthermore, bunch-to-bunch closed-orbit differences can be observed, providing a direct measurement of the long-range beam-beam PACMAN effects.

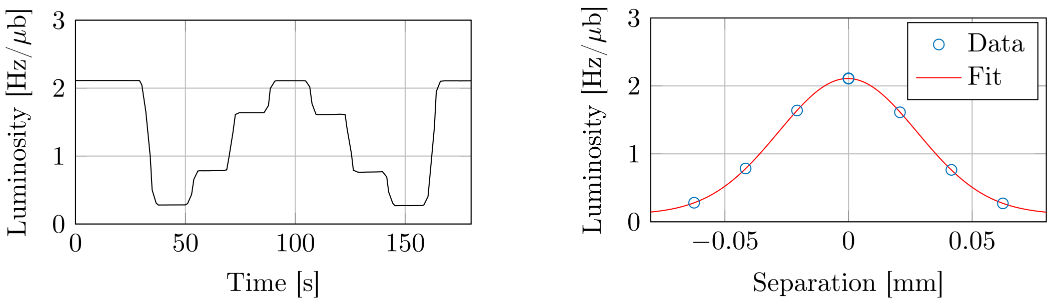

(a) Luminosity evolution over time.

(b) Fitted beam profile.

Figure 5.1: Luminosity during a typical emittance scan in one plane.

5.2 Method

Using a closed orbit bump around an IP, the beam separation d_u is scanned in each transverse plane. Then, following equation (3.14), a Gaussian of the form

$$\mathcal{L}(d_u) = \mathcal{L}_{\text{peak}} \exp\left(\frac{-(d_u - d_{u0})^2}{2\Sigma_u^2}\right) \tag{5.1}$$

is fitted to the measured luminosity, giving both the convoluted beam sizes Σ_u and the initial (parasitic) separations d_{u0} are found. An example luminosity measurement during a scan with 7 separation steps is shown in figure 5.1. As the luminosity can be measured bunch-by-bunch, this fitting can be done independently for each colliding bunch pair.

To derive the transverse emittances ε_u from the convoluted beam sizes Σ_u, it is assumed that the beam sizes of Beam 1 and Beam 2 are equal[1]. Following equation (3.9), the transverse beam size σ_{sep} in the separation plane (no crossing angle) is given by

$$\Sigma_{\text{sep}} = \sqrt{2}\sigma_{\text{sep}} \tag{5.2}$$

In the crossing plane, the longitudinal bunch profile is projected onto the convoluted beam size Σ_{xing} through the crossing angle. For a Gaussian longitudinal beam profile with bunch length σ_s and half crossing angle α, and still assuming equal beam sizes in beam 1 and beam 2, the convoluted beam size following equation (3.12) is given by

$$\Sigma_{\text{xing}} = \sqrt{2\sigma_{\text{xing}}^2 \cos^2(\alpha) + 2\sigma_s^2 \sin^2(\alpha)} \tag{5.3}$$

The normalized transverse emittances ε_u is then derived from the transverse beam sizes σ_u, assuming a negligible dispersion at the IPs and $\beta^* = \beta_x^* = \beta_y^*$:

$$\varepsilon_u = \frac{\gamma_{\text{rel}}}{\beta^*}\sigma_u^2 \tag{5.4}$$

where γ_{rel} is the relativistic gamma factor and β^* is the value of the β-function at the IP.

[1]The sizes of the two LHC beams are typically within $5 - 10\,\%$ according to independent measurements, e.g. from the synchrotron light monitors.

5.3 Impact of the Longitudinal Distribution

In the nominal LHC cycle, the bunch length is increased artificially during the energy ramp by injecting RF phase noise to guarantee beam stability [39]. This changes the longitudinal bunch profile from a Gaussian to a distribution with a larger core and lighter tails, which can be represented as a q-Gaussian with $q \approx 0.85$ [40]. Over the course of several hours in collisions, the longitudinal distribution resumes a more Gaussian shape as shown in figure 5.2.

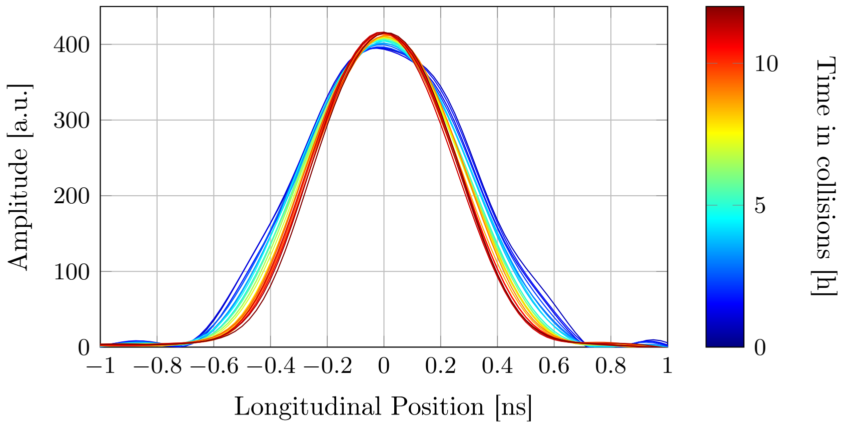

Figure 5.2: Longitudinal profiles in LHC fill 4964 for beam 1, bunch 142 over 12 h in collisions, measured by the 40 GS/s on-demand oscilloscope in SR4. Other bunches (including those in beam 2) behave similarly.

If the longitudinal distribution is not Gaussian, the simple factorization of the convoluted beam size in equation (5.3) is no longer valid for the crossing plane. While still assuming a Gaussian distribution in the transverse dimensions, the luminosity during a scan of the separation d in the crossing plane fulfills

$$\mathcal{L} \propto \exp\left(\frac{-d^2}{4\sigma_{\text{xing}}^2}\right) \cdot \mathcal{C}(d) \tag{5.5}$$

$$\mathcal{C}(d) = \int_{-\infty}^{\infty} \mathrm{d}s\, (f_1 * f_2)(2s)\, \exp\left(\frac{-s\sin(\alpha)(s\,\sin(\alpha) - d)}{\sigma_{\text{xing}}^2}\right) \tag{5.6}$$

where α is the half crossing angle, $f_{1,2}$ are the longitudinal distribution functions for beam 1 and beam 2 respectively, and $*$ denotes a convolution.

It is worth noting that $\alpha = 0$ yields $\mathcal{C}(d) = \text{const.}$ and the separation dependency reduces to the separation factor for Gaussian beams shown in equations (5.1) and (5.2). This proves that the longitudinal distribution only affects scans in the crossing plane.

If the longitudinal distribution is measured for both beams, the convolution and the integration in equation (5.6) can be done numerically, and a non-linear fit of (5.5) to the measured scan points yields σ_{xing}.

To benchmark this approach, emittance scans were simulated for bunches with a transverse Gaussian and longitudinal q-Gaussian distribution with $q = 0.85$. This longitudinal distribution is typical for bunches in the LHC at the start of collisions.

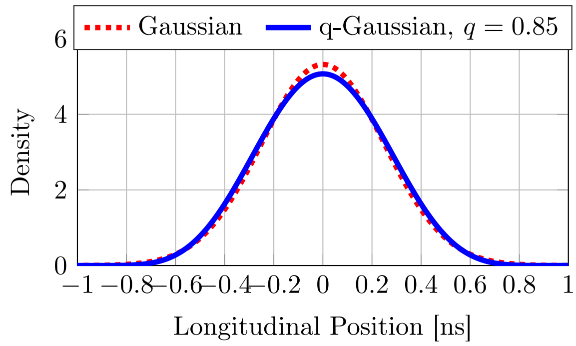

Figure 5.3: Longitudinal q-Gaussian profile used in the simulations, compared to a Gaussian profile of the same r.m.s. width.

The simulations were done for two different sets of machine parameters and emittances, corresponding to the configurations used during the LHC 2015 and 2016 run periods. The simulated data was then analyzed using the method proposed above, and compared to the Gaussian approximation based on full-width half-maximum (FWHM) and true r.m.s. bunch length measurements. The operational bunch length measurement in the LHC, the Beam Quality Monitor, is based on the FWHM method.

Results are compiled in table 5.1. While commissioning emittance scans during 2015 LHC operation, the Gaussian approach using the operational FWHM bunch length measurement was used and a systematic error of up to 15% in the crossing plane due to the non-Gaussian longitudinal distribution was assumed [36]. For measuring low-emittance beams at a large crossing angle, as at the LHC in 2016 and 2017, using measured longitudinal bunch profiles or a true r.m.s. bunch length measurement becomes crucial. For the machine parameters considered, the error on the emittance measurement using a Gaussian fit knowing the true r.m.s. bunch length is below 5 %.

The impact of a changing longitudinal distribution on emittance scan was directly probed at the LHC in 2016 during the commissioning of the longitudinal bunch flattening at the LHC [28]. During these tests, a sinusoidal RF phase modulation was applied to compensate for the shrinking caused by synchrotron radiation damping by flattening the bunch profiles. The distribution was measured before and after, the change is shown in figure 5.4a.

Before and after the longitudinal bunch flattening, an emittance scan was done in the crossing plane. Results are shown in figure 5.4b, and compared to the yield of equation (5.5) considering the change in longitudinal distribution and intensity, but assuming no change in transverse emittance. The agreement shows that the change of longitudinal distribution is well accounted for by using direct fit considering the measured longitudinal profile.

Table 5.1: Reconstruction of emittances from simulated emittance scans in the crossing plane with a longitudinal q-Gaussian ($q = 0.85$) distribution for LHC machine parameters of 2015, 2016 and 2017.

2015 parameters: $\beta^* = 45\,\text{cm}$, $\alpha = \pm 145\,\mu\text{rad}$, $\sigma_s = 7.5\,\text{cm}$		
	Reconstructed Emittances	
	$\varepsilon_{\text{ref}} = 2\,\mu\text{m}$	$\varepsilon_{\text{ref}} = 3.5\,\mu\text{m}$
Gaussian fit, FWHM bunch length	$1.85\,\mu\text{m}$	$3.34\,\mu\text{m}$
Gaussian fit, r.m.s. bunch length	$2.02\,\mu\text{m}$	$3.51\,\mu\text{m}$
Direct fit using longitudinal profile	$2.00\,\mu\text{m}$	$3.50\,\mu\text{m}$
2016 parameters: $\beta^* = 40\,\text{cm}$, $\alpha = \pm 185\,\mu\text{rad}$, $\sigma_s = 7.5\,\text{cm}$		
	Reconstructed Emittances	
	$\varepsilon_{\text{ref}} = 2\,\mu\text{m}$	$\varepsilon_{\text{ref}} = 3.5\,\mu\text{m}$
Gaussian fit, FWHM bunch length	$1.55\,\mu\text{m}$	$3.02\,\mu\text{m}$
Gaussian fit, r.m.s. bunch length	$2.09\,\mu\text{m}$	$3.57\,\mu\text{m}$
Direct fit using longitudinal profile	$2.00\,\mu\text{m}$	$3.50\,\mu\text{m}$
2017 parameters: $\beta^* = 30\,\text{cm}$, $\alpha = \pm 150\,\mu\text{rad}$, $\sigma_s = 7.5\,\text{cm}$		
	Reconstructed Emittances	
	$\varepsilon_{\text{ref}} = 2\,\mu\text{m}$	$\varepsilon_{\text{ref}} = 3.5\,\mu\text{m}$
Gaussian fit, FWHM bunch length	$1.60\,\mu\text{m}$	$3.08\,\mu\text{m}$
Gaussian fit, r.m.s. bunch length	$2.08\,\mu\text{m}$	$3.56\,\mu\text{m}$
Direct fit using longitudinal profile	$2.00\,\mu\text{m}$	$3.50\,\mu\text{m}$

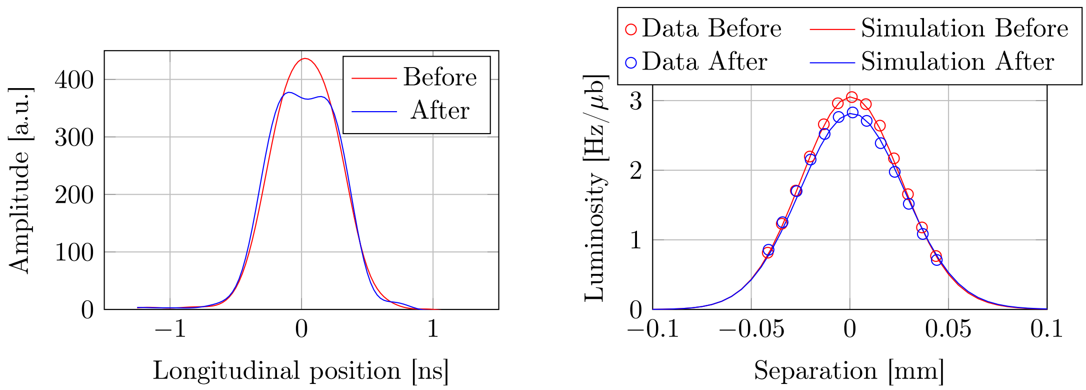

(a) Measured longitudinal profiles.

(b) Scan data and simulation based on the measured profiles.

Figure 5.4: Emittance Scans before and after the “bunch flattening” to change the longitudinal distribution.

As of 2017, neither an operational measurement of the longitudinal bunch profiles nor a true r.m.s. bunch length measurement is available at the LHC. Longitudinal bunch profiles are only acquired and stored by manual intervention on demand. However, since the machine parameters and the longitudinal distribution in the LHC are reproducible, the difference between the reconstruction based on a FWHM measurement and the real emittance reduces to a constant offset at the first order. This allows using the operational FWHM bunch length measurement and correcting the resulting emittance by an offset to be determined by a reference profile measurement and numerical simulations.

Simulations done for emittance scans at the start of Stable Beams yield e.g. a correction of $+0.5 \pm 0.1\,\mu$m for the 2016 LHC parameters and the bunch lengths of $\sigma_\mathrm{s} = 7.8 \pm 0.5\,$cm used operationally, as shown in figure 5.5. The residual error is at the 5 % level, which could be further reduced by applying a bunch length dependent correction. In the following, this is considered the systematic error originating from the longitudinal distribution in the crossing plane scans.

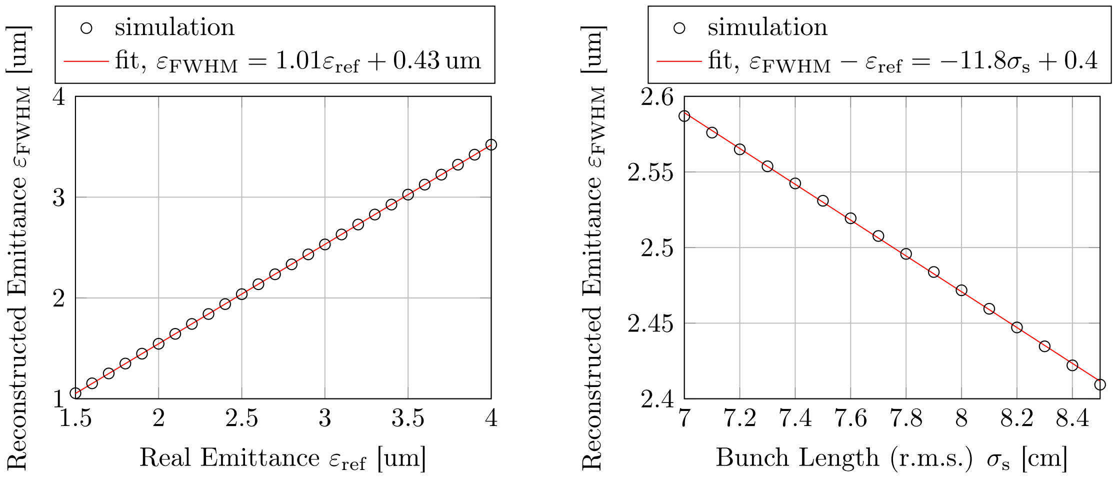

(a) $\sigma_s = 7.5\,$cm

(b) $\varepsilon_\mathrm{ref} = 2.5\,\mu$m

Figure 5.5: Simulation of the emittances reconstructed using a Gaussian fit and the FWHM bunch length compared to the real emittance, for 2016 LHC parameters and a longitudinal q-Gaussian ($q = 0.85$) distribution.

5.4 Emittance Scan Parameters

The scan parameters (the separation steps to scan and the time to acquire data at each step) were optimized to provide an adequate measurement while keeping the scans as short as possible. The steps were chosen to be equidistant and symmetric with respect to the initial position of the beams. A total effective separation of $\sim 3\sigma$ is aimed for, corresponding to a separation factor of $\mathcal{S} \approx 0.1$ and hence a change in luminosity by a factor of $\sim$10. Half of this total separation is applied symmetrically to the orbit of both beams with an opposite sign. Due to the crossing angle, Σ_{xing} is larger and hence the effect of the separation is smaller in the crossing plane. This is accounted for by applying a larger separation in the crossing plane to achieve the same separation factor in both planes.

The operational scan parameters established for the LHC based on these constraints are given in table 5.2. The scan range is given in units of the "LHC nominal" beam size σ_{nom} (assuming a nominal emittance of $\varepsilon_{\text{nom}} = 3.5\,\mu\text{m}$), while the actual beam emittance was $\varepsilon \approx 2.5\,\mu\text{m}$ as of 2016. Scans were done regularly at the CMS experiment which delivered bunch-by-bunch luminosity readings at a high rate (~0.5 Hz). A scan of both planes with these parameters takes ~5 min to complete.

Table 5.2: Emittance scan parameters

Number of separation steps	7
Integration time per step	10 s
Beam separation (crossing plane)	$2.5\,\sigma_{\text{nom}}$
Beam separation (separation plane)	$3.5\,\sigma_{\text{nom}}$

5.5 Error Considerations

At high luminosity, systematic errors are largely predominant, while statistical errors on the luminosity measurement are negligible. The error sources considered and their contribution to the total error of the derived emittances are given in table 5.3. As explained in section 5.3, it is assumed that the impact of the longitudinal distribution on the crossing plane scans is taken into account up to a residual error of 5 %.

Table 5.3: Errors on emittances derived from emittance scans

Error source	**separation plane**	**crossing plane**
luminosity non-linearity	5%	5%
steering magnet calibration	5%	5%
β^* (optics uncertainty)	3%	3%
head-on beam-beam kick	2%	2%
dynamic β^* (beam-beam)	2%	2%
transverse distribution	5%	5%
crossing angle	-	5%
longitudinal distribution	-	5%
combined error	9.6 %	11.9%

It is to be noted that most systematic error sources only affect the absolute scales of the derived beam sizes and emittances. Relative bunch-to-bunch differences and other fit parameters (in particular the parasitic separations d_0) are not affected. Also, only the luminosity non-linearity due to pile-up effects and the longitudinal bunch shape are expected to change over the course of a fill, and therefore to possibly affect the time evolution of the emittance.

5.6 Emittance Measurements

As a benchmark of the emittance scan method, the measured emittances are compared to those of other instruments, in particular the BSRT and the emittances derived from the absolute luminosity measurements in the ATLAS and CMS experiments. Using these measurements, the evolution of emittance in the LHC and its impact on luminosity performance will be discussed in greater detail in chapter 6.

Bunch-by-bunch emittances for a typical LHC fill of 2017 are shown for all bunches in figure 5.6, and details for 6 typical bunch trains in figure 5.7. The measurements from the BSRT agree well within the expected systematic and statistical uncertainties and all bunch-by-bunch features are well represented. As explained in the previous section, the systematic uncertainties affecting the absolute emittance scale for all bunches are strongly dominant over per-bunch statistical or systematic errors.

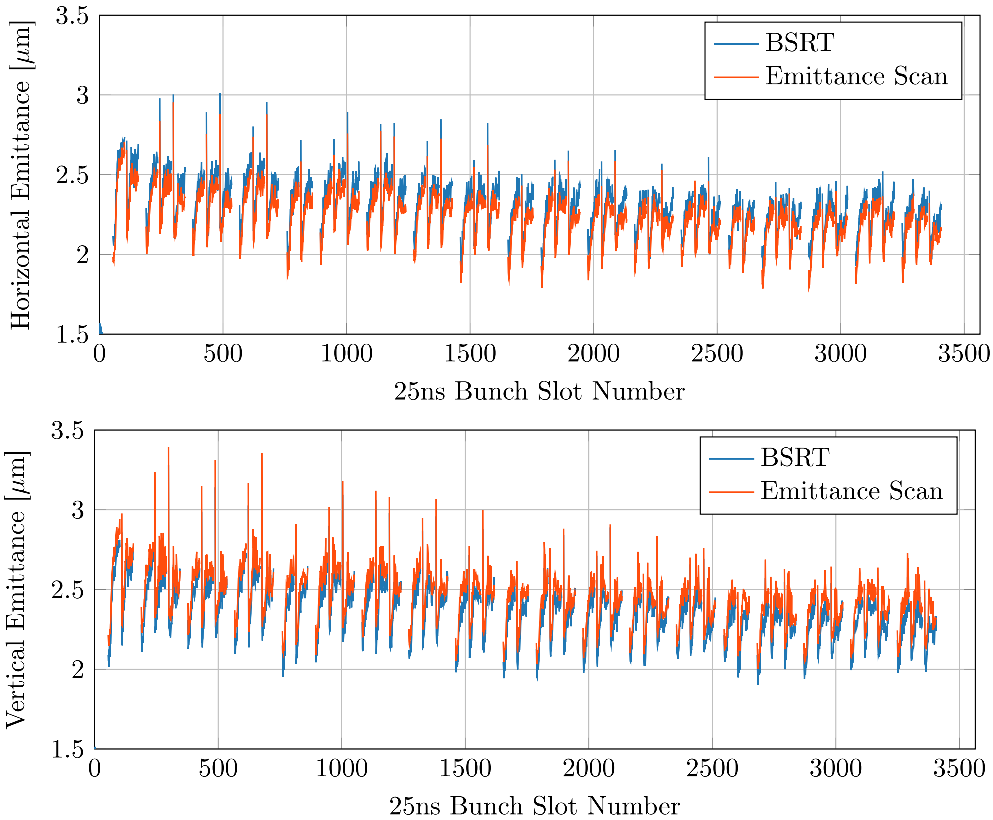

Figure 5.6: Bunch-by-bunch emittances of all 2556 bunches in the LHC (fill 6054, 2017).

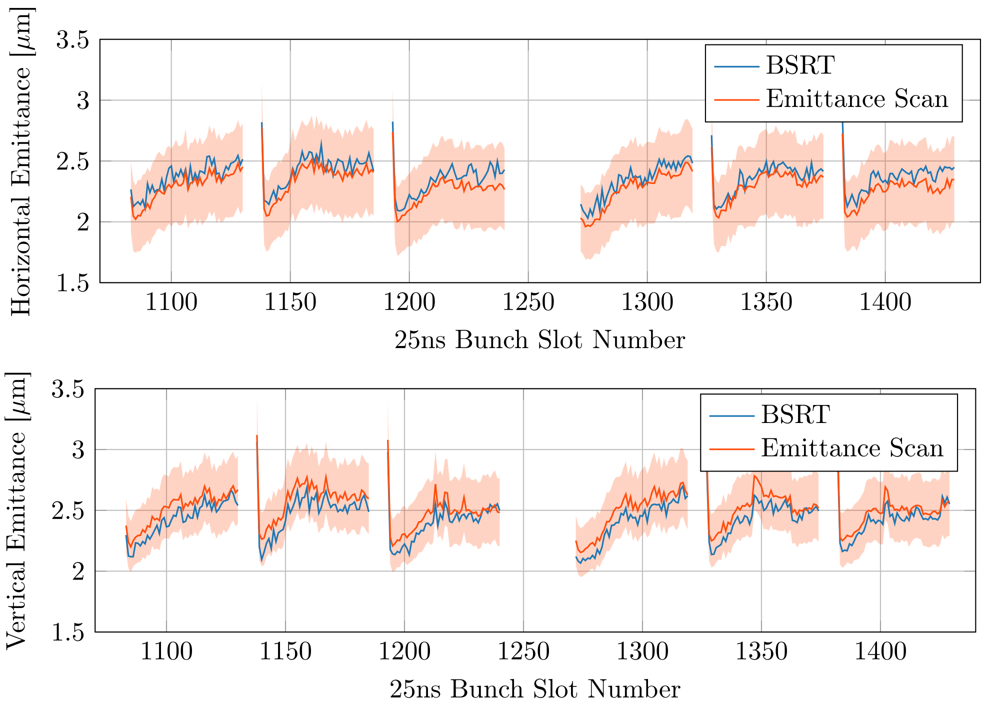

Figure 5.7: Bunch-by-bunch emittances of 6 typical bunch trains (fill 6054, 2017). The shaded error region for the emittance scan data includes both systematic and statistical uncertainties.

The bunch-by-bunch convoluted transverse emittances $\varepsilon_c = \sqrt{\varepsilon_x \varepsilon_y}$ can be compared to the emittances obtained from the bunch-by-bunch luminosity measurements from ATLAS and CMS by inverting equation (3.8). This was the first benchmark done during the commissioning of the emittance scans in 2015 [36]. A comparison for a typical LHC fill of 2015 is shown in figure 5.8.

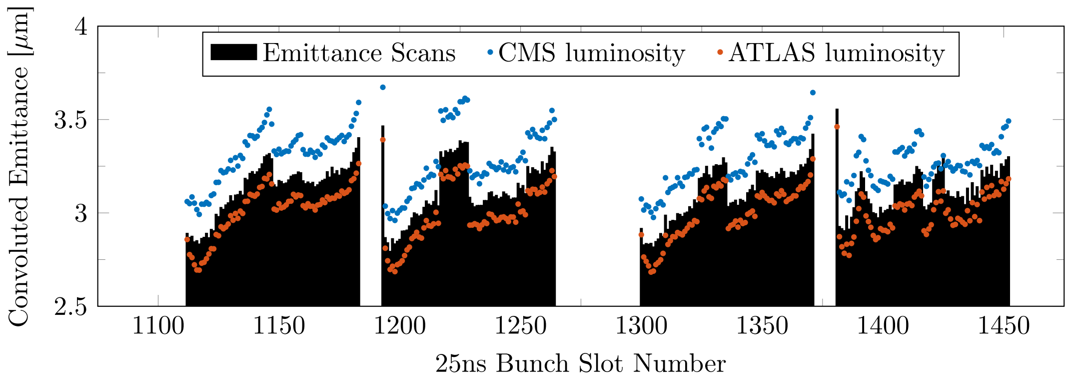

Figure 5.8: Bunch-by-bunch convoluted emittances of 4 bunch trains (fill 4400, 2015).

On a higher level, one can compare the measured emittances, averaged over all bunches, for all fills when Stable Beams is declared and the first emittance scan is done. This allows for a cross-check of all available emittance measurements over longer periods of time, e.g. over a year of LHC running. In figure 5.9, this comparison is shown for the 2017 LHC run.

While all measurements agree well within the expected uncertainties for most of the year, it is worth noting that the measurements from the BSRT diverged from the others in three periods (around fills 6025, 6180 and 6300), caused by the degradation of the BSRT image intensifier [26]. This stresses the importance of complementary, independent transverse emittance measurements. By comparing the measured emittances from the BSRT to measurements from emittance scans, and by cross-checking the plausibility against the luminosity measured by the experiments, this degradation was spotted timely and the synchrotron light spot on the BSRT image intensifier was moved to avoid the degraded area.

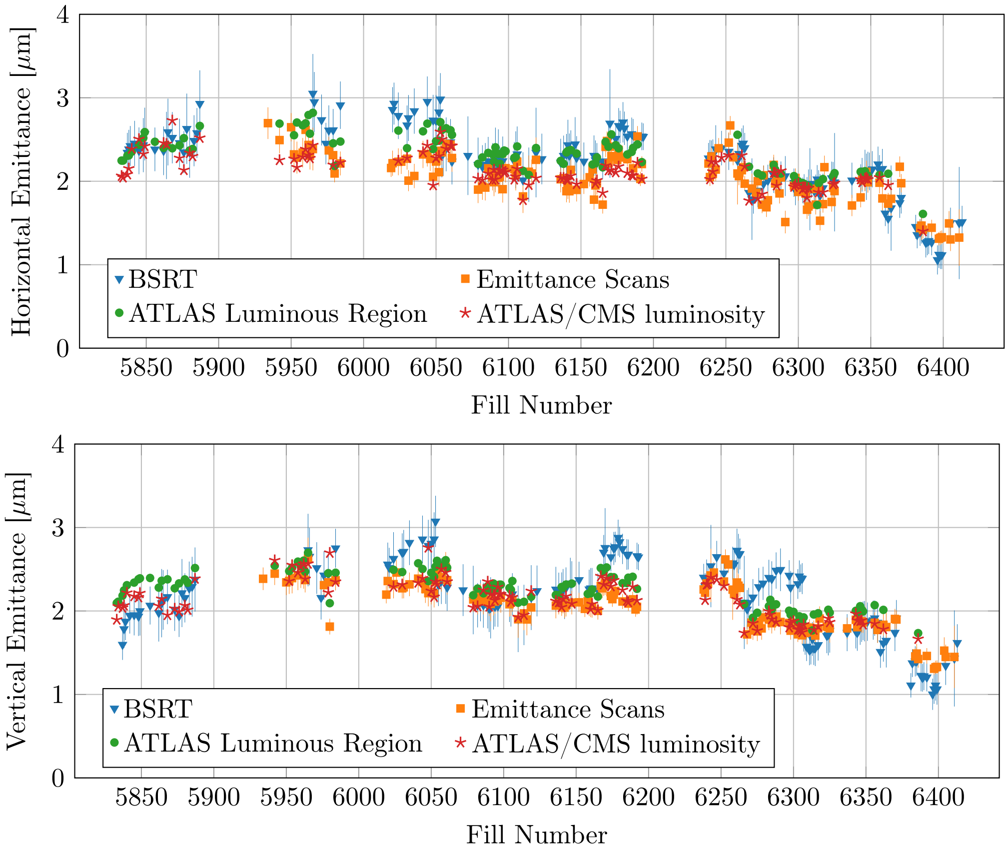

Figure 5.9: Average emittances at the start of Stable Beams of each fill in the 2017 LHC run. This plot includes all emittance measurements currently available at the LHC.

5.7 Closed-Orbit Offsets and PACMAN effects

The measured parasitic separations $d_{u,0}$ show a reproducible pattern dependent on the LHC filling scheme. This is explained by long-range beam-beam and "PACMAN" effects (see section 3.3). As explained in section 4.3, there are 15 possible long-range beam-beam encounters at each side of each IP in the LHC between the IP and the D1 separation dipole magnet [21], where the two beams are separated by the crossing angle, and up to 45 between the D1 and the D2 until the design orbits of the two rings reach their full separation and enter separate beam pipes. Due to the gaps in the filling scheme, some bunches have a different number of long-range beam-beam encounters, which leads to a difference in beam-beam kicks, and hence, to a different closed orbit. The number of long-range encounters per IP up to the D1 for a typical LHC filling scheme is given in figure 5.10.

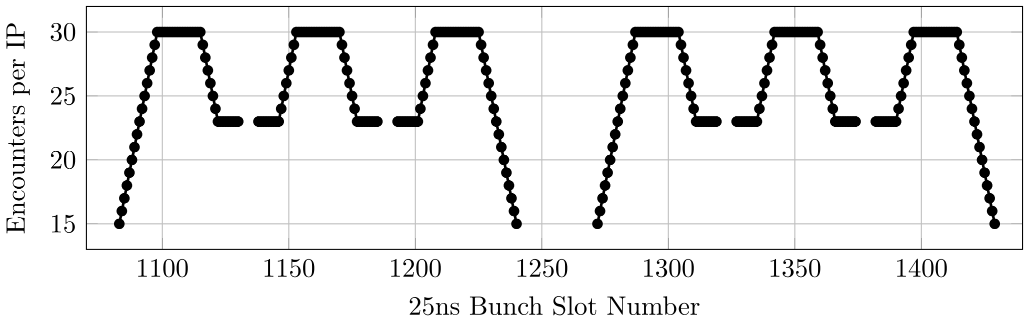

Figure 5.10: Number of long-range beam-beam encounters between the IP and the D1 for 6 bunch trains in LHC fill 5976 (2017).

TRAIN [41] is a FORTRAN code to simulate self-consistent closed orbits on a bunch-by-bunch level under the presence of long-range beam-beam effects originally developed at the LEP. It acts as a post-processor to the accelerator design code MAD-X [42]. First, the first-order transport matrices (linear optics) and second-order maps for the accelerator are generated in MAD-X, and markers are added where the beam-beam long-range interactions happen around the IPs. These are then used by the TRAIN code to add the kicks from long-range beam-beam interactions and find the closed orbits for all bunches. Since the kicks depend on the separation of the two beams around the IP, which depends on the closed orbits, this is done iteratively until convergence is achieved.

The simulations take into account the filling scheme as well as differences in intensities and emittances. The simulation gives self-consistent bunch-by-bunch closed orbits for either beam, as well as bunch-by-bunch tunes and chromaticities. The difference between the closed orbits of the two beams at the IP yields a separation, which can be observed during emittance scans.

In earlier versions, the TRAIN simulation only took into account the 15 long-range encounters from the IP up to the D1. This approximation only had a small effect for the LHC nominal design optics and its variations used at the LHC up to 2016. However for the "ATS" [32] optics used in the LHC as of 2017, the additional encounters between

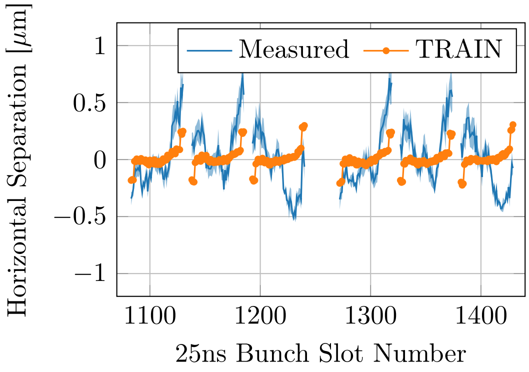

(a) IP5 horizontal

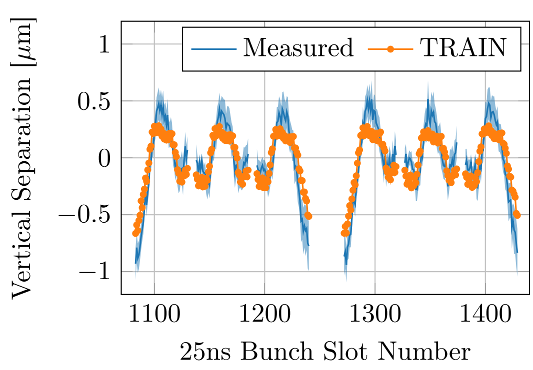

(b) IP5 vertical

Figure 5.11: Bunch-by-bunch separation at the IP for typical bunch trains (fill 5976, ATS optics, 2017). This TRAIN simulation considered only the 15 long-range beam-beam encounters up to the D1.

the D1 and D2 sum up to non-negligible orbit kick. This led to a discrepancy between measurements and simulation results. Since the design orbit separation is horizontal in both IP1 and IP5, but the crossing angle is vertical in IP1 and horizontal in IP5, this mostly affected IP5, where both the design orbit separation and the crossing angle are in the horizontal plane. An example is given in figure 5.11.

Following these results, I extended the TRAIN code to support an arbitrary number of long-range beam-beam encounters, including those with design orbit separation. As the first and second order maps provide no information on the design orbit, this requires passing it separately from MAD-X to TRAIN using a MAD-X “survey” file. In TRAIN, the separations from the closed orbit and the design orbit are then added to get the full beam-beam separation. For the LHC, the contributions of the crossing angle (closed orbit) and the design orbit separation are shown in figure 5.12. It should be noted that these contributions are both in the horizontal plane for IP5, but perpendicular for IP1.

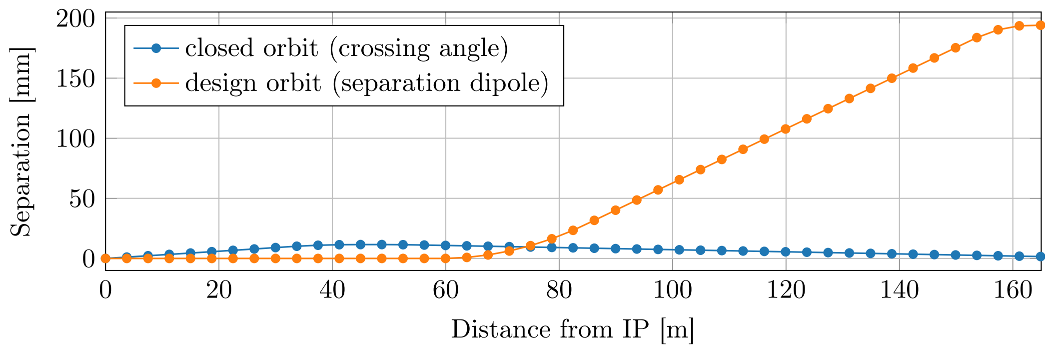

Figure 5.12: Beam-beam separation given by the crossing angle and the design orbit. The dots indicate the positions of possible long-range beam-beam encounters.

With this improvement in place, the full 45 beam-beam encounters between the IP and the D2 in the LHC can be considered by the simulation. As shown in figure 5.13, the extra encounters significantly change the orbit pattern in the horizontal plane, mainly in IP5. The simulation considering all encounters is then compared to the emittance scan measurements for an LHC fill with "ATS" optics in figure 5.14. The improved simulation agrees with the measurement within the expected uncertainties in both planes. These results are also reproducible for different LHC filling schemes: an example for an "8b4e" filling scheme consisting of short trains of 8 bunches spaced by 4 empty bunch slots [43] is given in figure 5.15.

The agreement obtained qualifies both the emittance scans as a method to measure these effects and the improved version of the TRAIN code to accurately simulate them, which can in turn be applied to future accelerators like the HL-LHC [44] or the FCC-hh [45].

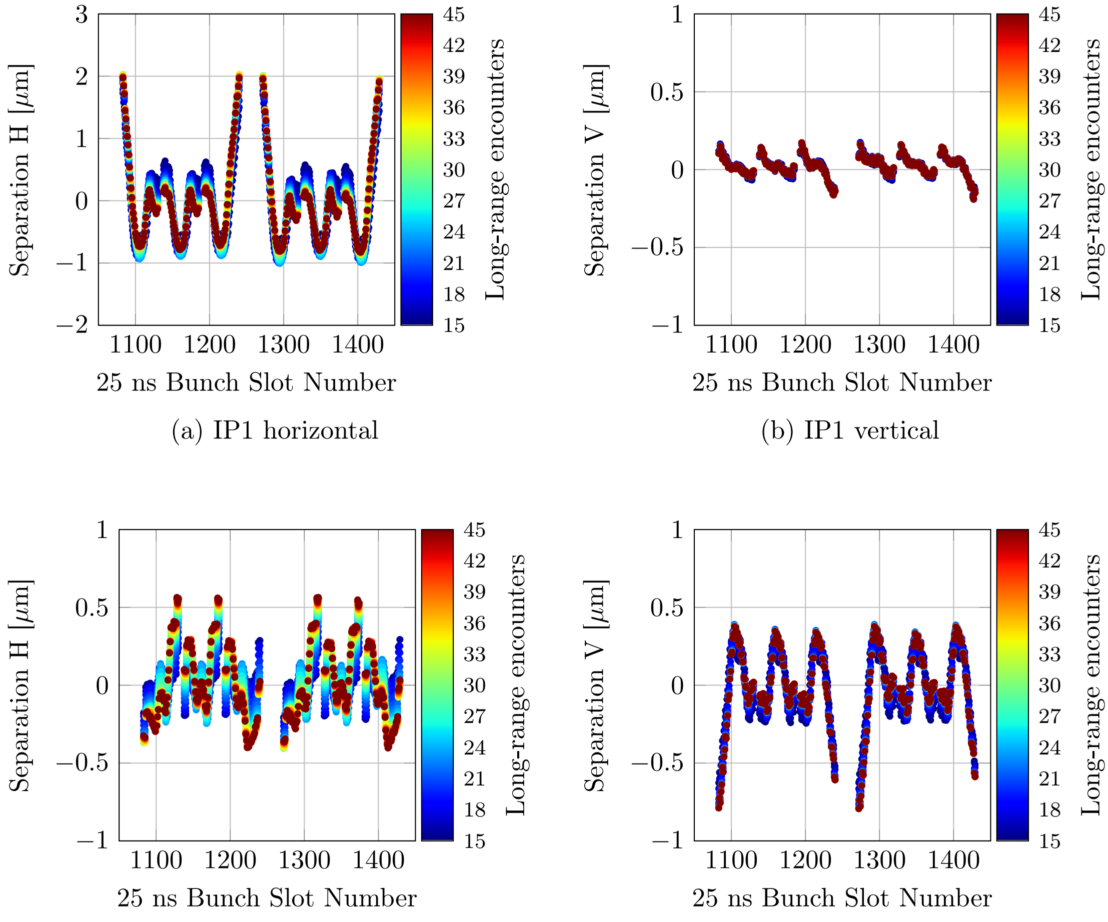

(a) IP1 horizontal

(b) IP1 vertical

(c) IP5 horizontal

(d) IP5 vertical

Figure 5.13: Simulated bunch-by-bunch orbit separations considering a varying number of beam-beam long-range encounters around each IP (ATS optics).

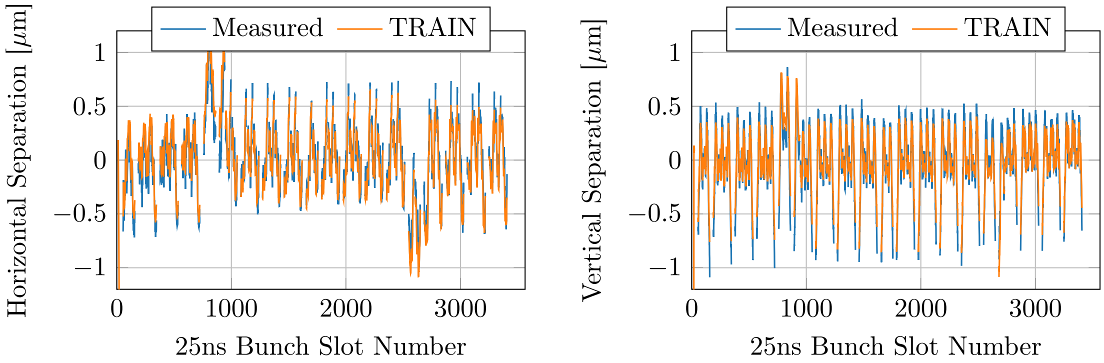

(a) IP5 horizontal (full beam) (b) IP5 vertical (full beam)

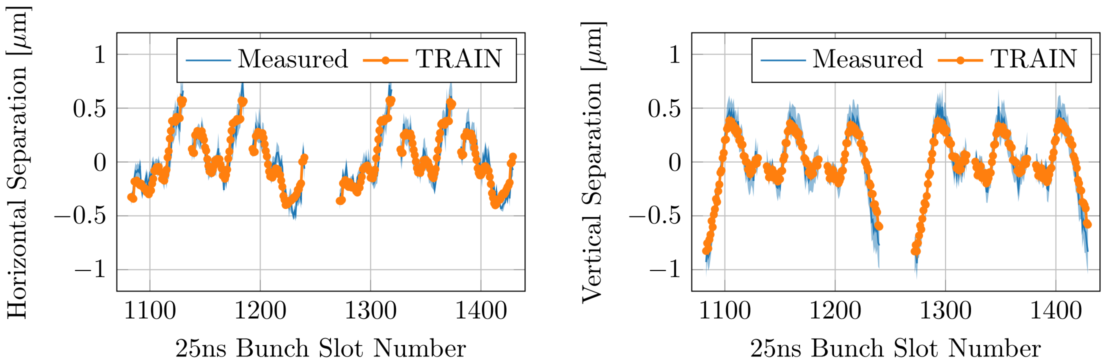

(c) IP5 horizontal (typical trains) (d) IP5 vertical (typical trains)

Figure 5.14: Bunch-by-bunch separation (fill 5976, ATS optics, 2017). The TRAIN simulation considers 45 long-range beam-beam encounters per IP side.

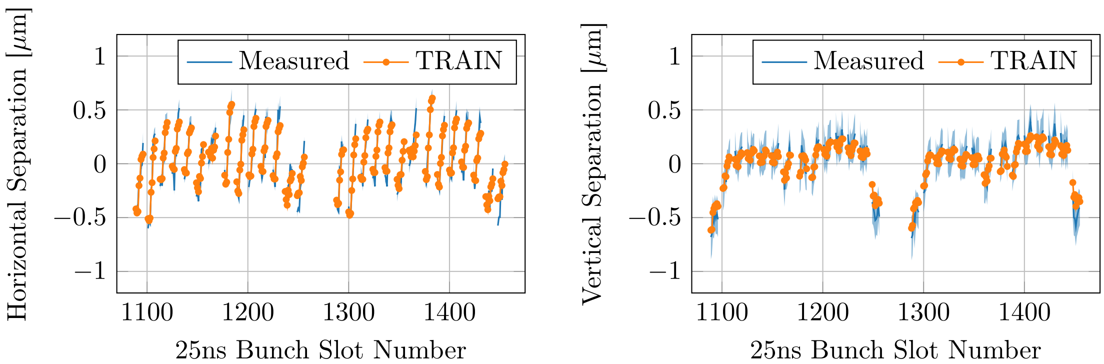

(a) IP5 horizontal (b) IP5 vertical

Figure 5.15: Bunch-by-bunch separation for an “8b4e” filling scheme (fill 6169, ATS optics, 2017). The simulation considers 45 long-range beam-beam encounters.

5.8 Other Applications

In addition to the usage of the emittance scans for deriving beam parameters on the LHC side, the CMS collaboration used them independently for tracking the calibration and performance of their luminosity monitors as of 2017 [46, 47]. The analysis is following that of the "van der Meer" luminosity calibration scans [35]. The calibration constant linking the measured rate R during the crossing of a bunch pair to the luminosity of that bunch pair is the *visible cross-section* σ_{vis} [48]

$$\mathcal{L}_{bb} = \frac{R}{\sigma_{vis}} \tag{5.7}$$

Following equation (3.8), $\mathcal{L}_{bb}$ can be calculated from the bunch intensities, measured by the FBCT, and the convoluted beam sizes $\Sigma_{x/y}$ measured by the emittance scans. From the measured rate and the expected luminosity, σ_{vis} is derived [48].

As the emittance scans were done regularly in every LHC fill, this allowed tracking the visible cross-section over the course of 2017, while previously there were only measurements during the dedicated "van der Meer" calibration sessions. An example is given in figure 5.16. While this was not a precision measurement, it allowed spotting e.g. degradations of the PLT luminosity monitor efficiency quickly, which were then mitigated by increasing the high voltage [46]. This improved the overall quality of the online luminosity measurement.

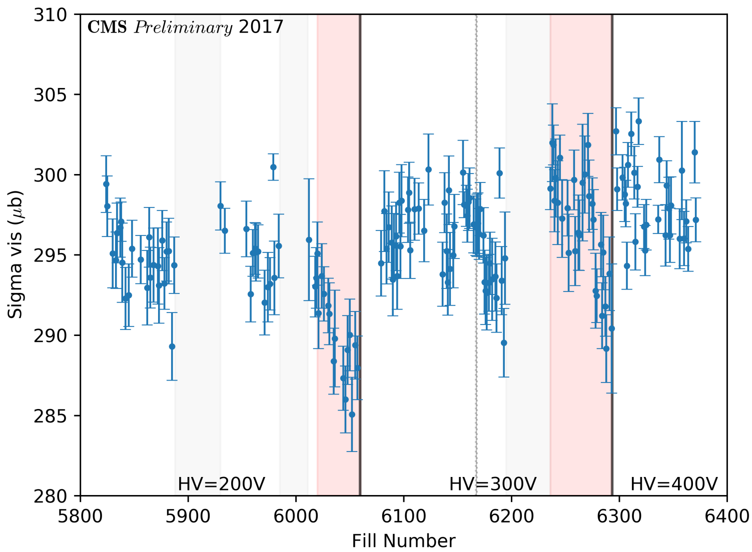

Figure 5.16: The visible cross-section of the CMS PLT luminosity monitor in 2017 operation, measured by emittance scans at the start of Stable Beams [49]. Over the periods highlighted in red, a decrease of efficiency was observed which was mitigated by increasing the operational high voltage [46].

6 LHC Luminosity and Optimization

6.1 Motivation

Delivering the maximum integrated luminosity to the experiments is one of the most important challenges in LHC operation. The LHC delivered the first luminosity to the experiments in 2010. After an initial operation for three years, the LHC was shut down in 2013-2014 for major upgrades and maintenance work ("Long Shutdown 1"). The machine was re-commissioned in 2015, and the 2016 and 2017 runs were dedicated to luminosity production. In the period of 2015-2017, the LHC delivered more than 90 fb^{-1} of integrated luminosity to the ATLAS and CMS experiments (figure 6.1). For 2018, an integrated luminosity of $\sim$70 fb^{-1} is expected [50].

To keep improving the LHC luminosity, it is crucial to understand the reach and evolution of the key beam parameters: the intensity and the emittance. Also, as shown in section 3.2, the luminosity decays once the beams are in collisions. The integrated luminosity reach hence depends both on the *peak luminosity* once the beams are brought into collisions, and on the luminosity time evolution once in collisions, the rate of decay.

In 2015-2018 LHC operation, I was part of the LHC luminosity studies group to model, understand and predict the LHC luminosity performance. An overview and some key results are given in the following.

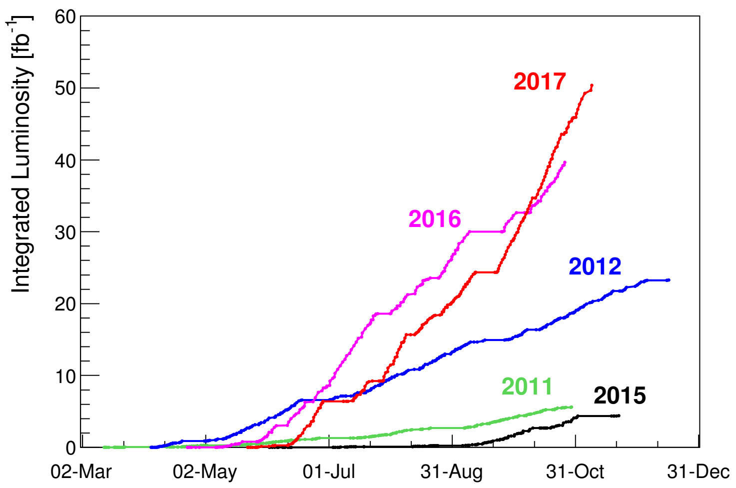

Figure 6.1: Integrated luminosity delivered to ATLAS and CMS in the past years of LHC operation [51]. The main production years were 2012, 2016 and 2017.

6.2 Peak Luminosity

Following equation (3.8), the number of bunches and the intensity and emittance of each bunch are the key parameters defining the luminosity. As the intensity decays while in collisions[1], the *peak luminosity* of the LHC is achieved at the beginning of Stable Beams, directly after bringing the beams in collisions.

The peak luminosities achieved in the two high-luminosity experiments during 2015-2017 LHC operation are shown in figure 6.2. A year always starts with a commissioning phase followed by an *intensity ramp-up*. During these periods, for machine protection reasons, the number of bunches is gradually increased until the maximum number of bunches is reached. Once conditions are stable, the bunch intensity is gently increased to maximize the luminosity yield.

In the past years of LHC operation, the luminosity reach was limited by technical difficulties. In 2016, beam-induced outgasing of one of the injection kicker magnets limited the maximum intensity [52]. In 2017, a major vacuum problem ocurred in a magnet interconnection of cell 16L2 [53] which led to frequent unplanned beam dumps and significantly impacted the luminosity (fills 6070 to 6200). For mitigation, two approaches were followed with the goal to reduce the beam-induced heating [54] at the affected cell: First, the "8b4e" filling schemes were introduced; these beams consist of short trains of 8 bunches separated by 4 empty bunch slots [43, 53]. Secondly, a solenoid was installed around the affected magnet interconnection [55]. With these measures in place, the original performance could be recovered. However, the "8b4e" filling schemes reduce the total number of bunches that can be injected into the LHC. While this was compensated for by higher intensity and lower emittances, it significantly increased the pile-up, which required luminosity levelling in ATLAS and CMS to $1.5 \cdot 10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$ (see chapter 7).

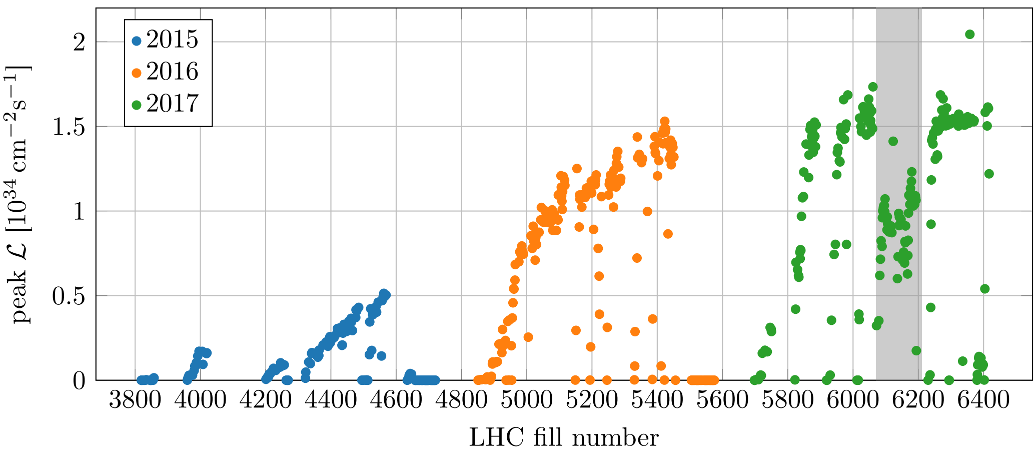

Figure 6.2: Measured peak luminosity at the CMS experiment in 2015-2017 operation. During the grey shaded period, the luminosity was limited by the "16L2" issue.

[1]The number of bunches remains constant throughout a fill, and the emittance only changes slightly, leaving the intensity decay as the predominant effect.

The LHC design luminosity of $10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$ was first surpassed in 2016. An important factor to make this possible was the reduction of the transverse emittance in the LHC injector chain while keeping the bunch intensity constant. The emittances measured in the LHC at the start of Stable Beams are shown in figure 6.3.

The emittance of the nominal LHC beam is $3.75\,\mu\mathrm{m}$ [6]. Beam of this type was used in 2015 and early 2016. In July 2016 (fill 5079) a new beam production scheme was made operational in the LHC injector complex, the "Batch Compression Merging and Splitting" (BCMS, [56, 57]), which reduced the emittance to $\sim 2\,\mu\mathrm{m}$ at the start of Stable Beams.

The "8b4e" beam, used in 2017 to mitigate the 16L2 vacuum issues, was originally produced in the injector chain using the "nominal" production scheme, yielding emittances of $\sim 3\,\mu\mathrm{m}$ at the start of Stable Beams. As of fill 6200, the production mechanism was changed to the BCMS-like "BCS" [58] production scheme. This allowed for emittances as low as $\sim 1.75\,\mu\mathrm{m}$ at the Start of Stable Beams, which lead to a new luminosity world record of $2.2 \cdot 10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$ before levelling (the previous record was $2.1 \cdot 10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$, reached in 2009 at KEKB [59]).

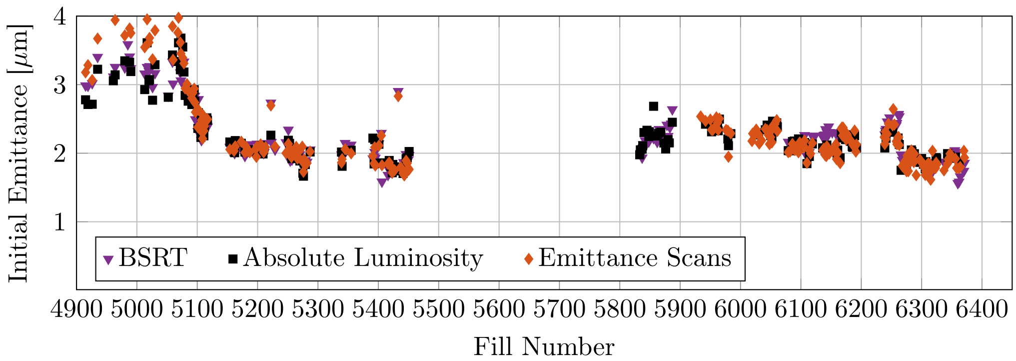

Figure 6.3: Convoluted emittances in 2016 and 2017 LHC operation at the start of Stable Beams. Measurements of known bad quality are excluded (e.g. periods when the BSRT intensifier was damaged).

6.3 Luminosity Time Evolution

The luminosity evolution for typical LHC fills of 2015, 2016 and 2017 are shown in figure 6.4. The decay of the luminosity is dominated by the beam intensity decay due to luminosity burn-off in the long term. Therefore, the fills with higher peak luminosity show a faster decay.

The total intensity losses can be quantified in terms of the *effective cross-section* following (3.17) as

$$\sigma_{\mathrm{eff}}(t) = \frac{\mathrm{d}\mathcal{N}(t)}{\mathrm{d}t} \frac{1}{\sum\limits_{\mathrm{IPs}} \mathcal{L}(t)} \tag{6.1}$$

where $\mathcal{N}$ is the total beam intensity of either LHC beam.

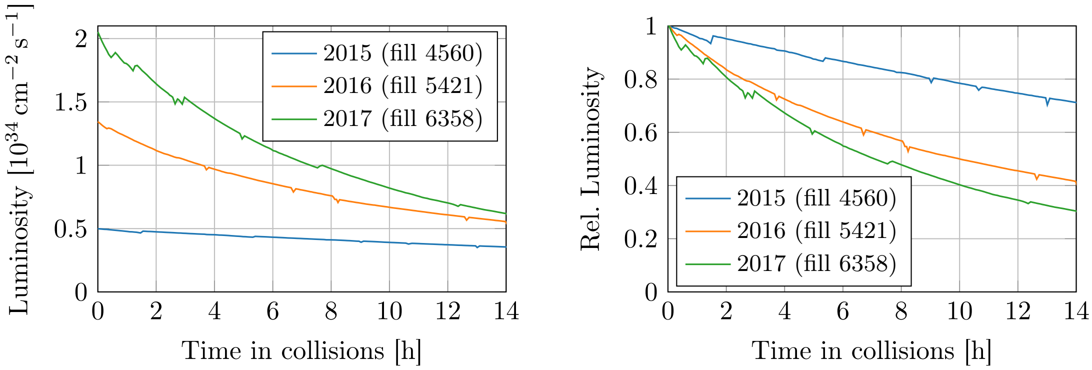

(a) absolute (b) relative to peak

Figure 6.4: Measured luminosity evolution at the ATLAS experiment over typical fills of 2015, 2016 and 2017[2] LHC operation.

If σ_{eff} is close to the proton-proton inelastic scattering cross-section of $\sigma_{\text{inel}} \approx 80\,\text{mb}$ [60] losses are dominated by burn-off (and therefore the intensity is optimally converted into collisions and luminosity), while a larger effective cross-section is a result of particles lost through other mechanisms.

The evolution of the effective cross-section in 2016 LHC operation is shown in figure 6.5. In the first few hours of Stable Beams, additional losses beyond burn-off are observed, which were typically higher for Beam 1 than for Beam 2 [61]. Afterwards the effective cross-section becomes compatible with inelastic scattering cross-section for both beams, indicating that losses are strongly dominated by burn-off [62].

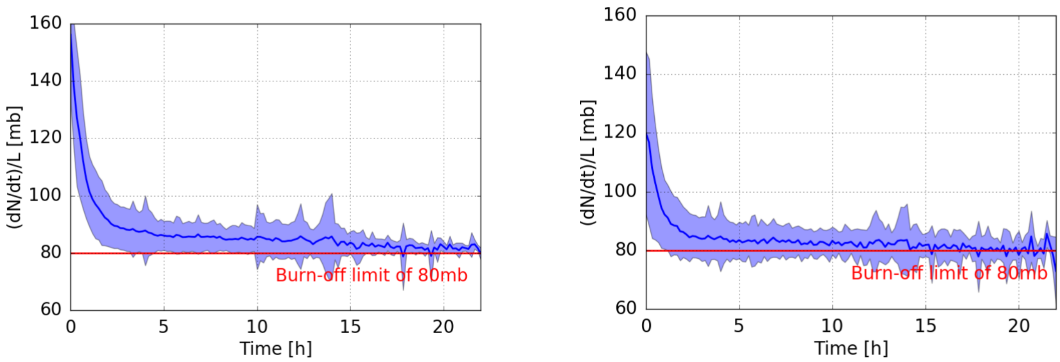

(a) Beam 1 (b) Beam 2

Figure 6.5: Effective cross-section, statistics of all physics fills in 2016 LHC operation [62].

[2]The 2017 fill shown is with "8b4e-BCS" type beams and without luminosity levelling. Earlier 2017 fills with "BCMS" type beams evolve like the 2016 fills. Luminosity levelling and its impact on the luminosity evolution will be discussed in chapter 7.

The evolution of the emittance in collisions depends on the initial emittances, beam parameters and machine settings. For LHC nominal beam with initial emittances of ∼3.5 um as in 2015, the emittances were shrinking at a rate of $-0.03\,\mu$m/h due to synchrotron radiation damping, a behaviour observed for the first time ever in a hadron accelerator [36]. With the lower initial emittances of ∼2 um of BCMS beams and more aggressive machine settings in 2016 and 2017, an emittance growth of $0.05\,\mu$m/h was observed [63]. An emittance model including synchrotron radiation, intra-beam scattering and beam-beam elastic scattering does not fully explain the observed emittance growth, indicating that additional sources, e.g. external noise, contribute to the emittance evolution [13, 62, 63].

6.4 Optimum Fill Time

Due to the LHC operational cycle (see section 4.8), each period of luminosity production is followed by a period during which no luminosity is produced (turnaround time). In order to optimize the overall integrated luminosity delivered to the experiments, the average luminosity produced over a full machine cycle must be optimized. It is given by

$$\langle \mathcal{L} \rangle_{\text{fill}} = \frac{\int_0^{T_{\text{SB}}} \mathcal{L}(t)\,\mathrm{d}t}{T_{\text{SB}} + T_{\text{TA}}} \tag{6.2}$$

where T_{SB} is the time spent for luminosity production in Stable Beams and T_{TA} is the turnaround overhead time. While T_{TA} is given by the minimal cycle time and possibly the effect of technical faults which delay the cycle, the operators can choose when to dump a fill, defining T_{SB} (unless a fill is dumped prematurely by a machine protection interlock).

Using luminosity measurements from previous fills as a reference, the fill time T_{SB} can be optimized to maximize the integrated luminosity with minimal assumptions. For $\mathcal{L}(t)$, any functional form fitted to the measured luminosity evolution can be used provided that fit quality is good. Possible approaches are a double-exponential [64] or a generalized form of the fractional power law (3.17) that would result from pure intensity burn-off [65]

$$\mathcal{L}(t) = \frac{\mathcal{L}_0}{(1 + t/\tau)^{\mu}} \tag{6.3}$$

where $\mathcal{L}_0$, μ and τ are fit parameters.

Fits to the measured luminosity over an example LHC fill using these two forms are shown in figure 6.6; other fills studied are similar in fit quality. For the fills studied, the fit residuals are smaller than 3% and dominated by machine effects, e.g. drifts of the orbit and luminosity optimization scans. The difference between the two fits is smaller than 0.5%. In the following, the fractional power law fits are used.

Figure 6.7 shows the optimum fill times with respect to the turnaround time, obtained through univariate optimization of (6.2) inserting a fit of (6.3) to LHC fills of the past years of LHC running. Note that the optimum fill time depends significantly on the peak luminosity and the luminosity evolution. While the LHC in general has a good fill-to-fill

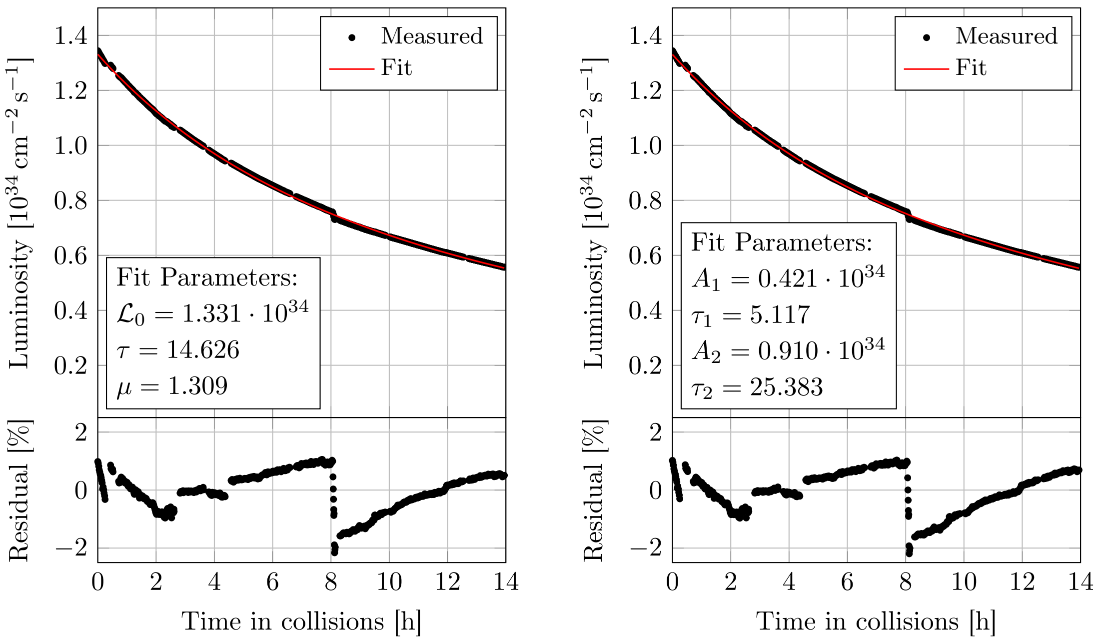

(a) fractional power law (b) double exponential

Figure 6.6: Fits to the luminosity evolution of LHC fill 5421 at the ATLAS experiment. Fit residuals are dominated by orbit drifts and the increase of bunch length after ~8 h.

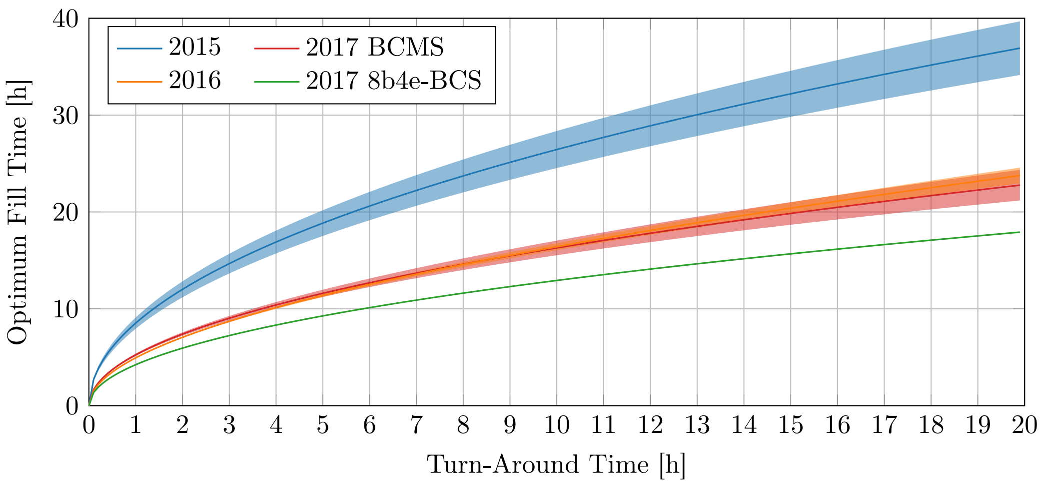

Figure 6.7: Optimum fill length for 2015, 2016 and 2017 luminosity evolution. The shaded error regions show the fill-to-fill spread. The 2017 “8b4e-BCS” case has low statistics due to the use of luminosity levelling for most fills.

reproducibility of the luminosity, the optimization must be redone whenever the beam parameters or machine settings are changed significantly, e.g. for different beam types, different number of bunches and changes of β^* or the crossing angle.

A histogram of the turnaround times observed in 2016 operation [66] is given in figure 6.8a. The theoretical minimum turnaround time is ∼2 h. In 40% of the cases, the turnaround time is 2.5-4 h, with small delays introduced by routine operations which are not fully automated [66] (e.g. setting up the beam production in the injectors or measuring and correcting beam parameters at injection energy). Minor technical problems during the re-filling (another 40% of the cases) typically yield turnaround times of 4-10 h, e.g. if the beams are dumped by an interlock before the new fill reaches Stable Beams or a short intervention is needed to repair faulty equipment. In the remaining 20% of the fills, major technical problems introduced machine downtimes longer than 10 h, e.g. when malfunctioning equipment required an extended intervention in the LHC tunnel.

Following the expected turnaround time, the integrated luminosity yield per week of operation is extrapolated in figure 6.8b. Around the optimum, the luminosity yield is only moderately sensitive to small deviations of the fill time. This allows keeping fills slightly longer than the calculated optimal fil time without losing more than 1-2% of integrated luminosity in ATLAS and CMS. Doing so increases the integrated luminosity delivered to the levelled low-luminosity experiments (ALICE and LHCb) and allows for an operational contingency, e.g. to avoid re-filling during shift handover periods. For example, following this analyis, a fill length of 15-18 h was targeted in 2016 (for a calculated optimum of 15 h for the average turnaround time) [67].

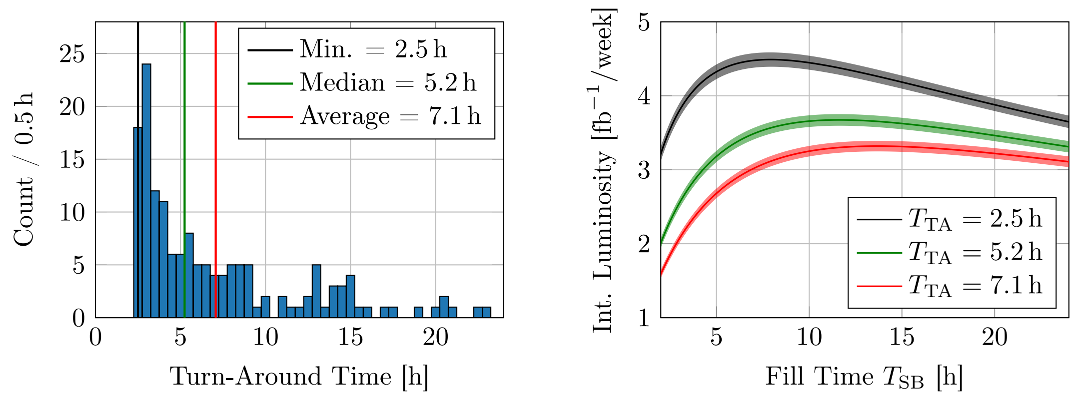

(a) turnaround times

(b) extrapolated integrated luminosity

Figure 6.8: Statistics of turnaround times and extrapolated integrated luminosity per week of 2016 LHC operation.

6.5 ATLAS/CMS Luminosity Ratio

During 2016 operation, the luminosities measured by the ATLAS and CMS experiments were different by up to 10% [68]. Delivering different integrated luminosity to the two high-luminosity experiments poses a problem as it leads to a difference in statistics, and hence a difference in discovery potential, for the two collaborations.

Since IP1 (ATLAS) and IP5 (CMS) are diametrically opposed, and the same bunch pairs collide in both IPs, one would expect the same luminosity in both experiments. However, this symmetry is broken since the crossing angle is vertical in IP1 and horizontal in IP5. Following equation (3.13), if the beams are not round ($\varepsilon_x \neq \varepsilon_y$), this leads a luminosity ratio equal to the ratio of the geometric factors[3]:

$$\mathcal{R} = \frac{\mathcal{L}_{\text{ATLAS}}}{\mathcal{L}_{\text{CMS}}} = \frac{\mathcal{F}_{\text{IP1}}}{\mathcal{F}_{\text{IP5}}} = \frac{\sqrt{1 + \left(\alpha \frac{\sigma_s}{\sigma_y}\right)^2}}{\sqrt{1 + \left(\alpha \frac{\sigma_s}{\sigma_x}\right)^2}} \tag{6.4}$$

The measured peak luminosities and luminosity ratios during 2016 are shown in figure 6.9. The ratio of the total integrated luminosities over the run was $\mathcal{R} = 0.94$. The trends over the year were introduced by machine setting changes. As explained in section 6.2, the "BCMS" beam production scheme was used as of fill 5079, reducing the emittance. Starting from fill 5300, the nominal crossing angle was reduced from $\pm 185\,\mu$rad to $\pm 140\,\mu$rad.

The systematic error on the observed ATLAS/CMS luminosity ratio was 4.2 % from uncorrelated errors of 3.4 % [69] on the ATLAS luminosity and 2.5 % [48] on the CMS luminosity.

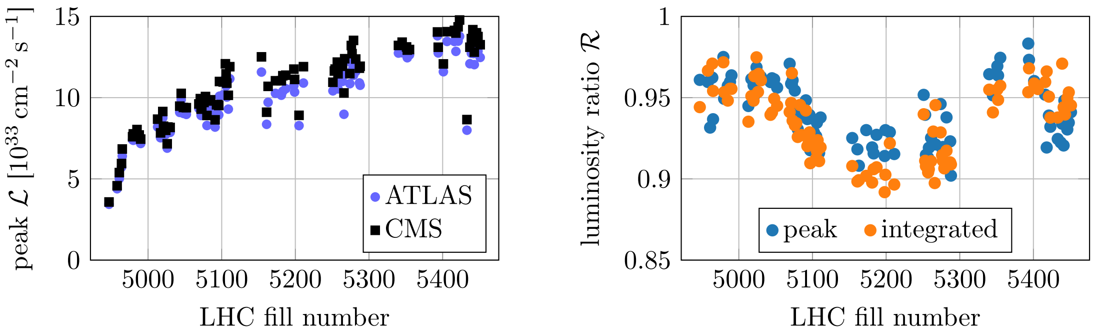

(a) peak luminosities (b) luminosity ratio

Figure 6.9: Peak Luminosity and ATLAS/CMS luminosity ratios in 2016 LHC operation.

The measured emittances in 2016 are shown in figure 6.10. After the introduction of the "BCMS" scheme, typical emittances were $\varepsilon_x \approx 2.5\,\mu$m in the horizontal and $\varepsilon_y \approx 2.0\,\mu$m in the vertical plane. As illustrated in figure 6.11, the expected luminosity ratio following equation (6.4) for these emittances is $\mathcal{R} \approx 0.95$. The predicted values are compared to the measured ones in figure 6.12, and a good agreement is observed.

[3] under the assumption of equal β^* and equal crossing angles in IP1 and IP5

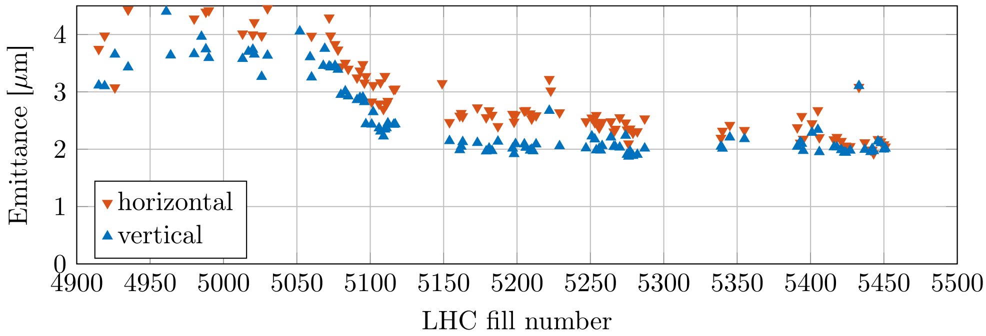

Figure 6.10: Measured emittances at the start of Stable Beams in 2016. Shown values are from emittance scans. BSRT values are not shown, but are in agreement.

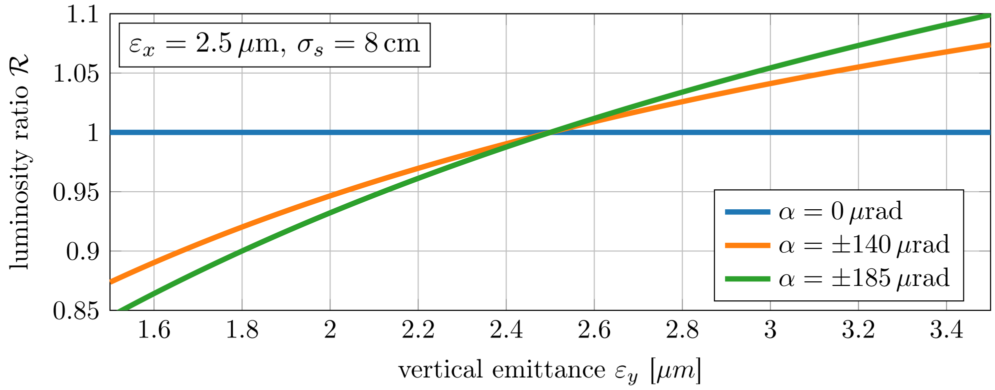

Figure 6.11: Predicted luminosity ratio for the two nominal crossing angles used in 2016 operation and typical beam parameters.

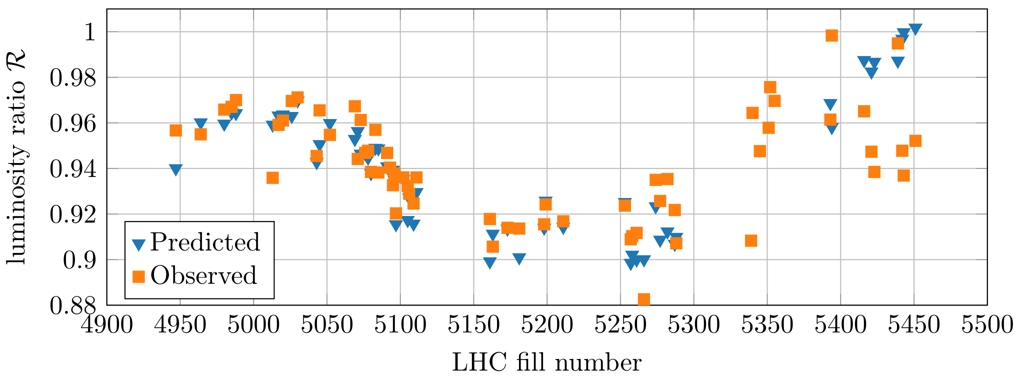

Figure 6.12: Predicted and observed luminosity ratio in 2016 operation.

To directly probe the effect of the crossing angle on the luminosity ratio, a special zero crossing angle test fill was carried out. For this test, the LHC was filled with 4 longitudinally spaced colliding bunch pairs of different emittances. Once in collisions, the crossing angle was reduced to zero.

In figure 6.13 the bunch-by-bunch luminosity ratio before and after the crossing angle reduction is shown. At the nominal crossing angle of ±140 μrad, the four bunches exhibited different luminosity ratios due to their different transverse emittances. After the crossing angle was reduced to zero, this difference vanished as the luminosity ratio was decreased by ∼7 % for the low-emittance bunches and ∼4 % for the high-emittance bunches. The residual luminosity ratio of ∼1 % is within the expected systematic uncertainties.

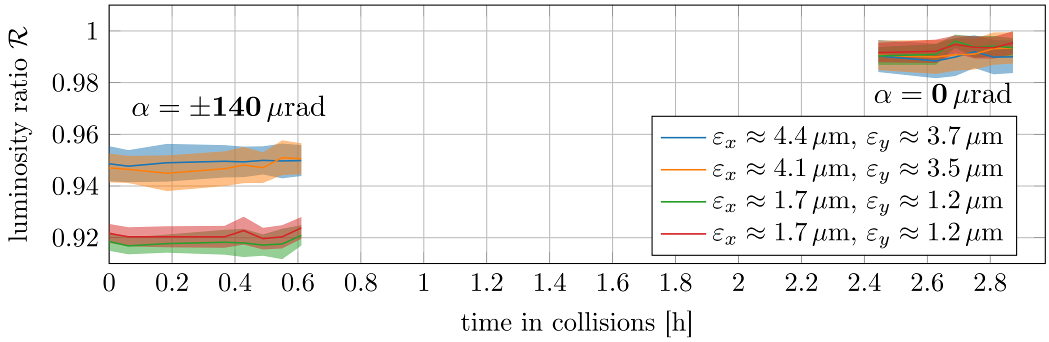

Figure 6.13: Observed luminosity ratios during the zero crossing angle test with statistical error bands (fill 5422). The period with no data is during the crossing angle adjustment.

It is worth noting that in 2017 operation, the beams were round at the start of Stable Beams ($\varepsilon_x \approx \varepsilon_y$); the measurement is shown in figure 6.14. The root cause for this different behavior is still under study. In line with the round beams, no significant luminosity imbalance between the ATLAS and CMS experiments was observed anymore ($\mathcal{R} \approx 1$).

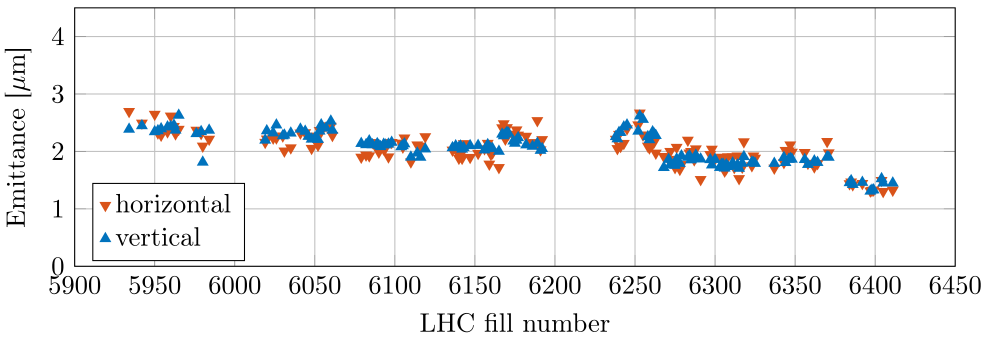

Figure 6.14: Measured emittances at the start of Stable Beams in 2017. Shown values are from emittance scans.

7 Luminosity Levelling

7.1 Motivation

With increasing instantaneous luminosity, the pile-up – the average number of interactions per crossing of a bunch pair – increases. According to equation (3.11), the pile-up μ is directly proportional to the instantaneous luminosity per colliding bunch pair.

The maximum acceptable pile-up is limited by the capabilities of the experimental detectors; with more simultaneous interactions, it becomes more challenging to identify all primary vertexes and to assign the detected secondary particles to the correct primary event. This increases the computing resource consumption per reconstructed event [70], and, above a certain limit, degrades the data quality up to an extent where the recorded data becomes unusable for certain physics analyses.

The two high-luminosity experiments at the LHC, ATLAS and CMS, were designed for a pile-up of $\mu \approx 20$ (at the design luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) [6]. In the past years of operation, they were improved to be capable of handling a pile-up of up to $\mu \approx 65$. The two specialized low-luminosity experiments, ALICE and LHCb, must run at a maximum pile-up of $\mu \approx 0.1$ and $\mu \approx 1.5$, respectively.

In order to fulfill these needs for all experiments, it is necessary to actively control the instantaneous luminosity. Such a *luminosity levelling* algorithm takes the measured luminosity of an experiment, and uses one or more machine parameter to reduce the instantaneous luminosity to the target communicated by the experiment, reducing the pile-up accordingly. Possible machine parameters used for this purpose are closed orbit separations (separation factor), the crossing angles (geometric factor) or the β^* (beam size at the IP).

In LHC operation, initially levelling by separation was used to satisfy the pile-up requirements of ALICE and LHCb. The luminosity in ATLAS and CMS was not levelled until the introduction of the low-emittance "8be4-BCS" beams (see section 6.2), which would have led to a pile-up of $\mu \approx 85$ without levelling. Hence, separation levelling was used to level the luminosity to $1.5 \cdot 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, limiting the average pile-up to $\langle\mu\rangle \approx 55$. The distribution of pile-up over the past years of LHC operation is shown in figure 7.1, with a clearly visible peak around the levelling target ($\langle\mu\rangle \approx 55$) for the 2017 data set.

A similar approach to luminosity levelling can also be used to increase the integrated luminosity by exploiting machine margins. Similarly to the pile-up, beam-beam effects limiting the performance of a collider get weaker over the course of a fill as the intensity decays. This allows decreasing the crossing angle or β^* over the course of a fill, beyond the limits set by the high intensity beams at the start of collisions. To distinguish this approach of levelling based on beam stability margins from the levelling on experiment request, and since it effectively allows levelling the luminosity "up", it is referred to as

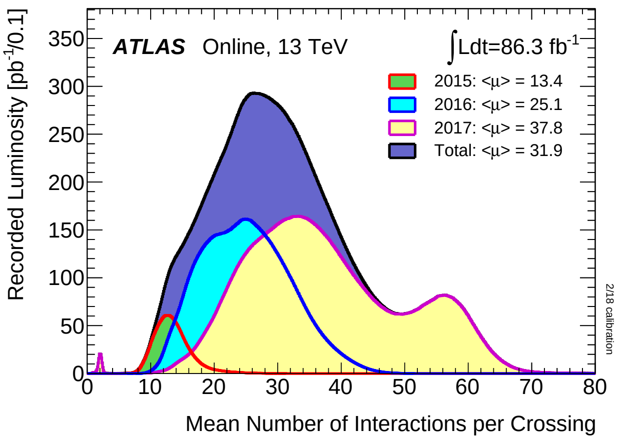

Figure 7.1: Observed pile-up at the ATLAS experiments in 2015-2017 operation [71].

anti-levelling. This was first implemented at the LHC for the crossing angle in 2017 as a part of this work, and later extended to allow changing both the crossing angle and β^*, which will be used in 2018 operation.

7.2 Levelled Luminosity Evolution

While levelling, the burn-off rate decreases due to the lower luminosity. As shown in figure 7.2, the typical evolution of luminosity in a fill with levelling consists of a first part, where the levelling is used to reduce the instantaneous luminosity to a certain target value, followed by a second part where the luminosity is left to decay without levelling.

If the emittance is constant and beam losses are dominated by burn-off while levelling, the bunch intensity decays linearly, as given by equation (3.17) with a constant luminosity. Note that while the actual luminosity is kept constant by the levelling algoritm, the *virtual luminosity* (theoretical luminosity without the levelling reduction factors) decays with the intensity.

Once the virtual luminosity reaches the levelling target, the levelling is not effective anymore and the luminosity decays naturally; the time in the fill until this happens is referred to as the *levelling time*. Considering only burn-off, it is given by

$$T_{\mathrm{lvl}} = \left(1 - \sqrt{\frac{\mathcal{L}_{\mathrm{lvl}}}{\mathcal{L}_{\mathrm{peak}}}}\right) \frac{k\, N_i}{n_{\mathrm{ip}} \sigma_{\mathrm{inel}} \mathcal{L}_{\mathrm{lvl}}} \tag{7.1}$$

where k is the number of bunches with initial intensity N_i, $\mathcal{L}_{\mathrm{peak}}$ is the theoretical peak luminosity, $\mathcal{L}_{\mathrm{lvl}}$ is the levelled luminosity, n_{ip} is the number of IPs[1], and σ_{inel} is the inelastic scattering cross-section ($\sigma_{\mathrm{inel}} \approx 80\,\mathrm{mb}$ for protons at 6.5 GeV [60]).

[1]Here it is assumed that all IPs are levelled to the same target luminosity, and that the same bunch pairs collide in all IPs. If this is not the case, the contribution of each IP to the total burn-off of each bunch must be considered individually on a bunch-by-bunch level.

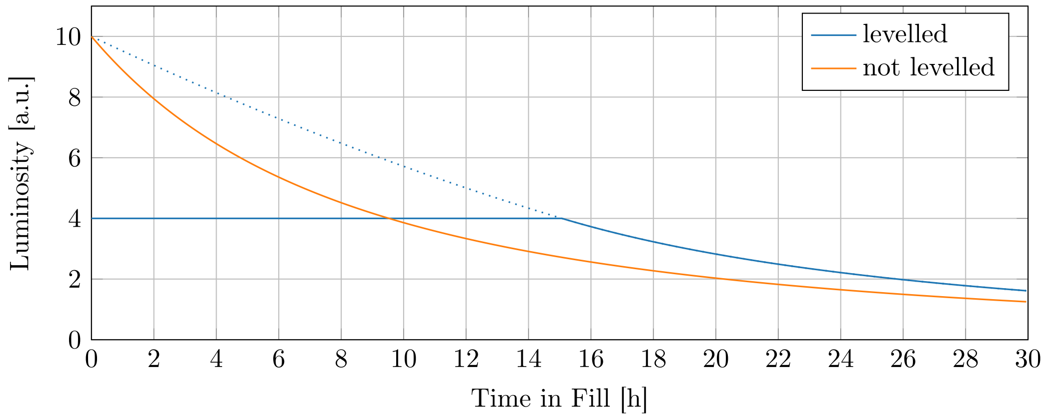

Figure 7.2: Luminosity evolution in a burn-off dominated scenario with and without levelling. The dotted line shows the virtual luminosity which is reduced to the target through levelling. The levelling time is $T_{\mathrm{lvl}} = 15\mathrm{h}$.

Since the instantaneous luminosity is reduced to the target initially, the integrated luminosity over the first hours of a fill is lower when levelling. In long fills, a significant part of the integrated luminosity originally given up can be recovered due to the slower luminosity decay.

The ratio of the instantaneous and integrated luminosities for a burn-off dominated scenario is shown in figure 7.3. Note that the reduction of the integrated luminosity in the first hours is largely dependent on the initial levelling factor; after twice the levelling time has passed, it is typically less than 10%.

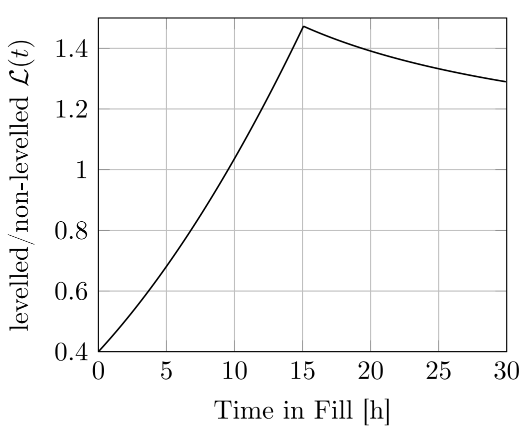

(a) instantaneous luminosity ratio

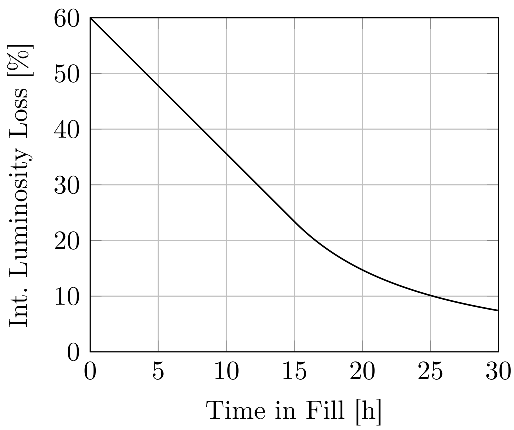

(b) integrated luminosity loss

Figure 7.3: Ratio of the instantaneous and integrated luminosities with and without levelling. The beam parameters and luminosities are the same as in figure 7.2.

7.3 Levelling by Separation

Levelling by separation uses local 4-corrector closed orbit bumps to introduce a beam separation at the IPs, introducing a separation factor as given by (3.14). The separation is introduced perpendicular to the crossing angle. To minimize the displacement of the beams with respect to the physical aperture, the separation is applied to both beams symmetrically with opposite sign (figure 7.4a). The effect of a separation on the luminosity is shown in figure 7.4b.

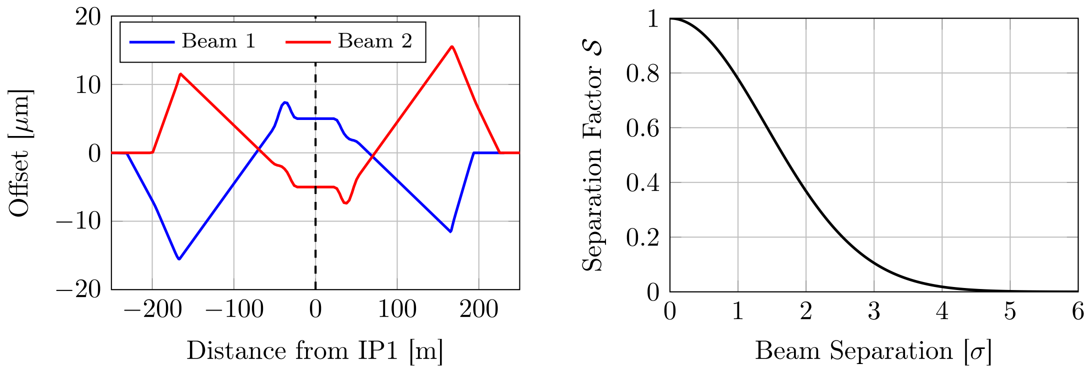

(a) closed orbit bumps (b) separation factor

Figure 7.4: Closed orbit bumps around LHC IP1 (ATLAS) for a total separation of $10\,\mu\text{m} \approx 1\sigma$ at the IP (with $\beta^* = 30\,\text{cm}$ optics). This reduces the luminosity by a separation factor of $\mathcal{S} \approx 0.8$.

Levelling by separation is operationally easiest at the LHC as the closed orbit deviations are small and local to each IP. This also allows levelling the luminosity in each IP independently to possibly different targets. For this reason, it has been used since the startup of the LHC to control the luminosity in the low-luminosity experiments, ALICE and LHCb.

However, the resulting offset collision reduces the stabilizing tune spread of the head-on beam-beam force and possibly drives additional resonances, which can lead to a loss of Landau damping and beam instabilities if applied to all IPs simultaneously (the region of least stability is around a separation of $\sim 1.4\sigma$ [72]). Such instabilities will deteriorate the beam quality due to emittance growth or intensity losses, or even lead to protective beam dumps in case the losses exceed a certain safety threshold. Such a behaviour was observed for example in 2012 while bringing very high intensity beams initially into collisions after the end of the squeeze [73, 74].

Starting from the introduction of the “8b4e-BCS” beams in mid-2017, levelling was also necessary for ATLAS and CMS to limit the pile-up at the beginning of each fill. Separation levelling was used for this purpose with initial separation factors of $\mathcal{S} \approx 0.65$, equivalent to an initial beam separation of $d \approx 1.3\sigma$. The levelling time was $\sim 3\,\text{h}$, after which the luminosity was left in natural decay. No signs of instabilities were observed with separation levelling of all four experiments in 2017 operation.

For reference, a test fill was done where the beams were left head-on. The luminosity of this test fill is compared to a levelled fill with the same peak luminosity in figure 7.5. As expected, the instantaneous luminosity after the end of the levelling is slightly higher in the levelled fill. In terms of integrated luminosity, ~4 % are lost to the levelling for fills with a typical length of 12-14h.

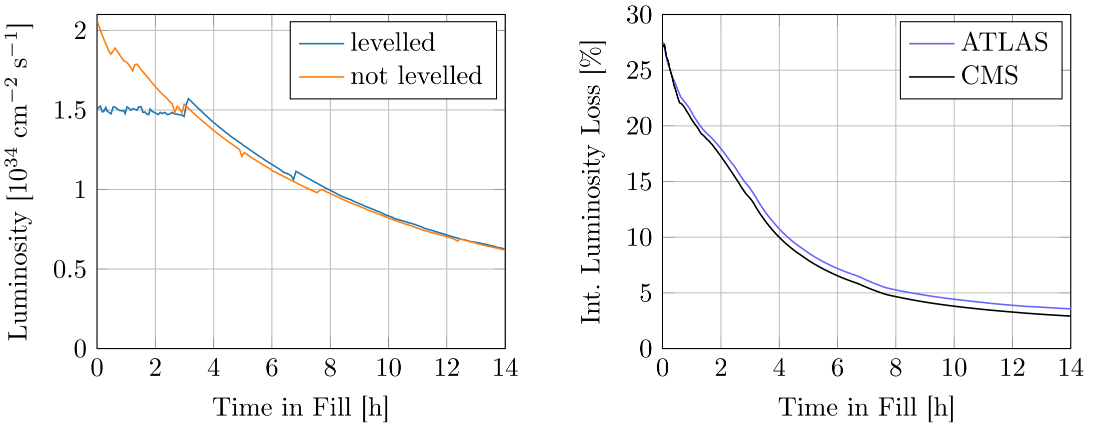

(a) instantaneous luminosity

(b) integrated luminosity ratio

Figure 7.5: Comparison of LHC fill 6358 (non-levelled test fill) to LHC fill 6360 (levelled).

7.4 Levelling by Crossing Angle

The crossing angle can also be used for luminosity levelling or anti-levelling. While the crossing angle bumps are closed orbit bumps around the IPs like the bumps for separation levelling, they are larger by two orders of magnitude for the same change in luminosity (figure 7.6). Therefore, changing the crossing angle with beams in collisions and Stable Beams is operationally more challenging than introducing (small) separation bumps.

In particular, the change of the closed orbit due to the change in crossing angle can not be treated as a small perturbation by the orbit feedback system and the collimators around the IP. If the reference the orbit feedback corrects against was not updated when the crossing angle bump is changed, the feedback loop would try to "correct" the change of closed orbit, which would lead to an unpredictable behaviour around the IP and possibly introduce orbit perturbations around the machine. Furthermore, the orbit feedback is needed to correct for orbit imperfections that could appear during the process of changing the crossing angle due to a not perfectly closed orbit bump ("leakage").

Likewise, the center of the collimators around the IP (TCTs within the crossing angle bump, see figure 7.6a), must be moved along with the orbit to prevent the beams from approaching the collimator jaws, which could lead to beam losses and increased background in the experiments.

Levelling by crossing angle changes the total pile-up along with the change in luminosity per bunch pair. However, due to the change in geometric factor, the luminous region is

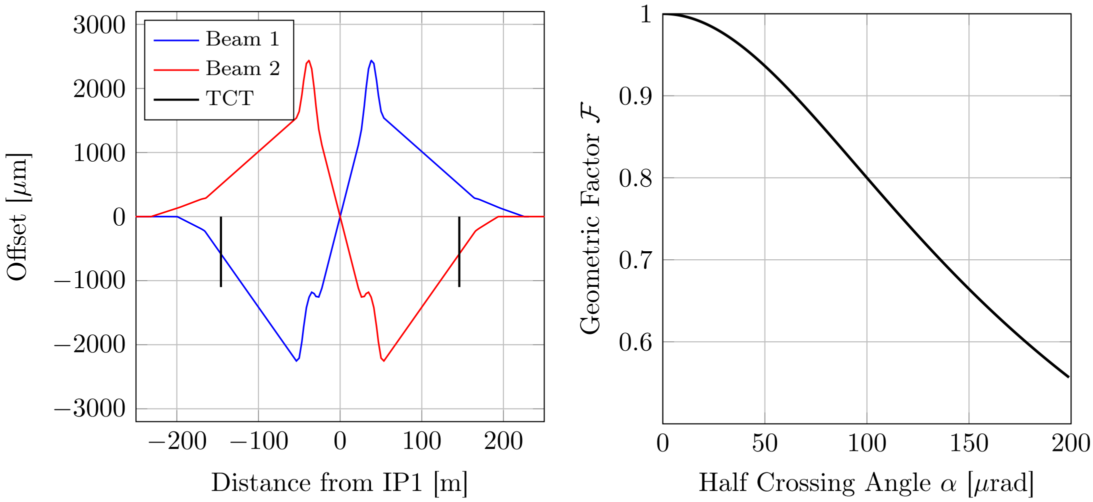

(a) closed orbit bumps (b) geometric factor

Figure 7.6: Closed orbit bumps around LHC IP1 (ATLAS) for a 50 μrad change of half crossing angle at the IP. Starting from a half crossing angle of 150 μrad, this changes the luminosity by a factor of ∼0.8. Note that the bump is two orders of magnitude larger with respect to figure 7.4.

shortened [11], and hence the pile-up density over the luminous region remains constant. If crossing angle levelling was used to reduce the pile-up, the unchanged pile-up density might still impact the performance of the experiments. Also, the physical aperture in the focusing triplets around the IPs limits the maximum crossing angle, especially at small β^*. Therefore, changing the crossing angle is particularly interesting for anti-levelling.

In the LHC, I implemented crossing angle anti-levelling in 2017 together with a team following my concepts and ideas. As the intensity decreases over the course of a fill, the long-range beam-beam effects due to the parasitic encounters next to the IP (see section 3.3) get weaker. This allows decreasing the crossing angle. The deteriorating effect of the parasitic encounters can be quantified in terms of *dynamic aperture* (DA), i.e. the area of the beam which remains stable for a high number of turns.

Figure 7.7 shows the dynamic aperture of the LHC with $\beta^* = 30\,\mathrm{cm}$ optics. A clear dependence on the crossing angle and the intensity is observed. In order to keep the beam losses due to long-range beam-beam effects in the shadow of the burn-off, a minimum dynamic aperture of 5σ should be kept. Therefore, the anti-levelling was set up to follow the iso-DA curve around a DA of 5σ [75].

The crossing angle anti-levelling was used in all physics fills of 2017 LHC operation. As a first step of gaining operational experience with this novel technique, and due to limitations of some of the experiments, fixed steps of 10 μrad in half crossing angle were done. In total, four different half crossing angles were used per fill, following the expected intensity decay: $\pm 150\,\mu$rad (initial), $\pm 140\,\mu$rad (after 2 h in collisions), $\pm 130\,\mu$rad (after 4 h) and $\pm 120\,\mu$rad (after 8 h, final value).

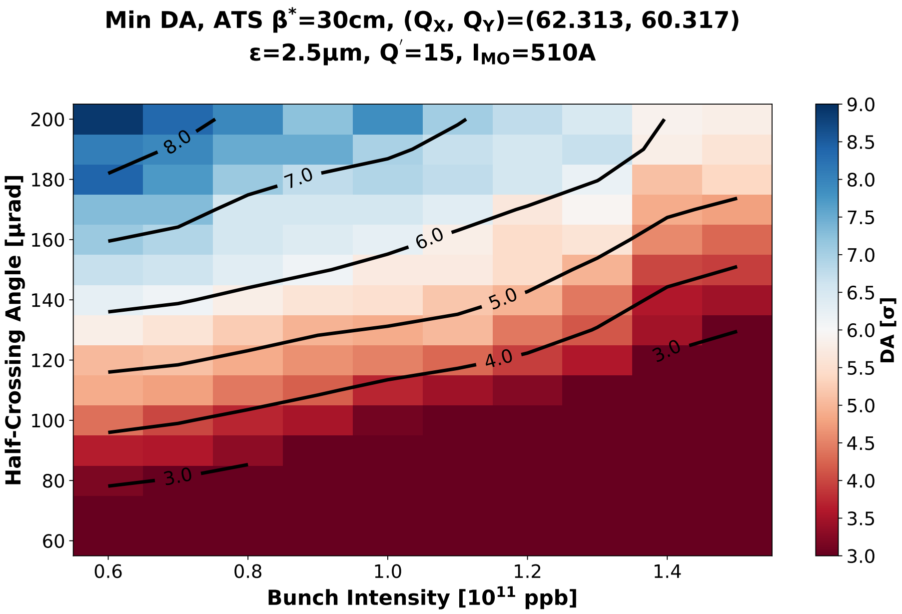

Figure 7.7: Dynamic Aperture simulation for the LHC with 30cm-ATS optics [75], showing the dependence of the DA on the intensity and crossing angle.

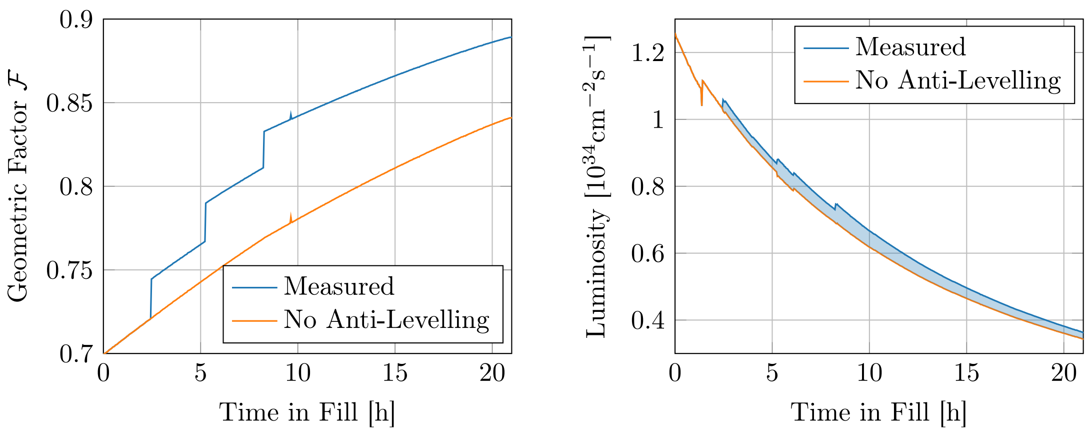

(a) geometric factor

(b) luminosity

Figure 7.8: Crossing angle anti-levelling in LHC fill 5849. The geometric factor is increased with each crossing angle step. For reference, the calculated evolution without the crossing angle anti-levelling is also shown.

Figure 7.8 shows the evolution of the geometric factor and the luminosity in a LHC fill due to these crossing angle changes. The gain in integrated luminosity due to these steps depends on the amount of extra beam losses introduced by the crossing angle changes and by running at lower crossing angles. Considering only the geometric factor, neglecting any extra losses, the gain is ∼5 % for fills lasting 12h in Stable Beams [76]. On the other hand, if all extra losses beyond burn-off were caused by the crossing angle anti-levelling, the gain would be ∼3 % for 12h fills [75]. Considering this most optimistic and most pessimistic estimates, the gain in integrated luminosity can be estimated as $4 \pm 1\,\%$, or in absolute terms $2 \pm 0.5\,\mathrm{fb}^{-1}$ out of the total of $50\,\mathrm{fb}^{-1}$ produced during 2017 LHC operation.

As a further improvement, the crossing angle steps can be reduced but applied more often, leading to a quasi-continuous crossing angle anti-levelling. This was tested in one fill in 2017 with promising results, shown in figure 7.9: The smaller steps allow following the iso-DA curve more precisely, increasing the gain due to the geometric factor to ∼8 % in 12h. Furthermore, the extra losses are decreased, as big steps in crossing angle (and hence DA) are avoided and the transition is smoother. This quasi-continuous crossing angle anti-levelling will be applied in 2018 LHC operation, and is expected to gain 6-7% in integrated luminosity. Also, it will use the measured intensity to follow the iso-DA curve automatically, whereas in 2017 the steps were done at pre-set times in a fill.

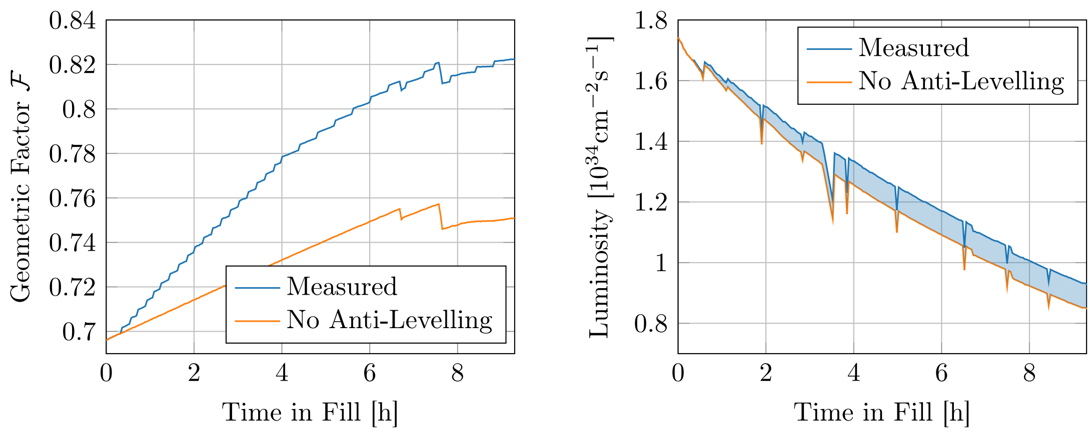

(a) geometric factor

(b) luminosity

Figure 7.9: Quasi-continuous crossing angle anti-levelling test in LHC fill 6061. The crossing angle is gradually reduced in small steps of $1\,\mu$rad. The overall rate of crossing angle reduction over the course of the fill is the same as for the regular physics fills (figure 7.8), i.e. $\pm 140\,\mu$rad, $\pm 130\,\mu$rad and $\pm 120\,\mu$rad are reached after 2 h, 4 h and 8 h into the fill, respectively. The two drops in the geometric factor after ∼7 h are not related to this test.

7.5 Levelling by β^*

The most advanced technique of controlling the luminosity for levelling is through the β function at the IP, β^*, effectively changing the beam size at the IP. Unlike a separation or the crossing angle, this implies a change of the machine optics which is in general a non-local process, and hence operationally most challenging.

Compared to levelling by separation, β^* levelling has the advantage of providing the full stabilizing effect of the head-on beam-beam tune spread independently of the luminosity. Furthermore, offset collisions which lead to a coherent beam-beam kick and could drive resonances are avoided. Over levelling by crossing angle, it has the advantage that the physical machine aperture does not limit[2] the maximum β^*. Therefore, β^* levelling is a baseline concept for the high-luminosity LHC upgrade [44] and future hadron colliders [45].

When changing the optics for β^* levelling, any closed orbit bumps in the machine (e.g. crossing angles) change their shape due to the different β function. Also, additional perturbations might be introduced on the orbit due to quadruple misalignments. Therefore, the reference for the orbit feedback system must be updated to follow the expected change of orbit, and the feedback loop must be operated with settings suitable to correct the unwanted perturbations.

Furthermore, changing the optics alters the β function and hence the beam size not only at the IP, but also in the adjacent low-β insertion, matching section and, in case of ATS-type [32] optics, in the adjacent arcs. The collimators in these regions need to follow both the change of orbit (displacing their center), as well as the change of beam size (changing the opening of the jaws). Figure 7.10 shows the β function around an LHC IP and the corresponding beam size at $\beta^* = 40\,\text{cm}$ and $\beta^* = 30\,\text{cm}$ optics.

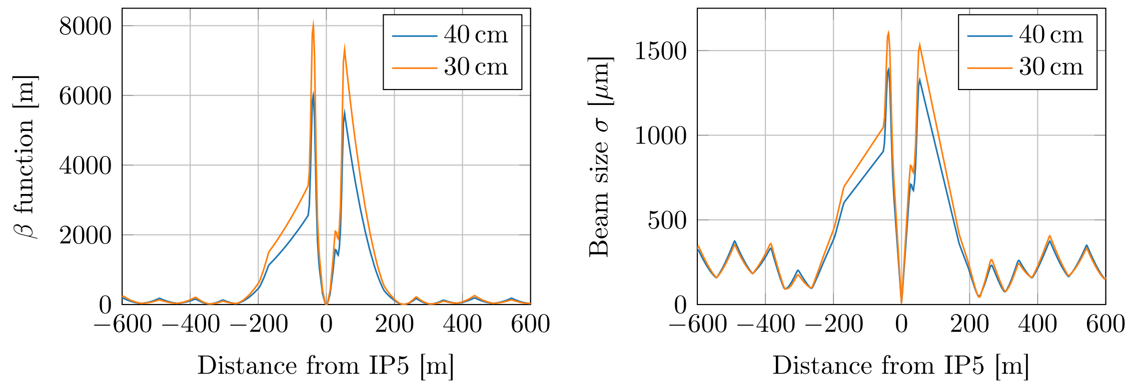

Figure 7.10: β function and beam size σ around LHC IP5 for $\beta^* = 40\,\text{cm}$ and $\beta^* = 30\,\text{cm}$ optics. Horizontal plane of beam 1 shown; the vertical plane and beam 2 show a similar behaviour.

[2] There is no upper β^* limit in the range that would be used for levelling. For example, there were special LHC runs with a β^* of up to 2.5 km. The injection β^* at the LHC is 11 m. Compared to the LHC collision β^* of 30 cm, this provides several orders of magnitude range for the levelling.

First LHC Machine Development sessions (MDs) to explore the feasibility of running through the existing betatronic squeeze in steps with the beams in collisions were done in 2012, 2015 [77, 78, 79] and 2017 [80]. In these MDs, existing squeeze segments were run using pre-programmed sequences, in the same way the LHC operational squeeze is played. This approach requires a static management of all machine settings and orbit references, limits the optics transitions to a single direction (it is not possible to "un-squeeze" the beams) and prevents all parallel operations in the LHC. Thus, while these tests showed that changing the optics while keeping the beams in collision is physically possible, the technique employed has limited use for operational β^* levelling.

Over the course of 2017, I guided a collaboration to implement a solution for operational β^* levelling at the LHC based on the same software stack and concepts used for crossing angle anti-levelling. The implementation is described in detail in chapter 8. It allows, based on a set of "repositories" of LHC squeeze functions, to move freely between any matched optics points provided in the repositories, automatically orchestrating the changes of magnet strengths, the orbit feedback system and collimators.

This new technique was then successfully tested in a near-operational scenario during an MD session at the end of 2017 [81]; the results are presented in the following. For 2018 LHC operation, it is foreseen to commission this technique for operation in anti-levelling mode [82], with the possibility of using it also for levelling at the beginning of fills instead of separation levelling later in the year.

β^* levelling test results: 40cm to 30cm

For testing our new β^* levelling technique we demonstrated multiple squeeze cycles between 40 cm and 30 cm β^* in IP 1 and 5 in a first test fill (6424). Due to the telescopic nature [32] of this squeeze segment, the collimators were not moved. Since our new mechanics allow running the squeeze backwards, we were able to un-squeeze the beams back to 40 cm once we reached $\beta^* = 30$ cm for multiple squeeze cycles in a single fill. The luminosities in IP 1 and 5, as shown in figure 7.11, follow the change in β^* closely.

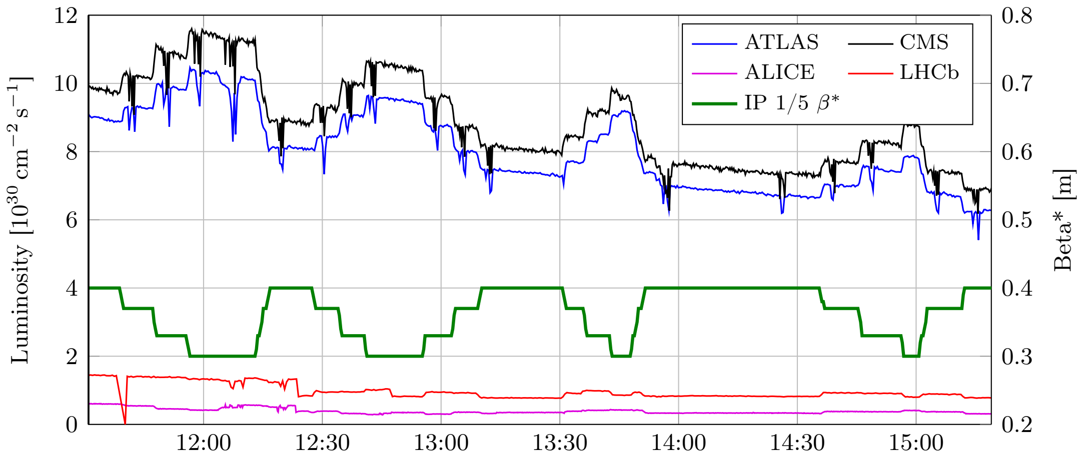

Figure 7.11: Luminosity from the LHC experiments during β^* levelling test fill 6424.

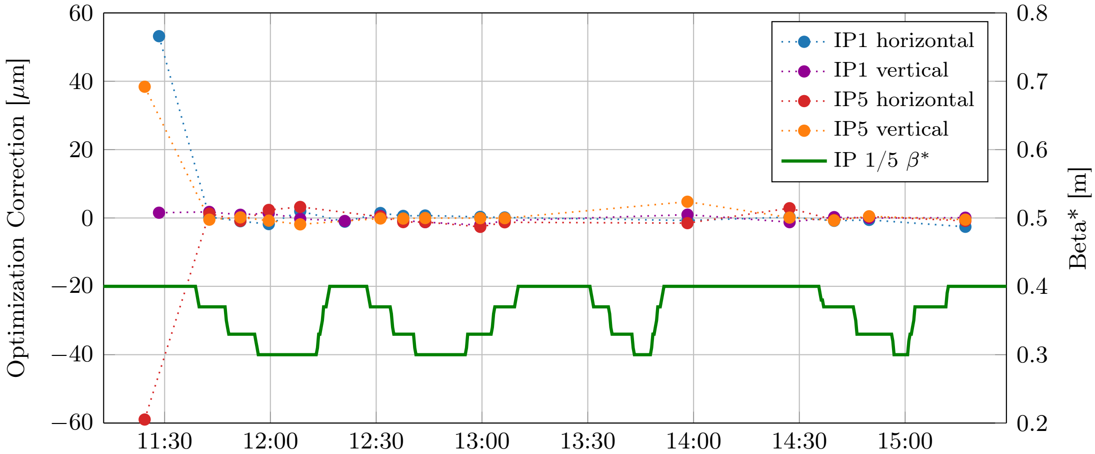

Figure 7.12: Corrections from luminosity optimization scans during the first β^* levelling test (fill 6424). Apart from the initial optimization to bring the beams into collisions, the corrections were lower than $2\,\mu$m.

The squeeze repository used had intermediate optics for $\beta^* = 37$ cm and $\beta^* = 33$ cm in addition to the two end points (40 cm and 30 cm). After each β^* step, the beams were re-centered in all IPs by scanning the separation for the optimum instantaneous luminosity. The largest changes with respect to the known feed-forward corrections found during a prevous test [80] were smaller than $2\,\mu$m, as shown in figure 7.12. This corresponds to $\sim 0.2\sigma$ at $\beta^* = 30$ cm, and is well correctable by luminosity optimization scans.

To probe the feasibility of using levelling by separation (for ALICE and LHCb) and levelling by β^* (for ATLAS and CMS) in parallel, IP 2 and 8 were separated by 1σ in the second cycle to 30 cm. This maximizes the sensitivity to changes in the beam separation, assessing the luminosity stability with separation levelling during β^* transitions. As shown in figure 7.13, the luminosity stability was $\sim$10% per β^* step without any feed-forward corrections. When arriving at $\beta^* = 30$ cm, the separation had drifted by $\sim 0.1\sigma$ in IP2 (separating the beams) and by $\sim 0.3\sigma$ in IP8 (bringing the beams closer together). The automatic separation levelling was then used to bring the luminosities back to their initial values.

In a third β^* cycle to 30 cm, we reduced the crossing angle at $\beta^* = 33$ cm and at $\beta^* = 30$ cm using the operational crossing angle anti-levelling tool, showing that there is no interference.

Over the course of these β^* cycles, we probed the settings (gain, eigenvalues) for the orbit feedback system (OFB) to improve the orbit control while avoiding overly aggressive corrections. Initially, the OFB was set up with 40 eigenvalues and a gain of 0.5; this is the standard configuration used at the LHC in Stable Beams. As shown in figure 7.14, increasing the gain to 1.0 eliminates the transient orbit excursions after each step, and using 200 eigenvalues in both planes allows for an orbit control at the level of $\sim 30\,\mu$m r.m.s. for all steps. This is equivalent to the level of orbit control in Stable Beams.

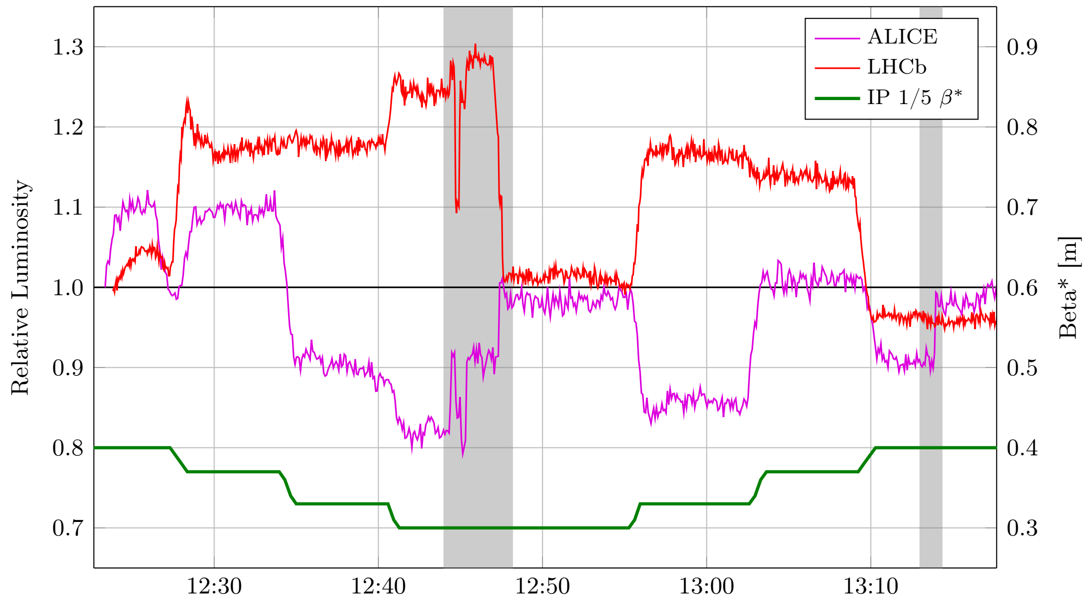

Figure 7.13: Relative Luminosity of IP 2 and 8 during a β^* cycle with an initial separation of 1σ. The gray shaded areas indicate periods where the automatic separation levelling was enabled to bring the luminosity to the initial value ($\pm 5\,\%$).

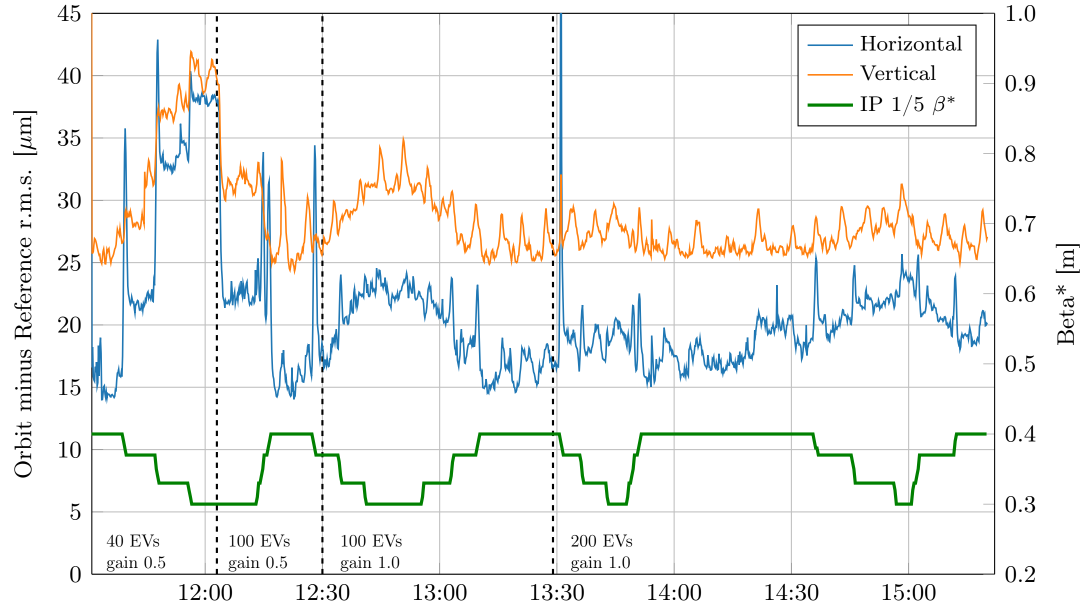

Figure 7.14: The orbit deviation from the feedback reference over the course of β^* levelling test fill 6424 for different settings of the OFB gain and eigenvalues (EVs).

β^* levelling test results: 40cm to 1m

In a second test fill (6425), we used the new β^* levelling mechanics to run backwards through the non-telescopic LHC pre-squeeze, de-squeezing the beams from 40 cm to 1 m β^* in IP 1 and 5. While this is not immediately foreseen as an operational scenario in the next years of LHC operation, it proves the flexibility of our technique.

Since during the LHC pre-squeeze the triplet strengths change, using this squeeze to change the β^* requires moving both the center and the gap of the tertiary collimators (TCTs) around IP 1 and 5. Also, the impact on the closed orbit around IP 1 and 5 is larger in this case. The luminosities and the β^* in IP 1 and 5 are shown in figure 7.15.

For the orbit feedback system, we used the same settings as established during the first fill: a gain of 1.0 and 200 eigenvalues per plane. No transients were observed during the steps, but a small residual orbit effect was visible at the last steps around $\beta^* = 1\,\mathrm{m}$ (figure 7.16). This perturbation vanished when squeezing back towards 40 cm and could be further improved by an orbit correction with a high number of eigenvalues. Overall, the orbit control was at the level of $\sim 50\,\mu\mathrm{m}$ r.m.s. throughout the whole process, which is sufficient for operation and machine protection.

Since the squeeze segment to 1 m was never commissioned in collisions before our test, no prior feed-forward for the beam positions at the IPs was available. At every matched optics point, a luminosity optimization scan was done to re-center the beams at all IPs and the results were incorporated into the repository to be used as a future feed-forward correction. In IP 2 and 8, where the nominal β^* did not change, no significant beam

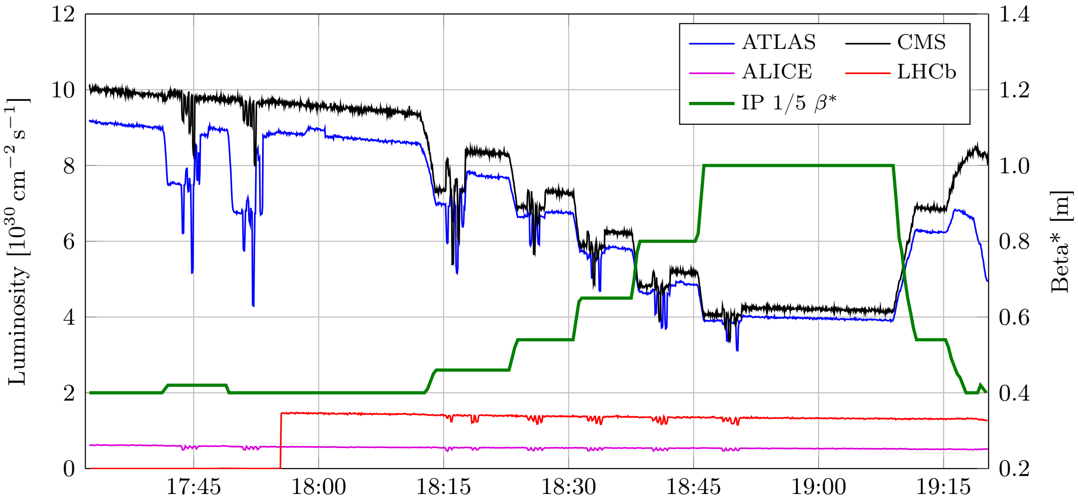

Figure 7.15: Luminosity from the LHC experiments during β^* levelling test fill 6425. The first segment until 18:10 is a special optics optimization for the forward physics experiments around ATLAS and CMS [83], re-distributing the strength among the quadrupoles at constant β^*. LHCb did not publish luminosity data before 17:50.

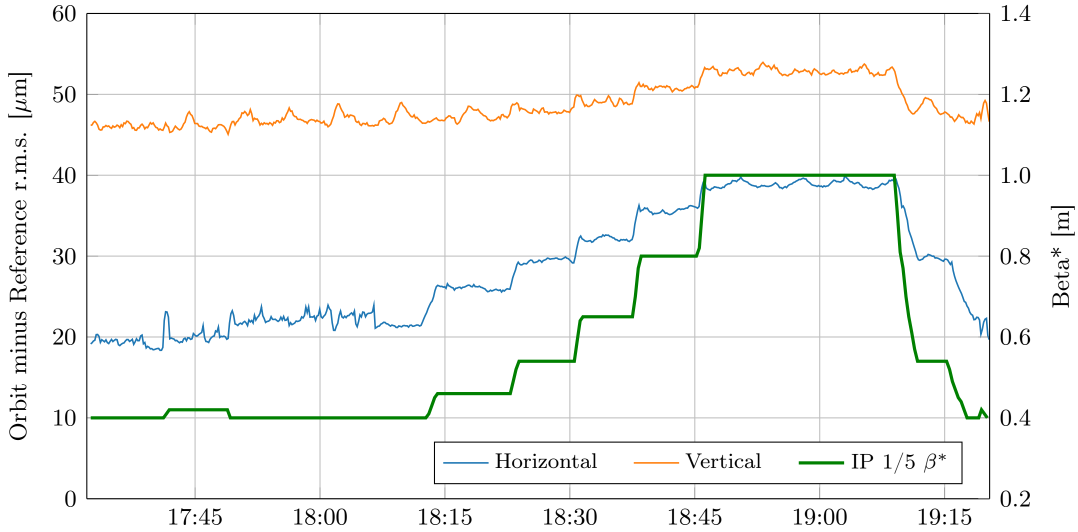

Figure 7.16: Orbit deviation from the feedback reference over the course of β^* levelling test (fill 6425).

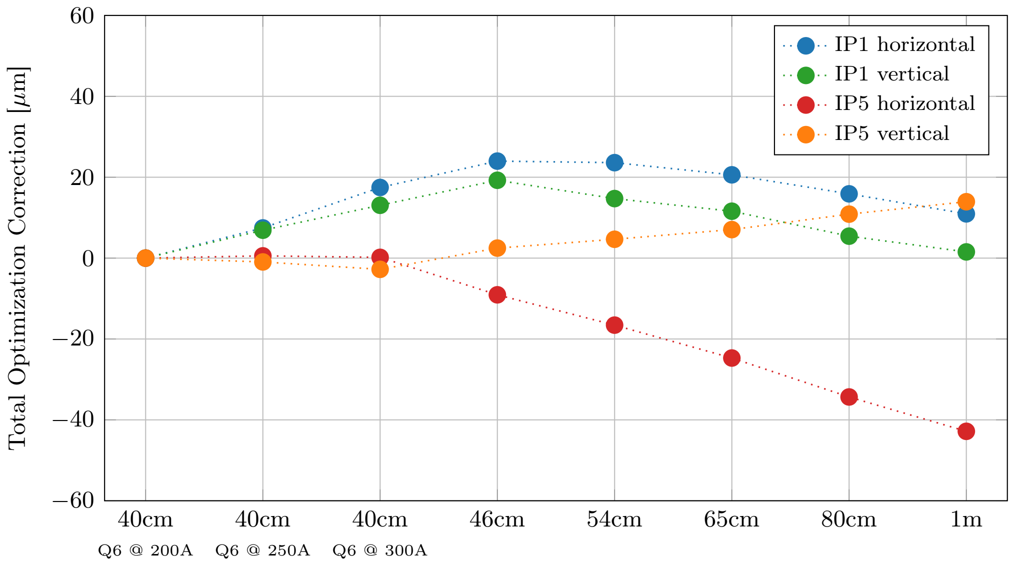

Figure 7.17: Total corrections from luminosity optimization scans during the de-squeeze to 1 m in fill 6425. The first 3 steps do not change the nominal β^* but re-distribute strengths among quadrupoles [83].

displacement was observed and the luminosity was constant. As shown in figure 7.17, the largest correction (~40 μm) was needed in the horizontal (crossing) plane of IP5. In this plane, the DOROS BPM measurements also indicate a shift of both beams with respect to the beam pipe as shown in figure 7.18. Since moving both beams in the same direction does not affect the luminosity, this can not be corrected by luminosity optimization scans. However, the DOROS results could be considered for a future feed-forward correction.

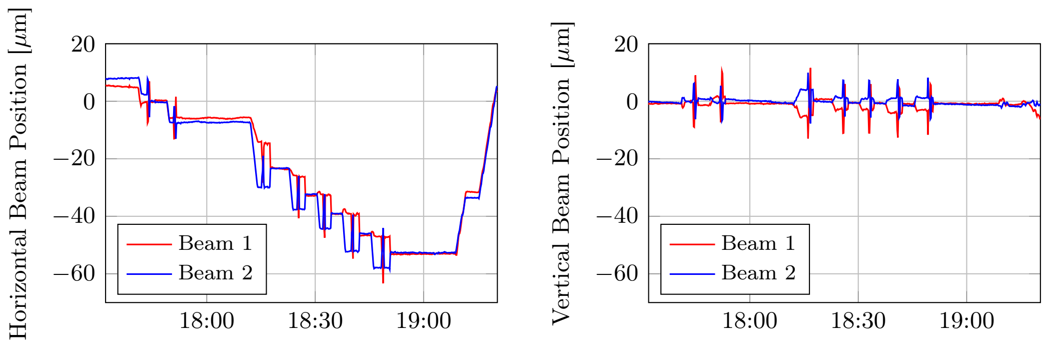

Figure 7.18: Beam positions at IP5 as extrapolated from the DOROS BPMs around the IP during the de-squeeze to 1 m in fill 6425. The zero positions are arbitrary. The data for the vertical plane has a 3rd order polynomial subtracted to correct for an observed instrumental drift.

The collimator center movements done by the automatic β^* levelling orchestration agreed with the measured orbit at the respective collimators within ~20 μm, which is at the level of uncertainty expected from the observed orbit perturbation. The change of collimator gaps, measured by the difference of the left and the right jaw motor positions, matched the pre-programmed settings in the repository within the accuracy of the collimation system (5 μm). An example is given in figure 7.19.

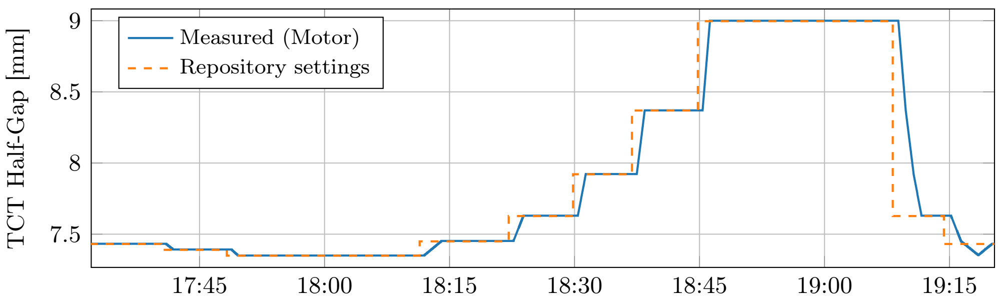

Figure 7.19: Half-gap of the tertiary collimator left of IP1 (TCTPH.4L1.B1) during the de-squeeze to 1 m in fill 6425. The respective settings in the squeeze repository are shown for comparison. Other TCTs show the same level of agreement.

7.6 Combined Levelling Scenarios

Two or more of the levelling and anti-levelling methods described above can be combined to a full levelling scenario for a collider to yield the optimal integrated luminosity within the given constraints of pile-up (which can significantly differ between experiments), beam-beam effects and machine limitations (e.g. aperture).

For example, in the second half of 2017 LHC operation (with “8b4e-BCS” beams), the luminosity of the ATLAS and CMS experiments was initially levelled to $1.5 \cdot 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ by separation. Once this was not effective anymore, the intensity had decreased due to burn-off, allowing to reduce the crossing angle by $\pm 20\,\mu$rad from $\pm 150\,\mu$rad to $\pm 130\,\mu$rad. A few hours later in the fill, the crossing angle was further reduced to $\pm 120\,\mu$rad. In parallel to all these manipulations in IP 1 and 5, the luminosities of the low-luminosity experiments in IP 2 and 8 were controlled via separation levelling (figure 7.20).

For 2018 LHC operation, a combination of levelling by separation and anti-levelling by both crossing angle and β^* is foreseen [82]: Initially, the pile-up will be controlled by separation levelling (if requested by the experiments). After the levelling is not effective anymore, the crossing angle will be reduced in small steps using quasi-continuous crossing angle anti-levelling between $\pm 160\,\mu$rad to $\pm 130\,\mu$rad. Towards the end of a fill, the β^* in IP 1 and 5 will be reduced from 30 cm to 27 cm and later to 25 cm.

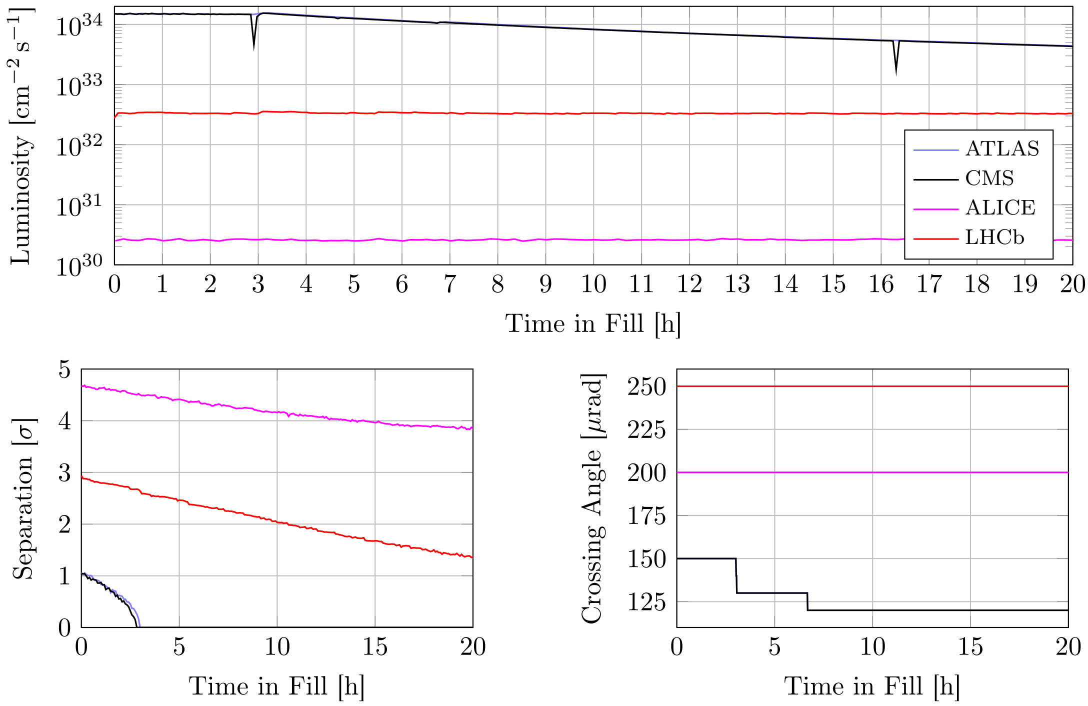

Figure 7.20: Luminosity and levelling parameter evolution in LHC fill 6360. A combination of separation levelling and crossing angle anti-levelling is used.

8 LHC Luminosity Control Software

8.1 Motivation

An important part of my work was to renovate the LHC luminosity control software stack. For better reliability and maintainability, I re-designed the system from scratch following a client-server architecture. The central component of the new system controlling beam displacements and other manipulations affecting the luminosity is the *LHC Luminosity Server* [84]. This new design also allowed implementing many of the techniques described in the previous chapters to make them available for day-to-day LHC operation.

For 2018 LHC operation, the operational use cases for the LHC luminosity control software stack are:

- **IP steering,** to let the operator directly modify the beam positions at any IP, e.g. for machine tests or at the request of the experiments.
- **Optimization scans,** probing the beam separation while acquiring the luminosity signal to find the beam head-on position giving maximum luminosity [85].
- **Luminosity calibration scans,** automatically synchronized with the experiments, to calibrate the absolute luminosity measurement using the van-der-Meer method [35, 86].
- **Emittance scans,** to measure the bunch-by-bunch emittances and beam separations by scanning the beam separation, as described in chapter 5.
- **Separation levelling,** to keep the luminosity at one or more IPs constant around a target value, provided by the experiment or by the operator, by adjusting the beam separation, as described in section 7.3.
- **Crossing angle anti-levelling,** to gradually decrease the crossing angle at IP 1 and 5 throughout a fill while the beams are in collision to optimally use margins while improving the luminosity, as described in section 7.4.
- β^* **anti-levelling,** to decrease the β^* at IP 1 and 5 some hours into the fill with the beams in collisions, as described in section 7.5, to improve the luminosity yield and to gain operational experience with this novel levelling technique.

Out of these functionalities, the previous LHC luminosity control software only allowed IP steering, optimization scans, levelling by separation and one particular type of luminosity calibration scans. In particular, automated emittance scans, calibration scans with arbitrary scanning patterns, and the novel levelling techniques by crossing angle or β^* are exclusive to the new LHC Luminosity Server.

8.2 Architecture

As most other top-level control software for the LHC, the new LHC luminosity control software is entirely written in Java. When I started re-designing the system in 2015, several fundamental choices were taken to make the new software extensible and to ensure its reliability. An overview of the architecture is given in figure 8.1.

All the logic and processing code is separated from the user interface (GUI) and encapsulated in a dedicated server (LHC Luminosity Server). The graphical control room application (LHC Luminosity Client) communicates remotely with this server over the network using the Java Message Service (JMS, [87]) and Remote Method Invocation (RMI, [88]) protocols.

This client-server separation allows the server to keep full control of all the functionality, all changes made to machine settings and all publications made to the LHC experiments, making them independent of any open GUI. Therefore, multiple instances of the GUI can be open on the operational consoles and closed at any time, while still guaranteeing a coherent behaviour (e.g. the accidental closure of a GUI does not affect an ongoing scan).

All operations in the Luminosity Server are implemented as *Tasks*. A task does one thing once (e.g. optimizing a selected set of IPs). Once the task completed, the results can be reviewed for reference, but can not be restarted or reused. If the same operation needs to be done again, a new task must be created.

The framework of the server allows tasks to be implemented in a transactional way, so the settings are reverted in case the task fails due to a problem or is aborted by the operator. This behaviour can be changed based on the needs of the actual task implementation. For example, an automated luminosity scan task is transactional, while a luminosity levelling or IP steering task is not.

The server was designed to contain an abstraction layer for accessing the lower layers of the LHC control system, both for data acquisition and controlling the beams. The LHC Luminosity Client acquires the data to display exclusively through the LHC Luminosity Server it is connected to, and never subscribes to any data source directly. This allows replacing data sources or the acquisition logic at the server level transparently without modifying the client code, e.g. by simulation loops for testing purposes.

In the operational configuration, the LHC Luminosity Server acquires data e.g. from the power converters, the LHC experiments, and the timing systems through the CERN Controls Middleware (CMW, [89]) framework. Publications to the experiments are done over the Data Interchange Protocol (DIP, [90]), accessed through CMW.

The beam displacements at the IPs are controlled through *trims* (setting changes) on *knobs* defined in the LHC Software Architecture (LSA, [91]) framework. The knobs are defined to generate 4-corrector closed orbit bumps around the IP concerned. In general the two beams are moved symmetrically in opposite directions to reduce the absolute displacement with respect to the common beam pipe. Likewise, changes of the crossing angle and the β^* are applied through trims on LSA, while additionally orchestrating the orbit feedback system and the collimators as needed (see section 8.5).

To ensure that only one task modifies the settings of a particular IP at a time, the Luminosity Server implements a reservation system for the IPs. A task can request a set

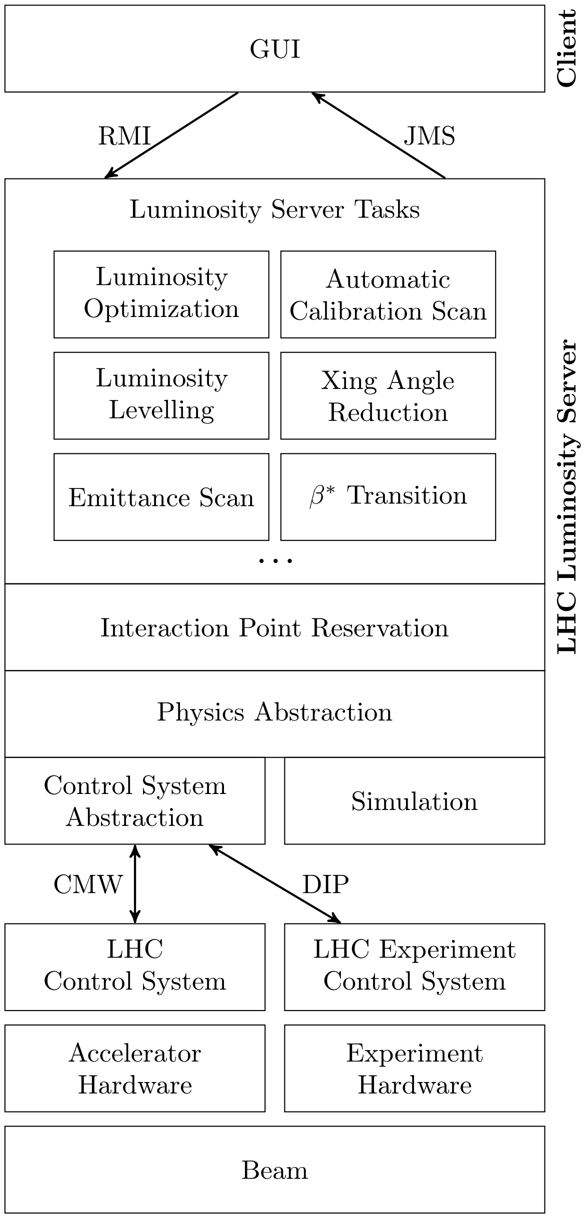

Figure 8.1: The basic architecture of the new LHC luminosity control software stack.

of IP reservations, and it is guaranteed that only one task will get a reservation for a particular IP at any given time.

If a task requests reservations for multiple IPs, they will be granted in IP order to prevent a deadlock situation (an example is given in figure 8.2). A task will not execute until all the reservations it requested are granted, and when a task finishes, fails, is cancelled or paused, the reservations it held are freed.

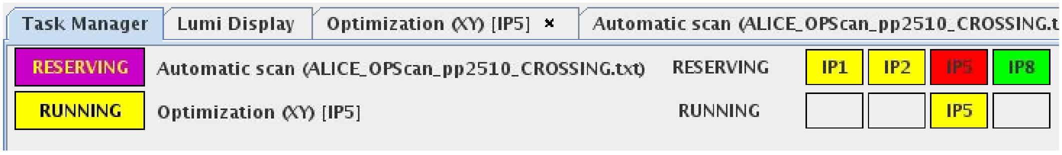

Figure 8.2: The “Task Manager” in the LHC Luminosity Client, with an overview of all active tasks along with their reservations. Granted reservations are shown in yellow; requested reservations are shown in green if they are available or in red if they are in use by another task.

By replacing the interface layer between the LHC control system and the Luminosity Server by a simulation loop, the server can run in simulation mode without accessing the LHC control system. In this mode, the server can even run outside the LHC controls network, which minimizes the risk of accidentally accessing operational devices during testing and development.

In simulation mode, the changes in separation, crossing angle or β^* any running task does are recorded, and an appropriate luminosity signal is simulated (according to chapter 3). As shown in figure 8.3, the beam conditions and various inputs to the simulation can be controlled by the user through a set of simulation control panels and any publication a task would make to the LHC experiments is displayed instead.

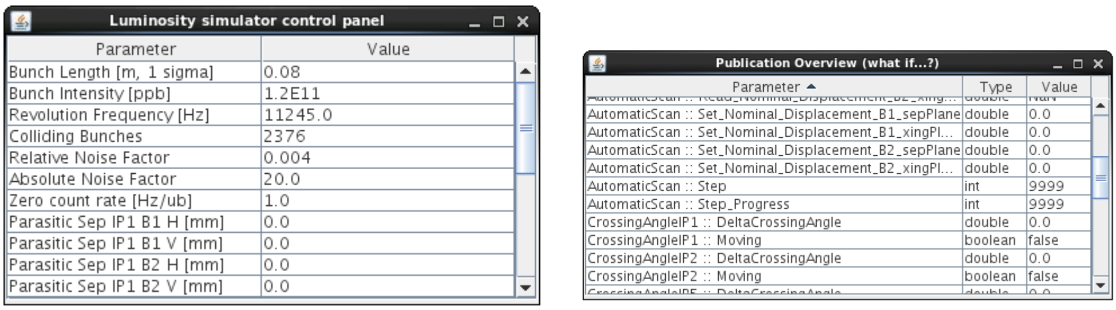

(a) luminosity simulator control (b) experiment publication overview

Figure 8.3: The control panel of the luminosity simulation loop and the display of publications of the LHC Luminosity Server in simulation mode.

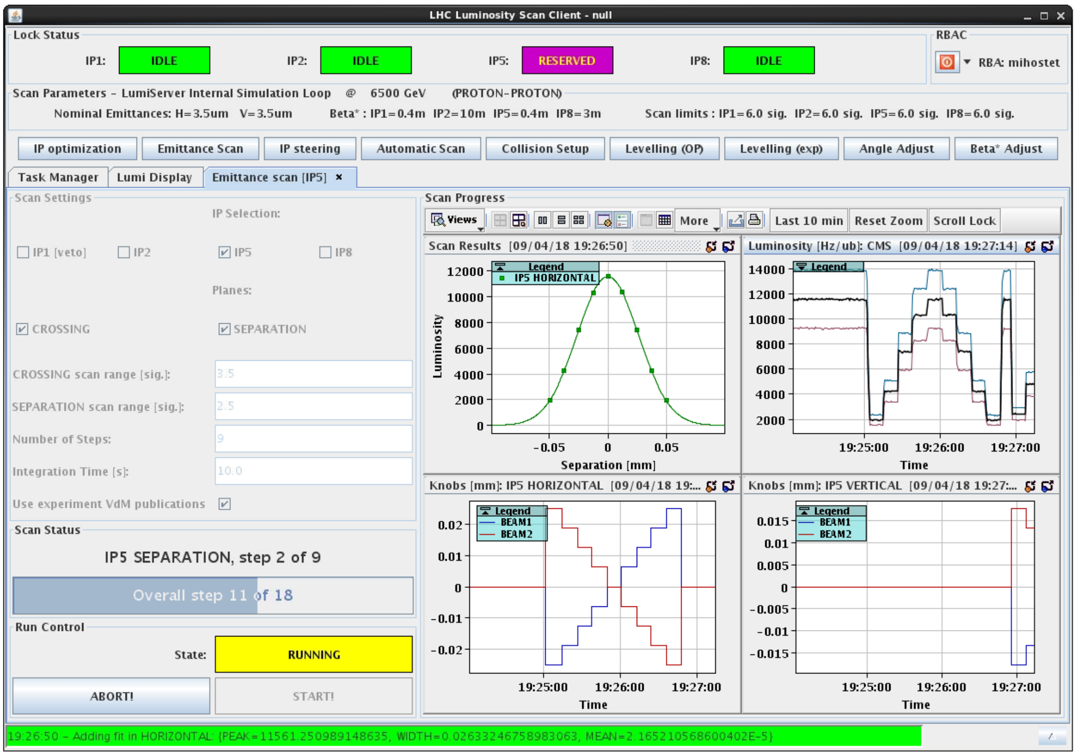

Figure 8.4: The LHC Luminosity Client, running an Emittance Scan while the server is in simulation mode. Note that the beam displacements (bottom charts) are reflected by the simulated luminosity (top charts).

The simulation mode allows to test any task and almost all aspects of the Luminosity Client and Server without accessing the LHC [92]. It is also used to dry-run various scan sequences and to predict the resulting luminosity patterns. This saves significant LHC beam time for developing and testing new functionality of the software, as tests in the LHC are only required during the final stages of commissioning. An example of an emittance scan task running in simulation mode is given in figure 8.4.

Furthermore, it is also possible to start the server in partial simulation mode (called *wonderland mode*). This allows to select either the simulated or the real implementation of every sub-system of the Luminosity Server at startup, and provides a special subscription multiplexer to dynamically switch every input signal between the real acquisition and the simulated counterpart (figure 8.5).

This mode can be used to test certain interactions of the server with the control system with the least impact on operations. For example, it is possible to test the communication with the LHC experiments over DIP without actually moving the beams or changing any settings in the LHC. Therefore, this can be done in parallel to LHC operation with no need for dedicated machine time.

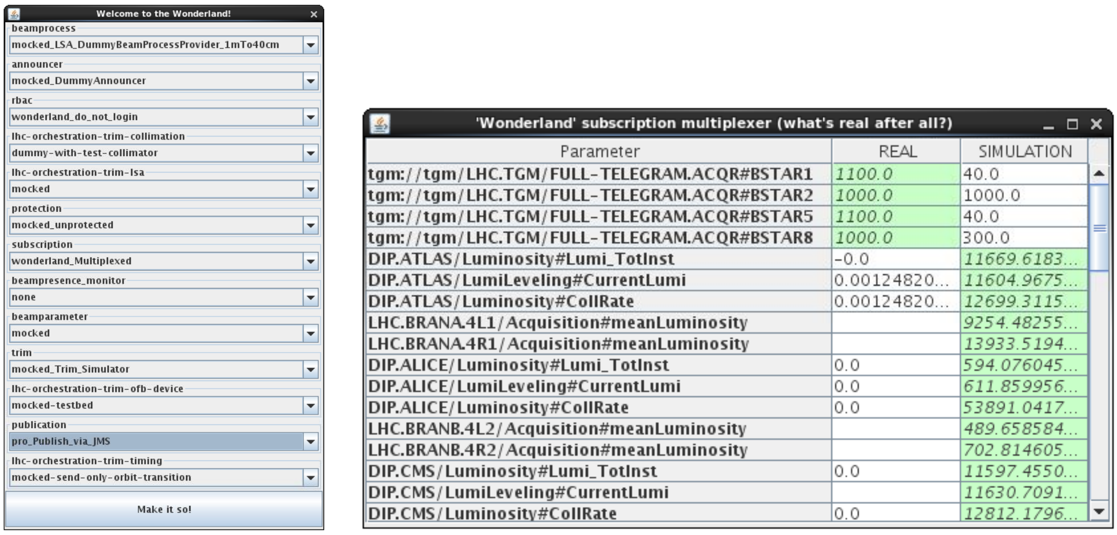

(a) profile selector (b) subscription data source selector

Figure 8.5: The interface of the LHC Luminosity Server running in *wonderland mode*. At startup, the profile selector allows to select a real implementation or a simulator for each sub-system (a). At runtime, every data subscription can be switched between the real acquisition and the internal simulation loop (b).

8.3 Beam Separation Scans

A basic functionality of the LHC Luminosity Server is to perform automated scans of the beam positions and separations at the IPs for experiment luminosity calibration, emittance scans or luminosity optimization.

At the start of each LHC fill, and regularly while in collisions, the positions of the beams at the IPs need to be optimized to ensure the beams are colliding head-on, so to deliver the maximum luminosity to the experiments (when levelling by separation, the optimization is only performed in the crossing plane, perpendicular to the separation).

The *Optimization* task achieves this in the following way: Independently, for each plane, the beam separation is scanned while acquiring a luminosity signal (increase in luminosity followed by a decrease, considering statistical errors) is found. Thereafter, a parabola is fitted to the top 3 points acquired during the scan to refine the peak position; an example is given in figure 8.6. Multiple IPs can be scanned in parallel, and the luminosity monitor to use for each IP can be selected by the operator.

Sanity checks are applied automatically both on the selected luminosity signal before a scan is launched, and on the results after a scan finishes. If potential problems are found for the selected luminosity signal (e.g. not publishing or publishing zeroes), a warning is issued before the scan is started. In case a problem is found with the result (e.g. a bad fit, non-reproducible luminosity at the initial point or an unusually large correction), a warning is issued along with the scan result.

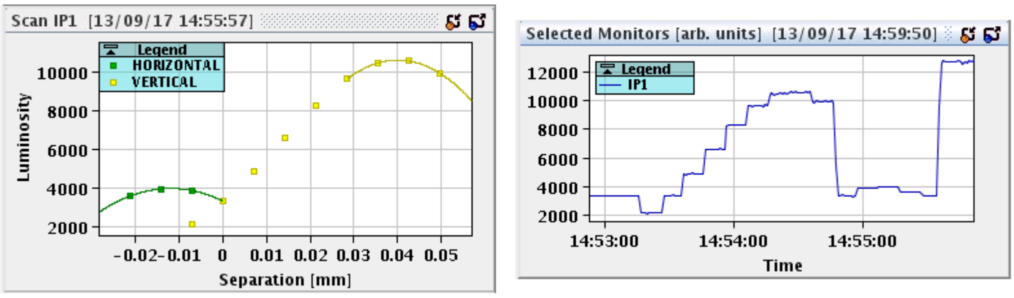

Figure 8.6: An optimization scan of IP1 (ATLAS). Beam separations of ~10 μm (horizontal) and ~40 μm (vertical) were found and corrected for.

The results for all scanned IPs are shown to the operator via the LHC Luminosity Client along with the acquired luminosity points and the used fits; the operator is given the choice of which results to apply. After applying, the optimization can be undone within a short time in case the luminosity did not increase.

The automatic optimization task has been used in LHC operation as of 2016. Compared to the software used in the previous years, the improved diagnostics and sanity checks allowed to quickly track down problems caused e.g. by missing or delayed luminosity data publications from the experiments.

Figure 8.7: Luminosity Client during an optimization scan in all IPs in LHC fill 5736. Due to the noise on the signal of IP2 (ALICE), the luminosity at the initial separation was not reproducible and a warning was issued.

Another application for beam separation scans is to calibrate the absolute luminosity measurements of the LHC experiments using the van-der-Meer method [35]. Since the first such scans in the LHC [86], the method has evolved from using simple beam separation scans to using a combiation of various scanning patterns to measure the beam width, to calibrate the length scale of the closed orbit bumps used against the experiment detectors, and to measure any deviation of the bunch distribution from a pure Gaussian [48, 93]. This allowed to reduce the systematic errors on the absolute luminosity calibration to 5% and below.

To automate these calibration sessions, the *Automatic Scan* task reads arbitrary scan sequences specified in a Domain Specific Language (DSL) from external text files [94]. While running the scan, it publishes the scan progress and actual beam separation to the experiments via DIP. This automatic scanning functionality greatly simplified the luminosity calibration scan sessions in 2016 and 2017 compared to the past. In particular, the experiment specific scan sequences for length-scale calibration (an example is given in figure 8.8) had to be followed manually by the operators in previous years. The new Luminosity Server can run them in fully automated way, which saves ∼20% of beam time per calibration session [95], and lowers the risk of operational mistakes.

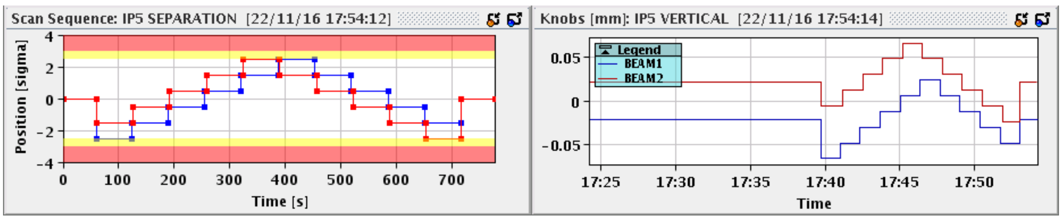

Figure 8.8: The CMS length-scale calibration scan run during the 2016 luminosity calibration session. The left chart shows the sequence as read from the file. The right chart shows the beam displacements sent to the control system by the Luminosity Server while running the scan (the zero position is arbitrary).

A third type of separation scan task automates the emittance scans described in chapter 5. In principle, this *Emittance Scan* task is a special case of the automatic scan task used for luminosity calibration. Since emittance scans are done once to twice per fill, a dedicated task was implemented to simplify the operational procedures by generating and running a pre-defined scanning sequence, providing only a limited set of input parameters with reasonable default values.

After each scan, the emittance scan task also provides average emittances estimates from a preliminary online analysis of the scan data. This analysis uses a Gaussian fit to the online luminosity measurement of the LHC experiment where the scan is executed. In the crossing plane, the derived beam width is corrected for the crossing angle and the correction for longitudinally non-Gaussian bunches as derived in chapter 5. While these estimates give a first indication of the beam emittances for operation, the full bunch-by-bunch luminosity data is stored for later offline analysis.

8.4 Separation Levelling

To keep the luminosity at one or more IPs stable within a given tolerance, the separation levelling tasks implement a feedback loop on the luminosity using the beam separation. The levelling target can either be sent by the experiments over DIP or be provided by the operator.

In a loop, the separation levelling task acquires the luminosity published by the experiment to be levelled for a certain integration time (15-60 s depending on the experiment and luminosity regime). After each acquisition, it is checked whether the measured luminosity is within the tolerance band of the target. If this is not the case, the beams are either brought closer together (if the luminosity below the target) or further separated (if the luminosity above the target).

Since the beam intensity decays the course of a fill the maximum achievable luminosity decreases ("virtual luminosity", see chapter 7). The luminosity levelling follows this and gradually decreases the beam separation to increase the luminosity back to the target. In case the separation reaches zero, the levelling is not effective anymore and the levelling task is stopped. An example is given in figure 8.9.

The separation levelling algorithm was first implemented in the previous LHC luminosity scan software, and has been used for levelling the luminosity in the ALICE and LHCb experiment since 2010. The new implementation in LHC Luminosity Server has been extended in 2017 to support separation levelling in ATLAS and CMS.

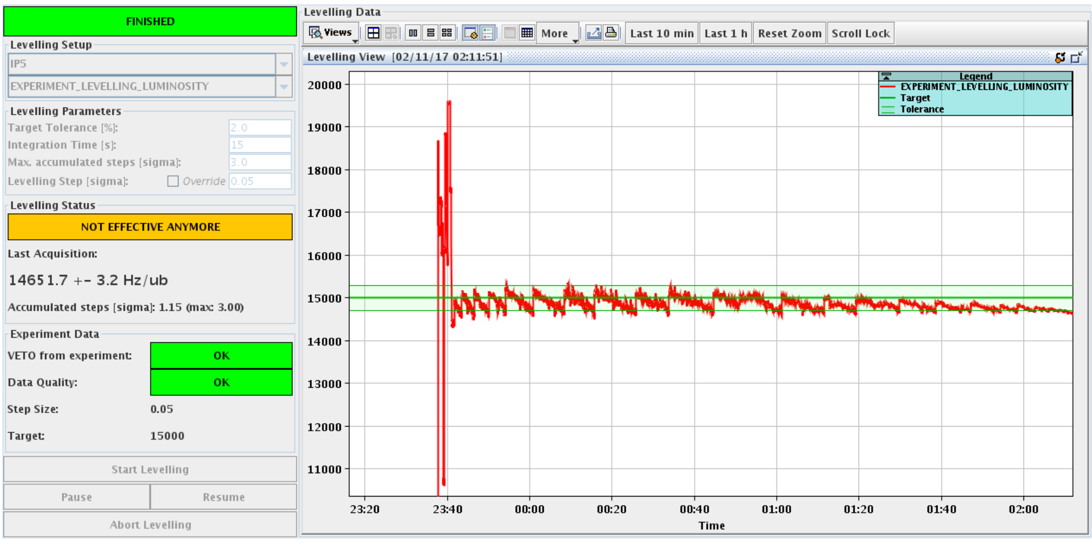

Figure 8.9: Separation levelling of IP5 (CMS) in LHC fill 6356, after the levelling task had stopped because the levelling was not effective anymore (beams head-on).

8.5 Equipment Orchestration

To allow for bigger settings changes, in particular for a change of crossing angle or β^* with the beams in collisions, the LHC Luminosity Server provides the funtionality of *orchestrating* the magnet power converters, the orbit feedback system, the collimators and the power converter software interlock system [96].

The orchestration acts on settings change (trim) requests generated by the tasks of the server. Instead of sending the trim request to LSA directly, it is passed through the orchestration layer. To apply an orchestrated trim, the following steps are taken:

1. Calculate the expected orbit response from the trim by considering the magnet kicks changed by the trim and the response matrix of the active optics (see section 2.2). If the trim contains functions to follow (e.g. for β^* changes), this calculation is done for each point of the function.

2. Get the change of all closed orbit bumps since the active feedback reference orbit was set. If there are any, build an updated reference orbit and send it to the feedback system to make sure all bumps present in the machine are included in the reference (this steps accounts for other applications which create or change orbit bumps without updating the orbit feedback reference).

3. Upload the calculated orbit response evolution to the orbit feedback system. Afterwards, arm the orbit reference change, so it will be played once a timing event to start the orbit reference transition is sent.

4. From the calculated orbit response, build functions to move the center of the affected collimators around the IP along with the beam orbit. Verify that the new positions do not violate any interlock limits. Then, arm the collimators with these functions so they will move accordingly once a timing event is sent.

5. If the trim includes an optics change (β^* changes), update the optics table in LSA to the target optics. This will make sure LSA and all other LHC applications receive the updated optics and response matrices for orbit calculations.

6. Send the trim to LSA, requesting to arm (but not to run) the affected magnet power converters with the functions generated by the luminosity server. They will start running the functions on a timing event.

7. Calculate and send the new feedback response matrix to the orbit feedback system with the requested number of eigenvalues and change the feedback gain. This allows to adjust the strength of the feedback during the change of orbit and/or optics.

8. On the LHC timing system, create a timing table to run the reference orbit changes, the collimators, the magnet power converters and the power converter interlock system synchronously. If the trim includes optics changes (β^* changes), prepare to send the optics change timing events for every matched optics point.

9. Run the timing table and wait for the power converters to reach their new currents. Since all systems involved are moved synchronously, they will reach their targets at the same time.

10. Calculate the feedback response matrix with the original number of eigenvalues and restore the previous feedback gain. In case the optics had changed due to the trim, this calculation is done for the new optics.

If any of the above steps fails or the new settings would violate the operational limits of collimator positions or power converter strength, the orchestration is cancelled and all changes are reverted. The error is propagated to the task which initiated the orchestrated trim, which in turn reports it to the operator.

8.6 Crossing Angle Anti-Levelling

Using the equipment orchestration functionality, the crossing angle anti-levelling tasks allow changing the crossing angle with beams in collisions, automatically updating the orbit feedback reference and moving the collimators appropriately.

For crossing angle levelling in 2017, a simple crossing angle steering task was used, which allows to change the crossing angle according to operator input. At certain points in a fill (typically after 2 h, 4 h and 8 h in Stable Beams), the operators used this task to decrease the half crossing angle by 10 μrad. As studied in section 7.4, this increased the LHC integrated luminosity of 2017 by $\sim$4%.

To allow for a quasi-contiuous crossing angle-levelling in 2018, a new task implementing a feedback loop on the measured bunch intensity was created (figure 8.10). This contiuous crossing angle levelling task acquires the bunch intensity, compares it against a pre-programmed iso-DA curve (see section 7.4), and performs steps of 1-5 μrad if dynamic aperture allows.

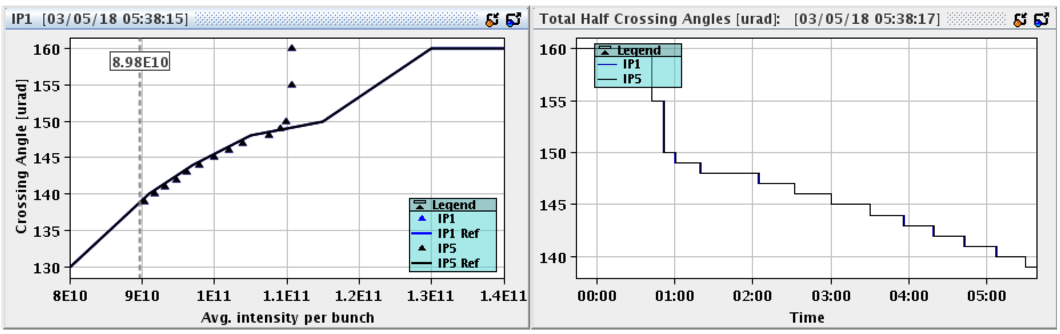

Figure 8.10: Automatic, quasi-continuos crossing angle levelling during LHC fill 6638. The left chart shows the reference iso-DA curve along with the steps done by the levelling task (triangles). The right chart shows the evolution of the crossing angle over time.

8.7 β^* Anti-Levelling

To change the β^* at one or more IPs for levelling and anti-levelling, the machine optics have to be changed in a smooth way. This is done with separated beams in every LHC cycle during the ramp and the squeeze step afterwards, where optics are changed from injection ($\beta^* = 10$-$11\,\mathrm{m}$) to physics settings. However, the established way to play the operational squeeze using pre-programmed sequences requires a static management of all machine settings (in particular any orbit corrections and closed orbit bumps), which makes it unsuitable for use in physics conditions with colliding beams and frequently changing corrections e.g. from luminosity optimization scans.

To overcome this limitation, I lead the project to allow playing any segment of a β^* squeeze[1] through the LHC Luminosity Server and its equipment orchestration functionality. While the optics settings (quadrupole strengths) and corrections are pre-matched and are driven to the hardware in absolute, settings which change regularly (e.g. orbit corrections, tune corrections or closed orbit bumps around the IPs) can be applied relative to the actual settings in the machine. The full list of parameters in either group are given in table 8.1. This more flexible implementation also allows to run squeeze segments "backwards" by inverting the time axis of the squeeze functions, which allows to squeeze and un-squeeze the beams as needed. Furthermore, the reference orbit for the orbit feedback system is dynamically calculated from the expected change of orbit. This also allows to use an orbit measured at the start of collisions instead of a static reference for the orbit feedback system to correct to, improving the orbit control significantly.

Table 8.1: LSA Parameter Groups driven by the LHC Luminosity Server for β^* changes. KNOB-level parameters are LSA knobs to control a certain physical quantity, e.g. a closed orbit offset, the tune, or the chromaticity, typically involving multiple magnets. K-level parameters are individual magnet strengths.

Absolute Parameters		**Relative Parameters**	
BETA-STAR	Beta	IP_ANGLE	KNOB
BETA-BEATING	KNOB	IP_CROSSING	KNOB
MATCHING QUADRUPOLE	K	IP_ANGLE	KNOB
TRIPLETS	K	IP_OFFSET	KNOB
TRIPLET CORRECTION	K	IP_SEPARATION	KNOB
WARM QUADRUPOLE	K	IP_SPECTROMETER	KNOB
		LUMI-SCAN	KNOB
		TUNE_TRIM	KNOB, K
		COUPLING	KNOB, K
		CHROMATICITY	KNOB, K
		ORBIT-H	K
		ORBIT-V	K
		ORBIT-TRIPLET	K
		WARM MAGNET	K
		SKEW SEXTUPOLES	K

[1] In principle, the same mechanics can be used for any optics change; however, optics changes are limited to β^* squeezes in regular LHC operation.

The first luminosity server task using this functionality just allows to change the optics by playing squeeze segments chosen by the operators. The interface of this task is shown in figure 8.11. Using this task, this new approach of β^* levelling was first validated in machine tests (the results are given in section 7.5). It will be used for anti-levelling in 2018 LHC operation [82]. In the future, a second task using the same underlying mechanics can be implemented to use the β^* changes for levelling according to the requests of the experiments, implementing the same feedback loop as the separation levelling tasks. Levelling the luminosity by β^* is a baseline scenario for the HL-LHC upgrade [44].

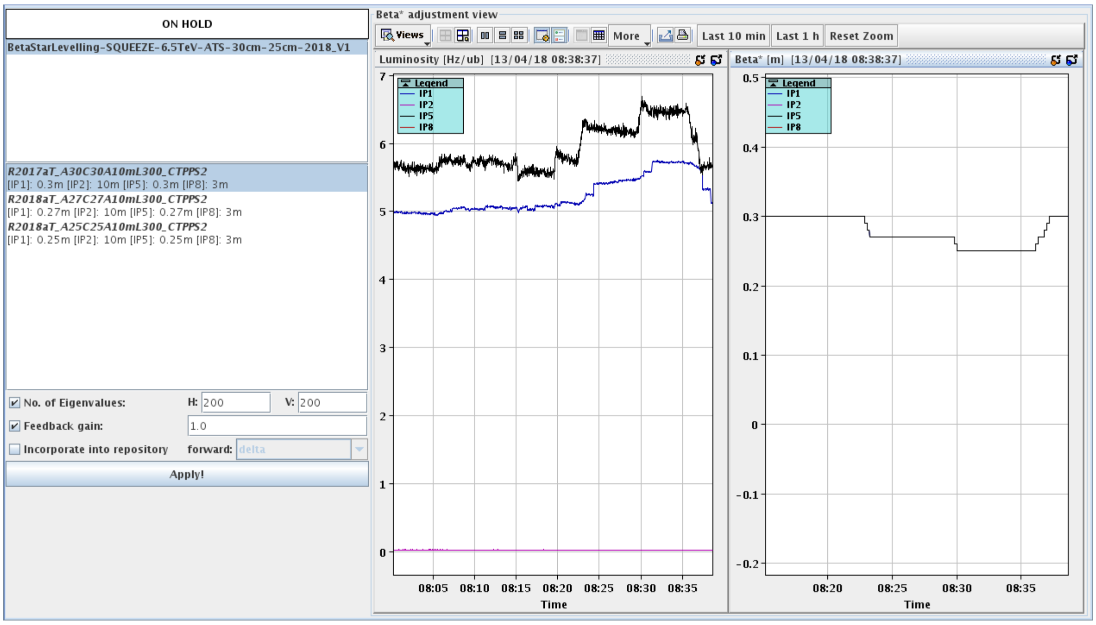

Figure 8.11: Screenshot of the task used for β^* changes, taken during 2018 LHC commissioning. The task has the squeeze segment from $\beta^* = 30\,\text{cm}$ to $\beta^* = 25\,\text{cm}$ in IP 1 and 5 loaded, which will be used in 2018 LHC operation.

9 Conclusions

This work was focused on a set of topics with the common goal of improving the overall luminosity performance of the LHC, and to introduce technologies vital for its upcoming luminosity upgrade, the HL-LHC.

A new method of using beam separation scans in regular LHC physics fills to derive bunch-by-bunch emittance and closed-orbit separations from the luminosity measurements was introduced. This allowed to better understand the emittance preservation (and hence the peak luminosity reach), and the beam-beam long-range effects, which lead to bunch-to-bunch differences in closed orbit, tune, and chromaticity. The direct measurement of bunch-by-bunch closed orbit separations was compared to simulations, and the existing TRAIN simulation code was significantly improved to consider also previously-neglected effects. This improved version of the TRAIN code not only reproduces all observations within the expected accuracy, but can also be used to predict beam-beam long-range effects in future colliders.

Using the data from these measurements, the luminosity measured by the LHC experiments, and measurements from beam instrumentation (intensity, bunch length), the peak luminosity reach and the time evolution of the LHC luminosity was studied. In particular, a method of predicting the optimum fill length (yielding the maximum integrated luminosity) from the measured luminosity evolution with minimum assumptions was established. The predictions made were followed by the LHC machine committee and physics coordinators in 2015-2018 LHC operation. Furthermore, it was shown that due to the alternating crossing scheme of the LHC, non-round beams introduce an asymmetry in the luminosities delivered to the ATLAS and the CMS experiments; such an asymmetry was observed in 2016 operation.

The last part of my work was focused on luminosity levelling (to control the event pile-up) and anti-levelling (to increase the integrated luminosity), in particular by changing the crossing angle or β^* with colliding beams. While the principle of these methods was demonstrated in dedicated machine studies before, they could not be used in regular operation due to various limitations of the LHC control systems. The new luminosity control software stack (the LHC Luminosity Server) developed as a part of this work allows to "orchestrate" all accelerator equipment needed to make smooth transitions between crossing angles or IP optics (β^*): the magnet power converters, the orbit feedback system, and the collimators around the IPs. This novel concept was initially applied for crossing-angle anti-levelling throughout 2017 LHC operation, gaining $\sim$4% ($2\,\text{fb}^{-1}$) of integrated luminosity for the ATLAS and CMS experiments. In the following, the method was extended for β^* transitions, which will be used for anti-levelling in 2018. For the first time in LHC, this will allow changing the β^* with the beams in collisions and Stable Beams declared, which is a prerequisite for β^* levelling in the HL-LHC.

Bibliography

[1] Klaus Wille. *The physics of particle accelerators: an introduction*. Oxford: Oxford Univ. Press, 2000. URL: https://cds.cern.ch/record/560708.

[2] Helmut Wiedemann. *Particle accelerator physics; 3rd ed.* Berlin: Springer, 2007. URL: https://cds.cern.ch/record/1083415.

[3] Shu-Yin Lee. *Accelerator Physics; 2nd ed.* London: World Scientific, 2004. URL: https://cds.cern.ch/record/939203.

[4] Alexander Wu Chao et al. *Handbook of accelerator physics and engineering; 2nd ed.* Singapore: World Scientific, 2013. URL: https://cds.cern.ch/record/1490001.

[5] Donald A. Edwards and Michael J. Syphers. *An introduction to the Physics of High Energy Accelerators*. Wiley Beam Phys. Accel. Technol. New York, NY: Wiley, 1993. URL: https://cds.cern.ch/record/249993.

[6] Oliver Sim Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. URL: https://cds.cern.ch/record/782076.

[7] The CERN LEP group. *LEP design report*. Geneva: CERN, 1984. URL: https://cds.cern.ch/record/102083.

[8] Lyn Evans and Vincent Hatton. *Operational aspects of the SPS collider*. Tech. rep. CERN-SPS-85-3-AOP. Geneva: CERN, 1985. URL: https://cds.cern.ch/record/158762.

[9] The CERN Study Group on New Accelerators. *Report on the design study of Intersecting Storage Rings (ISR) for the CERN proton synchrotron*. Geneva: CERN, 1964. URL: https://cds.cern.ch/record/110073.

[10] Werner Herr. *Particle Colliders and Concept of Luminosity*. CERN Accelerator School. Oct. 2012. URL: http://cas.web.cern.ch/sites/cas.web.cern.ch/files/lectures/granada-2012/herr.pdf.

[11] Massimiliano Ferro-Luzzi. *Luminosity and luminous region shape for pure Gaussian bunches*. Tech. rep. LHCb-PUB-2012-016. Geneva: CERN, Sept. 2012. URL: https://cds.cern.ch/record/1481670.

[12] Massimo Giovannozzi. "Simple models describing the time-evolution of luminosity in hadron colliders". In: Proc. of IPAC'14, pp. 1017 (June 2014). URL: https://cds.cern.ch/record/1742039.

[13] Fanouria Antoniou et al. "Building a Luminosity Model for the LHC and HL-LHC". In: Proc. of IPAC'15, pp. 2042 (2015). URL: https://cds.cern.ch/record/2141831.

[14] Maria Kuhn. "Emittance Preservation at the LHC". PhD thesis. Mar. 2013. URL: https://cds.cern.ch/record/1544573.

[15] Werner Herr and Tatiana Pieloni. *Beam-Beam Effects*. Proc. of CERN Accelerator School: Advanced Accelerator Physics. Aug. 2013. URL: https://cds.cern.ch/record/1982430.

[16] Esma Mobs. *The CERN accelerator complex*. CERN, OPEN-PHO-ACCEL-2016-009. Aug. 2017. URL: https://cds.cern.ch/record/2197559.

[17] The ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[18] The CMS Collaboration. "The CMS Experiment at the CERN LHC". In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[19] The ALICE Collaboration. "The ALICE experiment at the CERN LHC". In: *JINST* 3 (2008), S08002. DOI: 10.1088/1748-0221/3/08/S08002.

[20] The LHCb Collaboration. "The LHCb Detector at the LHC". In: *JINST* 3 (2008), S08005. DOI: 10.1088/1748-0221/3/08/S08005.

[21] Werner Herr. *Features and implications of different LHC crossing schemes*. Tech. rep. LHC-Project-Report-628. Geneva: CERN, Jan. 2003. URL: https://cds.cern.ch/record/604005.

[22] Colin Barschel et al. "Results of the LHC DCCT Calibration Studies". In: (Feb. 2012). URL: https://cds.cern.ch/record/1425904.

[23] Jacques Bosser and Claude Bovet. *Wire Scanners For LHC*. Tech. rep. Geneva: CERN, Sept. 1997. URL: https://cds.cern.ch/record/692054.

[24] Georges Trad et al. "Status of the beam profile measurements at the LHC". In: Proc. of Evian 2015, pp. 163 (2016). URL: https://cds.cern.ch/record/2294529.

[25] Georges Trad. "Development and Optimisation of the SPS and LHC beam diagnostics based on Synchrotron Radiation monitors". PhD thesis. Aug. 2014. URL: https://cds.cern.ch/record/2266055.

[26] Georges Trad. *Instrumentation performance and limitations*. Presented at the 8th LHC Operations Evian Workshop. Dec. 2017. URL: https://indico.cern.ch/event/663598/contributions/2782456/attachments/1565006/2486346/Evian2017_Emittances_GT.pptx.

[27] Giulia Papotti et al. "Longitudinal Beam measurements at the LHC: The LHC Beam Quality Monitor". In: Proc. of IPAC'11, pp. 1852 (Nov. 2011). URL: https://cds.cern.ch/record/1399087.

[28] Juan Esteban Muller. "Longitudinal intensity effects in the CERN Large Hadron Collider". PhD thesis. Apr. 2016. URL: https://cds.cern.ch/record/2196930.

[29] Marek Gasior, Jakub Olexa, and Ralph Steinhagen. "BPM Electronics based on Compensated Diode Detectors – Results from development Systems". In: Proc. of 2012 Beam Instrumentation Workshop, pp. 44 (Apr. 2012). URL: https://cds.cern.ch/record/1476070.

[30] Ralph Steinhagen. "LHC Beam Stability and Feedback Control - Orbit and Energy". PhD thesis. 2007. URL: https://cds.cern.ch/record/1054849.

[31] Roderik Bruce et al. "Collimator hierarchy limits: assumptions and impact on machine protection and performance". In: Proc. of MPP Workshop March 2013, pp. 99 (2013). URL: https://cds.cern.ch/record/1695409.

[32] Stephane Fartoukh. "An Achromatic Telescopic Squeezing (ATS) Scheme For The LHC Upgrade". In: Proc. of IPAC'11, pp. 2088 (Sept. 2011). URL: https://cds.cern.ch/record/1382077.

[33] *LHC Page 1*. URL: https://op-webtools.web.cern.ch/Vistar/vistars.php.

[34] Matteo Solfaroli Camillocci et al. "Combined Ramp and Squeeze to 6.5 TeV in the LHC". In: Proc. of IPAC'16, pp. 1509 (2016). URL: https://cds.cern.ch/record/2207383.

[35] Simon van der Meer. *Calibration of the effective beam height in the ISR*. Tech. rep. CERN-ISR-PO-68-31. Geneva: CERN, 1968. URL: https://cds.cern.ch/record/296752.

[36] Michael Hostettler and Giulia Papotti. "Beam Size Estimation from Luminosity Scans at the LHC During 2015 Proton Physics Operation". In: Proc. of IPAC'16, pp. 290 (2016). URL: https://cds.cern.ch/record/2207311.

[37] Michael Hostettler, Kajetan Fuchsberger, and Giulia Papotti. "Operational Experience with Luminosity Scans for Beam Size Estimation in 2016 LHC Proton Physics Operation". In: Proc. of IPAC'17, pp. 374 (2017). URL: https://cds.cern.ch/record/2287340.

[38] Michael Hostettler et al. "Luminosity Scans for Beam Diagnostics". Submitted to *Phys. Rev. ST Accel. Beams*. Preprint: arXiv:1804.10099 [physics.acc-ph]. 2018. URL: https://arxiv.org/abs/1804.10099.

[39] Philippe Baudrenghien and Themis Mastoridis. "Longitudinal emittance blowup in the large hadron collider". In: *Nucl. Instrum. Methods Phys. Res., A* 726 (2013), pp. 181–190. URL: https://cds.cern.ch/record/1709247.

[40] Stefania Papadopoulou et al. "Modelling and measurements of bunch profiles at the LHC". In: *Journal of Physics: Conference Series* 874.1 (2017), p. 012008. URL: https://cds.cern.ch/record/2207375.

[41] Hans Grote. *Self-consistent Orbits for Beam-beam Interactions in the LHC*. Tech. rep. LHC-PROJECT-NOTE-216. Geneva: CERN, Feb. 2000. URL: https://cds.cern.ch/record/691837.

[42] Laurent Deniau et al. *The MAD-X Program: User's Reference Manual*. Tech. rep. Geneva: CERN, Oct. 2017. URL: http://madx.web.cern.ch/madx/.

[43] Heiko Damerau et al. "LIU: exploring alternative ideas". In: Proc. of RLIUP Workshop (2014). DOI: 10.5170/CERN-2014-006.127. URL: https://cds.cern.ch/record/1977365.

[44] Giorgio Apollinari et al. *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2116337.

[45] Austin Ball et al. *Future Circular Collider Study: Hadron Collider Parameters*. Tech. rep. FCC-ACC-SPC-0001. Geneva: CERN, Feb. 2014. URL: https://fcc.web.cern.ch/.

[46] Olena Karacheban et al. "Emittance scans for CMS luminosity calibration". Presented at the LHCC students poster session. Feb. 2018. URL: https://indico.cern.ch/event/705390/contributions/2894890/.

[47] Olena Karacheban and the CMS collaboration. "CMS emittance scans for luminosity calibration in 2017". Presented at the XXII International Scientific Conference of Young Scientists and Specialists (AYSS). 2018. URL: https://indico.jinr.ru/contributionDisplay.py?sessionId=12&contribId=562&confId=436.

[48] *CMS Luminosity Measurements for the 2016 Data Taking Period*. Tech. rep. CMS-PAS-LUM-17-001. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2257069.

[49] Andres Delannoy. Plot prepared for [46] using vdm-toolkit by P. Tsrunchev. 2018.

[50] Jorg Wenninger. *Operational Scenario for 2018: Pulling it all together*. Presented at the LHC Performance Workshop, Chamonix. 2018. URL: https://indico.cern.ch/event/676124/contributions/2767843/attachments/1590637/2517642/LHC2018.PullingIt2gather.Jan18.pdf.

[51] Jorg Wenninger. Private communication. 2018.

[52] Mike Barnes et al. "MKI performance during 2016 and plans for EYETS". In: Proc. of Evian 2015, pp. 175 (2017). URL: https://cds.cern.ch/record/2293529.

[53] Lotta Mether. *16L2: Operation, observations and physics aspects*. Presented at the 8th LHC Operations Evian Workshop. Dec. 2017. URL: https://indico.cern.ch/event/663598/contributions/2782390/attachments/1574471/2485998/004_Evian2017_16L2_Lotta.pdf.

[54] Giovanni Iadarola and Giovanni Rumolo. "Electron cloud in the CERN accelerators (PS, SPS, LHC)". In: *AIP Conf. Proc. C* arXiv:1309.6795 (Sept. 2013), pp. 19–26. URL: https://cds.cern.ch/record/1603635.

[55] Gianluigi Arduini, Attilio Milanese, and Diego Perini. *Resistive solenoid in the 16L2 interconnect*. Presented at the 319th meeting of the LHC Machine Committee. URL: https://indico.cern.ch/event/665902/.

[56] Heiko Damerau et al. "RF Manipulations for Higher Brightness LHC-Type Beams". In: Proc. of IPAC'13, pp. 2600 (2013). URL: https://cds.cern.ch/record/1595719.

[57] Hannes Bartosik and Giovanni Rumolo. "Beams from the injectors". In: Proc. of Evian 2016, pp. 233 (2016). URL: https://cds.cern.ch/record/2293681.

[58] Alexandre Lasheen. *Overview of the beams from the injectors*. Presented at the 8th LHC Operations Evian Workshop. Dec. 2017. URL: https://indico.cern.ch/event/663598/contributions/2782426/attachments/1574425/2487466/Evian_overview_2017.pdf.

[59] Tetsuo Abe et al. "Achievements of KEKB". In: *Prog. Theor. Exp. Phys.* 2013 (2013), 03A001. URL: https://cds.cern.ch/record/1709394.

[60] The TOTEM Collaboration. *First measurement of elastic, inelastic and total cross-section at $\sqrt{s} = 13$ TeV by TOTEM and overview of cross-section data at LHC energies*. 2017. eprint: arXiv:1712.06153. URL: https://arxiv.org/abs/1712.06153.

[61] Michi Hostettler. *Analysis of lumi lifetime*. Presented at the 269th meeting of the LHC Machine Committee. URL: https://indico.cern.ch/event/547490/.

[62] Fanouria Antoniou et al. "Can we Predict Luminosity?" In: Proc. of Evian 2016, pp. 152 (2016). URL: https://cds.cern.ch/record/2293678.

[63] Stefania Papadopoulou. *Emittance, Intensity and Luminosity modeling and evolution*. Presented at the 8th LHC Operations Evian Workshop. Dec. 2017. URL: https://indico.cern.ch/event/663598/contributions/2782463/attachments/1574684/2493677/evian2017_stef.pdf.

[64] Giulia Papotti. *Observations on luminosity curves*. Presented at the 38th Meeting of LHC Beam Operation Committee. May 2012. URL: http://lhc-beam-operation-committee.web.cern.ch/lhc-beam-operation-committee/minutes/Meeting38-22_05_2012/lumi.GP.22May12.pdf.

[65] Vladimir Shiltsev and Erin McCrory. "Characterizing Luminosity Evolution in the Tevatron". In: *Proc. of PAC'2005*. May 2005, pp. 2536–2537. DOI: 10.1109/PAC.2005.1591170.

[66] Kajetan Fuchsberger. "Turnaround - Analysis and possible Improvements". In: Proc. of Evian 2016, pp. 33 (2017). URL: https://cds.cern.ch/record/2293508.

[67] Michi Hostettler. *Luminosity Evolution and Optimum Fill Length: an update for BCMS beams*. Presented at the 272th meeting of the LHC Machine Committee. URL: https://indico.cern.ch/event/559439/.

[68] Michael Hostettler et al. "Impact of the Crossing Angle on Luminosity Asymmetries at the LHC in 2016 Proton Physics Operation". In: Proc. of IPAC'17, pp. 2035 (2017). DOI: 10.18429/JACoW-IPAC2017-TUPVA005. URL: http://inspirehep.net/record/1626446/files/tupva005.pdf.

[69] ATLAS Collaboration. Preliminary uncertainty, derived from a calibration of the luminosity scale using x-y beam-separation scans performed in May 2016, following a methodology similar to that used in LHC run 1 and detailed in [93].

[70] Zachary Marshall for The ATLAS Collaboration. "Simulation of Pile-up in the ATLAS Experiment". In: *Journal of Physics: Conference Series* 513.2 (2014), p. 022024. URL: http://stacks.iop.org/1742-6596/513/i=2/a=022024.

[71] The ATLAS Collaboration. *Luminosity Public Results, Run 2*. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[72] Xavier Buffat et al. “Stability diagrams of colliding beams in the Large Hadron Collider”. In: *Phys. Rev. Spec. Top. Accel. Beams* 17.11 (2014), 111002. 9 p. URL: https://cds.cern.ch/record/2011653.

[73] Tatiana Pieloni et al. “Observations of Two-beam Instabilities during the 2012 LHC Physics Run”. In: Proc. of IPAC’13, pp. 1418 (2013). URL: http://cds.cern.ch/record/1581728.

[74] Michi Hostettler and Giulia Papotti. “Luminosity Lifetime at the LHC in 2012 Proton Physics”. In: Proc. of IPAC’13, pp. 1403 (2013). URL: http://cds.cern.ch/record/1567672.

[75] Nikos Karastathis et al. “Crossing Angle Anti-Levelling at the LHC in 2017”. To be presented at IPAC’18, Vancouver, Canada.

[76] Michi Hostettler. *Crossing Angle Anti-Levelling: Strategy and Status*. Presented at the 309th meeting of the LHC Machine Committee. URL: https://indico.cern.ch/event/649726/.

[77] Arkadiusz Gorzawski et al. *Collide and Squeeze MD*. CERN-ACC-NOTE-2015-0043. Dec. 2015. URL: http://cds.cern.ch/record/2115437.

[78] Arkadiusz Gorzawski et al. “Experimental Demonstration of β^* Leveling at the LHC”. In: Proc. of IPAC’16, pp. 1442 (2016). URL: http://cds.cern.ch/record/2207366.

[79] Arkadiusz Gorzawski. “Luminosity control and beam orbit stability with beta star leveling at LHC and HL-LHC”. PhD thesis. 2016. URL: http://cds.cern.ch/record/2238309.

[80] Jorg Wenninger et al. *beta* leveling with telescopic ATS squeeze (MD 2410)*. CERN-ACC-NOTE-2017-0052. Sept. 2017. URL: http://cds.cern.ch/record/2285184.

[81] Michi Hostettler et al. *β^* levelling using the LHC Lumi Server (MD 2427)*. CERN-ACC-NOTE-2018-0001. Jan. 2018. URL: https://cds.cern.ch/record/2300212.

[82] Michi Hostettler and Andrea Calia. *Beta* levelling implementation proposal*. Presented at the 95th Meeting of the LHC Beam Operation Committee. Mar. 2018. URL: https://indico.cern.ch/event/712834/.

[83] Stephane Fartoukh. *ATS optics for 2017*. Presented at the 74th Meeting of LHC Beam Operation Committee. Mar. 2017. URL: https://indico.cern.ch/event/625342/.

[84] Michael Hostettler et al. “Online luminosity control and steering at the LHC”. In: Proc. of ICALEPCS’17 (2017). URL: https://cds.cern.ch/record/2305966.

[85] Reyes Alemany-Fernandez, Mike Lamont, and Simon White. *Separation Scan Application Requirements*. EDMS specification. URL: https://edms.cern.ch/document/894460/1.

[86] Simon Mathieu White. "Determination of the Absolute Luminosity at the LHC". PhD thesis. 2010. URL: http://cds.cern.ch/record/1308187.

[87] *Java Message Service API*. JSR 914. URL: https://www.jcp.org/en/jsr/detail?id=914.

[88] *Java Remote Method Invocation*. URL: http://www.oracle.com/technetwork/java/javase/tech/index-jsp-136424.html.

[89] J Andersson et al. "The Controls Middleware (CMW) at CERN Status and Usage". In: Proc. of ICALEPCS'03 (2003). URL: https://cds.cern.ch/record/693152.

[90] The CERN EN-ICE group. *Data Interchange Protocol*. URL: https://cern.ch/dip/.

[91] Mike Lamont and Lionel Mestre. *LHC Era Core Control Application Software*. Presented at ICALEPCS'05. URL: https://icalepcs2005.web.cern.ch/icalepcs2005/Presentations/PostDocs/TH1_4-8O-2.pdf.

[92] Andrea Calia, Kajetan Fuchsberger, and Michi Hostettler. "Testing the Untestable: A Realistic Vision of Fearlessly Testing (Almost) Every Single Accelerator Component Without Beam and Continuous Deployment Thereof". In: *Proc. of IBIC'16*, pp. 400–403. URL: http://jacow.org/ibic2016/papers/tupg30.pdf.

[93] The ATLAS Collaboration. "Luminosity determination in pp collisions at $\sqrt{s}$ = 8 TeV using the ATLAS detector at the LHC". In: *Eur. Phys. J.* C76.12 (2016), p. 653. DOI: 10.1140/epjc/s10052-016-4466-1. arXiv: 1608.03953 [hep-ex]. URL: https://cds.cern.ch/record/2208146.

[94] Reyes Alemany-Fernandez et al. *Automatic scan application around the LHC interaction points*. EDMS specification. URL: https://edms.cern.ch/document/1367878/7.

[95] Reyes Alemany-Fernandez and Michi Hostettler. *Report about new lumi calibration program*. Presented at the LPC Meeting of 23 October 2015.

[96] Kajetan Fuchsberger et al. "LHC Orbit Correction Reproducibility and Related Machine Protection". In: Proc. of IPAC'12 (2012). URL: https://cds.cern.ch/record/1470556.

List of Figures

M. Hostettler

Curriculum Vitae

✉ *michi.hostettler@cern.ch*

Education

May 2015 – Jun 2018 **PhD studies**, *CERN, BE-OP-LHC and University of Bern*.
LHC Luminosity-related data processing, analysis and modeling; operational luminosity levelling using crossing angle and β^*; observation and simulation of beam-beam effects.
Supervisors: Prof. Dr. Antonio Ereditato, Dr. Giulia Papotti, Dr. Kajetan Fuchsberger.

Sep 2013 – Mar 2015 **Master of Science**, *University of Bern*, Summa Cum Laude.
Experimental/Applied Physics. Running ATLAS detector simulations on HPC systems.
Supervisors: PD Dr. Sigve Haug, Prof. Dr. Antonio Ereditato, Prof. Dr. Michele Weber

Aug 2012 – Jul 2013 **CERN Technical Studentship**, *CERN, BE-OP-LHC*.
LHC Luminosity-related data processing, analysis and modeling in the BE-OP-LHC section.
Supervisor: Dr. Giulia Papotti.

Sep 2008 – Jul 2012 **Bachelor of Science**, *University of Bern*, Insigni Cum Laude.
Major Physics (120 ECTS)
Minor Mathematics (30 ECTS)
Minor Computer Science (30 ECTS)

Aug 2003 – Jul 2007 **Matura**, *Gymnasium Bern-Kirchenfeld*, Bern, GPA: 5.15/6.
Schwerpunktfach (*Major*): Physics and Applied Mathematics
Ergänzungsfach (*Minor*): Biology and Chemistry
Maturaarbeit (*Matura thesis*): Simulation von chemisch-dynamischen Gleichgewichten (*Developing software to simulate dynamic equilibria in chemistry*)

Aug 1995 – Jul 2003 **Primary School**, *Pädagogisches Ausbildungszentrum NMS*, Bern.

PhD thesis

title *LHC Luminosity Performance*

supervisors Prof. Dr. A. Ereditato, Dr. Giulia Papotti, Dr. Kajetan Fuchsberger

description The work was globally aimed towards improving the luminosity yield of the LHC. This involved studies and modelling of the LHC luminosity evolution, the underlying beam parameters (emittance, intensity), and making proposals for optimization (e.g. optimum fill length). A new method for measuring the emittances and beam-beam long-range effect was put in place, and results were compared to simulations from the TRAIN code, which then was improved based on the results. Also, the LHC luminosity control software stack was re-designed to be more robust and allow for new features like fully-automated calibration scans and operational crossing angle and β^* luminosity levelling. These improvements were used in daily operation since and improved the LHC luminosity in 2017 by a few fb^{-1}.

Master thesis

title	*Enabling ATLAS for HPC*
supervisors	Prof. Dr. A. Ereditato, Prof. Dr. Michele Weber, PD Dr. Sigve Haug
description	Working with the ATLAS HPC working group, possibilities for running ATLAS production on current and future HPC resources were assessed. In particular, event generation and detector simulation workloads were tested on a Cray XK7 machine located at the Swiss National Supercomputing Centre (CSCS), and a solution for running multi-core detector simulation jobs continuously on HPC systems in a fully automated manner has been implemented and tested. Possibilities for using GPUs as accelerators on hybrid systems have been investigated.
grade	6 of 6, prize for the best thesis in applied physics at the University of Bern in 2015.
available at	`https://cds.cern.ch/record/2253719`

Bachelor thesis

title	*Entwicklung einer Steuer- und Messsoftware für Nahfeld-THz-Mikroskopie*
supervisors	Prof. Dr. T. Feurer, F. Enderli
description	For an existing near-field THz microscope at the Institute for Applied Physics in Bern, a new measurement and control software was developed in NI LabVIEW. The microscope does 3-D time-domain scans using THz waves. The new software is well structured and increases the time measurement precision by using an external trigger signal to sync the probe beam control to the lock-in-amplifier used for measurement.
grade	6 of 6

Publications

2018	M. Hostettler *et al.*: "Beam-Beam Orbit Effects in the LHC", ICFA Workshop on Beam-Beam Effects in Circular Colliders (BB2018)
2018	M. Hostettler *et al.*: "Luminosity Scans for Beam Diagnostics", in preparation, to be submitted to Phys. Rev. ST Accelerators and Beams.
2018	M. Hostettler *et al.*: "β^* levelling using the LHC Lumi Server (MD 2427)", CERN MD Note, CERN-ACC-NOTE-2018-0001.
2017	M. Hostettler *et al.*: "Online Luminosity Control and Steering at the LHC", 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)
2017	A.A. Gorzawski, T. Pieloni, K. Fuchsberger, M. Hostettler, J. Wenninger: "Long-Range Beam-Beam Orbit Effects in LHC, Simulations and Observations from Machine Operation in 2016", 8th International Particle Accelerator Conference (IPAC'17)
2017	M. Hostettler, K. Fuchsberger, G. Papotti: "Impact of the Crossing Angle on Luminosity Asymmetries at the LHC in 2016 Proton Physics Operation", 8th International Particle Accelerator Conference (IPAC'17)
2017	M. Hostettler *et al.*: "Operational Experience with Luminosity Scans for Beam Size Estimation in 2016 LHC Proton Physics Operation", 8th International Particle Accelerator Conference (IPAC'17)

2016 F. Antoniou, G. Arduini, M. Hostettler, M. Lamont, S. Papadopoulou, Y. Papaphilippou, G. Papotti, B. Salvachua: "LHC luminosity modeling for Run II", 7th International Particle Accelerator Conference (IPAC'16)

2016 M. Hostettler, G. Papotti: "Beam Size Estimation from Luminosity Scans at the LHC during 2015 Proton Physics Operation", 7th International Particle Accelerator Conference (IPAC'16)

2015 S. Haug, M. Hostettler, G. Sciacca, M. Weber : "The ATLAS ARC backend to HPC", 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP'15)

2015 M. Hostettler, A. Filipcic, S. Haug, R. Walker, M. Weber : "ATLAS computing on CSCS HPC", 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP'15)

2014 M. Hostettler, S. Haug: "Using CSCS HPC Resources for LHC ATLAS Data Processing", Platform for Advanced Scientific Computing (PASC'14)

2013 M. Hostettler, G. Papotti: "Luminosity Lifetime at the LHC in 2012 Proton Physics Operation", 4th International Particle Accelerator Conference (IPAC'13)

2013 G. Papotti, A. Gorzawski, M. Hostettler, R. Schmidt, "Beam Losses Through the LHC Operational Cycle in 2012", 4th International Particle Accelerator Conference (IPAC'13)

2013 M. Hostettler, G. Papotti: "Observations from LHC Operation", ICFA Mini-Workshop on Beam-Beam Effects in Hadron Colliders (BB2013)

Erklärung

gemäss Art. 28 Abs. 2 RSL 05

Name/Vorname: Hostettler Michael Artur

Matrikelnummer: 08-112-856

Studiengang: Doktorat Phil.-Nat., Physik

Bachelor ☐ Master ☐ Dissertation ☑

Titel der Arbeit: LHC Luminosity Performance

LeiterIn der Arbeit: Prof. Dr. Antonio Ereditato

Ich erkläre hiermit, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen benutzt habe. Alle Stellen, die wörtlich oder sinngemäss aus Quellen entnommen wurden, habe ich als solche gekennzeichnet. Mir ist bekannt, dass andernfalls der Senat gemäss Artikel 36 Absatz 1 Buchstabe r des Gesetzes vom 5. September 1996 über die Universität zum Entzug des auf Grund dieser Arbeit verliehenen Titels berechtigt ist. Ich gewähre hiermit Einsicht in diese Arbeit.

Ort/Datum

Bern, 21. Juni 2018

Unterschrift