



Impact of beam–beam effects on absolute luminosity calibrations at the CERN Large Hadron Collider

A. Babaev¹, T. Barklow², O. Karacheban^{3,4}, W. Kozanecki⁵, I. Kralik⁶, A. Mehta⁷, G. Pasztor⁷, T. Pieloni^{8,a}, D. Stickland⁹, C. Tambasco⁸, R. Tomas⁴, J. Wańczyk^{4,8}

¹ Tomsk Polytechnic University, Tomsk, Russia

² SLAC National Accelerator Laboratory, Stanford, CA, USA

³ Rutgers, The State University of New Jersey, Piscataway, NJ, USA

⁴ CERN, Geneva, Switzerland

⁵ IRFU-CEA, Université Paris-Saclay, 91191 Gif-sur-Yvette, France

⁶ Institute of Experimental Physics SAS, Watsonova 47, 04001 Kosice, Slovak Republic

⁷ MTA-ELTE Lendület CMS Particle and Nuclear Physics Group, Eötvös Loránd University (ELTE), Budapest, Hungary

⁸ Laboratory of Particle Accelerator Physics, EPFL, Lausanne, Switzerland

⁹ Princeton University, Princeton, NJ, USA

Received: 7 July 2023 / Accepted: 26 October 2023

© The Author(s) 2024

Abstract At the Large Hadron Collider (LHC), absolute luminosity calibrations obtained by the van der Meer (*vdM*) method are affected by the mutual electromagnetic interaction of the two beams. The colliding bunches experience relative orbit shifts, as well as optical distortions akin to the dynamic- β effect, that both depend on the transverse beam separation and must therefore be corrected for when deriving the absolute luminosity scale. In the *vdM* regime, the beam–beam parameter is small enough that the orbit shift can be calculated analytically. The dynamic- β corrections to the luminometer calibrations, however, had until the end of Run 2 been estimated in the linear approximation only. In this report, the influence of beam–beam effects on the *vdM*-based luminosity scale is quantified, together with the associated systematic uncertainties, by means of simulations that fully take into account the non-linearity of the beam–beam force, as well as the resulting non-Gaussian distortions of the transverse beam distributions. Two independent multi-particle simulations, one limited to the weak-strong approximation and one that models strong-strong effects in a self-consistent manner, are found in excellent agreement; both predict a percent-level shift of the absolute *pp*-luminosity values with respect to those assumed until recently in the physics publications of the LHC experiments. These results also provide guidance regarding further studies aimed at reducing the beam–beam-related systematic uncertainty on

beam–beam corrections to absolute luminosity calibrations by the van der Meer method.

Contents

1	Introduction
2	Luminosity-calibration methodology at the LHC
2.1	Absolute luminosity scale from measured beam parameters
2.2	Beam conditions during van der Meer scans
2.3	Beam–beam-induced biases and their correction
2.3.1	Orbit shift
2.3.2	Optical distortions
2.3.3	Beam–beam correction strategy
3	Beam–beam simulation codes
3.1	Linear approximation with MAD-X
3.2	Analytical estimate of optical distortions at zero beam separation
3.3	Weak-strong limit: B*B
3.4	Strong-strong model: COMBI
3.5	Cross-validation of simulation codes
3.5.1	Beam–beam-induced orbit shift
3.5.2	Impact of amplitude-dependent β -beating effects in the Gaussian-bunch approximation
3.5.3	Impact of optical distortions on the beam–beam overlap integral
3.5.4	Beam-parameter dependence of beam–beam corrections

^a e-mail: Tatiana.Pieloni@epfl.ch (corresponding author)

4	Calculated impact of beam–beam dynamics on luminosity calibrations
4.1	Methodology
4.2	Beam–beam correction procedure in the fully symmetric Gaussian-beam configuration
4.2.1	Validation of the scaling hypothesis
4.2.2	Beam–beam parameter and tune dependence of the optical-distortion correction
4.2.3	Practical implementation of parameterized beam–beam corrections
4.2.4	Separation-integrated estimates of beam–beam corrections to vdM calibrations
4.3	Impact of non-Gaussian unperturbed transverse beam profiles
4.3.1	Single-bunch models
4.3.2	Beam–beam-correction formalism in the presence of non-factorization
4.3.3	Impact of non-Gaussian tails on beam–beam-induced biases
4.4	Impact of beam ellipticity at the interaction point
4.5	Impact of non-zero crossing angle
4.6	Impact of multiple interaction points
4.6.1	Motivation and methodology
4.6.2	Impact of multiple interaction points on the tune spectra
4.6.3	Comparison of reference “no beam–beam” configurations in the presence of one non-scanning IP
4.6.4	Sensitivity of beam–beam corrections to the phase advance between IPs
4.6.5	Parameterization of multi-IP effects
4.6.6	Beam–beam correction strategies in the presence of multi-IP collisions
4.7	Impact of beam–beam asymmetry
4.7.1	Beam–beam asymmetry scenarios
4.7.2	Parametric characterization of beam–beam asymmetries
4.7.3	Predicted impact of transverse beam-size asymmetries
5	Systematic uncertainties
5.1	Optical configuration
5.1.1	β^* uncertainty at the scanning IP
5.1.2	Nominal collision tunes
5.2	Deviations from the fully symmetric Gaussian-beam configuration
5.2.1	Non-Gaussian unperturbed transverse-density profiles
5.2.2	Beam ellipticity at the scanning IP
5.2.3	Non-zero crossing angle
5.2.4	Beam–beam imbalance
5.3	Multiple interaction points
5.4	Other potential sources of systematic uncertainty
5.4.1	Long-range encounters

5.4.2	Lattice non-linearities
5.4.3	Numerical accuracy of polynomial parameterization
5.5	Summary of systematic uncertainties
6	Summary and outlook
Appendix A:	Impact of non-Gaussian distortions of the source field
Appendix B:	Luminosity integration in COMBI
Appendix C:	Parameterization of optical-distortion corrections
References

1 Introduction

The determination of the absolute scale of the luminosity delivered to the ALICE, ATLAS, CMS and LHCb experiments at the LHC relies almost entirely [1] on the van der Meer (*vdM*) method [2,3]. Luminosity calibrations, which must be performed under specially tailored beam conditions, use beam-separation scans, also known as *vdM* scans, to relate the collision rate measured by a given luminometer to the absolute luminosity inferred from directly measured beam parameters. The proportionality between the measured rate and the absolute luminosity is expressed in terms of a “visible” cross-section denoted σ_{vis} , which is specific to the luminometer and the counting method considered.

The accuracy requirements on σ_{vis} are driven by the physics program. At the LHC, the comparison of the most precisely measured cross-sections with the corresponding theoretical predictions of the Standard model provide some of the most stringent tests of higher-order calculations; they also put strong constraints on parton-density distributions. In particular, the experimental uncertainty affecting the fiducial cross-sections for inclusive vector-boson production ($pp \rightarrow Z + X$, $pp \rightarrow W + X$) is totally dominated [4–6] by the systematic uncertainty in the integrated luminosity, that in these three publications ranges from 1.8 to 2.5%; reducing that uncertainty by a factor of three would make it comparable to the combination of all other experimental uncertainties, and significantly improve the sensitivity of the associated Standard-Model tests. For more experimentally challenging processes such as $t\bar{t}$ production, the overall measurement uncertainty remains significantly impacted by that on the luminosity [7,8], even with the recently achieved sub-percent uncertainty on the Run-2 integrated luminosity [9]. At the HL-LHC, the measurement of the absolute Higgs couplings drives the accuracy specifications. Projections of the experimental and theoretical uncertainties that should be achievable by then led to set a 1% goal on the overall integrated-luminosity uncertainty [10,11]. Since the latter is known to receive comparable contributions from *vdM* calibrations, from rate-related instrumental non-linearities, and

from long-term luminometer stability [9], the corresponding target for vdM uncertainties at the HL-LHC has to remain significantly smaller than this goal; it has been set to approximately 0.6% by both the ATLAS and the CMS [11] Collaboration.

The beam parameters used to determine the absolute luminosity during vdM scans are the particle populations and transverse sizes associated with each colliding-bunch pair. During a beam-separation scan, the mutual electromagnetic interaction between two opposing bunches induces separation-dependent orbit shifts, as well as variations in the transverse size and shape of each bunch, thereby distorting the luminosity-scan curves from which the beam-overlap integrals, or equivalently the convolved transverse beam sizes, are extracted. Depending on the beam conditions under which the scans are performed, the magnitude of the resulting calibration bias – if left uncorrected – represents a significant fraction of, or can even exceed, the σ_{vis} systematic-uncertainty budget.

The methodology first developed in Refs. [12, 13] to quantify the impact of beam–beam effects on the absolute luminosity scale has been adopted since 2013 by all LHC Collaborations. With as input the beam separation dialed-in at each scan step, plus the measured bunch currents and convolved beam sizes, the orbit shift induced by beam–beam deflections was calculated analytically [14], and the optical distortion associated with the dynamic- β effect was evaluated in the linear approximation [13] using the MAD-X package [15]. The two effects impacted the luminosity scale in opposite ways, resulting at the time in a net upwards correction to σ_{vis} of 1–1.5% for vdM calibrations at $\sqrt{s} = 13$ TeV.

A recent reevaluation of this methodology, using a new beam–beam simulation package specifically developed for vdM -scan studies, revealed that the linear approximation used in the MAD-X simulation results in a significant underestimate of the optical-distortion correction [16].

These findings motivated the studies reported in the present paper, which is organized as follows.

After a brief overview of the vdM -calibration methodology at the LHC and of the impact thereon of beam–beam effects (Sect. 2), the relevant simulation tools are presented in Sect. 3: the MAD-X package used in the original implementation [12, 13]; the B*B package presented in Ref. [16], that models beam–beam dynamics in the transverse plane under the weak-strong approximation; and the long-established COMBI [17] multiparticle code, that is more CPU-intensive but can simulate beam–beam effects in the strong-strong regime with the optional inclusion of longitudinal dynamics. Following a systematic cross-validation of the latter two packages, B*B is used in Sect. 4 to develop an easy-to-use parameterization of beam–beam corrections to vdM scans in the limit of round, initially Gaussian bunches of equal bright-

ness¹ that collide at a single interaction point (IP) with zero crossing angle. Deviations of the colliding-bunch configuration from this idealized limit: non-Gaussian tails, elliptical transverse bunch profiles, non-zero crossing angle, collisions at multiple IPs, or unequal-brightness beams, are then either accounted for using simulation-based adjustments to the idealized parameterization, or found to be small enough to be treated as a systematic uncertainty. These and other sources of systematic uncertainty, such as tune or β^* settings, that may affect the beam–beam correction in the vdM regime are consolidated in Sect. 5. An overall summary and a brief outlook are offered in Sect. 6.

2 Luminosity-calibration methodology at the LHC

The vdM -scan formalism [2, 3] that underpins the determination of the absolute luminosity scale at the CERN ISR, RHIC and the LHC is summarized, for the simplest case, in Sect. 2.1 below; generalizations of this formalism can be found in Refs. [16, 19, 20]. At the LHC, because of both instrumental and accelerator-physics reasons [1], vdM scans are not performed during normal physics operation, but rather under dedicated beam conditions (Sect. 2.2). Their two fundamental ingredients, the transverse beam separation and the measured collision rate, are both affected by the beam–beam interaction (Sect. 2.3), at a level that is significant on the scale of the precision goals outlined in Sect. 1.

2.1 Absolute luminosity scale from measured beam parameters

In terms of colliding-beam parameters, the bunch luminosity \mathcal{L}_b is given by

$$\mathcal{L}_b = f_r n_1 n_2 \int \hat{\rho}_1(x, y) \hat{\rho}_2(x, y) dx dy, \quad (1)$$

where the opposing bunches are assumed to collide head-on and with zero crossing angle, f_r is the LHC revolution frequency, $n_1 n_2$ is the bunch-population product and $\hat{\rho}_{1(2)}(x, y)$ is the normalized particle density in the transverse (x - y) plane of beam 1 (2) at the IP.² With the standard assumption that the particle densities can be factorized into

¹ There exist multiple definitions of the concept of *beam brightness*, depending on the type of beam and the application considered [18]. In the context of the present paper, the term *brightness* refers to the particle density in transverse phase space. For round beams, it is defined as $\mathcal{B} = n_p / \epsilon_N$, where n_p is the bunch population and ϵ_N the normalized transverse emittance.

² This paper uses a right-handed coordinate system with its origin at the nominal IP, and the z -axis along the direction of LHC beam 1; the latter circulates in the clockwise direction when the LHC rings are viewed from above. The x -axis points from the center of the LHC ring to the IP, and the y -axis points upwards.

independent horizontal and vertical component distributions, $\hat{\rho}(x, y) = \rho_x(x) \rho_y(y)$, Eq. (1) can be rewritten as

$$\mathcal{L}_b = f_r n_1 n_2 \Omega_x(\rho_{x1}, \rho_{x2}) \Omega_y(\rho_{y1}, \rho_{y2}), \tag{2}$$

where

$$\Omega_x(\rho_{x1}, \rho_{x2}) = \int \rho_{x1}(x) \rho_{x2}(x) dx$$

is the beam-overlap integral in the x direction (with an analogous definition in the y direction). In the method proposed by van der Meer [2] at the ISR (Fig. 1, top), the overlap integral (for example in the y direction) can be calculated as

$$\Omega_y(\rho_{y1}, \rho_{y2}) = \frac{R_y(0)}{\int R_y(\delta_y) d\delta_y}, \tag{3}$$

where $R_y(\delta_y)$ is the collision rate, or equivalently the luminosity in arbitrary units, measured during a vertical scan at the time the two beams are separated vertically by the distance δ_y . Because the collision rate $R_y(\delta_y)$ is normalized to that at zero separation $R_y(0)$, any quantity proportional to the luminosity can be substituted in Eq. (3) in place of R .

Defining the vertical convolved bunch size Σ_y [1] as

$$\Sigma_y = \frac{1}{\sqrt{2\pi}} \frac{\int R_y(\delta_y) d\delta_y}{R_y(0)}, \tag{4}$$

and similarly for Σ_x , the bunch luminosity in Eq. (2) can be rewritten as

$$\mathcal{L}_b = \frac{f_r n_1 n_2}{2\pi \Sigma_x \Sigma_y}, \tag{5}$$

which allows the absolute bunch luminosity at zero separation to be determined from the revolution frequency f_r , the bunch-population product $n_1 n_2$, and the product $\Sigma_x \Sigma_y$ which is measured directly during a pair of orthogonal scans.

If the transverse density profile of each beam B ($B = 1, 2$) can be described by a single Gaussian of width σ_{iB} ($i = x, y$), the convolved widths are given by

$$\Sigma_i = \sqrt{\sigma_{i1}^2 + \sigma_{i2}^2} = \sqrt{\beta_{i1}^* \epsilon_{i1} + \beta_{i2}^* \epsilon_{i2}} \tag{6}$$

where ϵ_{iB} is the geometrical emittance of beam B in plane i , and β_{iB}^* is the corresponding value of the β function at the IP.³ In such a case, the beam-separation dependence of the collision rate is given by

$$R_i(\delta_i) = R_i(0) e^{-\frac{\delta_i^2}{2\Sigma_i^2}}. \tag{7}$$

The luminosity curve $R_i(\delta_i)$ is also Gaussian, and Σ_i coincides with the standard deviation of that distribution. It is important to note, however, that the vdM method does not

³ Throughout most of the present paper, and unless explicitly specified otherwise, the IP β -function is implicitly assumed to be the same for the two beams and in the two planes: $\beta_{x1}^* = \beta_{y1}^* = \beta_{x2}^* = \beta_{y2}^* = \beta^*$.





Fig. 1 Top: van der Meer method at the CERN ISR. Shown is the monitor rate $R_y(\delta_y)$ as a function of the relative vertical separation δ_y of the two beams (Figure reproduced from Ref. [21], © CERN). Bottom: beam-separation dependence of the visible interaction rate measured at the ATLAS IP during a horizontal vdM scan in pp collisions at $\sqrt{s} = 8$ TeV (LHC fill 3311), before (red circles) and after (purple squares) noise and background subtraction. The subtracted contributions are shown as triangles. The scan curve is fitted to a Gaussian function multiplied by a sixth-order polynomial, plus a constant. (Figure reproduced from Ref. [22], © CERN, CC-BY-4.0 license)

rely on any particular functional form of $R_i(\delta_i)$: the quantities Σ_x and Σ_y can be determined for any observed luminosity curve from Eq. (4) and used with Eq. (5) to determine the absolute luminosity at $\delta_x = \delta_y = 0$.

In the more general case where the factorization assumption breaks down, i.e. when the particle densities cannot be factorized into a product of uncorrelated x and y components,

Eq. (2) no longer holds, and a single pair of horizontal and vertical scans is no longer sufficient to measure the overlap integral in Eq. (1). One must then generalize the formalism to the two-dimensional case [3], and scan over a grid in the (δ_x, δ_y) beam-separation space to measure the product of the convolved bunch widths [1, 16]:

$$[\Sigma_x \Sigma_y] = \frac{1}{2\pi} \frac{\int R_{x,y}(\delta_x, \delta_y) d\delta_x d\delta_y}{R_{x,y}(0, 0)}. \tag{8}$$

Here the square brackets highlight the fact that in the presence of non-factorization, the quantity $[\Sigma_x \Sigma_y]$ can no longer be broken down into a product of two independent quantities. Equation (5), however, remains formally unaffected, as do Eqs. (9)–(10) below.

In terms of luminometer observables, the bunch luminosity can be written as

$$\mathcal{L}_b = \frac{\mu_{\text{vis}} f_{\text{r}}}{\sigma_{\text{vis}}}, \tag{9}$$

where μ_{vis} is the average number of inelastic collisions per bunch crossing detected by the luminometer considered, and σ_{vis} is the associated visible cross-section. Since μ_{vis} is a directly measurable quantity, the calibration of the absolute luminosity scale amounts to determining the visible cross-section σ_{vis} . Equating the absolute luminosity computed from beam parameters using Eq. (5) to that measured according to Eq. (9), yields:

$$\sigma_{\text{vis}} = \mu_{\text{vis,pk}} \frac{2\pi \Sigma_x \Sigma_y}{n_1 n_2}, \tag{10}$$

where $\mu_{\text{vis,pk}}$ is the visible interaction rate per bunch crossing reported at the peak of the scan curve (Fig. 1, bottom). Equation (10) provides a direct calibration of the visible cross-section σ_{vis} in terms of the peak visible interaction rate $\mu_{\text{vis,pk}}$, the product of the convolved bunch widths $\Sigma_x \Sigma_y$, and the bunch-population product $n_1 n_2$.

In the presence of a significant crossing angle, the formalism becomes more involved [16, 19, 23]. A non-zero crossing angle in either the horizontal or the vertical plane widens the corresponding luminosity-scan curve by the so-called geometrical factor F :

$$F = \frac{\sqrt{1 + \tan^2 \theta_c/2 (\sigma_{z1}^2 + \sigma_{z2}^2)/(\sigma_{c1}^2 + \sigma_{c2}^2)}}{\cos \theta_c/2}. \tag{11}$$

Here θ_c is the full crossing angle, σ_{zB} ($B = 1, 2$) are the RMS bunch lengths of beams 1 and 2, and σ_{cB} the transverse single-beam sizes in the crossing plane. The peak luminosity is reduced by the same factor. The corresponding increase in the measured value of Σ_x or Σ_y is exactly compensated by the decrease in $\mu_{\text{vis,pk}}$, so that Eqs. (4)–(10) remain valid, and no correction for the crossing angle is needed in the determination of σ_{vis} .

2.2 Beam conditions during van der Meer scans

The strength of the beam–beam interaction is traditionally quantified by the *linear beam–beam parameter*, defined as [18, 24]:

$$\xi_{x2} = \frac{n_1 r_0 Z_1 Z_2 \beta_{x2}^*}{2\pi A_{\text{ion},2} \gamma_2 \sigma_{x1} (\sigma_{x1} + \sigma_{y1})} \tag{12}$$

Here ξ_{x2} is the horizontal beam–beam parameter experienced by beam 2 (B2), the “witness beam”; n_1 is the bunch population of beam 1 (B1), the “source beam”; $r_0 = e^2/4\pi\epsilon_0 m_p c^2$ is the classical radius of the proton; $A_{\text{ion},B}$ and Z_B ($B = 1, 2$) are the atomic mass number and charge number of the beam- B particle type (proton or fully stripped ion); β_{x2}^* is the value of the B2 horizontal β function at the IP; γ_2 is the relativistic factor of the B2 particles, and σ_{x1} (σ_{y1}) is the horizontal (vertical) transverse RMS size of B1. Formulas for the other beam and the other plane are obtained by interchanging B1 and B2, and/or x and y .

For the most frequent case of pp collisions, Eq. (12) takes the more familiar form:

$$\xi_{x2} = \frac{n_1 r_0 \beta_{x2}^*}{2\pi \gamma_2 \sigma_{x1} (\sigma_{x1} + \sigma_{y1})}.$$

In the case of equally populated, equally sized round beams, this expression becomes much simpler:

$$\xi = \frac{n r_0 \beta^*}{4\pi \gamma \sigma_0^2} \tag{13}$$

where n is the bunch population, and $\sigma_0 = \sqrt{\epsilon \beta^*}$ is the nominal RMS beam size at the IP.

In 2018, during high-luminosity physics running in proton-proton (pp) mode, the LHC collided up to 2544 bunches with typical initial intensities of 1.1×10^{11} p /bunch, grouped in trains of 36 to 144 bunches with a minimum interbunch spacing of 25 ns. Beams crossed with a half-angle $\theta_c/2$ of $\pm 130 \mu\text{rad}$ in order to mitigate the impact of the long-range beam–beam interaction at parasitic crossings. At the start of stable beams, the emittance was typically $2 \mu\text{m}\cdot\text{rad}$ [25], the single-bunch luminosity around $7.8 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ at IP1 and IP5, and the total luminosity close to $1.9 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at each of these two IPs. These values correspond to a pile-up parameter μ of around 55 inelastic pp collisions per bunch crossing, and to a bunch-averaged, head-on beam–beam parameter $\langle \xi \rangle$ of approximately 0.005. The brightness, however, varied significantly along the bunch string, occasionally resulting in ξ values as high as 0.007 for some of the bunches.

In contrast, during pp vdM scans (Table 1), the injected emittance is deliberately blown up and the bunch population significantly lowered, in order to reduce the impact of beam–beam effects as well as minimize the unbunched-beam fraction and the intensity of satellite bunches. The bunches

are isolated rather than in trains, and their number is limited to 152 at most, in order to eliminate parasitic crossings and collide with zero crossing angle in the interaction regions where the beam-line layout so permits. The β function at the IP is increased such as to bring the pile-up parameter μ down to around 0.5, i.e. in a regime where luminometers are free of instrumental non-linearities; this carries the additional advantage that it significantly increases the transverse luminous size, allowing a more precise measurement of its beam-separation dependence [1,22].

2.3 Beam–beam-induced biases and their correction

The mutual electromagnetic interaction between colliding bunches shifts their orbits, and therefore modifies their transverse separation at the IP (Sect. 2.3.1); it also distorts their transverse density distributions (Sect. 2.3.2). These two effects depend on the *nominal separation* Δ dialed-in at each scan step. Their combination impacts both the normalized integrals (Σ_x, Σ_y) and the peak ($\mu_{\text{vis,pk}}$) of the luminosity-scan curves used in determining the absolute luminosity scale. The strategy for correcting the resulting biases is outlined in Sect. 2.3.3; its detailed implementation is developed in later chapters, in particular in Sects. 4.2.3 and 4.6.5.

2.3.1 Orbit shift

When two positively charged bunches collide with a non-zero impact parameter, they experience a mutually repulsive angular kick equivalent to that of a dipole located at the collision point, the strength of which depends on the beam separation. In the round-beam limit and for pp collisions, the angular kick experienced by a B2 bunch during a horizontal beam-separation scan is given by [26]:

$$\theta_{2x} = \frac{2r_0 n_1}{\gamma_2} \frac{\delta_x}{\delta} \left[\frac{1 - e^{-\delta^2/2\Sigma_R^2}}{\delta} \right], \tag{14}$$

and similarly for a vertical scan. Here δ_x (resp. δ) is the horizontal (resp. total) beam separation, and $\Sigma_R = \Sigma_x = \Sigma_y$ is the transverse convolved beam size. This formalism has been extended by Bassetti and Erskine [14] and by Ziemann [27] to the case of elliptical beams.

This angular deflection, first observed at the Stanford Linear Collider in e^+e^- collisions [26], can be measured using beam-position monitors (BPMs) installed both upstream and downstream of the IP, as illustrated in Fig. 2 for the LHC [28].

In a circular collider, the beam–beam angular kick experienced by each beam B ($B = 1, 2$) in the i -plane ($i = x, y$) results in a shift of its position at the IP, given by [29]

$$\delta_{iB}^{bb} = \frac{\beta_{iB}^*}{2 \tan(\pi Q_i)} \theta_{iB}, \tag{15}$$




Fig. 2 Total beam–beam deflection angle (B1–B2) as a function of the nominal separation Δ (denoted here by “knob setting”), during a vertical vdM scan at the ATLAS IP in pp collisions at $\sqrt{s} = 8\text{TeV}$ (LHC fill 3316). For each beam separately, the deflection angle is obtained from the difference between the outgoing- and incoming-beam angles measured by BPMs located in the LHC arcs outside the closed-orbit bump used for the scans. The zero of the horizontal axis is arbitrary, since only relative beam displacements matter. The vertical convolved beam size, denoted here by Σ , is extracted from a fit to Eq. (14), shown by the red curve, with as input the bunch-averaged population N . (Figure reproduced from Ref. [28], © CERN)

where β^* is the value of the β function at the IP and Q_i is the betatron tune. The actual beam separation at the IP δ_i therefore differs slightly from the nominal separation Δ_i :

$$\delta_i = \Delta_i + \delta_i^{bb}, \tag{16}$$

where the beam–beam-induced change in beam separation, hereafter denoted by “orbit shift”, is given by $\delta_i^{bb} = \delta_{i1}^{bb} - \delta_{i2}^{bb}$. Since δ_{iB}^{bb} is, to first order, proportional to the corresponding beam–beam parameter ξ_{iB} (Eqs. (12) and (14)), the orbit shift varies from one bunch to the next. The Δ -dependence of δ_i^{bb} mirrors that displayed in Fig. 2 for the deflection angle, with a peak-to-peak swing of $\pm(1-2)\mu\text{m}$ under typical vdM-scan conditions, to be compared to typical Σ_i values in the 110–160 μm range. The value of δ_i^{bb} changes from scan step to scan step, thereby expanding in a non-linear fashion the beam-separation scale, i.e. the horizontal axis of scan curves such as that illustrated in Fig. 1b. As a result, the overlap integral of Eq. (4) increases by typically 0.7–1.4% per plane, corresponding to a positive correction to σ_{vis} in the 1.4–2.8% range.

In practice, the correction for the beam–beam orbit shift is implemented as follows. At each scan step and for each colliding-bunch pair separately, δ_i^{bb} is calculated using the Bassetti–Erskine formula [14], with as input the measured bunch populations (n_1, n_2) and uncorrected convolved bunch widths (Σ_x, Σ_y), as well as the beam energy, the β^* setting and the tunes. The separation-dependent collision rate $R_i(\delta_i)$ is then integrated according to Eq. (4) to obtain the beam–beam corrected values $\Sigma_x^{c,\text{Orb}}$ and $\Sigma_y^{c,\text{Orb}}$, using the beam–




Fig. 3 Beam–beam force exerted by the B1 (B2) source bunch as a whole (red curve), as a function of the betatron amplitude of a B2 (B1) test particle, for equally sized round beams. The amplitude is in units of the RMS beam size. The green and brown lines correspond to the linear component of the force at amplitudes of zero and 4σ respectively, and are akin to the effect of a quadrupole. The short (long) double arrows illustrate the range of transverse kicks experienced by a small (large) amplitude particle for beams either in head-on collision (solid blue lines), or transversely separated by 4σ (dashed magenta lines)

beam corrected separation δ (Eq. (16)) instead of the nominal separation Δ . The agreement of this simple analytical procedure with the predictions of self-consistent multi-particle simulations will be addressed in Sect. 3.5.

2.3.2 Optical distortions

Not only does the electromagnetic field of the B1 bunch deflect the B2 bunch as a whole: it also acts as a non-linear

lens that perturbs the trajectory of the individual particles in that bunch, thereby modifying the transverse density distribution of both bunches in a separation-dependent manner.

For small-amplitude particles and beams in head-on collision (Fig. 3, short blue arrows), the force is rather linear and resembles that of a quadrupole (dotted green line), resulting in a tune shift proportional to the beam–beam parameter ξ and in the subsequent dynamic- β effect [13]. This “quadrupole strength” is proportional to the derivative of the beam–beam force; it is largest, and repulsive, for small-amplitude particles, changes sign around 1.6σ , and becomes weakly attractive at larger amplitude (dotted brown line).

If for simplicity one assumes that for a given beam separation, all particles are subject to the same quadrupolar-like force (the strength and sign of which depend on the beam separation), then the value of β^* at the scanning IP is modulated by the linear component of the beam–beam force. This results in a modulation of the transverse beam size, and therefore in a beam–separation-dependent modulation of the actual luminosity; however the actual shapes of the transverse density distributions projected on the x and y axes remain unaffected by the quadrupolar-like force. This is the approximation that was adopted in the first implementation of the optical-distortion correction [12, 13], and that will be further discussed in Sect. 3.1.

While for small amplitudes (short arrows) the force remains approximately linear, at amplitudes larger than 1σ (long arrows) it includes significant non-linear contributions. Large-amplitude particles, therefore, experience a tune shift and a β -beating that depend both on the particle amplitude (short vs. long arrows) [30], and on the beam separation (blue vs. magenta arrows). The resulting optical distortions include not only a change in optical magnification as in Ref. [13], but

Table 1 Beam conditions during vdM scans at the LHC. Middle column: typical parameters during the 2015–2018 pp scan sessions at $\sqrt{s} = 13$ TeV. Right column: reference parameters used for the cross-validation of the beam–beam codes described in Sect. 3; the parameter values are chosen to match those used in the initial, MAD-X based simulations of Ref. [13], which were representative of 2011 vdM scans at IP1 and IP5. In the absence of the beam–beam interaction, the simulated bunches are assumed round in the transverse plane, of equal radius, and strictly Gaussian in all three dimensions

Parameter	Typical scans (LHC Run 2)	Reference parameter set
Beam energy E_B [TeV]	6.5	3.5
Nominal tune settings Q_x/Q_y	64.31/59.32	64.31/59.32
Normalized emittance ϵ_N [$\mu\text{m}\cdot\text{rad}$]	2.2–3.5	4.00
IP1,5: β^* [m]	19.2	1.50
IP2/8: β^* [m]	19.2/24.0	–
IP1,5: $\theta_c/2$ [μrad]	0	0
IP2: $\theta_c/2$ [μrad]	70–195 (y)	–
IP8: $\theta_c/2$ [μrad]	450–550 (x)	–
Transverse convolved bunch sizes $\Sigma_x \approx \Sigma_y$ [μm]	110–160	56.7
Bunch spacing [ns]	≥ 525	–
Number of colliding bunches	30–124	1
Bunch population n_p [10^{11}]	0.7–1.0	0.85
Linear beam–beam parameter ξ [per IP]	0.0022–0.0056	0.0026
Typical bunch luminosity @ IP1, 5 [$\text{cm}^{-2}\text{s}^{-1}$]	$3\text{--}8 \times 10^{28}$	4×10^{29}
Typical total luminosity @ IP1, 5 [$\text{cm}^{-2}\text{s}^{-1}$]	$2.5\text{--}9 \times 10^{30}$	4×10^{29}

also distortions of the shape of the transverse density distributions. Describing their beam-separation dependence requires numerical simulations, that are detailed in Sect. 3.

In practice, the correction for optical distortions is implemented as follows. Separately for each scan step in a horizontal and vertical vdM -scan pair, the *luminosity-bias factor* $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with beam-beam-induced optical distortions is extracted, as a function of the nominal separation Δ_i , from one of the multiparticle simulations described in Sect. 3. Here \mathcal{L} refers to the luminosity that would be measured in the presence of beam-beam optical-distortion effects, and \mathcal{L}^0 is the corresponding luminosity if beam-beam effects were turned off altogether, all other conditions remaining unchanged. The quantity \mathcal{L}^0 is dubbed the *nominal luminosity*; it is akin to “Monte Carlo truth”, and is accessible in the simulation only. The beam-beam corrected collision rate $R_i^c(\delta_i)$ is then computed by dividing the measured collision rate by this simulation-based luminosity-bias factor:

$$R_i^c(\delta_i) = \frac{R_i(\Delta_i)}{[\mathcal{L}(\Delta_i)/\mathcal{L}^0(\Delta_i)]_{\text{Opt}}} \quad (17)$$

and used instead of $R_i(\delta_i)$ in computing the beam-beam corrected convolved bunch sizes $\Sigma_x^{c,\text{Opt}}$ and $\Sigma_y^{c,\text{Opt}}$ (Eq. (4)), the peak rate $\mu^c = \mu_{\text{vis,pk}} R_i^c(0)/R_i(0)$, and from these quantities the visible cross-section $\sigma_{\text{vis}}^{c,\text{Opt}}$ (Eq. (10)).

2.3.3 Beam-beam correction strategy

Conceptually, the principle of the beam-beam correction to vdM calibrations is to determine the visible cross-section σ_{vis}^c from the convolved bunch sizes and peak collision rates corrected both for the orbit shift (Sect. 2.3.1) and for optical distortions (Sect. 2.3.2), i.e. corrected to the values $(\Sigma_x^c, \Sigma_y^c, \mu^c)$ that these observables would take if the beam-beam interaction could be turned off during the scan. The quantity σ_{vis}^c is then the proportionality constant that translates a measured visible interaction rate μ_{vis} into the corresponding bunch luminosity \mathcal{L}_b . It is important to note that even though the actual luminosity is always modified by the beam-beam interaction, including during head-on collisions typical of routine physics running, beam-beam *corrections* are only needed *during scans*, basically because the beam-separation dependence of the beam-beam effects distorts the vdM -scan curves. Once σ_{vis}^c has been determined as specified above, it can always be used to translate the measured collision rate into luminosity units, irrespective of the extent to which this collision rate has been enhanced by the beam-beam interaction.

In practice, orbit-shift and optical-distortion corrections must be applied on a bunch-by-bunch basis, with as inputs the measured bunch populations (n_1, n_2) and uncorrected convolved bunch widths (Σ_x, Σ_y) , as well as the beam energy,

the β^* setting and the tunes. In many cases, these corrections can be extracted from a one-time parameterization of the simulation results in terms of the bunch-specific beam-beam parameter value and of the nominal tunes Q_x, Q_y . This procedure avoids the repeated use of CPU-intensive, time-consuming multiparticle simulations; it is detailed in Sect. 4.

3 Beam-beam simulation codes

Beam-beam corrections to vdM calibrations were originally based on MAD-X (Sect. 3.1). Since then, only the zero-separation case has proven amenable to analytical treatment (Sect. 3.2). Beam-separation-dependent effects have been investigated using two independent multiparticle codes dubbed B*B (Sect. 3.3) and COMBI (Sect. 3.4), that have been extensively cross-validated (Sect. 3.5).

3.1 Linear approximation with MAD-X

Since the linear part of the beam-beam force is similar to a quadrupolar field (Fig. 3), one expects the beam-beam interaction to contribute additional focusing or defocusing, thereby shifting the tunes and affecting the optical functions all around the ring, including at the IP itself: this is known as the “dynamic- β ” effect. In the specialized case of equally sized round beams colliding head-on at a single location, the resulting change in the β function at the IP is given by [13]:

$$\frac{\beta^*}{\beta_0^*} = \frac{1}{\sqrt{1 - 4\pi\xi \cot(2\pi Q) - 4\pi^2\xi^2}} \quad (18)$$

where β_0^* is the value of the unperturbed β function at the IP, β^* its value in the presence of the beam-beam interaction and Q the tune. The beam-beam parameter ξ is proportional to the derivative of the beam-beam force.⁴ Equation (18) implies that the beam-beam induced change in the IP β function, and therefore in the IP beam-size squared and in the luminosity, depends only on ξ and on Q . In addition, this dynamic- β effect is, to first order, proportional to ξ . This is the physical motivation underlying the parameterized-correction approach developed in Sect. 4.

In Ref. [13], the general-purpose optics code MAD-X [15] was used to model the dynamic- β effect during simulated vdM scans. In this software package, beam-beam elements can be inserted at one or several IPs, and their impact on the tunes and on the single-particle optical functions computed

⁴ The sign of the linear term in the denominator of Eq. (18) is flipped compared to that in Eq. (3) of Ref. [13]. This is because the latter follows the convention that ξ is negative (positive) for equal- (opposite-) charge beams. In the present paper, in contrast, ξ is positive by definition (Eq. (12)), and therefore Eq. (18) is adjusted so as to be applicable to equal-charge beams.

as a function of the beam separation Δ . The procedure effectively assumes that for a given beam separation, all particles in the bunch experience the same beam–beam kick, equal to that applied to a zero-amplitude particle for that particular value of Δ . Unperturbed bunches are implicitly supposed to be strictly Gaussian, and to remain so in the presence of the beam–beam interaction: only the change in optical magnification between the LHC arcs and the IP is accounted for in this method.

This study was carried out for the reference parameter set listed in Table 1. The corrections were then adapted to the beam conditions of different *vdM* sessions, using the assumption that $\Delta\beta/\beta_0^* = \beta^*/\beta_0^* - 1$ scales linearly with the value of ξ inferred from the beam parameters measured during each scan. The simulated beam-separation dependence of the transverse beam size squared, i.e. the value of $(\mathcal{L}/\mathcal{L}^0)^{-1}$, is illustrated in Fig. 4. Under head-on conditions ($\Delta = 0$), the intrinsically *defocusing* beam–beam force results in a $\sim 0.7\%$ reduction of the beam size squared at the IP, i.e. to an *increase* of the luminosity \mathcal{L} . This apparent contradiction results from the numerical value of the fractional tunes the LHC optics was designed for (Eq. (18)). As Δ increases, the tune shift and the dynamic- β effect weaken, change sign (in the scanning plane only) around $\Delta/\sigma^0 \sim 1.6$, then peak and finally vanish asymptotically at very large separation. The change in optical magnification at $\Delta = 0$ is slightly different in the x and y planes, because the corresponding fractional tunes q_x and q_y differ by design; the difference in beam-separation dependence also reflects the fact that the vertical kick never changes sign during a horizontal scan, while the horizontal kick does.

Historically, the optical-distortion correction to σ_{vis} predicted by MAD-X under typical Run-1 and Run-2 *vdM* conditions lay in the 0.2–0.4% range, much smaller than that associated with the orbit effect. For a long time, therefore, it was considered small enough that even if imperfect, it remained sufficiently accurate in view of the systematic uncertainty assigned at the time to the overall beam–beam correction, as documented e.g. in Refs. [12,22].

In hindsight, however, the limitations of applying the MAD-X approach to beam–beam corrections may not have been fully appreciated. Since at zero beam separation, the slope of the beam–beam force is steepest for zero-amplitude particles (Fig. 3, dotted green line), and since in MAD-X all particles in the witness bunch are assumed to experience the same linearized force, the dynamic- β effect predicted by MAD-X at $\Delta = 0$ is likely to be an overestimate; this is confirmed analytically below. The situation is reversed at large beam separation. When the beams are separated by (for instance) $\Delta = 4\sigma_0$, the derivative of the force at large amplitude (say 3 to 4 σ_0) is typically larger (dotted green line) than at small amplitude (dotted brown line), suggesting that MAD-X underestimates the optical distortions at large beam




Fig. 4 Beam-separation dependence of the transverse RMS beam sizes during a simulated horizontal *vdM* scan. The unperturbed beams are assumed to be round and perfectly Gaussian; their parameters are listed in the right column of Table 1. The vertical axis is the ratio squared of the actual beam size σ to its unperturbed value σ_0 ; the horizontal axis is the nominal beam separation in units of σ_0 . The MAD-X calculation of the single-particle dynamic- β effect (blue curves) is compared to the results of the B*B (square markers) and COMBI (red curves) multiparticle codes. The dark- and light-colored curves and markers illustrate the evolution, during the horizontal scan, of the horizontal and the vertical beam size respectively. The green circular and triangular markers display the predictions of Eq. (22), that only apply at zero beam separation

separation. This conjecture, however, can only be confirmed using multiparticle simulations.

3.2 Analytical estimate of optical distortions at zero beam separation

The particle phase space at any location s along a storage ring is described by an ellipse, the shape of which depends on s in a manner described by the well known Courant-Snyder parameters α , β and γ [18]:

$$J_u = \frac{1}{2}(\gamma_u u^2 + 2\alpha_u u p_u + \beta_u p_u^2).$$

Here $u = x, y$ is the deviation of the single-particle orbit from the reference trajectory, $p_u = du/ds$, and J_u is the action variable that represents the invariant of the motion for each single particle when the reference energy is not changing. The particle position in transverse phase space is fully described by the action J_u and by the corresponding phase variable ϕ_u defined as:

$$\tan \phi_u = -\beta_u \frac{p_u}{u} - \alpha_u. \tag{19}$$

In addition to distorting the closed orbit, the beam–beam interaction acts as a non-linear electromagnetic lens that includes a quadrupolar term; the latter affects the optical functions β and the betatron tunes of the individual par-

ticles [30,31]. During a νdM scan and for each colliding bunch, the quadrupolar component of the electromagnetic field of the opposing bunch changes as a function of the relative transverse separations δ_i ($i = x, y$) between the beams centroids, and of the single-particle transverse actions J_x and J_y [31,32]. This results in a shift in the betatron tune, the so-called detuning with amplitude, given by [33]:

$$\begin{aligned} \Delta Q_x(J_x, J_y) = & -\xi_x \int_0^\infty \frac{1}{(1+t)^2} e^{-\frac{J_x+J_y}{2\epsilon(1+t)}} \\ & \times \left[I_0\left(\frac{J_x}{2\epsilon(1+t)}\right) - I_1\left(\frac{J_x}{2\epsilon(1+t)}\right) \right] \\ & \times I_0\left(\frac{J_y}{2\epsilon(1+t)}\right) dt \end{aligned}$$

for the horizontal dimension. Here ξ_x is the beam-beam parameter, I_0 and I_1 are modified Bessel functions of the first kind, ϵ is the emittance (assumed equal in x and y), and t is a bound variable. This integral can only be solved when one action is zero, obtaining (for $J_y = 0$):

$$\begin{aligned} \Delta Q_x(J_x) = & -\xi_x \frac{2\epsilon}{J_x} \left[1 - I_0\left(\frac{J_x}{2\epsilon}\right) e^{-\frac{J_x}{2\epsilon}} \right] \\ \Delta Q_y(J_x) = & -\xi_x \left[I_0\left(\frac{J_x}{2\epsilon}\right) + I_1\left(\frac{J_x}{2\epsilon}\right) \right] e^{-\frac{J_x}{2\epsilon}}. \end{aligned} \tag{20}$$

For a zero-amplitude particle ($J_x = J_y = 0$) and beams in head-on collision ($\delta_x = \delta_y = 0$), the change in tune is equal to the beam-beam parameter ξ_x .

Head-on beam-beam β -beating can be derived from the tune shift⁵ in Eq. (20) as:

$$\frac{\Delta\beta_x}{\beta_x}(J_x, J_y) = -\frac{\Delta Q_x(J_x, J_y)}{\xi_x} \cdot \frac{\Delta\beta_0}{\beta_x} \tag{21}$$

where

$$\frac{\Delta\beta_0}{\beta_x} = \frac{\Delta\beta_x}{\beta_x}(0, 0)$$

is the linear β -beating or, equivalently, the β -beating at $J_x = J_y = 0$.

In the zero-separation case, the β -beating averaged over the particle distribution (assumed Gaussian) can be calculated as:

$$\begin{aligned} \frac{\overline{\Delta\beta}}{\beta_x} = & \int_0^\infty dJ_x \int_0^\infty dJ_y \frac{\Delta\beta_x}{\beta_x}(J_x, J_y) e^{-J_x-J_y} \\ \approx & 0.633 \frac{\Delta\beta_0}{\beta_x}. \end{aligned}$$

For bunches colliding with zero transverse separation, therefore, the particle-action distribution is modified such that the

⁵ The overall minus sign in Eqs. (20) and (21) expresses the fact that when the charges of the colliding bunches have the same sign, the tune shift is negative, even though the beam-beam parameter remains positive by definition (Eq. (12)).

average beam-beam beating is reduced to about 63% of the single-particle estimate computed in Ref. [13]. The corresponding impact on the head-on luminosity can be derived analytically, as follows.

The RMS single-beam size $\sigma_x = \sqrt{\beta_x \epsilon}$, including the action-dependent, beam-beam-induced β -beating given by Eq. (21), is computed using the following relation:

$$\begin{aligned} x^2 = & 2\beta_x J_x \left(1 + \frac{\Delta\beta_x}{\beta_x} \right) \cos^2 \phi_x \\ = & \beta_x J_x \left(1 + \frac{\Delta\beta_x}{\beta_x} \right) (1 + \cos 2\phi_x). \end{aligned}$$

Assuming Gaussian particle-density distributions, this yields

$$\begin{aligned} \sigma_x^2 = & \int_0^\infty dJ_x \int_0^\infty dJ_y \int_0^{2\pi} \frac{d\phi}{2\pi} x^2 e^{-J_x-J_y} \\ = & \beta_x \left(1 + \int_0^\infty dJ_x \int_0^\infty dJ_y e^{-J_x-J_y} J_x \frac{\Delta\beta_x}{\beta_x}(J_x, J_y) \right) \end{aligned}$$

where, for simplicity, we used an emittance value of $\epsilon = 1$. The triple integral has an exact solution:

$$\sigma_x^2 = \beta_x \left(1 + \frac{1}{2} \frac{\Delta\beta_0}{\beta_x} \right). \tag{22}$$

This calculation implies that in head-on collisions, beam-beam-induced linear β -beating of magnitude $\Delta\beta_0/\beta$ in both the horizontal and the vertical plane, causes a relative luminosity change of

$$\frac{\Delta L}{L} \approx -\frac{1}{2} \frac{\Delta\beta_0}{\beta}. \tag{23}$$

For collisions with zero transverse separation, in other words, the beam-beam interaction changes the luminosity by only half of what is expected from linear β -beating: this is consistent with the COMBI predictions displayed in Fig. 4.

Using the reference parameter set in Table 1, for instance, the beam-beam-induced linear β -beating amounts to about -0.6% , leading to a head-on luminosity enhancement of about 0.3% . During high-luminosity physics running typical of LHC Run 2, the effect is computed to be two to three times larger; it is both ξ - and tune-dependent (see Eq. (18)), and therefore its magnitude changes as beam conditions evolve.

When beams are transversely separated, analytical computations become impossible, forcing one to resort to numerical simulations such as those described below.

3.3 Weak-strong limit: B*B

The B*B package [16] is a multiparticle-simulation program developed specifically for assessing beam-beam biases in νdM scans, that was optimized for speed by adopting several simplifying assumptions. It aims at predicting the corresponding beam-beam corrections with better than 0.1%

accuracy (for a given set of input parameters), so as not to contribute significantly to the overall vdM -calibration uncertainty. The code is written in C++; it can be used as a standalone application, or as a library available in C, C++, Python and R.

The initial transverse particle density distributions $\hat{\rho}_B(x, y)$, where $B = s, w$ refers to either the source (s) or the witness (w) bunch, are modeled by a (linear combination of) two-dimensional elliptical Gaussian(s). Since during typical vdM scans at LHC, beam–beam-induced bunch-shape deformations remain small, they are taken as negligible when computing the electromagnetic field of the source bunch, which therefore remains unperturbed as a function of the transverse beam separation. This makes it possible to precompute this field over a two-dimensional grid in the transverse plane at initialization time, and to use only fast interpolations on all subsequent machine turns. The opposing, “witness” bunch is represented by a set of $\mathcal{O}(1000)$ macroparticles, that are transported around the ring using linear maps, with as input the nominal (x, y) phase advance between consecutive collision points. The B*B simulation, therefore, falls in the “weak-strong” category in that it models the transverse deformation of the density distribution of the witness bunch ($\hat{\rho}_w \rightarrow \hat{\rho}_w + \delta\hat{\rho}_w$) caused by the electromagnetic field of an unperturbed source bunch, the density distribution of which remains unaffected. Effects such as coherent bunch oscillations, or the distortion of the shape (and therefore of the field) of the source bunch induced by the deformation of the witness bunch, are therefore implicitly neglected.

The macroparticles are selected from a two-dimensional (x, y) grid of betatron amplitudes, and assigned weights that are precalculated from the initial Gaussian density distribution $\hat{\rho}_w$. Their initial betatron phase is chosen randomly; the phase then samples the full $[0, 2\pi]$ interval over the next $\mathcal{O}(100-1000)$ turns or so.

The overlap integral of the perturbed and unperturbed bunches (Eq. (1)) is computed as the sum over macroparticles, of the source-bunch density $\hat{\rho}_s(x_w, y_w)$ evaluated at the current location (x_w, y_w) of the macroparticle considered, and multiplied by the weight of that same macroparticle. If the populations and initial transverse-density distributions of the two colliding bunches are identical ($n_1 = n_2$ and $\hat{\rho}_1(x, y) = \hat{\rho}_2(x, y)$), the simulation needs to be run once only; in the presence of any initial B1–B2 asymmetry, it needs to be run twice, with the roles of the source and the witness bunch swapped between the two beams. The overall beam–beam bias affecting the overlap integral is approximated by

$$\begin{aligned} & \int (\hat{\rho}_1 + \delta\hat{\rho}_1) (\hat{\rho}_2 + \delta\hat{\rho}_2) dx dy - \int \hat{\rho}_1 \hat{\rho}_2 dx dy \\ & \approx \int (\delta\hat{\rho}_1 \cdot \hat{\rho}_2 + \hat{\rho}_1 \cdot \delta\hat{\rho}_2) dx dy, \end{aligned}$$

up to some physical constants and where the second-order term $\int \delta\hat{\rho}_1 \delta\hat{\rho}_2 dx dy$ is neglected.

The uncertainties arising from the manner in which the beam–beam force is switched on in the calculation (gradually or instantaneously), from the granularity of the simulation (finite number of macroparticles, largest sampled transverse amplitude, random choice of the initial phases), and from other simulation-control parameters such as the number of accelerator turns, are discussed in detail in Ref. [16] and found to lie well below 0.1%.

Finally, even though the density distributions $\hat{\rho}_{s,w}$ are only two-dimensional, and therefore represent the projection of the full six-dimensional distribution onto a plane perpendicular to the beam axis, the B*B package is capable of simulating the *geometrical* effects that arise in the presence of a non-zero crossing angle in a plane of arbitrary orientation [16]. Simulating *longitudinal dynamics*, however, and in particular the potential impact of a finite crossing angle on beam–beam corrections to the vdM calibration, requires a fully six-dimensional treatment such as that outlined below.

3.4 Strong-strong model: COMBI

The COMBI code [17] has been developed over the years for simulating, in a self-consistent manner, the coherent beam–beam interaction between multiple bunches coupled by head-on and/or long-range beam–beam encounters [34–36]. It includes a first level of parallelization based on the Message Passing Interface (MPI) [37], and a second one sharing several CPUs per node using OpenMPI [38,39].

The code has been optimized to handle simultaneously multiple bunches at several interaction points, thereby allowing flexible collision patterns. The circumference of each accelerator ring is modeled by a number of equally spaced slots that define the possible bunch positions. At each location one can assign an action (e.g. head-on or long-range beam–beam interaction, linear or non-linear magnetic element, Landau octupole, linear or non-linear map, etc.) that will be executed when a bunch is present. The actions corresponding to head-on or long-range beam–beam interactions require one bunch from each beam in order to be performed. Each macroparticle is tracked individually under the effect of the preassigned actions, in either four or six dimensions. In the LHC arcs, the macroparticles are transported by applying a linear transfer map to their coordinates, using phase advances precomputed by MADX with, as input, the nominal optical configuration used during vdM sessions.

In B*B, the transverse density distribution of the source bunch, and therefore its electromagnetic field, remain unchanged from turn to turn; only the witness bunch contains macroparticles, that are tracked over thousands of turns in the transverse plane. Coherent effects, therefore, cannot be modeled, and neither can longitudinal dynamics. COMBI, in

contrast, describes the two partners in a colliding-bunch pair as independent sets of macroparticles. In this paper, the initial, unperturbed density distributions are uncoupled single Gaussians by default; however, arbitrary distributions, such as a linear combination of Gaussians, can be used instead.

The beam–beam interaction can be described by different models: a four-dimensional Gaussian lens [35,36], a six-dimensional Gaussian lens [40,41], or a field computed from the actual charge distributions by the HFMM method [42]. Multiple beam–beam encounters, and therefore the evolution of bunch parameters such as emittance or transverse barycenter position, as well as coherent beam–beam effects, are treated in a self-consistent manner: the particle trajectories affected by the beam–beam interaction modify the density distribution of the corresponding bunch, and the fields produced by both partners are updated turn by turn from these modified distributions. Whether these fields are estimated in the Gaussian approximation, or by the HFMM method, yields effectively identical results for the full range of beam–beam parameter values considered in this paper (see Appendix A).

At a given IP, the luminosity per colliding-bunch pair can be computed either analytically, using the Gaussian formalism of Ref. [23], or by evaluating numerically the actual overlap integral of the two colliding distributions. The first method evaluates the luminosity from Eqs. (5)–(7) by substituting the single-beam sizes σ_{iB} with the RMS transverse widths of the macroparticle distributions, and the separation δ with the distance between the barycenters of these distributions. In the second method, the overlap integral of the macroparticle distributions is computed using functionalities developed specifically for this purpose. This is the approach adopted throughout this paper, except where explicitly specified otherwise; the numerical-integration procedure and the associated convergence studies are detailed in Appendix B.

3.5 Cross-validation of simulation codes

The full impact of the beam–beam interaction on the beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}(\Delta)$ can be broken down as follows.

The impact of the orbit shift can be expressed either in terms of the beam–beam induced distortion of the *actual* beam separation δ_i , as discussed in Sect. 3.5.1 below, or in terms of an orbit-related luminosity-bias factor, denoted by $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}(\Delta)$ for a given *nominal* separation Δ . In physically intuitive terms, and in the context of a weak-strong model such as B*B, the orbit shift results from applying to all particles in the witness bunch the same electromagnetic kick, computed from the field produced by the source bunch as a whole and averaged over all particles in the witness bunch. Since this is tantamount to a dipole kick, only the orbit of the witness bunch is affected; the size and shape of the macroparticle distribution remain invariant as $|\Delta|$ increases.

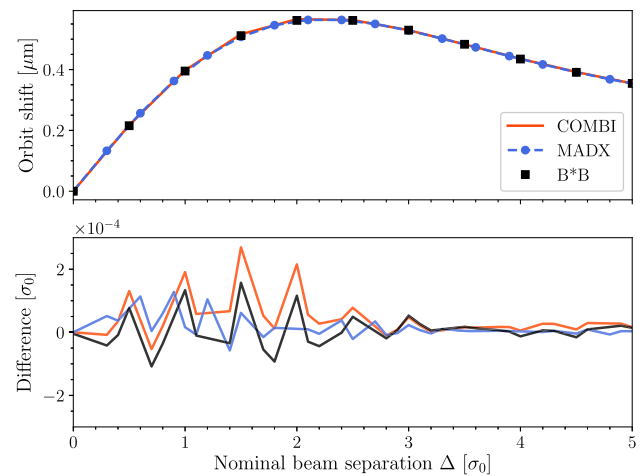


Fig. 5 Top: Bunch-centroid displacement δ_{iB}^{bb} due to the beam–beam deflection during a simulated vdM scan, as predicted by COMBI, MADX and B*B. The bunch parameters are listed in the right column of Table 1. The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . Bottom: difference between the simulated orbit shift and that calculated analytically using Eqs. (14)–(15), in units of σ_0 . The wiggles reflect the combination of a finite sampling step ($0.2\sigma_0$) and of numerical noise

The impact of the optical distortions is quantified in terms of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\Delta)$ first introduced in Sect. 2.3.2. It results from the combination of beam-separation and amplitude-dependent β -beating effects that modulate the RMS transverse beam size (Sect. 3.5.2), and of Δ -dependent bunch-shape distortions that affect the beam–beam overlap integral (Sect. 3.5.3).

The B*B and COMBI packages have been mutually benchmarked, and their results compared to those obtained either analytically (where possible) or using MADX. The cross-package comparisons of the beam–beam induced orbit shift and of the predicted optical distortions are based on the reference parameter set of Table 1, and assume that in the absence of the beam–beam interaction the transverse density distributions are strictly Gaussian. The consistency of B*B and COMBI results at higher beam–beam parameters is quantified in Sect. 3.5.4. All the results presented in this section assume in addition that the beams collide only at the IP where the vdM scan is taking place; multiple-IP effects will be discussed in Sect. 4.6.

3.5.1 Beam–beam-induced orbit shift

Figure 5 displays the beam–beam-induced, single-beam orbit shift (δ_{iB}^{bb}) at the IP, as simulated by MADX, B*B and COMBI, and compares it to the analytical prediction. The three models agree extremely well among themselves, and their results are indistinguishable from those of the analytical prediction.

During a luminosity-calibration session, the bunches typically collide at more than one IP (Sect. 4.6); however performing simultaneous beam-separation scans at different IPs is carefully avoided. Under these conditions and to an excellent approximation, the orbit shifts associated with slightly misaligned collisions at a non-scanning IP remain static during a beam-separation scan at another IP. Their impact on the beam separation at the scanning IP is therefore expected to remain constant during the scan, and has been neglected in the simulations reported in this paper.

3.5.2 Impact of amplitude-dependent β -beating effects in the Gaussian-bunch approximation

The beam-separation dependence of the transverse RMS beam-size ratio (squared for easier interpretation in terms of luminosity) is presented in Fig. 4. The results of B*B and COMBI agree to better than 2×10^{-4} .

At $\Delta = 0$, both predict a decrease in beam-size squared (or equivalently an increase in head-on luminosity) about half of that obtained using MAD-X. This is remarkably consistent with the analytical predictions of Eqs. (22) and (23): in MAD-X, all particles in the witness bunch are subject to the same quadrupole-like force as the zero-amplitude particle, while in B*B and COMBI, the slope of the beam-beam force decreases and even changes sign as the betatron amplitude increases (Fig. 3), resulting in a smaller overall change in optical demagnification at the IP.

Amplitude detuning also explains the different Δ -dependence of the beam sizes, with that in MAD-X being more pronounced than in B*B and COMBI. At very large beam separation, where the electromagnetic force becomes similar to that of a distant, point-like charge, all three curves tend towards a common asymptote. As for the different evolution of the horizontal and vertical beam sizes during a scan, it simply reflects the tune-dependence apparent in, for instance, Eq. (18).

If one assumes that the optical distortions discussed above modify the transverse RMS bunch sizes at the IP without significantly affecting their initial, purely Gaussian shape, Eqs. (5)–(7) apply. The combination of the beam-beam induced orbit shift and of the optical distortions then yields the following expression for the luminosity-bias factor during a beam-separation scan in plane i ($i = x, y$):

$$\frac{\mathcal{L}}{\mathcal{L}^0}(\Delta_i) = \frac{\Sigma_x^0 \Sigma_y^0}{\Sigma_x(\Delta_i) \Sigma_y(\Delta_i)} e^{-\frac{1}{2}[(\delta_i/\Sigma_i)^2 - (\Delta_i/\Sigma_i^0)^2]} \quad (24)$$

$$= \frac{(\sigma^0)^2}{\sigma_x(\Delta_i) \sigma_y(\Delta_i)} e^{-\frac{1}{4}[(\delta_i/\sigma_i)^2 - (\Delta_i/\sigma^0)^2]} \quad (25)$$

where $\sigma_{x,y}$, $\Sigma_{x,y}$, \mathcal{L} (resp. σ_0 , $\Sigma_{x,y}^0$, \mathcal{L}^0) are the RMS single-beam sizes, the inverse of the overlap integrals and the lumi-




Fig. 6 Beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ during a simulated horizontal vdM scan. The unperturbed beams are assumed to be round and perfectly Gaussian; their parameters are listed in the right column of Table 1. The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The MAD-X (dashed blue) and COMBI (dotted grey) curves are valid only in the Gaussian-bunch approximation (Sect. 3.5.2); the COMBI (red curve) and the B*B markers (black squares) are obtained from the overlap integrals of the multiparticle distributions (Sect. 3.5.3). The green circular marker displays the prediction of Eq. (23), that only applies at zero beam separation

nosity in the presence (resp. absence) of the beam-beam interaction.

Combining Eq. (25) with the beam-size ratios shown in Fig. 4 yields the beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0](\Delta)$ in the Gaussian-bunch approximation, as illustrated in Fig. 6 for MAD-X (dashed blue) and COMBI (dotted grey). These curves are not simply the inverse of the beam-size ratios displayed in Fig. 4, because they include in addition the impact, at a given nominal separation, of the beam-beam induced orbit shift. At zero and moderate separation ($\Delta/\sigma_0 < 1$), the luminosity bias $[\mathcal{L}/\mathcal{L}^0] - 1$ is positive and dominated by the dynamic- β effect; as the separation increases, the orbit shift, which is represented by the exponential term in Eq. (24), progressively takes over.

3.5.3 Impact of optical distortions on the beam-beam overlap integral

If the transverse-density distributions are sufficiently distorted by the non-linearity of the beam-beam force, the Gaussian-bunch approximation encapsulated in Eqs. (24)–(25) is no longer valid, and the luminosity bias must be calculated numerically from the beam-separation-dependent overlap integrals of the macroparticle distributions. To this effect, an integrator module, detailed in Appendix B, has been developed for COMBI; B*B provides an equivalent functionality [16]. The resulting beam-separation dependence of

the luminosity-bias factor is shown by the red curve and the black markers in Fig. 6. At each simulated scan step (indicated by the markers), B*B and COMBI agree to better than 2×10^{-4} ; the difference is not systematic, but fluctuates around zero. The difference between the Gaussian-bunch approximation and the numerically calculated overlap integral demonstrates that the non-Gaussian tails (or more precisely non-Gaussian deviations from the originally Gaussian shape), that are induced by the beam-beam interaction, have a significant impact on the luminosity as soon as $\Delta/\sigma_0 \geq 0.2$. The excellent agreement between the two multiparticle simulations also demonstrates that strong-strong beam-beam effects, that are modeled by COMBI but not by B*B, remain negligible in the low beam-beam parameter regime considered here.

Since the predicted orbit effect is identical in all simulations (Fig. 5), the differences between MAD-X on the one hand (Fig. 6, dashed blue curve), and B*B/COMBI on the other (black markers/red curve), is entirely associated with optical distortions. At zero separation, the difference is entirely explained by amplitude detuning (Eq. (23)); once the separation increases, beam-beam-induced non-Gaussian tails play a growing role, as illustrated by the difference between the red and grey curves.⁶

In the context of the beam-beam correction strategy outlined in Sect. 2.3, it is convenient to separate the full luminosity bias presented in Fig. 6 into its optical-distortion and orbit-shift components. The luminosity-bias factor associated with optical distortions and denoted by $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ is defined as the ratio, scan step by scan step, of the full luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ and of the bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ associated with the orbit shift:

$$[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\Delta) = \frac{[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}(\Delta)}{[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}(\Delta)}. \quad (26)$$

Its beam-separation dependence is presented in Fig. 7. While MAD-X overestimates the dynamic- β effect at zero separation, it strongly underestimates the optical distortions at medium and large beam separations. Since all models predict the same orbit effect, and since the optical distortions and the orbit effect impact the overlap integrals in opposite ways, their mutual cancellation is stronger in B*B and COMBI than in MAD-X, resulting in a net overall beam-beam correction of significantly smaller magnitude. This will be discussed quantitatively in Sect. 4.




Fig. 7 Beam-separation dependence, during a simulated horizontal *vdM* scan, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with the optical distortions, and of that induced by the beam-beam orbit effect ($[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$). The unperturbed beams are assumed to be round and perfectly Gaussian; their parameters are listed in the right column of Table 1. The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The MAD-X (dashed blue line with markers), weak-strong B*B (black triangles) and strong-strong COMBI (red curve) optical-distortion curves are obtained by dividing the corresponding full beam-beam bias factor shown in Fig. 6 by that predicted analytically for the orbit effect alone (green curve)

The beam-separation dependence of the optical-distortion factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ predicted by B*B and COMBI (Fig. 7) exhibits a characteristic S-shape or “wobble”, the amplitude of which increases with ξ (Fig. 8) and with the fractional tunes (Fig. 12). The beam separation where the wobble amplitude exhibits either a peak ($\Delta \sim 1.6 \sigma_0$) or a dip ($\Delta \sim 3 \sigma_0$), is determined by the shape of the particle-density distribution. This is because the latter dictates both the dependence of the beam-beam force on the distance to the center of the source bunch, and the transverse distribution of the witness particles in the opposing bunch; an example will be offered in Sect. 4.3 (Fig. 18).

Some intuitive insight into the beam-separation dependence of $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ can be gained by conceptually separating what occurs in the Gaussian core of the witness bunch from what happens to the large-amplitude particles in its non-Gaussian tails. The grey curve in Fig. 6 reflects the evolution, during the scan, of the transverse RMS bunch sizes (Fig. 4); in a sense, therefore, it represents the impact of beam-beam effects on the overlap of the Gaussian component of the particle distribution. The difference between the red and grey curves, in contrast, reflects the beam-separation dependence of the luminosity contribution of the non-Gaussian transverse tails; the latter dominate the overall upward trend of the red curve in Fig. 7, and are responsible for most of its S-shape. A quantitative understanding of these features, however, is accessible neither to analytical calculations nor to simple physical arguments. It can only be obtained from simula-

⁶ The very small difference, at zero separation, between the grey and red curves in Fig. 6 may grow when the bunches collide not only at the scanning IP, but also at additional IPs (Sect. 4.6). In this case, and depending on the phase advance between IPs, beam-beam-induced non-Gaussian tails at these additional IPs may contribute noticeably to the overlap integral at the scanning IP, even for zero beam separation.

tions, because it requires taking into account, at each scan step, the interplay between:

- the actual beam separation;
- the transverse bunch shapes;
- the non-linearity of the beam–beam force;
- the combined separation- and amplitude-dependence of the force exerted by the source bunch as a whole on those particles in the witness bunch that significantly contribute to the luminosity at the scan step considered.

3.5.4 Beam-parameter dependence of beam–beam corrections

The results in Figs. 6 and 7 were obtained using the reference parameter set in the right column of Table 1. This corresponds to a beam–beam parameter $\xi = 2.59 \times 10^{-3}$, at the low end of the ξ range explored during *vdM* scans at $\sqrt{s} = 13$ TeV. The cross-validation of B*B and COMBI was therefore repeated with a larger bunch current and a smaller emittance, both typical of routine physics running and corresponding to a beam–beam parameter value well beyond the *vdM* range. The beam-separation dependence of the optical-distortion luminosity-bias factor at these two ξ values is presented in Fig. 8. While in perfect agreement at low ξ , the two codes exhibit hints of a small but systematic difference in the high- ξ regime when the beam separation becomes large enough. This discrepancy is attributed to the fact that B*B is intrinsically a weak-strong model, and therefore cannot account for coherent bunch oscillations, contrarily to COMBI. This interpretation is confirmed by the comparison of the tune spectra predicted by the two packages: even in the *vdM* regime, a π -mode peak is apparent in the COMBI spectrum, but is missing (as it should be) from the B*B spectrum.

Quantitatively, it turns out that throughout the *vdM* regime and even in the high- ξ case considered here, the discrepancy is of no practical significance. Since it manifests itself only in the tails of the beams ($\Delta > 2\sigma_0$), where the particle density is low, the difference, between B*B and COMBI, in beam–beam-induced distortions has only a very small impact on the integral of the scan curves, i.e. on the perturbed transverse convolved beam sizes (Eq. (4)). This can be quantified using the methodology that will be presented in Sect. 4.2.4. For the high- ξ setting illustrated in Fig. 8, the difference between B*B and COMBI translates into a systematic uncertainty of 0.04% on the absolute luminosity scale. In typical *vdM* scans, where ξ is lower by a factor of about 1.5, the inconsistency becomes negligible; there the two simulation codes remain in more than adequate agreement, validating the use of the less resource-intensive B*B package for parameterizing beam–beam corrections to *vdM* calibrations.

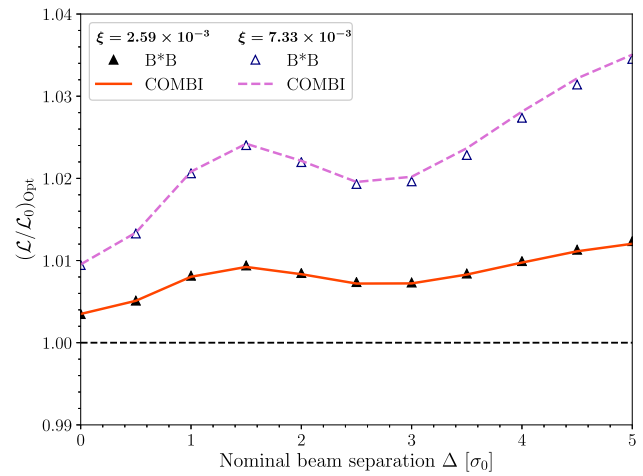


Fig. 8 Beam-separation dependence, during a simulated horizontal *vdM* scan, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with optical distortions, for the *vdM* reference parameter set ($\xi = 2.59 \times 10^{-3}$, solid red curve and black filled triangles), and for a beam–beam parameter typical of routine physics running ($\xi = 7.33 \times 10^{-3}$, dashed purple curve and blue open triangles). The COMBI (B*B) results are shown by the straight-line segments (markers). The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0

4 Calculated impact of beam–beam dynamics on luminosity calibrations

4.1 Methodology

The beam–beam biases affecting a *vdM* calibration are accounted for by correcting the luminosity-scan curves, one bunch pair at a time, according to the procedure outlined in Sect. 2.3. While the orbit shift (Sect. 2.3.1) can be calculated analytically, correcting for optical distortions (Sect. 2.3.2) requires the knowledge of the beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$, such as that illustrated in Fig. 8. The latter can be obtained by running B*B or COMBI with, as input:

- the interaction-region (IR) configuration (beam energy, nominal or measured β^* and crossing-angle values);
- the unperturbed tunes Q_x and Q_y (or equivalently the unperturbed fractional tunes q_x and q_y), i.e. the values of the horizontal and vertical tunes with the beam–beam interaction switched off at the IP where the simulated scans are taking place. Physically, these correspond to the tune values that would be measured, for the bunch pair under study, before the beams are brought into collision at the scanning IP considered;
- the measured parameters of the bunch pair under study: bunch population, horizontal and vertical beam sizes or emittances, as well as the bunch length in case of a non-zero nominal crossing angle.

Since vdM calibrations must be performed on a bunch-by-bunch basis, with possibly over 100 colliding-bunch pairs and typically five to ten x – y scan pairs per scan session and per IP, the approach sketched above can become unwieldy from the computing viewpoint. This motivated the development of a much lighter technique, based on the fact that at least in the vdM regime and under some simplifying assumptions, beam–beam biases, to a very good approximation, scale almost linearly with the beam–beam parameter ξ . This makes it possible to construct a simple polynomial parameterization of the luminosity-bias curves, that is extracted from B*B or COMBI simulations over a grid in beam-parameter space and is applicable to most cases of practical interest for pp vdM calibrations at the LHC. The impact of violating the underlying assumptions is accounted for either by a contribution to the systematic uncertainty that is associated with the beam–beam correction procedure, or, in the case of multi-IP effects, by a simulation-guided adjustment to the parameterized correction.

Some cases do not lend themselves to the parameterization technique, such as off-axis vdM scans, diagonal scans, or vdM calibrations performed with a large crossing angle (which is unavoidable at the LHCb IP).

An off-axis vdM scan is a horizontal (or vertical) beam-separation scan where the beams are partially separated in the non-scanning plane, i.e. in the vertical (or horizontal) direction. A diagonal scan is one in which the beams are scanned transversely along an inclined straight line in the x – y plane, rather than only along either the x or the y axis. Such “generalized”, one-dimensional vdM scans are sometimes used to measure and correct for non-factorization effects [22, 43, 44]. While the orbit-shift correction can still be calculated analytically using a more general form of the Bassetti–Erskine formula, the increased dimensionality of the parameter space makes it impractical to invest in general enough and precise enough a parameterization of optical-distortion effects. In such a case, deriving the corrections from a large set of fully simulated scans, labor- and CPU-intensive as it may be, appears more tractable by comparison.

Allowing non-zero crossing angles also increases the dimensionality of the parameter space, with similar implications. Therefore, except where explicitly stated otherwise, the results offered in the remainder of this report are restricted to the zero (or moderate) crossing-angle case, which covers the needs of the ATLAS and CMS experiments (as well as those of ALICE, at the cost of a slight increase in systematic uncertainty).

This section is organized as follows. The parameterization approach mentioned above is described in Sect. 4.2. The associated “fully symmetric Gaussian-beam configuration” assumes that in the absence of any beam–beam interaction, the colliding bunches in beam 1 and beam 2:

- can be modeled by factorizable transverse-density distributions, and exhibit a single-Gaussian profile in all three dimensions. The impact of violating this assumption is evaluated in Sect. 4.3;
- are round in the transverse plane ($\sigma_{xB} = \sigma_{yB}$, $B = 1, 2$). The impact of violating this assumption is evaluated in Sect. 4.4;
- intersect at zero crossing angle. The impact of violating this assumption is evaluated in Sect. 4.5;
- collide only at the IP where beam-separation scans are performed. The impact of violating this assumption is evaluated in Sect. 4.6;
- are beam–beam symmetric, i.e. equally populated ($n_1 = n_2$), of the same transverse size ($\sigma_{i,1} = \sigma_{i,2}$, $i = x, y$), and therefore subject to the same beam–beam parameter ($\xi_{i,1} = \xi_{i,2}$, $i = x, y$). The impact of violating these assumptions is evaluated in Sect. 4.7.

4.2 Beam–beam correction procedure in the fully symmetric Gaussian-beam configuration

The parametrization strategy relies on the fact that at fixed tunes, beam–beam induced biases to the visible cross-section depend only on the beam–beam parameter ξ (Sect. 4.2.1). The simulated ξ - and tune-dependence of the luminosity-bias functions (Sect. 4.2.2) leads to parameterizing them by second-order polynomials that can be used, in the context of the beam–beam correction procedure of vdM calibrations, as a proxy for full-fledged B*B or COMBI simulations (Sect. 4.2.3). Combining these parameterized bias functions with hypothetical vdM -scan curves devoid of fitting biases and of step-to-step fluctuations provides robust and intuitive insight into the magnitude and beam-conditions dependence of beam–beam corrections to luminosity calibrations (Sect. 4.2.4).

4.2.1 Validation of the scaling hypothesis

To characterize the scaling properties (or lack thereof) of beam–beam biases during vdM scans, pairs of horizontal and vertical beam-separation scans were generated under the assumptions above using B*B.⁷ The beam–beam parameter spanned the range $\xi \sim 0.002$ – 0.008 , thereby covering from vdM scans with low-brightness proton beams at the low end, to conditions slightly above routine physics running at the high end. The input beam energies, bunch populations, emittance and β^* values were chosen to be representative of vdM scans during LHC Runs 1 and 2 (Sect. 2.2); also included was a high- ξ setting representative of Run-2 high-luminosity

⁷ In view of the consistency of the B*B and COMBI results demonstrated in Sect. 3, the B*B package was chosen for practical reasons, first and foremost because it is less computationally expensive.




Fig. 9 Beam-separation dependence, during a simulated horizontal vdM scan, of the luminosity-bias factor associated with the beam–beam orbit shift only (green squares), the optical distortions only (red triangles), and their combination (black circles). The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The beams satisfy the assumptions listed in Sect. 4.1. The input unperturbed-tune and beam–beam parameter values are indicated at the top; the other relevant parameters are $E_B = 6.5$ TeV, $\beta^* = 19.2$ m and $\epsilon_N = 2.6$ $\mu\text{m-rad}$. The points are the results of the simulation; the lines are intended to guide the eye

physics running. The unperturbed fractional tunes were constrained to satisfy $q_y = q_x + 0.010$ so as to reflect routine LHC collision settings.

The primary output of the simulation is the beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$. An example is presented in Fig. 9, with a value of ξ chosen to lie roughly in the middle of the vdM range (Table 1). The orbit-shift bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ (green squares) is computed analytically:

$$[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}(\Delta) = e^{-\frac{1}{2}[(\delta_i/\Sigma_R)^2 - (\Delta/\Sigma_R)^2]} \quad (27)$$

where Δ is the nominal separation in the scanning plane, δ_i is computed using Eqs. (14) to (16), and $\Sigma_R = \sqrt{\Sigma_x^0 \Sigma_y^0}$ is the unperturbed, round beam-equivalent convolved beam size that in this particular case satisfies $\Sigma_R = \Sigma_x^0 = \Sigma_y^0 = \sqrt{2} \sigma_0$. The optical-distortion bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ (red triangles) can then be calculated using Eq. (26).

For given values of q_x and q_y , beam–beam-induced distortions scale with the beam–beam parameter, in the sense that they depend only on ξ . This can be proven mathematically, both for the orbit shift (Eqs. (14)–(15)) and for the dynamic- β effect at zero separation (Eq. (18)). That this remains true at any beam separation, within the constraints of the fully symmetric beam configuration and over the ξ range specified above, can only be demonstrated by simulation. To this effect, luminosity-bias curves such as those presented in Fig. 9 were




Fig. 10 Beam-separation dependence, during simulated horizontal vdM scans, of the difference between the luminosity-bias factor (full beam–beam effect) calculated for the collision parameters shown in the legend, and that computed using the reference parameters (Fig. 9, black circles) listed at the top of the present figure. For each parameter set, the bunch intensity is adjusted such that the value of ξ remains the same. The error bars are explained in the text

compared for different combinations of bunch populations, beam energies, β^* and emittance values that all correspond to a given value of ξ ; they were found to be identical within the statistics of the simulation.

An example is presented in Fig. 10, which displays the difference, at each scan step, between the luminosity-bias factor computed for three distinct sets of beam parameters, and that associated with the parameters used in Fig. 9; all four parameter sets correspond to $\xi = 0.004$. The error bars represent the statistical uncertainty associated with the randomization of initial conditions in the B*B code [16]; they are computed as the error on the mean over eight different runs, with different random seeds, for each beam separation and each parameter set. The differences in luminosity-bias values between the four beam-parameter sets and the reference set are statistically consistent with zero, and never exceed 10^{-4} . The exercise was repeated for a range of ξ values, leading to the conclusion that at fixed Q_x and Q_y , the luminosity-bias curves indeed depend only on ξ , to better than 10^{-4} on the absolute luminosity scale.

4.2.2 Beam–beam parameter and tune dependence of the optical-distortion correction

The combined beam-separation and ξ dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ is illustrated in Fig. 11. The simulations shown span the full beam–beam parameter range ($2.04 \times 10^{-3} \leq \xi \leq 7.83 \times 10^{-3}$), and are carried out at the nominal LHC tune settings ($Q_x = 64.31$, $Q_y = 59.32$). A similar ξ -dependence is observed for $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ and




Fig. 11 Beam-separation dependence, during simulated horizontal vdM scans, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with optical distortions, for several values of the beam–beam parameter ξ . The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The beams satisfy the assumptions listed in Sect. 4.1. The input unperturbed fractional-tune values are indicated at the top. The points are the results of the simulation; the lines are intended to guide the eye




Fig. 12 Beam-separation dependence, during simulated horizontal vdM scans, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with optical distortions, for several values of the unperturbed fractional tunes (q_x, q_y). The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The beams satisfy the assumptions listed in Sect. 4.1. The input ξ value is indicated at the top. For each curve, the value of the horizontal tune is indicated in the legend; the value of the vertical tune is set to $q_y = q_x + 0.01$. The points are the results of the simulation; the lines are intended to guide the eye

$[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ (not shown). The beam–beam parameter dependence of all three variables at fixed nominal separation is found to be well modeled by a second-order polynomial of ξ , the coefficients of which depend on the nominal separation.

Since both the orbit shift (Eq. (15)) and the dynamic- β effect (Eq. (18)) depend on the unperturbed tunes, and as part of investigating the systematic uncertainties associated with beam–beam corrections, their sensitivity to the input tune values was characterized by extending the above-described simulations to an area in the (q_x, q_y) plane that encompasses the full range of operational tune settings and of beam–beam tune shifts expected during vdM sessions. Horizontal and vertical beam-separation scans were simulated, for several values of ξ , over a grid⁸ bounded by $0.2975 < q_x < 0.3100$, with the vertical unperturbed fractional tune set to $q_y = q_x + \delta_q$ where $\delta_q = 0.0075, 0.0100, 0.0125$.

The combined beam-separation and tune dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ is illustrated in Fig. 12 for $\xi = 7.83 \times 10^{-3}$. For a given nominal separation, the lower the tunes, the more they approach the quarter integer, and therefore the smaller (i.e. the closer to unity) the optical-distortion bias factor becomes [13]. Similarly, the lower the

tunes, the smaller (i.e. the further away from unity) the orbit-shift bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ (not shown), albeit with a weaker tune dependence. As a result, the tune dependence of the full beam–beam-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ is slightly steeper than that associated with optical distortions only: the lower the tunes, the faster $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ drops in magnitude (Fig. 13), and therefore the larger the overall beam–beam correction to the luminosity scale. In addition, the lower the tunes, the less the optical distortions contribute to the overall beam–beam bias to σ_{vis} .

4.2.3 Practical implementation of parameterized beam–beam corrections

Concretely, given a measured vdM -scan curve such as that displayed in Fig. 1 (bottom), the beam–beam correction procedure can be implemented as follows.

1. At each scan step in the i scanning plane ($i = x, y$), correct the nominal separation Δ_i for the deflection-induced separation shift δ_i^{bb} , as detailed in Sect. 2.3.1, using Eqs. (14)–(16).
2. At each scan step in the i scanning plane, correct the measured rate $R_i(\Delta_i)$ for the optical-distortion bias as detailed in Sect. 2.3.2, using Eq. (17). The beam-separation dependence of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ is either extracted directly from a set of B*B

⁸ For technical reasons associated with the numerical evaluation of overlap integrals in B*B [16], the input tune values had to be chosen such that $\text{gcd}(10000 q_x, 10000 q_y) = 1$, where gcd refers to the greatest common divider. This avoids integration over closed curves that may lead to unexpected systematic effects.




Fig. 13 Beam-separation dependence, during simulated horizontal *vdM* scans, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$ associated with the full beam–beam effect, for several values of the unperturbed fractional tunes (q_x, q_y) . The horizontal axis is the nominal beam separation in units of the unperturbed transverse beam size σ_0 . The beams satisfy the assumptions listed in Sect. 4.1. The input ξ value is indicated at the top. For each curve, the value of the horizontal tune is indicated in the legend; the value of the vertical tune is set to $q_y = q_x + 0.01$. The points are the results of the simulation; the lines are intended to guide the eye

or COMBI simulations, or, if the assumptions listed in Sect. 4.1 can be considered valid, simply computed using the polynomial parameterization described below. The latter is valid only for on-axis, one-dimensional *vdM* scans with zero crossing angle; other configurations, and in particular off-axis one-dimensional scans as well as two-dimensional grid scans in the (δ_x, δ_y) beam-separation plane, require dedicated simulations.

The simulations described in Sect. 4.2.2 were used to parameterize the optical-distortion luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ in bins of normalized nominal separation over the range $0 \leq \Delta_i/\sigma_0 \leq 6$, separately for *x* and *y* scans, by polynomials of the form

$$\begin{aligned}
 [\mathcal{L}/\mathcal{L}^0(\xi_R, q_x, q_y)]_k &= p_{0k} + p_{1k} \xi_R + p_{2k} q_x + p_{3k} q_y \\
 &+ p_{4k} \xi_R^2 + p_{5k} q_x^2 + p_{6k} q_y^2 \\
 &+ p_{7k} \xi_R q_x + p_{8k} \xi_R q_y \\
 &+ p_{9k} q_x q_y.
 \end{aligned} \tag{28}$$

Here the index *k* ($k = 1, 25$) labels the nominal-separation bin in the scanning plane considered. The round-beam equivalent beam–beam parameter ξ_R is defined as

$$\xi_R = \frac{nr_0 \beta^*}{4\pi \gamma \sigma_R^2}, \tag{29}$$

where the single-beam transverse RMS beam size σ_0 in Eq. (13) is approximated by its measured⁹ round-beam equivalent $\sigma_R = \sqrt{\Sigma_x \Sigma_y}/2$, thereby partially accounting for a potential ellipticity of the beams at the IP. The input fractional-tune values (q_x, q_y) are those of the unperturbed tunes defined in Sect. 4.1. In the case where bunches collide only at the IP where the scan is taking place, i.e. remain fully separated, transversely and/or longitudinally, at the other three IPs, these unperturbed tunes are identical to the nominal LHC tunes, or equivalently to the values one would measure before the beams are put in collision at the scanning IP. If, however, some bunches also collide at one or more other IPs, the unperturbed tunes input to the parameterization for those bunches must take into account the beam–beam tune shift arising from collisions at non-scanning IPs; a prescription to this effect is offered in Sect. 4.6.

The parameterization above, that amounts to a second-order Taylor expansion in ξ_R, q_x and q_y , was found sufficient to approximate the exact simulation results to better than 10^{-3} on $[\mathcal{L}/\mathcal{L}^0]$ over most of the grid of simulated points in (ξ, q_x, q_y, Δ) space. The function defined by Eq. (28) is linear in all 10 parameters p_{lk} . In each nominal-separation bin *k*, the set of parameters p_{0k}, \dots, p_{9k} is determined by a weighted linear least-square fit. A smooth dependence of $[\mathcal{L}/\mathcal{L}^0]$ on the nominal separation, for given values of ξ_R, q_x and q_y , can be achieved by, for instance, cubic spline interpolation between adjacent separation bins. Parameterizations based on the same functional form were also constructed for the orbit-shift and full beam–beam bias factors $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ and $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$, and achieved similar numerical accuracy.

Tabulated values of the parameters p_{lk} in separation steps of 0.25 in Δ/σ_R , for both vertical and horizontal scans and over the full range of ξ_R and tune values defined earlier, have been made available to all LHC experimental Collaborations. The optical-distortion parameters, which are the most useful, are documented in Appendix C, and are publicly accessible in computer-readable form [45].

⁹ It will be shown in Sect. 4.2.4 that in the *vdM* regime, beam–beam effects, if left uncorrected, lead to a percent-level underestimate of Σ_x and Σ_y , and therefore of σ_R . The use of the measured transverse beam size – as opposed to that of the experimentally inaccessible unperturbed beam size – in estimating ξ_R therefore leads to a slight overestimate of the beam–beam parameter input to the simulation. However, since the overall beam–beam bias on the effective cross-section is typically less than 1% (Fig. 16, black curve), a sub-percent overestimate of ξ_R biases the overall beam–beam correction to σ_{vis} by 1–1.5% of itself. This can usually be neglected in view of the much larger fractional systematic uncertainties detailed in Sect. 5. Alternatively, the correction can be iterated upon, by using as input to the second iteration a ξ_R value based on the corrected effective single-beam size obtained from the first iteration.

4.2.4 Separation-integrated estimates of beam–beam corrections to vdM calibrations

The beam–beam correction procedure detailed in the preceding section is designed to be applied to measured vdM -scan curves, one scan step at a time. One then extracts the beam–beam corrected peak rate and convolved beam sizes needed to calculate the visible cross-section (Eq. (10)) using a carefully chosen fit function. More often than not, the latter must diverge from a perfect Gaussian in order to faithfully model the data over the full beam-separation range.

In order to provide consistency checks on this intricate analysis chain, as well as a physically intuitive, albeit approximate, breakdown of the impact of individual beam–beam effects on the luminosity calibration under study, the same procedure can be applied to an x – y pair of hypothetical beam-separation scan curves, that suffer from no statistical fluctuations and that, in the $\xi_R \rightarrow 0$ limit would be perfectly Gaussian:

$$\mathcal{L}^0(\Delta) \propto \frac{1}{4\pi\sigma_0^2} \exp\left(-\frac{\Delta^2}{4\sigma_0^2}\right).$$

Here and in the remainder of this paper, the “0” index indicates a “nominal” quantity, i.e. one for which beam–beam effects have been fully turned off in the simulation.

At a given nominal separation Δ_i in scanning plane i ($i = x, y$), and with the beams transversely centered on each other in the non-scanning plane, the luminosity in the presence of, say, the full beam–beam effect is then calculated as

$$\mathcal{L}(\Delta_i) = \mathcal{L}^0(\Delta_i) \times [\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}(\Delta_i) \tag{30}$$

with $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}(\Delta_i)$ provided by the parametrization detailed in Sect. 4.2.3, with input parameters representative of the actual beam conditions during the scans under study. The impact of either the orbit shift or the optical distortions alone can be evaluated by substituting $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}(\Delta_i)$ or $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\Delta_i)$ for $[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}(\Delta_i)$ in Eq. (30).

One then defines “figures of merit” (FoMs) that characterize the impact of beam–beam effects on vdM observables.

1. The peak-rate bias factor characterizes the dynamic- β effect at zero beam separation:

$$\mu_{pk}/\mu_{pk}^0 = \frac{\mathcal{L}(\Delta_i = 0)}{\mathcal{L}^0(\Delta_i = 0)}, \tag{31}$$

and is sensitive to optical distortions only.

2. The beam-size bias factors characterize, in each plane, the beam–beam impact on the convolved transverse beam size:

$$\Sigma_i/\Sigma_i^0 \quad (i = x, y), \tag{32}$$

where the overlap integrals

$$\Sigma_i = \frac{1}{\sqrt{2\pi}} \frac{\int \mathcal{L}(\Delta_i) d\Delta_i}{\mathcal{L}(\Delta_i = 0)} \tag{33}$$

$$\Sigma_i^0 = \frac{1}{\sqrt{2\pi}} \frac{\int \mathcal{L}^0(\Delta_i) d\Delta_i}{\mathcal{L}^0(\Delta_i = 0)} \tag{34}$$

are calculated numerically, with their integrands smoothed by cubic-spline interpolation.

3. The σ_{vis} -bias factor characterizes the beam–beam impact on the absolute luminosity scale, and is given by the product of the previous three FoMs:

$$\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 = (\mu_{pk}/\mu_{pk}^0) \times (\Sigma_x/\Sigma_x^0) \times (\Sigma_y/\Sigma_y^0). \tag{35}$$

As an illustration, Fig. 14 displays the beam–beam parameter dependence of the peak-rate bias factor for the full beam–beam effect (black circles), the orbit shift only (green squares), and the optical distortion only (red triangles). Since at zero separation there is no orbit shift, only optical distortions matter in this case.

The evolution of the horizontal beam-size bias factor is presented in Fig. 15. While orbit shifts, if left uncorrected, result in an underestimate of Σ_x and Σ_y , optical distortions have the opposite effect, resulting in a partial cancellation of the bias.

This is even more apparent in the visible-cross-section bias curves (Fig. 16). While the full beam–beam effect on the visible cross section remains below 1% over the whole ξ range (black circles), individual contributions from the orbit shift (green squares) and the optical distortion (red triangles) can each lead to a 3 – 4% bias, a clear demonstration of how much overcorrection can be expected from too approximate a treatment of optical distortions.

In Figs. 14, 15 and 16, the curves are fits to quadratic functions of ξ_R :

$$f(\xi_R)_j = p_{0j} + p_{1j}\xi_R + p_{2j}\xi_R^2, \tag{36}$$

where j labels the type of correction (orbit shift only, optical distortion only, or full beam–beam). Even though the linear term is clearly dominant, a quadratic term is needed to describe the evolution of all three FoMs to satisfactory precision over the full range of beam–beam parameter values.

The impact of beam–beam effects on the absolute-luminosity scale can be expressed equivalently by either the σ_{vis} bias factor $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0$, the visible cross-section bias $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1$ (typically a fraction of a percent), or a multiplicative correction factor $(\sigma_{\text{vis}}/\sigma_{\text{vis}}^0)^{-1}$ to the raw visible cross-section. This FoM approach, that lends itself to a simple polynomial parameterization in terms of ξ_R , q_x and q_y , offers the advantage that its results are easy to interpret, as




Fig. 14 Dependence of the peak-rate bias factor on the beam-beam parameter, with beams colliding at the scanning IP only. The unperturbed tunes are fixed at the nominal LHC values ($Q_x = 64.31$, $Q_y = 59.32$). The symbols display the simulation results; the curves are described in the text. Since the orbit shift does not impact the peak-rate bias factor, the red and black symbols and curves lie exactly on top of each other

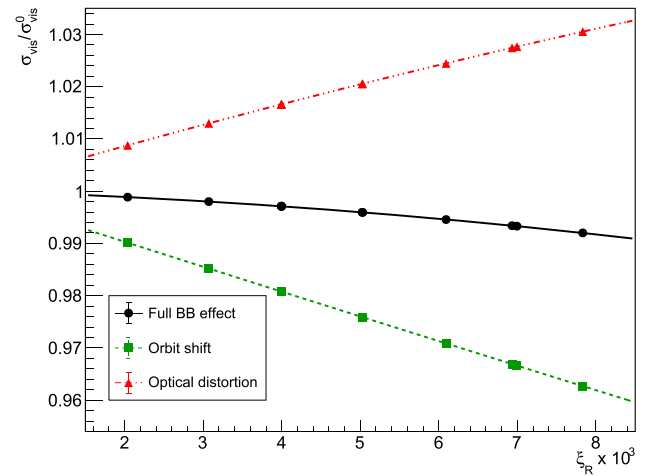


Fig. 16 Dependence of the σ_{vis} bias factor on the beam-beam parameter, with beams colliding at the scanning IP only. The unperturbed tunes are fixed at the nominal LHC values ($Q_x = 64.31$, $Q_y = 59.32$). The symbols display the simulation results; the curves are described in the text




Fig. 15 Dependence of the horizontal beam-size bias factor on the beam-beam parameter, with beams colliding at the scanning IP only. The unperturbed tunes are fixed at the nominal LHC values ($Q_x = 64.31$, $Q_y = 59.32$). The symbols display the simulation results; the curves are described in the text

well as insensitive both to the quality of the fit to the measured scan curves and to rate fluctuations, from one scan step to the next, due for instance to counting statistics or beam-position jitter. This technique also simplifies the evaluation of some of the systematic uncertainties discussed in the remainder of this chapter. It is, however, not recommended for determining the central value of the beam-beam correction on real data, since it ignores the deviations of the actual scan curves from a perfect Gaussian.

4.3 Impact of non-Gaussian unperturbed transverse beam profiles

All the simulation results presented in this report so far make the explicit assumption that in the absence of the beam-beam interaction at the scanning IP, the unperturbed transverse beam profiles, i.e. the particle density functions $\hat{\rho}_B(x, y)$ ($B = 1, 2$) in Eq. (1), can be perfectly modeled by the uncorrelated x - y product of two single, one-dimensional Gaussians. Should this not be the case, the beam-separation dependence of the luminosity-bias curves will deviate from that presented in Sect. 4.2. This is due both to the modified spatial dependence of the field generated by the source bunch, and to the fact that the fraction of particles in the witness bunch that experience a given electromagnetic kick is different from that in the pure-Gaussian case.

Both B*B and COMBI accept as input unperturbed transverse-density distributions, functional forms that are more general than a single Gaussian. In order to assess the potential impact, on beam-beam corrections, of non-Gaussian beams, a realistic model of $\hat{\rho}_B(x, y)$ is needed, or at least one that represents the “worst-case” deviation from the ideal Gaussian shape while remaining representative of actual beam conditions during vdM -calibration sessions at the LHC (Sect. 4.3.1). Taking into account potentially non-factorizable unperturbed density distributions requires a minor generalization of the beam-beam correction formalism (Sect. 4.3.2). While the influence of non-Gaussian tails on the beam-separation dependence of the luminosity-bias factors appears sizeable, the resulting bias on the beam-beam corrections to the visible cross-section remains moderate enough to be treated as a systematic uncertainty (Sect. 4.3.3).

4.3.1 Single-bunch models

In a hadron collider, transverse-density distributions cannot be calculated from first principles, and existing beam-profile monitors, such as wire scanners or synchrotron-light telescopes, have too limited a dynamic range to provide sufficiently precise beam-tail measurements. The only remaining experimental handle, therefore, is that provided by non-factorization analyses such as those described in, for instance, Refs. [9, 22, 46–48]. In this approach, the bunch-density distributions are modeled by the sum of two or three three-dimensional Gaussians. A simultaneous fit to the beam-separation dependence, during a νdM -scan pair, not only of the luminosity but also of the position, size, shape and orientation of the luminous region, makes it possible to estimate the single-bunch parameters of each colliding-bunch pair.

In order to quantify the largest plausible impact of non-Gaussian beam profiles on the beam–beam corrections calculated in Sect. 4.2, single-bunch parameters were extracted from the results of non-factorization analyses of a 2012, a 2017 and a 2018 νdM session. The former is representative of Run-1 νdM conditions at the ATLAS IP, up to and including July 2012 [22], the latter two of Run-2 νdM sessions at the CMS IP, for which “bunch tailoring” in the injector chain [22, 49] significantly reduced non-Gaussian tails. In both cases, a “worst-case” parameter set (in terms of deviations from a perfectly Gaussian shape) was selected from the fitted parameters of the analyzed colliding-bunch pairs. In order to separate the impact of non-Gaussian profiles discussed in this section, from that of elliptical beams (Sect. 4.4) and of beam–beam imbalance (Sect. 4.7), as well as to maximise the sensitivity of the study, in each parameter set the most non-Gaussian transverse profile is assigned to both planes and both beams.

In a first step, the particle density distribution input to B*B is chosen to be factorizable by construction, and modeled by the product of two uncorrelated double Gaussians:

$$\begin{aligned} \hat{\rho}_B(x, y) &= \rho_{xB}(x) \rho_{yB}(y) \\ &= \frac{1}{\sqrt{2\pi}} \left[\frac{w_n}{\sigma_n} \exp\left(-\frac{x^2}{2\sigma_n^2}\right) + \frac{w_w}{\sigma_w} \exp\left(-\frac{x^2}{2\sigma_w^2}\right) \right] \\ &\quad \times \frac{1}{\sqrt{2\pi}} \left[\frac{w_n}{\sigma_n} \exp\left(-\frac{y^2}{2\sigma_n^2}\right) + \frac{w_w}{\sigma_w} \exp\left(-\frac{y^2}{2\sigma_w^2}\right) \right] \end{aligned} \quad (37)$$

where $B = 1$ or 2 , the labels “ n ” and “ w ” refer to the narrow and wide components of the distribution, and their relative population is constrained by $w_w = 1 - w_n$. The functional form reflects the assumptions of unperturbed round beams and of equal particle-density distributions ($\hat{\rho}_1(x, y) = \hat{\rho}_2(x, y)$). The three parameter sets input to the B*B simulation are listed in Table 2. The 2012 (2017) parameter set

clearly results in the most (the least) non-Gaussian shape, as evidenced qualitatively by the combination of the largest (smallest) beam-size ratio σ_w/σ_n and the largest (smallest) weight of the wide component w_w , and as quantified by the single-beam kurtosis computed from the parameters listed in the Table.¹⁰

The functional form chosen for $\hat{\rho}_B(x, y)$ makes it possible to compute analytically the unperturbed convolved transverse beam size Σ^0 from the values of σ_n , σ_w , w_n and w_w , using Eq. (4): these are listed in the top half of Table 3 (rows 3 to 6). The Gaussian-equivalent single-beam size σ_{SG}^0 , i.e. the transverse R.M.S. width of perfectly Gaussian bunches that would yield the same unperturbed convolved beam size as the non-Gaussian bunches studied here, is given by $\sigma_{SG}^0 = \Sigma^0/\sqrt{2}$ (Eq. (6)). These definitions naturally lead to using the “Gaussian-equivalent” beam–beam parameter ξ , inferred from σ_{SG}^0 using Eq. (13), as the common metric in which to quantify the impact of non-Gaussian tails.

4.3.2 Beam–beam-correction formalism in the presence of non-factorization

In a second step, and in order to explore the potential impact of non-factorization on beam–beam corrections, the simplest two-dimensional, non-factorizable double Gaussian:

$$\begin{aligned} \hat{\rho}_B(x, y) &= \frac{w_n}{2\pi\sigma_n^2} \exp\left(-\frac{x^2 + y^2}{2\sigma_n^2}\right) \\ &\quad + \frac{w_w}{2\pi\sigma_w^2} \exp\left(-\frac{x^2 + y^2}{2\sigma_w^2}\right) \end{aligned} \quad (38)$$

is chosen as an alternative density-distribution model to be input to B*B. Since by construction such a distribution cannot be expressed as the product of uncorrelated horizontal and vertical components, the concept of one-dimensional convolved beam size defined by Eq. (4) ceases to be meaningful, and only the convolved–beam-size product $[\Sigma_x \Sigma_y]_{2D}$ introduced in Eq. (8) remains amenable to physical interpretation. Accordingly, the two beam-size bias factors Σ_i/Σ_i^0 ($i = x, y$) introduced in Eqs. (32)–(34) must be replaced by the single quantity

$$\begin{aligned} \frac{[\Sigma_x \Sigma_y]_{2D}}{[\Sigma_x \Sigma_y]_{2D}^0} &= \frac{\int \mathcal{L}(\Delta_x, \Delta_y) d\Delta_x d\Delta_y}{\int \mathcal{L}^0(\Delta_x, \Delta_y) d\Delta_x d\Delta_y} \\ &\quad \times \frac{\mathcal{L}^0(0, 0)}{\mathcal{L}(0, 0)}, \end{aligned} \quad (39)$$

¹⁰ Deviations of a distribution from the strictly Gaussian shape can be characterized by its kurtosis. This statistic is defined as $\Gamma_2 = m_4/m_2^2 - 3$, where m_4 and m_2 are, respectively, the fourth and the second moment of the distribution [50]. The kurtosis is zero for a Gaussian, positive for a leptokurtic distribution with longer tails, and negative for a platykurtic distribution with tails that fall off more quickly than those of a Gaussian.

Table 2 Single-bunch parameters representative of strongly non-Gaussian single-beam profiles during ATLAS and CMS pp vdM scans in 2012, 2018 and 2017. The bunches are assumed to be round and, within each parameter set separately, beam–beam symmetric, i.e. of equal brightness for beam 1 and beam 2

Date of vdM session LHC fill number	July 2012 2855	June 2018 6868	July 2017 6016
E_B [TeV]	4.0	6.5	6.5
n_p [10^{11} p/bunch]	1.1	0.85	0.84
β^* [m]	11.0	19.2	19.2
q_x, q_y	0.31, 0.32	0.31, 0.32	0.31, 0.32
σ_n [μm]	57.9	85.0	84.0
σ_w [μm]	115	125	116
σ_w/σ_n	1.99	1.47	1.38
w_n	0.634	0.670	0.840
$w_w = 1 - w_n$	0.366	0.330	0.160
Single-beam kurtosis Γ_2	1.40	0.47	0.25

Table 3 Scan parameters representative of strongly non-Gaussian beam profiles during ATLAS and CMS pp vdM scans in 2012, 2018 and 2017, calculated analytically from the single-bunch parameters listed in Table 2

Date of vdM session Single-beam kurtosis	July 2012 1.40	June 2018 0.47	July 2017 0.25
Factorizable density distribution: Eq. (37)			
Σ^0 (from Eq. (4)) [μm]	107.1	137.3	125.4
$\sigma_{\text{SG}}^0 = \Sigma^0/\sqrt{2}$ [μm]	75.7	97.1	88.7
Gaussian-equivalent ξ [10^{-3}]	6.04	3.04	3.61
Non-factorizable density distribution: Eq. (38)			
Σ_F^0 (from Eq. (41)) [μm]	100.4	134.9	124.4
$\sigma_{\text{SG}}^0 = \Sigma_F^0/\sqrt{2}$ [μm]	71.0	95.4	88.0
Gaussian-equivalent ξ [10^{-3}]	6.88	3.16	3.67

where the scan and the integration are carried out over a grid in the two-dimensional nominal-separation space (Δ_x, Δ_y) that extends over $\pm 5 \sigma_{\text{SG}}^0$ in both directions. The σ_{vis} -bias factor defined by Eq. (35) then takes the more general form

$$\sigma_{\text{vis}2D}/\sigma_{\text{vis}2D}^0 = \frac{\mu_{pk}}{\mu_{pk}^0} \times \frac{[\Sigma_x \Sigma_y]_{2D}}{[\Sigma_x \Sigma_y]_{2D}^0}, \tag{40}$$

where the subscript “2D” indicates that the two-dimensional formalism introduced in Eq. (8) is used throughout the analysis.

Using ξ as a common metric for both the factorizable and the non-factorizable case also requires a generalization of the formula that expresses, in terms of single-beam parameters, an observable that is akin to the transverse convolved bunch size and that can be measured in a standard, one-dimensional vdM scan. For instance, for a vertical scan with the beams centered on each other in the horizontal plane, the measured apparent convolved width is given by:

$$\Sigma_{yF} = \frac{1}{\sqrt{2\pi}} \frac{\int R_{x,y}(\delta_x = 0, \delta_y) d\delta_y}{R_{x,y}(\delta_x = 0, \delta_y = 0)}, \tag{41}$$

which reduces to Eq. (4) when $\hat{\rho}(x, y) = \rho_x(x) \rho_y(y)$. The subscript “F” emphasizes that while measurable, this observable cannot be used in Eqs. (5) and (10) unless the density distributions are factorizable. The importance of

this distinction is illustrated by the significant differences between the top half (rows 3–6) and the bottom half (rows 7–10) of Table 3: even though both sets of scan parameters are based on Table 2, their values are significantly different in the two cases, because the non-linear x – y correlation built into Eq. (38) affects the beam-overlap integral (Eq. (1)). These differences, the size of which appears correlated with the single-beam kurtosis, are indicative of the magnitude of the non-factorization biases that can be introduced, even in the absence of the beam–beam interaction, by applying the one-dimensional, factorizable FoM formalism of Sect. 4.2.4 to non-factorizable beams. In contrast, the analysis strategy encapsulated in Eq. (40) guarantees that non-factorizable correlations embedded in the unperturbed density distributions cancel in the $\sigma_{\text{vis}2D}$ ratio; the latter, therefore, reflects the impact of beam–beam biases only. The study of non-factorization biases proper lies beyond the scope of the present report.

4.3.3 Impact of non-Gaussian tails on beam–beam-induced biases

In the case of factorizable density distributions, the impact of non-Gaussian tails on beam–beam corrections can be




Fig. 17 Beam-separation dependence, during simulated vertical vdM scans, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ associated with the deflection-induced orbit shift, for single-Gaussian (red) and factorizable double-Gaussian (black) bunch-density distributions that yield the same unperturbed convolved beam size. The horizontal axis is the nominal vertical beam separation. The bunch parameters are listed in the second column of Table 2 (LHC fill 2855). The beams satisfy the assumptions listed in Sect. 4.1, except (in the case of the black curve) that regarding the single-Gaussian transverse profile

visualized by comparing the beam-separation dependence of the luminosity-bias factors $[\mathcal{L}/\mathcal{L}^0]$ for a given set of non-Gaussian beam parameters, to that computed for perfectly Gaussian beams that in the absence of beam–beam effects, would yield the same value of the unperturbed convolved beam size Σ^0 , and therefore the same head-on luminosity. This is illustrated in Fig. 17 for the orbit-shift effect. In the single-Gaussian case (red curve), the steeper drop of $[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$ with increasing separation, and its asymptotic behavior for $\Delta_y \geq 300 \mu\text{m}$, indicates that at such large separations ($\Delta_y > 4\sigma_{\text{SG}}^0$), the number of particles contributing to the source field becomes negligible, so that the latter increasingly appears to the witness bunch as a point source. In the non-Gaussian case (black curve), the charge distribution is more spread out, and the interplay between the source-field distribution and the density distribution in the witness bunch produces a less intuitive separation dependence.

The situation is reversed for optical distortions (Fig. 18): the Gaussian beams are subject to larger optical distortions at large beam separation (red curve), presumably because in such configurations the tail particles in the witness bunch experience a more intense source-field gradient when close to the center of the source bunch. The tails of the non-Gaussian beams, being subject to a more diffuse source-field distribution, experience weaker gradients and correspondingly smaller optical distortions (black curve).

In the case of non-factorizable distributions, a physically intuitive graphical representation of beam–beam biases becomes impractical, since $[\mathcal{L}/\mathcal{L}^0]$ depends on both Δ_x and

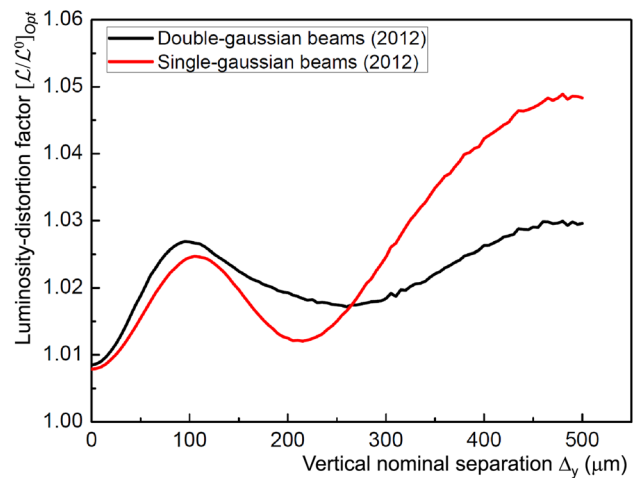


Fig. 18 Beam-separation dependence, during simulated vertical vdM scans, of the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ associated with optical distortions, for single-Gaussian (red) and factorizable double-Gaussian (black) bunch-density distributions that yield the same unperturbed convolved beam size. The horizontal axis is the nominal vertical beam separation. The bunch parameters are listed in the second column of Table 2 (LHC fill 2855). The beams satisfy the assumptions listed in Sect. 4.1, except (in the case of the black curve) that regarding the single-Gaussian transverse profile

Δ_y , and the curves in Figs. 17 and 18 must be replaced by surfaces.

The impact of non-Gaussian tails on the corresponding vdM calibrations is best summarized by the σ_{vis} -bias factors presented in Table 4. These are all computed using Eqs. (39)–(40), in order to ensure that factorizable and non-factorizable models are treated identically; it has been verified that for the factorizable models, the one-dimensional FoM formalism detailed in Sect. 4.2.4 yields the same results within 0.01% of σ_{vis} , well within the expected accuracy of the overlap-integral calculation (Sect. 3.3). Listed in this table is the full beam–beam bias for either factorizable or non-factorizable double-Gaussian models, for the corresponding single-Gaussian model (that is factorizable by construction), as well as for their difference DG–SG. This difference is also listed after scaling it linearly to a common, arbitrary reference value of $\xi = 5.6 \cdot 10^{-3}$, that lies at the upper end of the beam–beam parameter range covered by Run-2 vdM scans. This simple-minded procedure, inspired by the roughly linear dependence of the σ_{vis} bias factor on ξ apparent in Fig. 16, allows direct comparisons across configurations corresponding to different values of ξ .

Globally, non-Gaussian tails modify the σ_{vis} bias by an amount that – for the models chosen here – varies from less than 10% to about half of the bias itself. Predictably enough, the larger the single-beam kurtosis, the larger the relative impact of the tails. The non-factorizable models seemingly predict a smaller influence of the tails than their factorizable counterpart; this may however depend on the particular form

Table 4 Visible cross-section bias associated with the full beam–beam effect, calculated using B*B, for double- (DG) and single- (SG) Gaussian beams yielding the same Gaussian-equivalent beam–beam parameter ξ , and difference between the two (DG–SG). The unperturbed DG transverse-density distributions are described either by a factorizable or by a non-factorizable function; their parameters are listed in Table 2. The bunches are assumed to collide at the scanning IP only

Date of νdM session		July 2012	June 2018	July 2017
Single-beam kurtosis		1.40	0.47	0.25
Factorizable density distribution: Eq. (37)				
ξ [10^{-3}]		6.03	3.05	3.61
σ_{vis2D} bias [%]:	DG	−0.27	−0.15	−0.25
	SG	−0.57	−0.22	−0.27
	DG–SG	0.30	0.07	0.02
DG–SG scaled to $\xi = 5.6 \cdot 10^{-3}$		0.28	0.13	0.03
Non-factorizable density distribution: Eq. (38)				
ξ [10^{-3}]		6.88	3.16	3.67
σ_{vis2D} bias [%]:	DG	−0.50	−0.17	−0.23
	SG	−0.64	−0.21	−0.25
	DG–SG	0.14	0.04	0.02
DG–SG scaled to $\xi = 5.6 \cdot 10^{-3}$		0.11	0.07	0.03

of x – y correlation embedded into the specific functional form chosen for $\hat{\rho}_B(x, y)$. A practical proposal to translate the results listed in Table 4 into a realistic systematic uncertainty is offered in Sect. 5.2.1. Repeating the above studies using more refined single-beam models extracted from future νdM sessions may help refine the understanding of such biases.

4.4 Impact of beam ellipticity at the interaction point

In the fully symmetric beam configuration (Sect. 4.1), the beams are round at the IP, and the beam–beam interaction can be described by a single parameter ξ (Eq. (13)). In the more general case where the transverse beam profile is an upright¹¹ ellipse, for instance because of unequal horizontal and vertical emittances or β^* values, the horizontal and vertical beam–beam parameters ξ_x and ξ_y (Eq. (12)) no longer coincide, and the parametric approach described in Sect. 4.2.3 is, strictly speaking, no longer applicable.

During νdM scans at the LHC, however, the round-beam hypothesis is only moderately violated: the measured horizontal and vertical IP- β functions typically differ by 10% or less, the corresponding emittances by 30% at most, and the increase in projected transverse beam size associated with a non-zero crossing angle at the ALICE and LHCb IPs does not exceed 15%.

To assess the impact of beam ellipticity on beam–beam corrections, B*B is used to simulate νdM scans over a square grid in (Σ_x^0, Σ_y^0) space (Fig. 19), with as input unperturbed beams that satisfy the assumptions listed in Sect. 4.1 except for the round-beam hypothesis. The chosen Σ_x^0 range is

slightly wider than that observed during Run-2 νdM scans at $\sqrt{s} = 13$ TeV; the vertical/horizontal aspect ratio Σ_y^0/Σ_x^0 is varied from 0.7 to 1.4, corresponding to sampling the corresponding emittance ratio from 0.5 to 2.0, well beyond the range observed during νdM -calibration sessions. The horizontal and vertical beam–beam parameters cover the range $2.6 \times 10^{-3} < \xi_x, \xi_y < 5.1 \times 10^{-3}$ (Fig. 20). The grid is chosen such that for a constant value of the round-beam equivalent beam–beam parameter ξ_R , the horizontal and vertical beam–beam parameters sample the full span allowed by the range of chosen Σ_x^0 and aspect-ratio values. For example, the group of five points at $\xi_R \approx 3.6 \times 10^{-3}$ in Fig. 20 corresponds to the five diagonal points in Fig. 19; the residual curvature is due to the fact that the grid is built from equally spaced (Σ_x^0, Σ_y^0) , rather than (ξ_x, ξ_y) , values.

The resulting peak-rate, beam-size and visible cross-section bias factors (Eqs. (31)–(35)) for elliptical configurations are then compared to the same quantities calculated using the round-beam parameterization (Eq. (36)) at the corresponding value of ξ_R . The results are illustrated in Fig. 21 for the visible cross-section. On the scale shown here, the bias $(\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1)$, and therefore the corresponding beam–beam correction, seem to depend on ξ_R only; stated differently, the beam–beam-induced biases appear rather insensitive to the aspect ratio. Upon closer inspection however, all four FoMs do reveal an actual dependence on the Σ_y^0/Σ_x^0 aspect ratio, albeit a weak one. Since the orbit-shift effect can be accurately calculated, even for elliptical beams, using the Bassetti–Erskine formalism [14], only optical-distortion biases matter. The difference between the bias calculated by B*B for a given set of elliptical-beam parameters, and that extracted from the round-beam parameterization using the corresponding value of ξ_R , is displayed in Fig. 22. Its dependence on the vertical/horizontal aspect ratio is close to linear, and remains below 0.03% over the full range of

¹¹ In view of the minor impact predicted by the simulations described in this section for upright (horizontal or vertical) elliptical beams, and of the small residual transverse coupling measured at the LHC IPs, the fully general case of beam ellipses tilted in the x - y plane is not considered in the present paper.

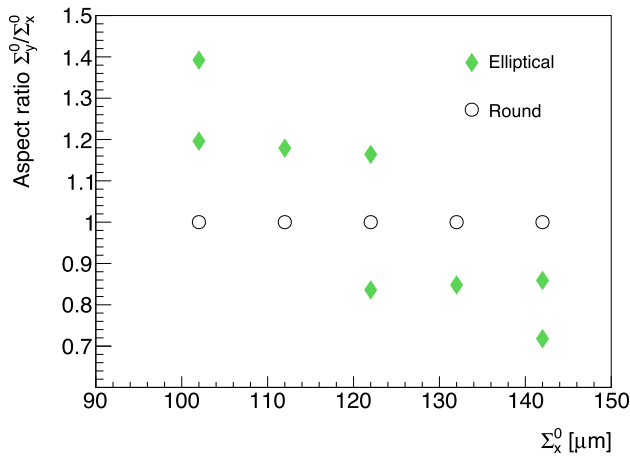


Fig. 19 Grid of unperturbed transverse convolved beam sizes used to characterize the impact of beam ellipticity on beam–beam corrections, and their vertical/horizontal aspect ratio. The diamonds (circles) correspond to elliptical (round) beams. The beam energy, bunch intensity, IP- β function and unperturbed tunes are fixed at, respectively, $E_B = 6.5$ TeV, $n_p = 0.777 \times 10^{11}$ protons/bunch, $\beta^* = 19.2$ m and $(Q_x, Q_y) = (64.31, 59.32)$

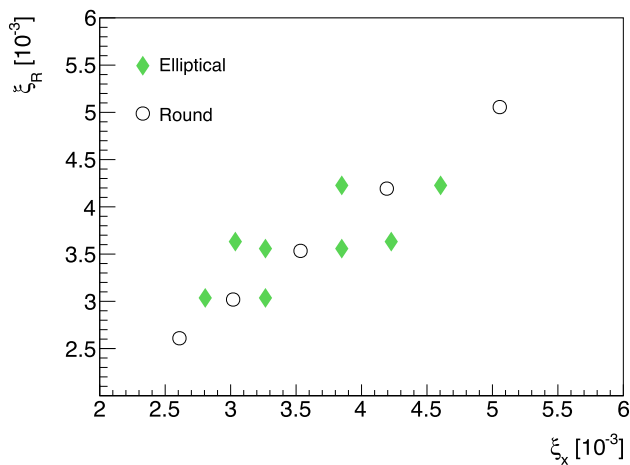


Fig. 20 Grid of beam–beam parameter values used to characterize the impact of beam ellipticity on beam–beam corrections. The diamonds (circles) correspond to elliptical (round) beams. The horizontal axis shows the value of ξ_x defined by Eq. (12), the vertical axis the corresponding round-beam equivalent parameter ξ_R defined by Eq. (29)

beam parameters considered in this study. Its dependence on ξ_R is illustrated by the spread over the three parameter sets clustered at $\Sigma_x^0/\Sigma_y^0 \approx 0.8$ and 1.2 . The beam–beam parameter values considered in this study, and that are associated with elliptical-beam configurations, span the range $3.0 \times 10^{-3} < \xi_R < 4.2 \times 10^{-3}$: this is representative of, but does not quite cover, the full ξ range listed in Table 1. Even if the mean value of ξ_R was doubled, however, the impact of potentially elliptical beam profiles on the beam–beam correction to σ_{vis} would still remain well below 1 permil, and can therefore be treated as a systematic uncertainty.




Fig. 21 Dependence of the visible cross-section bias on the round-beam equivalent beam–beam parameter ξ_R , for optical distortions only (red), orbit shift only (green), and full beam–beam effect (black). The filled (open) markers tag the results of individual B*B simulations assuming upright elliptical (round) beams; the orbit-shift bias is calculated analytically using the elliptical-beam formula [14]. The curves show the round-beam estimates based on the parameterization of Eqs. (35)–(36). The points clustered around each ξ_R value correspond to different settings of the Σ_y^0/Σ_x^0 aspect ratio




Fig. 22 Aspect-ratio dependence of the difference in σ_{vis} bias associated with optical distortions, between the result of B*B elliptical-beam simulations, and the corresponding round-beam estimate based on the parameterization of Eqs. (35)–(36). The filled (open) markers tag the results of individual simulations assuming upright elliptical (round) beams. The points clustered around each value of the aspect ratio correspond to different ξ_R settings. The horizontal dashed line corresponds to no difference in σ_{vis} bias between round and elliptical beams. The spread of the round-beam simulation results (open triangles) reflects the numerical accuracy of the subtracted parameterization

4.5 Impact of non-zero crossing angle

In the presence of a significant crossing angle, the analysis becomes considerably more involved, and a full exposition thereof would exceed the scope of the present report. The discussion below is therefore focussed on assessing the beam–

beam correction uncertainty associated with vdM scans at zero nominal crossing angle, during which the actual, residual crossing angle does not exceed few tens of microradians: this covers most luminosity-calibration sessions at IP1 (ATLAS) and IP5 (CMS). The same treatment, albeit at the cost of a small additional uncertainty, remains applicable to scans at IP2 (ALICE) (and to some non-standard sessions at IPs 1 and 5) that involve nominal crossing angles θ_c in the 140–400 μrad range. At IP8 however, the LHCb requirement for an unambiguous separation, in the crossing plane, of beam-gas vertices from beam 1 and beam 2, leads to θ_c settings of up to 1100 μrad . Such large crossing angles would require simulation studies more extensive than can be reported here.

In the absence of any beam–beam interaction, the geometric factor F defined in Eq. (11), and generalizations thereof [16, 19], provide an accurate description of the combined impact of a non-zero crossing angle and of a finite bunch length on the unperturbed apparent transverse beam size in the crossing plane, as well as of the beam-separation dependence of the unperturbed luminosity.

In the presence of the beam–beam interaction, the implications of a non-zero crossing-angle for the calculation of beam–beam biases fall in two conceptually distinct categories.

- *Geometrical effects in the transverse plane* that depend on the crossing angle and on the bunch length, such as modifications of the transverse beam–beam overlap and the calculation of the effective strength of the source field experienced by the particles in the witness beam, are adequately modeled [16] by the B*B package and well reproduced by COMBI simulations. B*B simulations, however, are intrinsically limited to four dimensions (horizontal and vertical positions and angles), since the longitudinal variables are “projected out” and the associated dynamics implicitly neglected.
- *Longitudinal dynamics*, such as synchrotron motion or beam–beam-induced transverse-shape variations along the bunch, can be modeled by COMBI using the 6D beam–beam strong-strong model [41]. This model divides particle distributions into slices along the length of each bunch. The kick experienced by each macroparticle in one beam is computed based on the statistical moments of the charge distributions in each slice of the opposing bunch; the computation is repeated for the macroparticles in the other beam.

In order to properly account for longitudinal effects and assess their importance, COMBI is chosen here to characterize the interplay of crossing-angle and beam–beam effects. This first study is exploratory in nature and restricted to the vdM regime.

As predicted analytically [16, 51], the magnitude of the orbit shift extracted from COMBI simulations for a purely horizontal crossing angle of several hundred microradians, is found consistent with that computed using the Bassetti–Erskine formula with, as input:

- in the crossing (horizontal) plane, the effective convolved beam size given by $\Sigma_{x,eff} = F \times \Sigma_x$, where F is the geometric factor defined in Eq. (11) and Σ_x is the horizontal convolved beam size at zero crossing angle, given by Eq. (6);
- in the non-crossing (vertical) plane, the vertical convolved beam size Σ_y .

The dependence of the beam–beam-induced visible cross-section bias on the full crossing angle is illustrated in Fig. 23 for a beam–beam parameter setting typical of Run-2 vdM scans. For the configurations shown in this figure, the geometric factor spans the range $1.0 < F < 1.09$. Since this factor impacts both σ_{vis}^0 and σ_{vis} in an almost identical manner,¹² it cancels in the crossing-angle dependence of the ratio $\sigma_{vis}/\sigma_{vis}^0$. The absolute value of the σ_{vis} bias increases with the angle, even though the “effective” beam–beam parameter in the crossing plane (i.e. that calculated from the effective transverse beam size) decreases. The difference between the horizontal- and vertical-crossing configurations is due to the difference between the x and y tunes. Also notable is the sensitivity to β^* : the crossing-angle dependence of the beam–beam bias is weaker for $\beta^* = 24\text{ m}$ than for $\beta^* = 19\text{ m}$. The results are insensitive to the assumed value of the synchrotron tune, as well as to the chromaticity of the ring lattice (reducing the chromaticity to zero from its nominal setting of 15 units impacts the results by less than 0.01%). This suggests that the effect is mostly geometrical in nature, as confirmed by the qualitative agreement of the COMBI results with those of a preliminary B*B study that covered a similar parameter space.

In practical terms, the impact of the crossing angle on the estimated beam–beam correction can be taken as the difference in visible cross-section bias between $\theta_c = 0$ and the actual crossing-angle value. For a nominally zero crossing angle, the actual angle, as measured by the BPM system or by K-modulation in the final-triplet quadrupoles, is of the order of 10 μrad : the associated shift in the beam–beam bias, as read off Fig. 23 at $\beta^* = 19\text{ m}$, does not exceed 0.01%. For a nominal crossing angle around 150 μrad , which is typical of ALICE vdM sessions, neglecting the crossing angle would underestimate the beam–beam correction by less than

¹² In the ratio $\sigma_{vis}/\sigma_{vis}^0$, both numerator and denominator include a geometric factor. In the case of σ_{vis} , the transverse convolved beam size is in addition modified by the dynamic- β effect, resulting in a 0.07% increase in F in the worst case ($\beta^* = 19\text{ m}$, $\theta_c = 1000\text{ }\mu\text{rad}$).

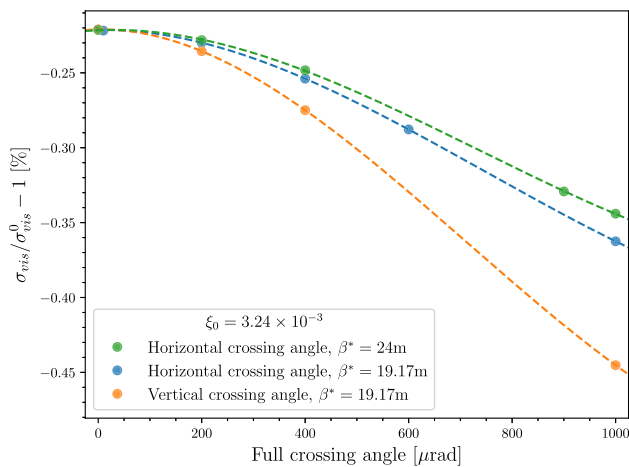


Fig. 23 Crossing-angle dependence of the visible cross-section bias predicted by COMBI strong-strong, 6-dimensional simulations, for a horizontal (blue) and a vertical (orange) crossing angle at $\beta^* = 19.17$ m, and for a horizontal crossing angle at $\beta^* = 24$ m (green). The beam-beam parameter value at zero crossing angle is indicated in the legend. The beam energy, bunch intensity, normalized emittance, bunch length and tunes are fixed at, respectively, $E_B = 6.5$ TeV, $n_p = 0.783 \times 10^{11}$ p/bunch, $\epsilon_N = 2.95 \mu\text{m}\cdot\text{rad}$, $\sigma_z = 7.5$ cm, $Q_x = 64.31$, $Q_y = 59.32$, $Q_s = 0.0023$. The bunches are assumed to collide only at the IP where the vdM scans are performed. The points are the results of the simulation; the lines are meant to guide the eye

0.02%, a value small enough to be neglected. Crossing angles in excess of $500 \mu\text{rad}$, however, induce still-to-be understood biases of 0.1% or more (Fig. 23), and may require further study.

4.6 Impact of multiple interaction points

4.6.1 Motivation and methodology

All results presented in this paper so far assume that the beams collide at the scanning IP only, i.e. that they are fully separated, either transversely or longitudinally, at the other three IPs of the LHC. In practice however, this is often not the case: to make optimal use of beam time during luminosity-calibration scans at, say, the CMS IP, the experimental detectors at some or all of the other IPs typically collect collision data at fixed (most often zero) beam separation, thereby significantly affecting the unperturbed-tune spectra and therefore the magnitude of beam-beam induced biases.

For a given colliding-bunch pair at a given scanning IP, the *number of non-scanning interaction points* N_{NSIP} is defined as the average number of head-on collisions experienced, at interaction points *other* than the scanning IP, by the two members of the pair. The parameter N_{NSIP} ranges from zero when beams are fully separated at all but the scanning IP, to $N_{\text{NSIP}} = 3$ when both bunches in the pair collide at three other IPs; it can assume half-integer values when the two




Fig. 24 Schematic layout of the LHC beams and interaction points. Interaction points 1, 2, 5 and 8 house, respectively, the ATLAS, ALICE, CMS and LHCb experiments. (Figure reproduced from Ref. [30], © CERN)

members of the pair experience different numbers of additional collisions.

Because IP1 (ATLAS) and IP5 (CMS) are separated by half a ring circumference, all bunches that collide in IP1 also collide in IP5, unless fully separated in the transverse plane at the non-scanning IP. Each member of such a colliding-bunch pair can optionally also collide in IP2 (ALICE) and/or IP8 (LHCb), both of which are located $1/8$ th of the ring away¹³ from IP1; however its collision “partner” at these IPs cannot be the same as that at IPs 1 and 5, because of the ring layout (Fig. 24). Injected-bunch patterns are optimized, on a fill-by-fill basis, to achieve the desired sharing of collisions (or absence thereof) among the four experimental detectors.

The parameter N_{NSIP} , therefore, depends on the scanning IP considered, and, for a given scanning IP, can vary along the bunch string in a manner that depends on the injected-bunch pattern. From 2010 to 2017, the most frequent configuration during pp vdM scans at the ATLAS and CMS IPs corresponded to $N_{\text{NSIP}} = 1$ at IPs 1 and 5, i.e. to bunches that collide in IP1 and IP5 only; from 2018 onwards, all four scanning IPs had to deal with multiple non-scanning IPs, for reasons of operational efficiency.

In order to characterize the impact of multiple collision points on beam-beam corrections to luminosity calibrations, COMBI was used to simulate vdM scans at one IP (e.g. IP1), with additional head-on collision points optionally inserted at one or more other IPs (in this example IP5, IP2 and/or

¹³ Up to a longitudinal shift of 11.24 m (30 RF buckets) at IP8, a requirement driven by the longitudinal asymmetry of the LHCb detector.

Table 5 Beam conditions assumed for the multi-IP simulations described in Sect. 4.6. The crossing angle θ_c is set to 0 at all IPs so as to fully decouple the multi-IP study from potential longitudinal-dynamics effects

Parameter	Value
Beam energy E_B [TeV]	6.5
Typical bunch population n_p [10^{11}]	0.78
Bunches colliding	1
ϵ_N [$\mu\text{m-rad}$]	2.9
IP1,5: β^* [m]	19.2
IP2/8: β^* [m]	19.2/24.0
θ_c [μrad]	0
$\Sigma_x^0 = \Sigma_y^0$ [μm]	128
Nominal tune settings Q_x/Q_y	64.31/59.32
Typical ξ [per IP]	0.0032

IP8). The bunch parameters are listed in Table 5. The beam transport between consecutive IPs is modeled by linear maps that reflect the nominal phase advance in the LHC rings during the Run-2 vdM sessions at $\sqrt{s} = 13$ TeV (Table 6), and head-on collisions are simulated at each non-scanning IP as described in Sect. 3.4. The unperturbed beams satisfy the assumptions listed in Sect. 4.1, except for the absence of additional collisions.

This section is organized as follows. Multi-IP collisions increase the beam–beam tune spread, thereby shifting the means of the unperturbed-tune spectra to lower values (Sect. 4.6.2). They also affect the peak-rate and beam-size bias factors in a manner that depends on the phase advance between consecutive collision points around the rings, but that leaves the σ_{vis} bias almost totally insensitive to these phase-advance values. The analysis can be greatly simplified by an astute choice of the reference “no beam–beam” configuration (Sect. 4.6.3), up to a small residual ambiguity that depends on the phase advance between the non-scanning IP(s) and that where the scans are taking place (Sect. 4.6.4). This approach naturally leads to an effective parameterization of the impact of multi-IP collisions on beam–beam corrections to vdM calibrations (Sect. 4.6.5). Candidate beam–beam correction strategies in the presence of multi-IP collisions are summarized in Sect. 4.6.6.

4.6.2 Impact of multiple interaction points on the tune spectra

A comparison of simulated vertical-tune spectra for $N_{\text{NSIP}} = 0$ and $N_{\text{NSIP}} = 1$ is presented in Fig. 25. The single-IP beam–beam parameter ξ is chosen to be typical of vdM scans in pp collisions. For beams colliding at the scanning IP only ($N_{\text{NSIP}} = 0$), two narrow peaks appear: one near the nominal tune ($q_y \approx 0.320$), and the other near $q_y - Y * \xi$,

where $Y \approx 1.1$ is the Yokoya factor in the soft-Gaussian approximation [52]. These peaks correspond to, respectively, the so-called σ and π modes of coherent bunch oscillation. When beams collide in addition at one non-scanning IP ($N_{\text{NSIP}} = 1$), the peak of the σ mode is unaffected, but the beam–beam tune spread roughly doubles and the π mode is shifted further down, near $q_y - 2Y * \xi$. In both cases, the mean tunes lie roughly half-way between the σ and π peaks, and therefore differ by roughly $-\xi/2$ between the two configurations. Adding a second non-scanning IP results in a further widening of the tune spread, to roughly three times that for $N_{\text{NSIP}} = 0$, and in a correspondingly larger downward shift of the mean tune.

4.6.3 Comparison of reference “no beam–beam” configurations in the presence of one non-scanning IP

In Sects. 2, 3 and 4.1 to 4.5, it is implicitly assumed that the beams collide at the scanning IP only. The beam–beam-induced peak-rate, beam-size and luminosity-bias factors defined in Sect. 4.2.4 are all expressed with respect to an unambiguous reference configuration that corresponds to “no beam–beam anywhere”, in which the “nominal” reference variables ($\sigma_0, \mu_{pk}^0, \Sigma_x^0, \Sigma_y^0, \mathcal{L}^0, \sigma_{\text{vis}}^0, \dots$) are identified by a “0” index.

In the presence of collisions at non-scanning IPs, however, an ambiguity arises: should beam–beam-induced biases be referenced to a configuration in which the beam–beam interaction is switched off *at every single IP*, or to one in which the beam–beam interaction is switched off *at the scanning IP only*, but remains active at the non-scanning IPs (if any)?

Confronting this question requires a more precise nomenclature than that used so far:

- when the reference configuration is one in which the beam–beam interaction is switched off at all IPs, i.e. if vdM observables are corrected to “no beam–beam anywhere”, the corresponding reference variables carry a “0” index as before. In particular, the reference luminosity $\mathcal{L}_{\text{no-bb}}$ at the scanning IP is equal to the nominal luminosity defined in Sect. 2.3.2 ($\mathcal{L}_{\text{no-bb}} = \mathcal{L}^0$). This is referred to below as the “ \mathcal{L}^0 normalization”;
- when the reference configuration is one in which vdM observables are corrected to “no beam–beam at the scanning IP only” even when the bunches considered also collide at one or more other IPs, the corresponding reference variables carry a “u” index instead. In particular, the reference luminosity $\mathcal{L}_{\text{no-bb}}$ at the scanning IP is given by the *unperturbed luminosity* \mathcal{L}^u ($\mathcal{L}_{\text{no-bb}} = \mathcal{L}^u$), i.e. that which would be measured at the scanning IP if the beam–beam interaction could be turned off at that IP only. This quantity is accessible only in simulations; its beam-separation dependence is in general notably dif-

Table 6 Horizontal (μ_x) and vertical (μ_y) betatron phase advance between consecutive collision points in the LHC Run-2 *vdM* configuration for *pp* collisions at $\sqrt{s} = 13$ TeV. Beam 1 moves from the top to the bottom row, beam 2 from the bottom to the top row

From-to	Beam 1		From-to	Beam 2	
	$\mu_x^{B1} [2\pi]$	$\mu_y^{B1} [2\pi]$		$\mu_x^{B2} [2\pi]$	$\mu_y^{B2} [2\pi]$
IP1–IP5	31.9757	29.6486	IP5–IP1	31.9844	29.7613
IP5–IP1	32.3343	29.6714	IP1–IP5	32.3256	29.5587
IP1–IP2	8.2960	7.6692	IP2–IP1	8.2728	7.9577
IP2–IP5	23.6797	21.9794	IP5–IP2	23.7116	21.8036
IP5–IP8	24.0891	21.3685	IP8–IP5	23.8146	21.9544
IP8–IP1	8.2452	8.3029	IP1–IP8	8.5110	7.6042
Full turn	64.3100	59.3200	Full turn	64.3100	59.3200




Fig. 25 Vertical coherent-tune spectra in collisions simulated using COMBI, for head-on collisions at the scanning IP only ($N_{NSIP} = 0$, blue curve), and in the presence of head-on collisions at one additional IP ($N_{NSIP} = 1$, green curve). The beams satisfy the assumptions listed in Sect. 4.1, except for the absence of additional collisions. The beam-beam parameter ξ is set to 3.24×10^{-3} per collision, and the tunes to their nominal value (Q_x, Q_y) = (64.31, 59.32)

ferent from that of \mathcal{L}^0 . The same notation, and physical meaning, are applied to the other reference variables. This is referred to below as the “ \mathcal{L}^u normalization”;

- in the limiting case of collisions at the scanning IP only ($N_{NSIP} = 0$), \mathcal{L}^u and \mathcal{L}^0 are one and the same, as are all the other reference variables: there is no difference between “nominal” and “unperturbed” quantities.

The beam-separation dependence of the $\mathcal{L}^u/\mathcal{L}^0$ ratio, during a vertical scan at IP1 and in the presence of head-on collisions at IP5, is presented in Fig. 26 (blue lozenges). Even though the unperturbed luminosity \mathcal{L}^u is by definition unaffected by beam-beam effects at the scanning IP, it exhibits a monotonic increase with beam separation when compared to the reference luminosity \mathcal{L}^0 . This seemingly counter-intuitive observation is explained by the fact that the horizontal axis in Fig. 26 is the *nominal* transverse beam size σ^0 , not the *actual* beam size during the scan; equivalently, the

actual relative separation, i.e. that normalized to the actual unperturbed beam size, differs from the nominal relative separation, which is normalized to the nominal beam size. The monotonic increase of the $\mathcal{L}^u/\mathcal{L}^0$ ratio therefore suggests that in this particular configuration, the collisions at IP5 result in an enlarged vertical convolved beam size at IP1 ($\Sigma_y^u > \Sigma_y^0$). This interpretation is confirmed by the evolution, during the scan, of the individual unperturbed single-beam sizes σ_{yB}^u ($B = 1, 2$) that is illustrated in Fig. 27 (open lozenges): both are larger than σ_0 , they remain constant as a function of the beam separation at the scanning IP, and their difference reflects the fact that the vertical phase-advance values from IP1 to IP5 differ significantly between the two beams (Table 6, top two rows).

Also shown in Fig. 26 is the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{Full BB}$ in the \mathcal{L}^0 normalization (red squares), that displays an *s*-like shape superposed on the monotonic beam-separation dependence of the $\mathcal{L}^u/\mathcal{L}^0$ curve. The ratio of these two curves:

$$[\mathcal{L}/\mathcal{L}^0]_{Full BB} / (\mathcal{L}^u/\mathcal{L}^0) = [\mathcal{L}/\mathcal{L}^u]_{Full BB}$$

represents the luminosity-bias factor in the \mathcal{L}^u normalization in the presence of one non-scanning IP ; it is displayed as filled green circles, and is highly reminiscent of the black curve in Fig. 9.

This same luminosity-bias factor that is presented in Fig. 26 for $N_{NSIP} = 1$, is compared in Fig. 28 to the equivalent quantity for $N_{NSIP} = 0$. Up to a small systematic difference that represents the impact of the additional head-on collisions that occur at the non-scanning IP, the two curves are essentially identical. This observation suggests that the beam-separation dependence of the beam-beam-induced orbit-shift and optical-distortion effects is encapsulated in that of $[\mathcal{L}/\mathcal{L}^u]_{Full BB}$, even in the presence of collisions at the non-scanning IP(s). This interpretation is confirmed by comparing the separation dependence of the actual single-beam transverse RMS sizes σ_{iB} ($i = x, y; B = 1, 2$) with, and without, additional collisions at the non-scanning IP. When $N_{NSIP} = 1$ (Fig. 27), $\sigma_{y1} > \sigma_{y2}$; these two quantities exhibit a very similar beam-separation dependence with




Fig. 26 Beam-separation dependence, during a simulated vertical *vdM* scan in the Run-2 configuration with one non-scanning IP, of the luminosity-bias factors at the scanning IP associated with the unperturbed luminosity ($\mathcal{L}^u/\mathcal{L}^0$, blue lozenges), with the full beam–beam effect in the \mathcal{L}^0 normalization ($[\mathcal{L}/\mathcal{L}^0]_{\text{Full BB}}$, red squares), and with the full beam–beam effect in the \mathcal{L}^u normalization ($[\mathcal{L}/\mathcal{L}^u]_{\text{Full BB}}$, filled green circles). The horizontal axis is the beam separation in units of the nominal transverse beam size σ_0




Fig. 28 Beam-separation dependence, during a simulated vertical *vdM* scan, of the luminosity-bias factors at the scanning IP associated with the full beam–beam effect, for $N_{\text{NSIP}} = 0$ using the \mathcal{L}^0 normalization (open circles), and for $N_{\text{NSIP}} = 1$ using the \mathcal{L}^u normalization (filled green circles). The horizontal axis is the beam separation in units of the nominal transverse beam size σ_0 . The span of the vertical scale is identical to that of Fig. 26




Fig. 27 Beam-separation dependence, during a simulated vertical *vdM* scan, of the beam–size-bias factors at the scanning IP associated with the full beam–beam effect (σ_{yB}/σ^0 , filled squares) and of the corresponding unperturbed quantities (σ_{yB}^u/σ^0 , open lozenges), separately for B = beam 1 (blue) and B = beam 2 (red), using the \mathcal{L}^0 normalization and in the presence of collisions at one non-scanning IP ($N_{\text{NSIP}} = 1$). The horizontal axis is the beam separation in units of the nominal transverse beam size σ_0




Fig. 29 Beam-separation dependence, during a simulated vertical *vdM* scan, of the beam–size-bias factors at the scanning IP associated with the full beam–beam effect for $N_{\text{NSIP}} = 0$ using the \mathcal{L}^0 normalization (σ_{y1}/σ^0 , open squares), and for $N_{\text{NSIP}} = 1$ and using the \mathcal{L}^u normalization ($\sigma_{y1}/\sigma_{y1}^u$, filled green lozenges). For the sake of clarity, only beam-1 results are displayed; the beam-2 points would be indistinguishable. The horizontal axis is the beam separation in units of the nominal transverse beam size σ_0 . The span of the vertical scale is identical to that of Fig. 27

respect to the corresponding unperturbed beam size σ_{yB}^u . A direct comparison of the $\sigma_{yB}/\sigma_{yB}^u$ ratio with and without a non-scanning IP (Fig. 29) shows that the two configurations result in an almost identical beam-separation dependence, the shape of which closely parallels that computed [13] in the linear approximation using MADX, once the qualitative impact of amplitude detuning is taken into account.

The contrast in beam-separation dependence between the nominal (\mathcal{L}^0) and the unperturbed (\mathcal{L}^u) luminosity-bias curves apparent in Fig. 26, and the “universality” of the $[\mathcal{L}/\mathcal{L}^u]$ luminosity-bias curves (Fig. 28), can be further characterized by comparing the separation dependence of the unperturbed luminosity \mathcal{L}^u , computed numerically as the overlap integral over the unperturbed macroparticle distributions, with that of its Gaussian-equivalent counterpart $\mathcal{L}_{\text{Geq}}^u$.




Fig. 30 Beam-separation dependence, during *vdM* scans simulated in the Run-2 configuration with one non-scanning IP, of the luminosity-bias factors at the scanning IP associated with the unperturbed luminosity ($\mathcal{L}^u/\mathcal{L}^0$, lozenges), and with the Gaussian-equivalent unperturbed luminosity ($\mathcal{L}_{\text{Geq}}^u/\mathcal{L}^0$, circles). The horizontal axis is the beam separation in units of the nominal transverse beam size σ_0 . The filled (open) symbols corresponds to a vertical (horizontal) scan

The latter is defined as:¹⁴

$$\mathcal{L}_{\text{Geq}}^u(\delta_x) = \frac{1}{2\pi \Sigma_{x,\text{Geq}}^u \Sigma_{y,\text{Geq}}^u} \exp \left[-\frac{1}{2} \left(\frac{\delta_x}{\Sigma_{x,\text{Geq}}^u} \right)^2 \right],$$

and similarly for $\mathcal{L}_{\text{Geq}}^u(\delta_y)$. It is calculated under the assumption that the unperturbed macroparticle distributions are Gaussian, and adequately described (for this particular purpose) by the transverse RMS single-beam sizes σ_{iB}^u ($i = x, y; B = 1, 2$):

$$\Sigma_{i,\text{Geq}}^u = \sqrt{(\sigma_{i1}^u)^2 + (\sigma_{i2}^u)^2}.$$

The comparison between \mathcal{L}^u and $\mathcal{L}_{\text{Geq}}^u$ is presented in Fig. 30. During the vertical scan (filled symbols), the unperturbed luminosity (already shown in Fig. 26) and its Gaussian-equivalent counterpart are almost exactly equal, except at large separation ($\delta_y/\sigma^0 > 4$); both increase with separation relative to \mathcal{L}^0 . The same occurs during the horizontal scan (open symbols), except that here both curves drop with increasing separation, because the horizontal phase advances are such that the dynamic- β effect at IP5 translates into smaller horizontal beam sizes at IP1. As already documented [13] in the linear approximation, the dynamic- β effect at the non-scanning IP(s) propagates to the scanning IP, resulting in either larger or smaller unperturbed transverse beam-sizes depending on the phase advance.

The fundamental observable of the van der Meer method is the *beam-separation dependence* of the measured interaction

rate. Intuitively therefore, one can argue that a separation-independent change in unperturbed transverse beam size at the scanning IP (Fig. 27, open symbols) should not bias the σ_{vis} measurement, provided the shift in unperturbed tunes is properly accounted for. The same reasoning is offered in Ref. [13], where the argument is more straightforward: in the linear approximation (Sect. 3.1), additional collisions are tantamount to inserting a quadrupole-like perturbation at the non-scanning IP(s), that in turn results in a small tune shift accompanied by a β -beating wave around each ring. In the more general case treated here, the excellent agreement between \mathcal{L}^u and $\mathcal{L}_{\text{Geq}}^u$ over most of the scanning range (Fig. 30) indicates that the dominant effect is that of a beam- and plane-dependent shift of the effective β^* value, and of the corresponding unperturbed RMS single-beam sizes. Non-Gaussian tails associated with the non-linearity of the beam-beam force at the non-scanning IP(s) have but a very minor impact on the unperturbed transverse-density distributions at the scanning IP.

As a first illustration of the magnitude of multi-IP effects and of the quantitative impact, on beam-beam corrections, of the difference between the green circles and the red squares in Fig. 26, Table 7 compares the separation-integrated FoMs (Sect. 4.2.4) extracted from a simulated *vdM*-scan pair at IP1, with and without additional head-on collisions at IP5.

- In the \mathcal{L}^u normalization, and as expected from Fig. 28, the hierarchy of the FoMs is very similar with ($N_{\text{NSIP}} = 1$) and without ($N_{\text{NSIP}} = 0$) collisions at the non-scanning IP. The beam-beam bias is systematically more negative for $N_{\text{NSIP}} = 1$: this is qualitatively consistent with the combination of a decrease of the mean unperturbed tunes (Sect. 4.6.2), and of the tune-dependence of the luminosity-bias curves (Fig. 13). As a result, the magnitude of the overall beam-beam bias on σ_{vis} (row 7, columns 3 and 4) increases substantially in the presence of a non-scanning IP.
- For $N_{\text{NSIP}} = 1$, the hierarchy of the FoMs is very different in the \mathcal{L}^0 and \mathcal{L}^u normalizations, because the dynamic- β effect at the non-scanning IP has a significant impact on the transverse unperturbed beam sizes at the scanning IP (Fig. 27). The fact that the Σ_x (Σ_y) bias is more negative (positive) in the \mathcal{L}^0 normalization reflects the fact that the unperturbed horizontal (vertical) beam size at the scanning IP is smaller (larger) than its nominal value (Fig. 30).
- In contrast, the σ_{vis} -bias values for $N_{\text{NSIP}} = 1$ agree to better than 0.005% between the \mathcal{L}^0 and \mathcal{L}^u normalizations, indicating that in practice, the two normalizations yield equivalent beam-beam corrections. This can be understood by examining the bottom half of Table 7, that summarizes the impact of collisions at IP5 on the unperturbed scan variables at IP1: Σ_x^u decreases, Σ_y^u

¹⁴ Some of the constants appearing in the complete formulas (Eqs. (5)–(10)) are irrelevant here, and omitted for the sake of clarity.

Table 7 Fractional peak-rate, beam-size and σ_{vis} beam–beam-induced bias on measured (rows 3–7) and unperturbed (rows 8–12) scan variables, for simulated vdM scans at IP1, with (columns 2–3) and without (column 4) head-on collisions at IP5, using the nominal bunch parameters shown in Table 5 and the phase-advance settings listed in Table 6

	$N_{NSIP} = 1$		$N_{NSIP} = 0$
FoM	Normalization:		
	$\mathcal{L}_{no-bb} = \mathcal{L}^0$	\mathcal{L}^u	$\mathcal{L}_{no-bb} = \mathcal{L}^0$ ($\mathcal{L}^u = \mathcal{L}^0$)
Beam-beam bias on measured scan variables [%]			
$\mu_{pk} / \mu_{pk}^{no-bb} - 1$	0.346	0.410	0.431
$\Sigma_x / \Sigma_x^{no-bb} - 1$	-0.784	-0.454	-0.405
$\Sigma_y / \Sigma_y^{no-bb} - 1$	0.076	-0.323	-0.239
$\sigma_{vis} / \sigma_{vis}^{no-bb} - 1$	-0.365	-0.369	-0.215
Beam-beam bias on unperturbed scan variables [%]			
$\mu_{pk}^u / \mu_{pk}^{no-bb} - 1$	-0.064	-	0
$\Sigma_x^u / \Sigma_x^{no-bb} - 1$	-0.331	-	0
$\Sigma_y^u / \Sigma_y^{no-bb} - 1$	0.401	-	0
$\sigma_{vis}^u / \sigma_{vis}^{no-bb} - 1$	0.004	-	0

increases by a little more, resulting in a small decrease of the unperturbed peak interaction rate μ_{pk}^u ; the net change in σ_{vis} is at the 0.004% level. Physically, this expresses the fact that the changes in the unperturbed transverse convolved beam sizes and the unperturbed head-on luminosity are correlated such that the visible-cross-section remains invariant, at least in the configuration simulated here.

4.6.4 Sensitivity of beam–beam corrections to the phase advance between IPs

The results presented in the preceding paragraphs indicate that even though the separation dependence of the luminosity-bias factors are very different in the \mathcal{L}^0 and \mathcal{L}^u normalizations (Fig. 26), these two reference configurations appear to yield almost identical beam–beam corrections to σ_{vis} (Table 7). What remains to be quantified is how robust this conclusion is against large variations in the assumed phase advance between consecutive IPs, as well as in the presence of more than one non-scanning IP.

The study reported in Sect. 4.6.3 assumes the phase-advance values between IP1 and IP5 that are listed in the top two rows of Table 6; these correspond to one point in four-dimensional phase-advance space $(\mu_x^{B1}, \mu_y^{B1}, \mu_x^{B2}, \mu_y^{B2})$. To address the first question above, the analysis detailed in Sect. 4.6.3 was repeated while varying the horizontal and vertical phase advance between IP1 and IP5 over most of the $[0, 2\pi]$ range, while keeping the nominal tunes constant; in order to keep the parameter space manageable, the B1




Fig. 31 Phase-advance dependence of the difference, between the \mathcal{L}^0 and the \mathcal{L}^u normalization, of the beam–beam induced σ_{vis} bias. The horizontal and vertical axes represent the betatron phase advance between IP1 and IP5, here assumed to be the same for B1 and B2. The color scale quantifies the fractional inconsistency in σ_{vis} bias between the two reference configurations; it ranges from -0.06 to $+0.05\%$ of σ_{vis}

and B2 phase-advance values were in addition constrained to be the same, separately in the horizontal and the vertical plane ($\mu_x^{B1} = \mu_x^{B2} = \mu_x, \mu_y^{B1} = \mu_y^{B2} = \mu_y$). vdM scans were simulated at IP1 over a grid in (μ_x, μ_y) space, with the bunches colliding head-on at IP5, and using the bunch parameters listed in Table 5.

At each point in the (μ_x, μ_y) space, the luminosity-bias curves and the corresponding separation-integrated FoMs were computed in both the \mathcal{L}^0 and the \mathcal{L}^u normalization. The beam–beam-induced σ_{vis} bias, or equivalently the magnitude of the overall beam–beam correction to the vdM -based absolute luminosity scale, was found to depend slightly on the choice of the reference “no beam–beam” configuration, i.e. on whether the beam–beam correction is calculated in the \mathcal{L}^0 or in the \mathcal{L}^u normalization. The difference is phase-advance dependent, and exhibits a clear periodic structure (Fig. 31); periodic structures are also observed in such two-dimensional maps of the luminosity-bias factor (not shown), with different patterns for $[\mathcal{L}/\mathcal{L}^0]$ and $[\mathcal{L}/\mathcal{L}^u]$.

These periodicities have been verified to arise neither from statistical fluctuations, nor from numerical effects in the simulation.¹⁵ They are attributed to cross-talk between collision points: beam–beam kicks at one IP induce separation- and amplitude-dependent β -beating that propagates in a phase-dependent manner to the next IP, where it slightly modifies

¹⁵ It is interesting to note that the interplay between beam–beam effects at different IPs, and its sensitivity to the phase advance between IPs, was observed at LEP-II already, where they could strongly affect the instantaneous-luminosity balance between the four experimental detectors [53].

the transverse-density distributions and the particle-averaged β functions; the tunes are affected as well. The superposition of these effects across multiple IPs, in turn, modulates the associated beam–beam-induced biases. This interpretation is supported by the fact that the detailed shape of the simulated tune spectra, that of the transverse luminosity profiles at the witness IP, as well as that of the beam-separation dependence of the $\mathcal{L}/\mathcal{L}^u$ ratio at the scanning IP, are all observed to depend on the phase advance between IPs. A deeper understanding of these mechanisms may provide the key for resolving the ambiguity between the \mathcal{L}^0 and \mathcal{L}^u normalizations; as such, they remain under active study.

Fortunately, these subtle effects have only a minor impact on beam–beam corrections to vdM calibrations. The largest amplitude of the \mathcal{L}^0 - \mathcal{L}^u discrepancy over the entire plane amounts to 0.06% of σ_{vis} , which is smaller than, or at most comparable to, the absolute accuracy of the B*B or COMBI simulations. The reason why in Table 7 the \mathcal{L}^0 and \mathcal{L}^u normalizations yield essentially identical results, is that the Run-2 phase-advance between IP1 and IP5 happens to lie in a region of the $\mu_x - \mu_y$ plane where the inconsistency is the smallest. In this case therefore, and for bunches that collide only in these two IPs, the difference in σ_{vis} bias between the \mathcal{L}^0 and \mathcal{L}^u normalizations is negligible.

To assess the impact of more than one non-scanning IP, the above phase scans were repeated for collision patterns with up to three non-scanning IPs. To deal with the increased dimensionality of the problem, the fractional phase advance from IP1 to IP5 was fixed to $\Delta\mu_x = \Delta\mu_y = 0.405$, a choice that in Fig. 31 corresponds to the largest absolute difference in σ_{vis} bias between the \mathcal{L}^u and \mathcal{L}^0 normalizations. With that constraint, and with the overall tunes and the IP5-IP8 phase advances fixed at their nominal values, the horizontal and vertical fractional phase-advance values from IP1 to IP2 were scanned over the range $[0, 2\pi]$. The study was then repeated with the roles of IP2 and IP8 interchanged. In all these configurations, the difference in σ_{vis} bias between the \mathcal{L}^u and \mathcal{L}^0 normalizations displays the same periodic pattern as that in Fig. 31, and the largest absolute difference in σ_{vis} bias between the two reference configurations does not exceed 0.15% (for $\xi = 0.006$). While strictly speaking, the ensemble of these scans does not constitute an exhaustive sampling of the 12-dimensional phase-advance space (2 beams \times 2 planes \times 3 IP combinations), it suggests that the largest possible disagreement between the \mathcal{L}^0 and \mathcal{L}^u normalizations amounts to 0.1–0.2% of σ_{vis} when considering the full parameter space. Since the phase advance between LHC IPs can be measured to a few degrees ($\sim 0.015 \times 2\pi$) per beam and per plane, the “operating point” in phase-advance space is known to very good precision. In some cases therefore, the ambiguity on the σ_{vis} bias associated with the choice of reference configuration can be significantly reduced, as

already demonstrated for bunches that collide in IP1 and IP5 only.

4.6.5 Parameterization of multi-IP effects

To gain further insight into multi-IP effects and account, in a pragmatic fashion, for their impact on beam–beam corrections to luminosity calibrations, COMBI was used to simulate scans over the full range of beam–beam parameter values typical of vdM sessions at $\sqrt{s} = 13$ TeV, using the Run-2 phase-advance settings (Table 6), both without ($N_{\text{NSIP}} = 0$) and with ($1 \leq N_{\text{NSIP}} \leq 3$) head-on collisions at one or more non-scanning IP(s).

The results are summarized in Fig. 32 for three representative values of the single-IP beam–beam parameter ξ , using the \mathcal{L}^u normalization. Since both the orbit-shift and the optical-distortion bias are sensitive to the unperturbed-tune values (Sect. 4.2.2), the overall beam–beam bias on the visible cross-section associated with a given value of ξ is significantly affected by collisions at non-scanning IPs. Taking as an example the case of $\xi = 6 \times 10^{-3}$, the σ_{vis} bias when colliding only at the scanning IP amounts to -0.55% (orange circle). It grows to -1.0% for one non-scanning IP (pale-brown lozenge), to -1.3% for two non-scanning IPs (medium-brown lozenge), and to a little over -1.5% for the worst case of $N_{\text{NSIP}} = 3$ (darkest-brown lozenge).¹⁶ As N_{NSIP} increases by one unit, both the horizontal and the vertical mean unperturbed tune decrease by approximately $\xi/2$; as a consequence, the luminosity-bias curves shift downwards (Figs. 12, 13) and the σ_{vis} bias becomes accordingly more negative. The impact of this “multi-IP tune shift” is comparable in magnitude to the σ_{vis} bias for $N_{\text{NSIP}} = 0$; it varies from one bunch pair to another in a manner that depends on the single-bunch parameters and on the value of N_{NSIP} associated with each-colliding bunch pair at the IP where the scans are taking place. Fortunately however, this effect lends itself to a simple parameterization.

The diagonal curves in Fig. 32 display the visible cross-section bias for scans simulated at IP1 in the absence of collisions at any non-scanning IP, but with both the horizontal and the vertical unperturbed tune shifted by a quantity ΔQ_{mIP} that scales with ξ . For a given value of ξ , the intersections of the corresponding diagonal curve (orange in the example above) with the horizontal lines of the same hue corresponding to $N_{\text{NSIP}} = 1, 2, \dots$, determine the magnitude of the multi-IP equivalent tune shift ΔQ_{mIP} that needs to be applied simultaneously to the unperturbed horizontal and vertical tunes input to the COMBI simulation to mimic the effect of collisions

¹⁶ Half-integer values such as $N_{\text{NSIP}} = 2.5$ correspond to a bunch pair that collides at IP1 and in which one of the bunches experiences three additional collisions (IP5, plus both IP2 and IP8), while the other one is involved in only two (IP5, plus either IP2 or IP8).




Fig. 32 Horizontal lines, and markers at $\Delta Q_{\text{mIP}} = 0$: visible cross-section bias in the \mathcal{L}^u normalization, associated with vdM scans at IP1, for bunch pairs that in addition collide head-on at a fixed number of non-scanning IPs, and with the tunes set to their nominal value. The three sampled values of the single-IP beam-beam parameter ξ are identified by the blue triangles ($\xi = 2.7 \times 10^{-3}$), green squares ($\xi = 4.0 \times 10^{-3}$) and brown lozenges ($\xi = 6.0 \times 10^{-3}$). For a given value of ξ , the lighter, medium and darker colors correspond to additional collisions at $N_{\text{NSIP}} = 1, 2, 2.5$ and 3 non-scanning IPs respectively (see text). Inclined curves and circular markers: visible cross-section bias associated with vdM scans at IP1 with beams fully separated at all other IPs ($N_{\text{NSIP}} = 0$), as a function of the multi-IP equivalent mean tune shift ΔQ_{mIP} expressed in units of ξ . The circular markers at $\Delta Q_{\text{mIP}}/\xi = 0$ show the value of the σ_{vis} bias corresponding to beams colliding only at the scanning IP, with the unperturbed tunes set to their nominal value. The vertical dashed lines correspond, from right to left, to $\Delta Q_{\text{mIP}}/\xi = -0.5, -1.0, -1.25$ and -1.5 , and approximately correspond to the values of the effective tune shift ΔQ_{mIP} where an inclined curve associated with a given value of ξ intersects the corresponding horizontal lines

at non-scanning IPs. The corresponding values of $\Delta Q_{\text{mIP}}/\xi$ are indicated by the vertical dashed lines in the figure.

A more telling representation of the same results is offered in Fig. 33, which displays the multi-IP equivalent tune shift as a function of the number of non-scanning IPs. The results are consistent with a linear dependence that differs only slightly from the naive scaling law




Fig. 33 Dependence of the normalized multi-IP equivalent tune shift $\Delta Q_{\text{mIP}}/\xi$ on the number of non-scanning IPs, for several values of the single-IP beam-beam parameter ξ . The scanning IP is set to IP1, and the σ_{vis} bias is calculated using the \mathcal{L}^u normalization. The solid magenta line represents a two-parameter linear fit to the points; the black dashed line represents a naive parameterization of the form $\Delta Q_{\text{mIP}}/\xi = -0.5 \times N_{\text{NSIP}}$

$$\Delta Q_{\text{mIP}} = -0.5 \xi N_{\text{NSIP}} = -p_1 \xi N_{\text{NSIP}} \tag{42}$$

inspired by the qualitative evolution of the tune spectra. The purely empirical, two-parameter linear fit represented by the purple solid line:

$$\Delta Q_{\text{mIP}} = p_0 - p_1 \xi N_{\text{NSIP}}$$

provides a slightly better parameterization than Eq. (42), and appears accurate at the level of $\Delta Q_{\text{mIP}} \sim \pm 0.05 \xi$.

Figure 33 demonstrates that for a given scanning IP and for all values of ξ , the normalized multi-IP equivalent tune shift scales linearly with the number of non-scanning IPs. To what extent this scaling is “universal”, i.e. how strongly it depends on the choice of the IP where the scans are taking place, is addressed by Fig. 34. Irrespective of the scanning IP considered, the overall dependence of the σ_{vis} bias on the number of scanning IPs is reasonably well represented by the naive scaling law of Eq. (42), demonstrating that the dominant effect is the downward shift of the mean unperturbed tunes. There are, however, significant differences between marker families, most notably between the red crosses (IP5) and the cyan squares (IP8): these reflect the variety of phase-advance patterns between the scanning IP and the other collision points. Similarly, within a given family, there exist several bunch pairings between B1 and B2 that yield the same number of additional collisions, but that correspond to different phase-advance configurations. When IP5 is the scanning IP, for instance, the configuration $N_{\text{NSIP}} = 2$ can be achieved with additional collisions at IP1 and at either IP2 or IP8; these two cases can therefore yield slightly different σ_{vis} biases. To minimize modeling errors, therefore, it is advisable that scal-




Fig. 34 Visible cross-section bias in the \mathcal{L}^u normalization as a function of the number of non-scanning IPs (lower horizontal axis), for the four possible choices of scanning IP. The curve represents the dependence of the σ_{vis} bias in the $N_{\text{NSIP}} = 0$ case, but with the horizontal and vertical unperturbed tunes shifted downwards by $\Delta Q_{\text{mIP}}/\xi = -0.5 \times N_{\text{NSIP}}$ (upper horizontal axis). For a given scanning IP, there exist several B1-B2 bunch-collision patterns that yield the same value of N_{NSIP} . The beam-beam parameter and phase-advance values are listed in Tables 5 and 6 respectively




Fig. 35 Visible cross-section bias in the \mathcal{L}^0 normalization as a function of the number of non-scanning IPs, for the four possible choices of scanning IP. The results are extracted from the same simulations as those in Fig. 34; only the reference configuration differs. The curve, the axis scales and the legend are identical

ing parameterizations such as that illustrated by Fig. 33 be developed for each scanning IP separately, preferably taking into account the actual collision pattern experienced by each bunch that collides at the scanning IP under consideration.

Pursuing the same strategy in the \mathcal{L}^0 normalization meets with limited success (Fig. 35): for a given value of N_{NSIP} , the spread across different collision patterns is noticeably larger than in the \mathcal{L}^u normalization, even when restricted to one given scanning IP. This observation reflects the same

underlying physics as the periodic structure in Fig. 31; its implications are addressed below.

4.6.6 Beam-beam correction strategies in the presence of multi-IP collisions

In summary, collisions at non-scanning IPs manifest themselves in two main ways. The widening of the tune spectra and the resulting downward shift of the mean unperturbed tunes (Sect. 4.6.2) significantly increase the beam-beam-induced bias on the absolute luminosity scale (Fig. 34), mainly because of the reduced amplitude of the optical-distortion bias (Sect. 4.2.2). This tune shift is accompanied by a beam- and plane-dependent β -beating wave around each ring [13], that translates into a sub-percent rescaling of the unperturbed transverse single-beam sizes at the scanning IP (Fig. 27). The latter causes significant distortions of the luminosity-bias curves when using the \mathcal{L}^0 normalization (Fig. 26), that depend strongly on the phase advance between consecutive collision points.

The shift in unperturbed tunes (Sect. 4.6.2) has, by far, the dominant impact on beam-beam corrections to vdM calibrations. It can be accounted for by simulating collisions simultaneously at all relevant IPs, with the input tunes set to their nominal value, and by correcting the luminosity-scan curves to “no beam-beam anywhere”. This \mathcal{L}^0 normalization is the approach adopted in Ref. [16]. The alternative is to use the \mathcal{L}^u normalization. It is inspired by the fact that the fundamental observable of the vdM method is the beam-separation dependence of the measured interaction rate. What should matter, therefore, is the variation of the beam-beam force during the scan: underlying, separation-independent modifications of the unperturbed transverse-density distributions are expected not to affect the visible cross-section. In the \mathcal{L}^u approach, the luminosity-scan curves extracted from the same multi-IP simulation as above are corrected to “no beam-beam at the scanning IP only” (Sect. 4.6.3); the resulting luminosity-bias functions become much less sensitive to the presence of non-scanning IPs (Fig. 28).

The \mathcal{L}^0 and \mathcal{L}^u normalizations yield beam-beam corrections that, depending on the phase advance between consecutive IPs, either are indistinguishable (Table 7), or in the worst case differ by at most 0.1–0.2% of σ_{vis} (Sect. 4.6.4). These observations demonstrate that no phase-advance configuration exists in which beam-beam induced β -beating would significantly reduce the beam-beam bias on σ_{vis} .

The results of the full multi-IP simulations can be very closely approximated by those of the much faster single-IP simulation, with as input the same value of the beam-beam parameter ξ , but with the unperturbed tunes shifted downwards by an amount ΔQ_{mIP} that in the \mathcal{L}^u normalization and for a given scanning IP, scales with ξ and depends linearly on N_{NSIP} (Fig. 33). The scaling parameters are not fully uni-

versal, in the sense that they depend slightly on the choice of scanning IP (Fig. 34), or more fundamentally on the phase advance between the scanning and non-scanning IPs.

The scaling prescriptions offered in Sect. 4.6.5 in the \mathcal{L}^u normalization account for collisions at non-scanning IPs to better than 0.1–0.2% on σ_{vis} . It should be pointed out, however, that these simple-minded models reproduce only the σ_{vis} bias extracted from the multi-IP simulation (to an accuracy limited by the phase-advance-dependent inconsistency between the \mathcal{L}^0 and \mathcal{L}^u normalizations), but that they are not meant to reproduce the corresponding tune spectra. An additional minor limitation is that they assume that collisions at non-scanning IPs occur with zero transverse separation, and with zero crossing angle. These ad hoc recipes, therefore, offer but a partial substitute for full-fledged simulations that account in detail for every distinct multi-IP collision pattern, separately for each member of every colliding-bunch pair and for each scanning IP. For many purposes, however, these simple scaling laws appear sufficiently accurate to cover most beam conditions encountered so far during *vdM* sessions at the LHC.

On the practical side, the \mathcal{L}^u normalization offers several advantages. First, by reducing the multi-IP problem to that of a single collision point, it greatly simplifies the interpretation of simulation results and the estimate of systematic uncertainties. The \mathcal{L}^u approach also automatically accounts for the static (in the sense of separation-independent) change in unperturbed beam size associated with the dynamic- β effect at the non-scanning IPs. Finally, combining the parameterization of single-IP beam-beam biases offered in Sect. 4.2.3 with the scaling prescriptions of Sect. 4.6.5 provides a beam-beam-correction procedure that is straightforward to implement, requires no or little investment by individual users in complex, CPU-intensive simulations, and is in many cases sufficient to calculate absolute bunch-by-bunch corrections to *vdM* calibrations, or to estimate systematic uncertainties. This is particularly valuable when dealing with five to ten *vdM*-scan pairs that involve up to 150 colliding bunches with different populations and emittances, arranged in multiple collision patterns that each require a different multi-IP correction [9]: the computing overhead that would be associated with simulating each of those configurations in detail is considerable.

At the present level of understanding, there appears to be no factual argument in favor of the \mathcal{L}^0 normalization. Even if it eventually turns out that correcting to “no beam-beam anywhere” provides a more accurate correction in principle, the potential reduction in systematic uncertainty remains moderate ($< 0.2\%$) when compared in quadrature to the overall luminosity-calibration uncertainty (around 1%). In addition, and because of the larger phase-advance sensitivity of the \mathcal{L}^0 normalization, this hoped-for improvement could only be secured by simulating separately each and every collision

pattern associated with the scanning IP under consideration. Finally, the cumulative impact of the dynamic- β effect at the non-scanning IPs, that can shift (Sect. 4.6.3) the underlying unperturbed beam sizes σ_{iB}^u at the scanning IP by $\mathcal{O}(0.5\text{--}1\%)$ per plane, per beam and per non-scanning IP, totally blurs both the conceptual and the operational distinction between the orbit-shift and the optical-distortion correction. This is an additional reason why in the \mathcal{L}^0 normalization, a parameterization approach would be inapplicable: one must resort to a combined, “full beam-beam” multiplicative correction to the measured collision rate that can only be extracted from a dedicated simulation that has been tailored, separately for each colliding bunch, to the actual conditions of the *vdM*-calibration session.

In conclusion, clarifying the origin of periodic structures such as that in Fig. 31, and explaining why beam-beam corrections are less sensitive to phase-advance effects in the \mathcal{L}^u (Fig. 34) than in the \mathcal{L}^0 (Fig. 35) normalization, remain topics of active study. The ambiguity stands, therefore, as to which of the two reference “no beam-beam” configurations (i.e. which of the \mathcal{L}^0 or the \mathcal{L}^u normalization) provides the most accurate beam-beam correction in absolute terms. Pragmatically, the practical advantages of the \mathcal{L}^u normalization, and the larger complexity in implementing the \mathcal{L}^0 normalization in the *vdM*-data analysis, tilt the balance in favor of the former – at least for now. Until the \mathcal{L}^0 - \mathcal{L}^u ambiguity is lifted, and irrespective of the choice of normalization, it seems prudent to cover the difference by assigning a phase-dependence-related systematic uncertainty to the absolute magnitude of the beam-beam corrections; an explicit prescription to this effect will be offered in Sect. 5.3.

4.7 Impact of beam-beam asymmetry

The results presented so far in this chapter, with the exception of those in Sect. 4.4, assume beams that are round in the transverse plane ($\beta_{xB}^* = \beta_{yB}^* = \beta_B^*$ and $\sigma_{xB} = \sigma_{yB} = \sigma_B$, where $B = 1, 2$). In such a case, the distinction between horizontal and vertical beam-beam parameters disappears, and the beam-beam parameter of beam 2, defined in Eq. (12), simplifies to

$$\xi_2 = \frac{n_1 r_0 \beta_2^*}{4\pi \gamma_2 \sigma_1^2}, \quad (43)$$

where σ_1 is the transverse RMS size of beam 1. The same formula, *mutatis mutandis*, is used to define ξ_1 .

The parameterization described in Sect. 4.2, and the results derived therefrom, further assume that the beam-beam forces are exactly balanced between the two beams, i.e. that $\xi_1 = \xi_2 = \xi_R$ (Eq. (29)). Beam-beam asymmetry occurs when this assumption is violated, i.e. when $\xi_1 \neq \xi_2$, which can arise under three different scenarios (Sect. 4.7.1). Even though a parametric approach can provide physically intuitive insight

(Sect. 4.7.2), dedicated simulations are required to quantify the impact of such asymmetries on the accuracy of the beam–beam correction procedure (Sect. 4.7.3).

4.7.1 Beam–beam asymmetry scenarios

Consider first the case of a bunch-population asymmetry ($n_1 \neq n_2$, with $\beta_1^* = \beta_2^*$ and $\sigma_1 = \sigma_2$), which can be simulated in B*B by assigning different populations to the source beam and to the witness beam. Since the source beam in B*B remains unaffected by the beam–beam interaction, and since the two beams are subject to different beam–beam parameters, the simulation must be run twice, first with n_1 and n_2 assigned to, respectively, the source and the witness beam, and then with the inverse assignment. The predicted luminosity-bias curves, and the corresponding corrections to the FoMs, are then given by adding the two sets of results,¹⁷ implicitly assuming that strong-strong effects do not play a significant role even in the presence of a large bunch-population asymmetry (at least in the *vdM* regime). This procedure was validated by confronting the B*B results with those from COMBI for several asymmetric configurations: even in the worst case considered ($n_2/n_1 = 2$, $\xi_1 = 5.30 \times 10^{-3}$, $\xi_2 = 2.65 \times 10^{-3}$), the beam–beam corrections calculated by the two packages agreed to better than one part in 1600.

The B*B and COMBI results above are accurately reproduced by the parameterization detailed in Sect. 4.2.3, with as input the beam-averaged beam–beam parameter:

$$\langle \xi \rangle = \frac{\xi_1 + \xi_2}{2}. \tag{44}$$

This is a consequence of the fact that the ξ -dependence of the FoMs is close enough to linear (Figs. 14–16) for the averaging of the FoM values at two different ξ settings to be tantamount to computing the FoM once, with $\langle \xi \rangle$ as input.

The second scenario is that of a β -function asymmetry ($\beta_1^* \neq \beta_2^*$, with $n_1 = n_2$ and $\sigma_1 = \sigma_2$). This case is equivalent to that of the bunch-population asymmetry above, since only the product $n_k \beta_l^*$ ($k, l = 1, 2$ with $l \neq k$) plays a role in Eq. (43). The bunch population n_k of beam k controls the scale of the beam–beam force beam k exerts on beam l , while β_l^* determines how sensitive beam l is to this force.

In contrast to the above two scenarios, which involve only scale effects, a beam-size or emittance asymmetry ($\sigma_1 \neq \sigma_2$, with $n_1 = n_2 = n$ and $\beta_1^* = \beta_2^* = \beta^*$) modifies in different ways the spatial distributions of the source field and of the particles in the witness bunch, the trajectories of which are affected by this field.

¹⁷ When B*B is run on a symmetric configuration, the scan is simulated only once, and the luminosity-bias result is simply doubled to account for the conceptual swap between the source and the witness beam.

4.7.2 Parametric characterization of beam–beam asymmetries

A full characterization of beam–beam asymmetries requires the two parameters ξ_1 and ξ_2 . A physically intuitive parameterization can however be provided by a single, “effective” beam–beam parameter $\bar{\xi}$, defined as

$$\bar{\xi} = \frac{\xi_1 \sigma_1^2 + \xi_2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}, \tag{45}$$

on the basis of the following argument.

Under the naive assumption that the optical distortion of B1 induced by the field of B2 is proportional to ξ_1 , and that it can be adequately represented by a small fractional change λ_1 in the value of σ_1 :

$$\sigma_1 \rightarrow \sigma_1(1 + \lambda_1),$$

where $\lambda_1 \propto \xi_1$, one can show, using Eqs. (5) and (6) and after a bit of algebra, that to first order in ξ_1 , the luminosity-bias factor $[\mathcal{L}/\mathcal{L}^0]_{Opt,1}$ associated with the optical distortion of B1 alone obeys

$$[\mathcal{L}/\mathcal{L}^0]_{Opt,1} - 1 \propto \frac{\lambda_1 \sigma_1^2}{\sigma_1^2 + \sigma_2^2} \propto \frac{\xi_1 \sigma_1^2}{\sigma_1^2 + \sigma_2^2}.$$

Applying the same argument to the optical distortion of B2 induced by the field of B1, and combining the impact of the optical distortions of the two beams, it follows that

$$[\mathcal{L}/\mathcal{L}^0]_{Opt} - 1 \propto \frac{\xi_1 \sigma_1^2 + \xi_2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2} = \bar{\xi}.$$

In the limit of equal beam sizes ($\sigma_1 = \sigma_2$), $\bar{\xi}$ reduces to $\langle \xi \rangle$, and as such properly characterizes bunch-population or β^* asymmetries. If instead the bunch populations are balanced ($n_1 = n_2 = n$) but the transverse beam sizes differ ($\sigma_2 \neq \sigma_1$), then for a given value of $\Sigma_R = \sqrt{\sigma_1^2 + \sigma_2^2}$, the effective beam–beam parameter can be written as

$$\bar{\xi} = \frac{1}{2} \left[(\sigma_2/\sigma_1)^2 + (\sigma_1/\sigma_2)^2 \right] \frac{nr_0 \beta^*}{2\pi \gamma \Sigma_R^2}, \tag{46}$$

thereby incorporating the effect of beam-size asymmetries in a single-parameter, but approximate, characterization of the strength of the beam–beam interaction. In the limit of $\sigma_2/\sigma_1 = 1$, $\bar{\xi}$ reduces to the symmetric, round-beam equivalent beam–beam parameter ξ_R (Eq. (29)).

4.7.3 Predicted impact of transverse beam-size asymmetries

Since the beam–beam induced orbit shift depends only on the convolved beam size Σ_R (Sect. 2.3.1), but not on the σ_2/σ_1 ratio, only the beam–beam-induced optical distortion may be affected by a beam-size asymmetry. It is natural, therefore,

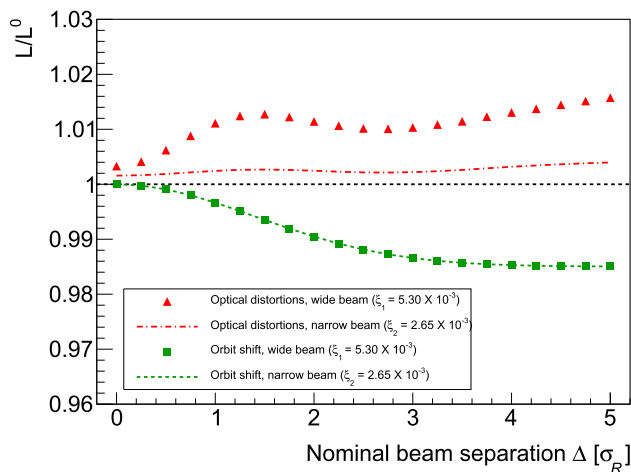


Fig. 36 Single-beam contributions to the beam-separation dependence, during a horizontal vdM scan simulated using B*B, of the luminosity-bias factor associated with the optical distortions ($[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$, red curve and triangles), and of that induced by the beam–beam orbit shift ($[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$, green curve and squares). The unperturbed beams are assumed to be round and perfectly Gaussian. Their bunch populations are assumed equal ($n_1 = n_2 = 0.78 \times 10^{11}$ p/bunch), and their transverse beam sizes differ by a factor of $\sqrt{2}$ ($\sigma_1 = 99.6 \mu\text{m}$, $\sigma_2 = 70.5 \mu\text{m}$). The wide-beam contributions are shown by the filled markers, the narrow-beam contributions by the dashed curves. The beams collide only at the IP where the scan is taking place. The horizontal axis is the nominal beam separation in units of the effective transverse beam size $\sigma_R = \Sigma_R/\sqrt{2}$

to characterize this effect by varying, in the simulation, the common bunch intensity and the beam-size ratio, while keeping Σ_R constant. When using COMBI, this can be done in a single step for each parameter set, since the spatial density distributions of both bunches evolve simultaneously. When using B*B, however, the source bunch remains unaffected. Therefore the simulation must be run twice, first with σ_1 and σ_2 assigned to, respectively, the source and the witness beam, and then with the inverse assignment, after which the contributions of the two passes are combined.

Such a study is presented in Fig. 36, for a highly asymmetric example in which $\xi_1/\xi_2 = 2$ and $\bar{\xi} = 4.42 \times 10^{-3}$. When the wider of the two bunches ($\xi_1 = 5.30 \times 10^{-3}$) plays the role of the witness beam (triangles and squares), the orbit-shift and optical-distortion contributions are of opposite sign and of comparable magnitude; when the narrower bunch ($\xi_2 = 2.65 \times 10^{-3}$) becomes the witness beam (dashed curves), the beam–beam parameter drops by a factor of two, the orbit-shift contribution remains unchanged as it should, but the optical-distortion effect becomes almost negligible. The combination of the two contributions, obtained by multiplying the luminosity-distortion factors associated with the two witness beams, is presented in Fig. 37. The two beams contribute equally to the orbit effect, but the optical-distortion is dominated by that experienced by the wider (i.e. the weaker) of the two beams.




Fig. 37 Beam-separation dependence, during a horizontal vdM scan simulated using B*B, of the luminosity-bias factor associated with the optical distortions ($[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$, red solid curve), and of that induced by the beam–beam orbit shift ($[\mathcal{L}/\mathcal{L}^0]_{\text{Orb}}$, green dot-dashed curve), in the presence of a beam-size asymmetry corresponding to $\xi_1/\xi_2 = 2$. The single-beam parameters are listed in the caption of Fig. 36. The beams collide only at the IP where the scan is taking place. The horizontal axis is the nominal beam separation in units of the effective transverse beam size $\sigma_R = \Sigma_R/\sqrt{2}$

To quantify the impact of a transverse beam-size asymmetry on beam–beam corrections to vdM calibrations, the simulation protocol detailed above was repeated for several configurations in which the bunch populations were symmetric ($n_1 = n_2 = n$), the beam-size ratio σ_2/σ_1 was varied from 1.0 down to 0.65, and the convolved beam size was kept constant at $\Sigma_R = 122 \mu\text{m}$. Two values of the common bunch population were considered ($n = 0.78 \times 10^{11}$ and $n = 1.21 \times 10^{11}$ p/bunch), near the lower end, and beyond the upper end, of the bunch-intensity range used during vdM scans; these translate into round-beam equivalent beam–beam parameter values of, respectively $\xi_R = 3.53 \times 10^{-3}$ and 5.50×10^{-3} . The visible cross-section bias $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1$ corresponding to each of these configurations, and computed from luminosity-bias curves such as those displayed in Fig. 37, is presented in Figs. 38 and 39.

As the beam-size ratio deviates more and more from unity, $\bar{\xi}$ grows (see Eq. (46)), and so does the σ_{vis} bias. Its dependence on $\bar{\xi}$ is almost linear (Fig. 38), with only a small quadratic component. The scaling law, however, is approximate only, as evidenced by the slight relative misalignment of the B*B curves at the interface between the two groups ($\xi_R \sim 5.2$). The dependence on the beam-size ratio (Fig. 39) is well modeled by a third-order polynomial of σ_2/σ_1 , the coefficients of which depend on the $n\beta^*$ product. A systematic difference between COMBI and B*B becomes apparent at large asymmetry, suggesting that coherent oscillations may start to play a role in that regime. Quantitatively however, this discrepancy is small enough in the asymmetry range of inter-




Fig. 38 Dependence on the effective beam–beam parameter $\bar{\xi}$ of the visible cross-section bias $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1$ associated with optical distortions, as predicted by B*B (diamonds) or COMBI (triangles) simulations, for low-intensity (open blue and red markers, bottom left) and high-intensity (filled green markers, top right) colliding-bunch configurations that cover a wide range of beam-size asymmetries at constant Σ_R (see text). The beams collide at the scanning IP only. The σ_{vis} bias in the fully symmetric configuration ($\sigma_1 = \sigma_2$) corresponds to the lowest value of $\bar{\xi}$ in each group, and is indicated by a horizontal black line. The curves are second-order polynomial fits to the points, and are intended to guide the eye. The arrows indicate the impact on the absolute luminosity scale of a 20% beam-size imbalance ($\sigma_2/\sigma_1 = 0.8$)




Fig. 39 Dependence on the transverse beam-size ratio σ_2/σ_1 of the visible cross-section bias $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1$ associated with optical distortions, as predicted by B*B (diamonds) or COMBI (triangles) simulations, for low-intensity (red and blue bottom curves) and high-intensity (green top curve) colliding-bunch configurations that cover a wide range of beam-size asymmetries at constant Σ_R (see text). The beams collide at the scanning IP only. The σ_{vis} bias in the fully symmetric configuration ($\sigma_1 = \sigma_2$) is indicated by a horizontal black line in each group. The curves are third-order polynomial fits to the points, and are intended to guide the eye. The arrows indicate the impact on the absolute luminosity scale of a 20% beam-size imbalance ($\sigma_2/\sigma_1 = 0.8$)




Fig. 40 Impact of the beam-size asymmetry on the visible cross-section bias $\sigma_{\text{vis}}/\sigma_{\text{vis}}^0 - 1$ associated with optical distortions, as simulated using either B*B (diamonds and circles) or COMBI (triangles). Shown is the change in σ_{vis} bias between an asymmetric beam–beam configuration characterized by a given setting of $\sigma_2/\sigma_1 < 1$, and the corresponding symmetric configuration ($\sigma_2 = \sigma_1$). The top three curves (green diamonds, red triangles and blue diamonds) correspond to the data displayed in Fig. 39, when the beams collide at the scanning IP only. The bottom two curves (circles) show the same results when the beams also collide head-on at two non-scanning IPs. The open (filled) symbols correspond to the low- (high-) intensity bunches (see text). The curves are third-order polynomial fits to the points, and are intended to guide the eye

est at LHC (at most 0.01% on σ_{vis} for $\sigma_2/\sigma_1 > 0.9$, i.e. for an emittance imbalance of 20%) that it can be absorbed in a systematic uncertainty.

The impact of the beam-size asymmetry on the absolute luminosity scale, i.e. the difference in the σ_{vis} bias between a given asymmetric configuration (in this example $\sigma_2/\sigma_1 = 0.8$) and the fully symmetric case ($\sigma_1 = \sigma_2$), is indicated by the arrows. The same results are presented in Fig. 40 (top three curves), on an expanded scale and after subtracting the σ_{vis} bias corresponding to the symmetric configuration.

The results displayed in Figs. 36–39 assume that the beams collide at the scanning IP only. Additional collisions at other IPs lower the unperturbed tunes (Sect. 4.6.2): as a result, the impact of the beam–beam-induced optical distortion is significantly reduced (Fig. 12), and so is the impact of the transverse beam-size asymmetry. This is illustrated by the two bottom curves in Fig. 40, for the case of bunches colliding at two non-scanning IPs ($N_{\text{NSIP}} = 2$), that was modeled using B*B to simulate single-IP collisions but with the unperturbed tunes shifted as prescribed by Eq. (42). At the largest asymmetry considered ($\sigma_2/\sigma_1 = 0.65$), the sensitivity of the σ_{vis} bias to the beam-size asymmetry drops by more than a factor of two. In addition, the bunch-current dependence of this bias, manifested by the difference between the filled and the open circles, is several times smaller for the $N_{\text{NSIP}} = 2$ case. This occurs because the shift ΔQ_{mIP} in unperturbed tunes

is proportional to both ξ and N_{NSIP} : when both are large enough, this tune shift more than compensates the increase in optical distortion at the scanning IP associated with the increased beam–beam parameter experienced by the weaker of the two beams.

During Run-2 pp νdM scans at $\sqrt{s} = 13$ TeV, the bunch population n_p rarely exceeded 10^{11} (Table 1), and the typical beam-size asymmetry σ_2/σ_1 hovered around 0.94, as extracted from one of the non-factorization analyses outlined in Sect. 4.3.1. In practice therefore, for these νdM sessions the impact of beam–beam asymmetries on beam–beam corrections to νdM scans does not exceed 0.02% (Fig. 40); such a small number can be safely accounted for as a systematic uncertainty in the absolute luminosity scale.

5 Systematic uncertainties

The sources of systematic uncertainty that may affect the absolute magnitude of beam–beam corrections during νdM scans, and that are detailed in Sect. 4, can be regrouped in three categories: optical configuration of the LHC rings, deviations from the fully symmetric Gaussian-beam configuration, and collisions at multiple IPs. The associated uncertainties are briefly elaborated upon in Sects. 5.1, 5.2 and 5.3 below; additional sources of potential bias are addressed in Sect. 5.4. All the uncertainties are then added in quadrature, and their combined magnitude compared in Sect. 5.5 to that of the overall beam–beam correction itself.

The quantitative impact of each uncertainty on the σ_{vis} bias, or equivalently on the magnitude of the beam–beam correction, is summarized in Table 8 for scans simulated using three different collision patterns ($N_{\text{NSIP}} = 0, 1, 2$). For each uncertainty source, the beam conditions input to the simulation are detailed in the chapters listed in the last column of the table. Not all simulations, however, use exactly the same beam conditions. In order to provide a self-consistent picture of the relative magnitude of the various uncertainties, and since the magnitude of the beam–beam bias, as well as most of the uncertainties, scale with ξ , the uncertainties listed in Table 8 are all expressed at a common value ξ_{sim} of the beam–beam parameter; the details of this procedure can be found either in the table itself, or in Sects. 5.1 to 5.3. Because the chosen value of ξ_{sim} lies at the upper end of the range covered during Run-2 νdM scans, Table 8 illustrates something of a “worst-case” scenario with respect to the magnitude of the beam–beam correction and of the associated systematic uncertainty. The numerical results listed in that table should therefore not be applied blindly to estimating actual νdM -calibration uncertainties, but rather used as guidance for a case-by-case and bunch-by-bunch error analysis of actual νdM -calibration data.

5.1 Optical configuration

5.1.1 β^* uncertainty at the scanning IP

Since the σ_{vis} bias scales with the beam–beam parameter (Sect. 4.2), any measurement error on ξ directly translates into an error of similar relative magnitude on beam–beam corrections to the luminosity scale. Existing accelerator instrumentation unfortunately does not allow accurate enough a determination of the transverse single-beam sizes that enter the definition of ξ (Eq. (12)). The horizontal and vertical beam–beam parameters must therefore be redefined in terms of two-beam observables that can be reliably measured:

$$\xi_{xA\nu} = \frac{n r_0 \beta_{xA\nu}^*}{\pi \gamma \Sigma_x (\Sigma_x + \Sigma_y)} \tag{47}$$

$$\xi_{yA\nu} = \frac{n r_0 \beta_{yA\nu}^*}{\pi \gamma \Sigma_y (\Sigma_x + \Sigma_y)}. \tag{48}$$

Here, $\beta_{xA\nu}^* = (\beta_{x1}^* + \beta_{x2}^*)/2$ is the beam-averaged horizontal β -function at the scanning IP, and similarly for $\beta_{yA\nu}^*$. The quantity $n = (n_1 + n_2)/2$ refers to the beam-averaged bunch population associated with the colliding-bunch pair under study, $\gamma = \gamma_1 = \gamma_2$ is the common relativistic factor (assumed here to be the same for B1 and B2), and Σ_x and Σ_y are the convolved transverse bunch sizes. During νdM scans, the observables n , Σ_x and Σ_y are measured on a bunch-by-bunch basis to sub-percent accuracy [9,46–48], and the uncertainty on the absolute beam energy is totally negligible for the purposes of this paper. The main uncertainty affecting the absolute scale of the beam–beam parameter, therefore, is that associated with potential deviations of β^* from its nominal or assumed value.

The facts that the ξ dependence of the σ_{vis} bias is roughly linear (Fig. 16), and that the measured value of ξ_{xj} (Eq. (12)) depends linearly on n_k and β_{xj}^* ($j, k = 1, 2; j \neq k$), justify the use of the beam-averaged variables defined in Eqs. (47)–(48) above. Furthermore, the very weak dependence of the σ_{vis} bias on the Σ_y^0/Σ_x^0 aspect ratio (Fig. 22) indicates that “averaging” over x and y , in the sense of substituting the round-beam equivalent parameter ξ_R (Eq. (29)) for the pair of horizontal and vertical beam–beam parameters ($\xi_{xA\nu}$, $\xi_{yA\nu}$), results in an acceptably small bias. What finally matters, therefore, is the uncertainty on the beam- and plane-averaged value of β^* at the scanning IP.

The β^* values at the LHC IPs are extracted from optical-function measurements that are carried out for every colliding-beam configuration as part of routine accelerator commissioning. During the first few years of LHC operation, these values were typically obtained using the phase-advance method; k-modulation is now considered a more promising approach [54]. It targets an absolute accuracy of a few per-

Table 8 Typical systematic uncertainties affecting beam–beam corrections to a hypothetical pp νdM calibration in a fully symmetric Gaussian-beam configuration, with the round–beam–equivalent beam–beam parameter set equal to ξ_{sim} , for three values of N_{NSIP} . For each source, the uncertainty is either evaluated at, or scaled linearly to, the value of ξ_{sim} indicated in the second column; if no value of ξ_{sim} is specified, the uncertainty listed covers the full range of ξ values encountered during pp νdM scans at the LHC. When an uncertainty is assumption-dependent, the value flagged by an asterisk is that used in computing the total uncertainty; the latter is compared to the overall beam–beam correction itself in the bottom two rows of the Table. The rightmost column indicates the chapter(s) where the corresponding issues are discussed in detail

Beam–beam (b–b) uncertainty source	$\xi_{sim} [10^{-3}]$	Uncertainty-determination procedure	σ_{vis} uncertainty [%] for $N_{NSIP} =$	Comments	See Sect.
			0 1 2		
Absolute ξ scale: β^* uncertainty at the scanning IP	5.60	Vary β^* by $\pm 10\%$ in the simulation or parameterization (Sect. 4.2.3), for each beam and in each plane	0.06 0.10 0.13	β^* uncertainty assumed uncorrelated between beams, correlated between planes	4.2.1 + 5.1.1
Nominal collision tunes	5.60	Vary q_x, q_y by ± 0.002 in the simulation or parameterization, for each beam	0.26 0.23 0.20	Tune uncertainty assumed correlated between beams and between planes	4.2.2 + 5.1.2
Non-Gaussian transverse-density distributions	5.60	B*B (or COMBI) simulations	0.13 0.22 0.30	Simulated for $N_{NSIP} = 0$, extrapolated to $N_{NSIP} \geq 1$ using Eq. (42)	4.3 + 5.2.1
Beam ellipticity at the scanning IP	5.60	B*B (or COMBI) simulations. Uncertainty scaled linearly from ξ_R to ξ_{sim}	0.03 (for all values of NNSIP)	Simulated for $\xi_R \leq 4.2 \times 10^{-3}$, $0.7 < \Sigma_y / \Sigma_x < 1.4$	4.4 + 5.2.2
Non-zero crossing angle	≤ 5.60	COMBI simulations	$< 0.01^*$ (for all values of NNSIP) < 0.02 (for all values of NNSIP)	For $\theta_c \leq 10 \mu\text{rad}^*$ For $\theta_c \leq 150 \mu\text{rad}$	4.5 + 5.2.3
Beam-beam imbalance	5.60	B*B and COMBI simulations	0.016* 0.059 0.136	For $\sigma_2 / \sigma_1 > 0.95^*$ For $\sigma_2 / \sigma_1 > 0.90$ For $\sigma_2 / \sigma_1 > 0.85$	4.7 +
Multiple IPs:					5.2.4
Phase advance	5.60	COMBI (or B*B) simulations	0 < 0.20 (for $N_{NSIP} > 0$)	Worst case: arbitrary phase advances between IPs	4.6.4 +
Multi-IP tune shift	5.60	Vary p_1 in Eq. (42) by $\pm 15\%$ in single-IP simulations. Ignore if using multi-IP simulation	0 0.05		4.6.5 + 5.3
Long-range encounters	–	None at the scanning IP during pp νdM scans at the LHC	–		5.4.1
Lattice non-linearities	–	COMBI simulations, with sextupoles and octupoles included	0.01* (for all values of NNSIP) 0.03 (for all values of NNSIP)	For $E_B \geq 6.5 \text{ TeV}^*$ at lower energies	5.4.2
Numerical accuracy of parameterization	–		< 0.10 (for all values of NNSIP)	Ignore if using simulation rather than parameterization	5.4.3
Total uncertainty	5.60	Uncertainties summed in quadrature	± 0.32	% of σ_{vis}	5.5
Total b-b correction	5.60	Parameterization (Sects. 4.2.3 and 4.6.5)	+0.52 ± 0.41 +0.86	+1.17 % of σ_{vis}	5.5

cent on β_{iB}^* ($i = x, y; B = 1, 2$), provided sufficient beam time can be dedicated to the measurement.

The absolute β^* -accuracy requirements associated with beam–beam-correction uncertainties are rather modest. When beams collide at the scanning IP only, and for $\xi = 5.6 \cdot 10^{-3}$, a $\pm 10\%$ uncertainty in β^* at the scanning IP, uncorrelated between beams (because the optics of the two rings are almost completely independent) but correlated between planes (because a focusing error presumably affects both the horizontal and the vertical β -function) results in $\pm 0.06\%$ uncertainty in σ_{vis} (Table 8). Treating this uncertainty as fully uncorrelated (correlated) between beams and between planes would decrease (increase) the corresponding σ_{vis} uncertainty by a factor of $\sqrt{2}$. The absolute uncertainty increases to 0.10% (0.13%) for $N_{\text{NSIP}} = 1(2)$, but its magnitude relative to the overall beam–beam correction remains about the same.

The $\pm 10\%$ β^* -uncertainty assumed in the example above is very conservative, and is representative of the typical difference between the nominal, i.e. the dialed-in, β^* setting and the measured values. Smaller uncertainties can be achieved by calculating ξ based on the measured β^* values themselves, while accounting rigorously for the systematic β^* -measurement uncertainties and for their correlation between beams and between planes.

5.1.2 Nominal collision tunes

Figures 12 and 13 illustrate the sensitivity of the optical distortion and of the full beam–beam bias to the unperturbed fractional tunes. Any deviation of (q_x, q_y) from their nominal setting results in an inaccurate prediction of the σ_{vis} bias.

For the purposes of the present discussion, a conservative uncertainty of ± 0.002 is assigned to each of q_x and q_y , and is assumed to be correlated between beams (as could be caused by, for instance, instrumental effects) and between planes (because a focusing error presumably affects both the horizontal and the vertical betatron frequency). The magnitude of this uncertainty is based on LHC operational experience,¹⁸ and is meant to cover setting errors, as well as the stability of the unperturbed tunes, over the course of the vdM -calibration fill(s). The corresponding uncertainty in σ_{vis} (Table 8) dominates the overall beam–beam uncertainty. It decreases slightly when the number of non-scanning IPs increases, because the additional collisions lower the unperturbed tunes, thereby reducing the sensitivity of the optical distortions to the exact values of the fractional tunes. The same effect is responsible for the N_{NSIP} -dependent sensitivity of σ_{vis} to beam–beam asymmetries (Fig. 40).

¹⁸ While tunes can often be measured to 0.001 or better before beams are brought into collision, the beam–beam tune spread makes it quite challenging to monitor the unperturbed LHC tunes while the beams are colliding.

5.2 Deviations from the fully symmetric Gaussian-beam configuration

Except where specifically stated otherwise, most of the results presented in this paper, and in particular the parameterizations detailed in Sects. 4.2.3 and 4.6.5, are based on the assumptions listed in Sect. 4.1. Each of these will be lifted, one at a time, in the following sections, and their individual impact on beam–beam-correction uncertainties estimated using the studies detailed in Sects. 4.3–4.7.

5.2.1 Non-Gaussian unperturbed transverse-density profiles

Of the three single-bunch models described in Table 2, the 2012 parameter set represents an extreme case, both because this is the vdM session that revealed the largest non-factorization biases at LHC ever [9, 22], and because the parameters in that Table are deliberately chosen to represent the worst case among all scans and all bunches in a given vdM session. The 2018 model emerges as the most appropriate to estimate the uncertainty associated with non-Gaussian tails in the unperturbed density distributions: not only is the kurtosis intermediate between that of the other two parameter sets, but it also reflects the performance of the injector chain in the fourth year of LHC Run 2, by which time procedures had long stabilized and beam conditions could be considered reasonably reproducible.

Since the factorizable double-Gaussian configuration systematically yields a larger deviation from the corresponding single-Gaussian model than its non-factorizable counterpart (Table 4), it is favored for the purpose of estimating systematic errors. The uncertainty listed in Table 8 for $N_{\text{NSIP}} = 0$ is therefore taken from the top half of Table 4. It is extrapolated to the cases of one and two non-scanning IPs by assuming that the fractional impact of additional collisions on the σ_{vis} bias is dominated by the multi-IP tune shift (Sect. 4.6.2), and thus insensitive to the shape of the unperturbed transverse-density distributions.

The choice of the 2018 factorizable model to estimate the impact of non-Gaussian tails is not devoid of some arbitrariness. The uncertainties listed in Table 8, therefore, should be considered as a first indication, albeit a prudent one. A refined evaluation would require a more extensive characterization of the transverse-density distributions extracted from non-factorization analyses, that could then be used to produce a wider palette of more realistic single-bunch models for input to B*B or COMBI. In order to explore potential correlations between systematic biases, one could in addition extend the existing studies to simulate bunches that are not purely Gaussian and that in addition collide at multiple interaction points.

5.2.2 Beam ellipticity at the scanning IP

The dependence of the optical-distortion bias on the Σ_y/Σ_x aspect ratio is summarized by Fig. 22, and does not exceed $\pm 0.025\%$ for $\xi_R \leq 4.2 \cdot 10^{-3}$. Scaling this linearly to the value of ξ_{sim} shown in Table 8 yields a beam–beam-correction uncertainty of 0.03%, under the condition that the orbit-shift correction take the actual ellipticity into account, i.e. that it be computed using the Bassetti–Erskine formalism [14] with, as input, the measured values of Σ_x and Σ_y .

5.2.3 Non-zero crossing angle

The crossing-angle dependence of the beam–beam bias in the *vdM* regime is summarized by Fig. 23. The associated systematic uncertainty is taken as the difference in σ_{vis} bias between $\theta_c = 0$ and the actual crossing-angle value; it is listed in Table 8 for two values of the full crossing angle.

It should be emphasized that this small uncertainty is applicable only at moderate crossing angles ($\theta_c \leq 150 \mu\text{rad}$), and only in the *vdM* regime ($\xi \leq 0.006$, $\beta^* \geq 10\text{m}$). For larger crossing angles in the *vdM* regime, additional effects may come into play. This may also be the case for scans performed in the physics regime, even at moderate crossing angle, when the beam–beam parameter becomes large enough, and/or when β^* becomes small enough ($\beta^* \leq 30\text{cm}$) for the hourglass effect to no longer be negligible. Of particular concern are the so-called “emittance scans”, that are *vdM*-like beam-separation scans that are used to monitor luminometer performance and transverse emittances during routine LHC operation [55]: the study of their sensitivity to beam–beam-induced biases remains mostly unexplored territory.

5.2.4 Beam–beam imbalance

Of the three beam–beam-asymmetry scenarios considered in Sect. 4.7, only the transverse-emittance imbalance (Sect. 4.7.3) requires special attention, since bunch-current and β^* -asymmetries are accounted for by the use of the beam-averaged beam–beam parameters defined in Eqs. (47)–(48). The emittance-asymmetry dependence of the σ_{vis} bias is controlled by the combination of ξ_R and of the transverse beam-size ratio σ_2/σ_1 (Eq. (46)); it also is sensitive to the number of non-scanning IPs (Fig. 40).

The beam–beam-imbalance uncertainty is taken as the difference in σ_{vis} bias between an asymmetric beam–beam configuration characterized by a given value of $\sigma_2/\sigma_1 < 1$, and the corresponding symmetric configuration ($\sigma_2 = \sigma_1$); it is listed in Table 8 for three values of the transverse beam-size ratio. The ratio closest to unity ($\sigma_2/\sigma_1 > 0.95$), that emerged as the most representative of typical beam conditions during the 2015–2018 *pp vdM* sessions [9], is the one used

when computing the total beam–beam uncertainty (penultimate row of Table 8). The N_{NSIP} dependence is discussed in Sect. 4.7.3, and is closely related to the tune-dependence of the optical-distortion bias (Sect. 5.1.2).

5.3 Multiple interaction points

Two sources of uncertainty affect the beam–beam correction strategies in the presence of multi-IP collisions, that are detailed in Sect. 4.6.6: the phase advance between consecutive IPs (Sect. 4.6.4), and the parameterization of the multi-IP equivalent tune shift ΔQ_{mIP} (Sect. 4.6.5). These uncertainties affect only those colliding-bunch pairs in which at least one of the two opposing bunches experiences collisions at one or more IPs other than the scanning IP ($N_{\text{NSIP}} > 0$).

The phase-advance uncertainty arises from the currently unresolved ambiguity in the choice of the reference “no beam–beam” configuration, or, in more technical terms, from the ambiguity between the \mathcal{L}^0 and the \mathcal{L}'' normalizations (Sect. 4.6.3). In the absence of any phase-advance information, e.g. for a bunch pattern that mixes different collision configurations that are all treated on the same footing, a blanket systematic uncertainty of at most 0.2% appears adequate (Sect. 4.6.4). This uncertainty can, in some cases, be significantly lowered by simulating the actual collision pattern associated with each colliding-bunch pair, with as input the actual IP-to-IP phase advance per beam and per plane, and by then comparing the results between the two reference no-beam configurations. An illustration of the potential gain provided by such a refined analysis is offered by the results presented in Table 7 for the $N_{\text{NSIP}} = 1$ case, a configuration that applies to the 13 TeV *pp vdM* scans at IP1 and IP5 in 2015, 2016 and 2017 [9].

The uncertainty associated with the equivalent multi-IP tune shift ΔQ_{mIP} needs to be considered only when invoking the parameterization of single-IP simulations proposed in Sect. 4.6.5 and encapsulated in Eq. (42), to compute beam–beam corrections in the presence of additional collisions. This parametrization, that is illustrated by the magenta curve in Fig. 34, is in good agreement with the simulation results when IP5 is the scanning IP. The discrepancies that are apparent for scans at some of the other IPs, reflect the variety of phase-advance patterns between the scanning IP and the other collision points. Varying the p_1 parameter in Eq. (42) by $\pm 15\%$ of its central value covers all the collision patterns shown in Fig. 34, thereby providing a measure of the uncertainty associated with this simplified procedure (Table 8). More precise results, and a smaller uncertainty, could again be obtained by simulating separately each of the collision patterns associated with the B1 and B2 bunch strings circulating in the LHC ring during the *vdM*-calibration session under analysis.

5.4 Other potential sources of systematic uncertainty

5.4.1 Long-range encounters

During pp vdM scans at the LHC, the bunches are longitudinally isolated, in the sense that consecutive bunches are separated by more than 500 ns, rather than grouped in trains with only 25 ns spacing. In addition, the fill pattern is systematically optimized to avoid parasitic crossings, at least at the scanning IP. This strategy eliminates separation-dependent, long-range beam–beam kicks that might distort the luminosity-scan curves.¹⁹

In contrast, during high-luminosity pp physics running, consecutive bunches are typically separated longitudinally by only 25 ns, the emittance is smaller than during vdM scans, and the bunch current significantly larger. As a result, and in spite of the non-zero crossing angle, long-range beam–beam effects are known to significantly impact emittance scans in several ways, most of which remain to be understood.

5.4.2 Lattice non-linearities

The strongest non-linearity is introduced by the beam–beam interaction itself. However, non-linear magnetic elements in the LHC rings also affect beam dynamics, potentially leading to additional distortions of the transverse particle-density distributions. Because the LHC relies on the combination of high chromaticity ($Q' \approx 10$ –15 units) and of high Landau-octupole currents to mitigate some strong impedance-driven instabilities, the next largest optical non-linearities at beam energies above 1 TeV are associated with the corresponding sextupoles and octupoles [56]. As demonstrated experimentally [57–59], machine imperfections contribute even less by comparison, and are therefore neglected in the present study.

In order to quantify the impact of these magnets on the σ_{vis} bias, both the chromaticity and the linear detuning with amplitude introduced by the Landau octupoles were added to the COMBI model. The chromaticity dependence of the beam–beam bias is presented in Fig. 41 for two different beam energies. As expected, the impact of the octupoles is more important in the 2011 case because of the lower beam energy: at zero chromaticity, powering the octupoles changes the beam–beam bias from -0.155% to -0.130% , a reduction in relative magnitude of about one sixth. In contrast, in the 2018 configuration the effect of the octupoles is three times

¹⁹ The situation is different during heavy-ion (HI) scans, which for operational reasons are carried out in bunch-train mode with a non-zero nominal crossing angle. However, the bunch charge is considerably smaller than during pp vdM scans, and the beam–beam parameter about an order of magnitude smaller. Parasitic crossings, therefore, are no issue from the viewpoint of beam–beam corrections to HI vdM calibrations.




Fig. 41 Chromaticity dependence of the σ_{vis} bias predicted by COMBI for pp vdM scans, with octupoles either under power (solid curves) or left at zero current (dashed curves). The red curves correspond to the 2011 configuration ($E_B = 3.5$ TeV), the blue curves to that in 2018 ($E_B = 6.5$ TeV); the corresponding octupole currents are indicated in the legend. The momentum spread and the synchrotron tune are set to, respectively, 1.12×10^{-4} and 2×10^{-3} . The beam–beam parameter is set to $\xi = 2.6 \times 10^{-3}$, and the beams collide at the scanning IP only

smaller, and almost negligible. As for the chromaticity, its impact is negligible in both configurations.

The very low sensitivity of beam–beam corrections to these non-linear elements can be understood as follows. Octupoles act more strongly on larger-amplitude particles. These barely contribute to the luminosity when the beams collide head-on; the fractional contribution of the tails to the overlap integral is highest at large beam separation, but there the luminosity is the lowest. This is why modifying the trajectory of the tail particles barely influences the separation-integrated FoMs, and therefore the σ_{vis} bias.

Similarly, high chromaticity potentially enhances synchrobetatron coupling; however since vdM scans are carried out with relaxed β^* settings and at zero crossing angle, this is of no concern.

Even though lattice non-linearities are almost negligible in vdM scans (Table 8), it remains important to check systematically the potential influence of the octupoles under operational scenarios where they are pushed to maximum power to ensure beam stability: their impact on beam–beam uncertainties may turn out to be no longer negligible, for instance during emittance scans.

5.4.3 Numerical accuracy of polynomial parameterization

Propagating to the σ_{vis} bias the parameterization uncertainties that affect the separation-dependent luminosity-bias factors (Eq. (28)) and that are discussed in Sect. 4.2.3, results in an upper limit of 0.1% on the associated beam–beam uncer-

tainty. This contribution should be ignored when resorting to a full B*B or COMBI simulation at each scan point, rather than to the parameterization.

5.5 Summary of systematic uncertainties

Adding in quadrature all the above uncertainties yields, for the example examined here ($\xi_{\text{sim}} = 5.6 \cdot 10^{-3}$), a total beam–beam uncertainty of 0.32–0.46%, depending on the number of non-scanning IPs (Table 8). The corresponding value of the beam–beam correction varies from 0.52 to 1.17%. The absolute magnitude of the total uncertainty grows when N_{NSIP} increases from 0 to 2, but its magnitude relative to the correction itself decreases from about 60% to roughly 40% through the mechanism discussed in Sects. 5.1.2 and 5.2.4.

Under operating conditions typical of actual *vdM*-calibration sessions, the number of non-scanning IPs is at least one. For $N_{\text{NSIP}} = 1$, the four dominant uncertainty contributions, in order of decreasing importance, are associated with the nominal tunes, non-Gaussian tails, multi-IP effects (in the absence of collision–pattern-specific studies), and the β^* uncertainty. Estimating the uncertainty associated with non-Gaussian tails, and verifying whether beam-size asymmetries might contribute at a significant level, imperatively requires input from the non-factorization analysis of luminous-region data.

It is worth noting that the larger N_{NSIP} , the larger the contribution of the orbit-shift correction relative to that of the optical distortion. Qualitatively therefore, estimating precisely the effects that contribute only to the optical-distortion uncertainty, viz. the beam ellipticity and the beam-size asymmetry, becomes less critical as the number of non-scanning IPs increases. In contrast, the β^* and tune uncertainties affect the magnitude of both the orbit shift and the optical distortion, in a fully correlated manner.

The uncertainties detailed in Table 8 have been deliberately estimated in a conservative fashion, and for the largest value of the beam–beam parameter encountered during Run-2 *vdM* scans in *pp* collisions at $\sqrt{s} = 13$ TeV. Most uncertainties scale with ξ , either linearly or with a slightly positive quadratic component. For the purposes of a rough estimate, therefore, they can be safely extrapolated linearly to smaller ξ values, that are more representative of typical *vdM* conditions. A more precise determination of these uncertainties requires a case-by-case analysis, possibly augmented by some of the additional simulation studies suggested in some of the preceding sections in this chapter.

6 Summary and outlook

Under conditions typical of *vdM* luminosity-calibration scans in *pp* collisions at the LHC, beam–beam-induced orbit shifts

and optical distortions, if left uncorrected, bias the absolute luminosity scale by an amount comparable to, or larger than, the total luminosity-uncertainty budget. The magnitude of the correction, which ranges from approximately 0.2 to 1.2%, mainly depends on the beam–beam parameter ξ and on the number of non-scanning IPs.

The contributions to the associated systematic uncertainty are listed in Table 8 and elaborated upon in Sect. 5. The total uncertainty amounts to roughly half of the correction itself, and is dominated by tune-related effects, either directly (accuracy and reproducibility of nominal-tune settings) or indirectly (beam–beam tune shift at non-scanning IPs, phase advance between consecutive IPs). The next most important source of uncertainty is the potential deviation of the unperturbed transverse bunch-density distributions from a perfectly Gaussian shape. Simulating the impact of non-Gaussian tails cannot be based on luminosity data alone: it requires, as input, single-bunch parameters that can only be extracted from the measured beam-separation dependence, during the scan, not only of the luminosity but also of the position, shape and orientation of the luminous region. These single-bunch parameters are also crucial to quantify transverse beam-size differences between the two beams, which, if large enough, may become a significant source of uncertainty.

The scope of the present paper is limited to beam-separation scans in the *vdM* regime ($\xi \leq 6 \cdot 10^{-3}$), under beam conditions that deviate only moderately from round, initially Gaussian bunches of equal brightness that collide with a zero or small nominal crossing angle. In the course of investigating the impact of actual departures from that idealized limit, several areas of further study have been identified, most notably:

- improved single-bunch models of non-Gaussian tails (Sect. 4.3), both factorizable and non-factorizable, that are more sophisticated than the simple double-Gaussian functions considered in Eqs. (37) and (38);
- the interplay between non-factorization, non-Gaussian tails, and beam–beam corrections, and in particular the application of beam–beam corrections to two-dimensional grid scans;
- the ambiguity affecting the choice of the reference “no beam–beam” configuration, i.e. the phase-dependence of the difference between the \mathcal{L}^0 and the \mathcal{L}^u normalizations (Sect. 4.6);
- the characterization of optical distortions at large nominal crossing angle (Sect. 4.5) in the *vdM* regime;
- the potential interplay between the impact of non-Gaussian beams (Sect. 4.3), beam ellipticity (Sect. 4.4), crossing angle (Sect. 4.5), and multi-IP effects (Sect. 4.6), that have so far been treated as fully uncorrelated;

- the extension of the ξ -scaling laws (Sect. 4.2) to the physics regime ($\xi \leq 10^{-2}$), with emittance scans as the primary use case;
- the combined impact, during emittance scans at low β^* in the physics regime, of higher-beam–beam parameter values and of the hourglass effect on the crossing-angle dependence of beam–beam biases;
- the impact of parasitic crossings on the absolute accuracy of emittance scans during routine physics running;
- the impact of lattice non-linearities on the absolute accuracy of emittance scans during routine physics running.

Acknowledgements We thank our colleagues in the LHC Luminosity Monitoring and Calibration Working Group (LLCMWG), and in particular V. Balagura, X. Buffat, R. Hawkings, W. Herr and J. Weninger, for many fruitful discussions and for valuable comments on part or all of the manuscript. We are also indebted to V. Balagura for teaching us how to run his B*B package, as well as for implementing additional functionalities needed by the studies reported above. Some of the authors have received support from either the Swiss Accelerator Research and Technology Institute (CHART), the U.S. Department of Energy (DOE Contract No. DE-AC02-76SF00515), or the Hungarian National Research, Development and Innovation Office (NKFIH Contract Nos. 124845 and 143460).

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: The parameterization outlined in Sec. 4.2.3 and detailed in Appendix C (Tables 10 and 11) is accessible in computer-readable form in the Zenodo repository, at the URL specified in Ref. [45].]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

Appendix A: Impact of non-Gaussian distortions of the source field

Because B*B is fundamentally a weak-strong model, beam–beam-induced shape-distortions of the witness bunch, i.e. of the multiparticle distribution, influence neither the shape of the source bunch nor the spatial dependence of the electromagnetic field it generates. The latter is computed once and for all at initialization time, and remains static both turn-by-turn and as a function of beam separation.

COMBI, in contrast, is based on a strong-strong model: both bunches are represented by macroparticles; both play,

with respect to their partner, the dual roles of source and witness bunch; and the electromagnetic fields they generate are updated on a turn-by-turn basis, and therefore also at each scan step. In order to estimate the impact of separation-dependent, non-Gaussian field distortions caused by the mutual interaction of the opposing bunches, the σ_{vis} bias (Eq. (35)) has been computed using COMBI in two different modes:

- for each witness bunch in turn, the field that its macroparticles are subjected to, is computed under the assumption that the transverse macroparticle distribution of the source bunch is a two-dimensional Gaussian, with the position of its centroid and its RMS widths computed on every turn from the actual spatial distribution of the “source” macroparticles. This procedure tracks the evolution of the relative position and of the width of the opposing multiparticle distributions, but the field calculation ignores potential deviations of the bunches from their initial, perfectly Gaussian shape. This is the mode in which all the COMBI results in this paper have been obtained.
- for each witness bunch, the field it is subjected to is computed by the HFMM method [42] from the actual charge distribution in the source bunch. In this mode, each macroparticle in the witness bunch is subjected to an electromagnetic field computed as the vector sum of the contributions of all the macroparticles in the source bunch. Since this computation must be repeated on every turn and for the opposing bunch, it is much more CPU-intensive than the Gaussian-field approximation above.

Figure 42 displays the beam–beam-parameter dependence of the σ_{vis} bias over a ξ range that extends all the way up to 25×10^{-3} , more than four times the highest ξ value encountered in νdM scans, and significantly higher than the largest per-IP beam–beam parameter that is sustained operationally in the actual LHC rings. The two calculations agree within 1% or better; for the ξ range and the nominal tunes considered in Table 1, this translates into a σ_{vis} difference of less than 0.005%. Computing the electromagnetic field as that of an equivalent Gaussian, the parameters of which evolve turn by turn, therefore constitutes a fully justified approximation whose impact on beam–beam corrections can be safely neglected.




Fig. 42 Top: ξ dependence of the beam–beam induced σ_{vis} bias, computed either in the Gaussian-kick approximation (solid black line), or with a self-consistent kick calculated using the HFMM field solver (filled orange circles). Bottom: ξ dependence of the ratio between the σ_{vis} -bias values computed by the two methods

Appendix B: Luminosity integration in COMBI

In COMBI, the calculation of the luminosity, or more precisely of the overlap integral (Eq. (1)), is based on a complete description of the macroparticle distribution in the two colliding bunches: no assumption is made about the shape of the density distributions when evaluating this integral. The latter is calculated on a turn-by-turn basis, both at the scanning IP and at the non-scanning IP(s), if any. The positions of all the macroparticles are updated at each IP on every consecutive turn, and each step in a beam-separation scan is initialized and simulated separately.

The transverse distribution of the macroparticles, called $H_{x,y}^B$, is discretized, separately for the two beams B ($B = 1, 2$), on a two-dimensional grid in the x - y plane. The number of cells in each grid is given by $N_{cells} = n \times m$, where n and m represent the number of bins in the x and y directions respectively. The grid boundaries are located at a distance $k_i \times \sigma_i^0$ ($i = x, y$) from the center of the grid, where σ_i^0 is the initial, unperturbed nominal transverse beam size inferred from the input emittance and β^* values in the i plane; the scale factor k_i is typically set to 12 in both the positive and the negative direction along the x and y axes. The cell area is thus given by

$$\Delta S = \Delta x \times \Delta y = 2k_x \sigma_x^0 / n \times 2k_y \sigma_y^0 / m.$$

The separation between the two beams is taken into account when filling two-dimensional histograms of the macroparticle distributions.




Fig. 43 Visible cross-section bias as a function of the number of macroparticles used to model the transverse-density distributions

Table 9 Parameter space used in the determination of the $[\mathcal{L}/\mathcal{L}^0]_{Opt}$ polynomial coefficients

Horizontal scans	Vertical scans
$0.002 \leq \xi_R \leq 0.007$	
$0.2975 \leq q_x \leq 0.3100$	$q_y - 0.0125 \leq q_x \leq q_y - 0.0075$
$q_x + 0.0075 \leq q_y \leq q_x + 0.0125$	$0.3075 \leq q_y \leq 0.3200$

With these definitions, the discretized macroparticle density distribution is given by

$$h_{x,y}^B = \frac{H_{x,y}^B}{N_{part} \Delta S},$$

where N_{part} is the total number of tracked macroparticles. The overlap density of the two bunches, i.e. the product $\hat{\rho}_1(x, y) \hat{\rho}_2(x, y)$ in Eq. (1), is therefore represented by:

$$\lambda_{x,y} = h_{x,y}^1 h_{x,y}^2.$$

The bunch luminosity can be calculated from the overlap integral I_{ovlp} as:

$$\mathcal{L}_b = f_r n_1 n_2 I_{ovlp}, \tag{49}$$

where n_1 and n_2 are the bunch populations and f_r the revolution frequency. The overlap integral is estimated by the two-dimensional trapezoidal method:

$$I_{ovlp} \simeq \frac{1}{4} \Delta x \Delta y \left(\lambda_{0,0} + \lambda_{m,0} + \lambda_{0,n} + \lambda_{m,n} + 2 \sum_{i=1}^{m-1} \lambda_{i,0} + 2 \sum_{i=1}^{m-1} \lambda_{i,n} + 2 \sum_{j=1}^{n-1} \lambda_{0,j} + 2 \sum_{j=1}^{n-1} \lambda_{m,j} + 4 \sum_{j=1}^{n-1} \left(\sum_{i=1}^{m-1} \lambda_{i,j} \right) \right).$$

Table 10 Parameterization of the $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ luminosity-bias factor as a function of the normalized nominal separation, for horizontal vdM scans. The functional form is defined in Eq. (50)

$\Delta_k [\sigma_0]$	p_{0k}	p_{1k}	p_{2k}	p_{3k}	p_{4k}	p_{5k}	p_{6k}	p_{7k}	p_{8k}	p_{9k}
0.00	1.00944	-9.9169	0.05163	-0.10946	-11.343	0.27849	0.51631	17.907	17.933	-0.70617
0.25	1.0965	-11.2335	-0.37109	-0.25749	-15.156	0.8075	0.59757	21.37	19.207	-0.38121
0.50	1.33057	-14.2614	-1.41804	-0.75044	-23.614	2.23701	1.09972	29.115	22.489	0.21938
0.75	1.6732	-18.353	-3.14234	-1.28652	-32.777	4.71339	1.58581	37.476	28.876	0.9846
1.00	2.09627	-23.4765	-4.63004	-2.55265	-38.221	3.68123	0.36501	50.199	34.146	7.71504
1.25	2.35234	-28.0037	-5.75125	-3.10402	-38.905	3.96085	-0.21768	57.845	41.93	10.74776
1.50	2.5998	-30.8986	-6.20145	-4.24446	-36.073	8.14783	4.83319	59.119	50.075	4.06889
1.75	2.71221	-32.1903	-5.80054	-5.33954	-33.326	6.80135	5.92624	54.751	58.075	5.3906
2.00	2.68384	-31.9924	-4.60368	-6.30519	-25.565	5.92361	8.49128	47.367	63.795	3.25386
2.25	2.72577	-30.6931	-4.15275	-7.00444	-15.931	4.41796	8.97022	42.853	63.188	4.6426
2.50	2.71357	-29.8333	-3.31831	-7.7255	-17.797	-0.48377	6.79757	35.712	67.066	11.47357
2.75	2.55328	-28.5073	-2.93845	-7.07616	-19.558	2.2039	8.83346	28.805	69.588	5.10042
3.00	2.63081	-28.984	-3.48143	-7.0339	-13.235	1.7513	7.59532	37.803	62.493	7.5962
3.25	2.51794	-29.2876	-3.37658	-6.41585	-9.579	2.73257	7.63313	38.021	63.687	5.42281
3.50	2.57552	-30.733	-3.87832	-6.30735	-5.79	3.30687	7.27443	49.126	58.156	5.87508
3.75	2.7132	-34.1604	-4.72201	-6.34866	0.329	4.88899	7.5148	62.801	56.443	5.47604
4.00	2.80473	-36.4992	-6.20828	-5.49987	-0.487	1.89557	1.09204	77.16	50.842	15.98952
4.25	2.84402	-40.437	-7.01717	-4.94963	5.671	5.45919	2.26449	89.531	51.974	11.6835
4.50	2.92941	-42.1662	-7.69945	-4.84568	7.414	12.10344	7.25898	100.023	47.916	1.02831
4.75	3.00531	-44.2824	-8.94959	-4.12019	4.043	8.3604	0.8071	119.399	36.595	12.11431
5.00	3.1174	-47.4622	-9.45872	-4.32462	10.289	7.49102	-0.57319	125.029	41.518	15.50293
5.25	3.08992	-48.0346	-9.55933	-4.06527	9.796	14.24273	5.29776	139.651	29.582	2.65662
5.50	3.22545	-51.1449	-10.7063	-3.79323	9.956	14.35667	3.04625	141.488	38.024	6.20924
5.75	3.25358	-51.7947	-11.62684	-3.08922	7.899	12.51096	-1.17893	151.272	30.874	12.68129
6.00	3.15344	-50.9016	-10.84895	-3.228	7.762	19.8657	7.03849	145.734	33.547	-3.89875

The instantaneous luminosity in Eq. (49) is calculated at each turn, and the result is averaged over the total number of selected turns. Typically a few hundred turns are necessary for the result to stabilize. The reliability of this method was confirmed by benchmarking it against analytical calculations, with the beam–beam effects turned off.

Figure 43 shows the evolution of the σ_{vis} bias as a function of the number of macroparticles used in the simulation. The result converges to within 0.01% of its asymptotic value for 5×10^6 macroparticles. Most of the simulation results presented in this paper are based on 10×10^6 macroparticles per bunch, implying that the results are numerically stable at the 0.001% level.

The uncertainty associated with statistical fluctuations in the discretization of the transverse-density distributions was evaluated separately, and typically cancels out when computing luminosity ratios such as $[\mathcal{L}/\mathcal{L}^0]$. It becomes significant only at large beam separation, when the overlap integral is computed from a small number of macroparticles in the tails. At these scan points, however, the luminosity values are very

small, and therefore have a negligible impact on the estimation of the beam–beam bias factors.

Appendix C: Parameterization of optical-distortion corrections

The polynomial parameters p_{0k}, \dots, p_{9k} from which to calculate the optical-distortion luminosity-bias factor

$$\begin{aligned}
 [\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\xi_R, q_x, q_y | \Delta_k) = & p_{0k} + p_{1k} \xi_R \\
 & + p_{2k} q_x + p_{3k} q_y + p_{4k} \xi_R^2 + p_{5k} q_x^2 + p_{6k} q_y^2 \\
 & + p_{7k} \xi_R q_x + p_{8k} \xi_R q_y + p_{9k} q_x q_y
 \end{aligned}
 \tag{50}$$

defined in Eq. (28) of Sect. 4.2.3, have been tabulated, for normalized nominal-separation values covering the range $0 < |\Delta_k/\sigma_0| < 6$ in steps of 0.25. Here the index k ($k = 1, 25$) labels the nominal-separation bin Δ_k in the scanning plane considered. This parameterization is valid in the fully symmetric Gaussian-beam configuration with zero

Table 11 Parameterization of the $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ luminosity-bias factor as a function of the normalized nominal separation, for vertical vdM scans. The functional form is defined in Eq. (50)

$\Delta_k [\sigma_0]$	p_{0k}	p_{1k}	p_{2k}	p_{3k}	p_{4k}	p_{5k}	p_{6k}	p_{7k}	p_{8k}	p_{9k}
0.00	1.00627	-9.8954	0.02224	-0.06076	-11.359	0.0568	0.18622	17.901	17.872	-0.18399
0.25	1.20592	-12.0782	-0.02827	-1.29247	-17.815	0.1942	2.21186	19.01	24.368	-0.28862
0.50	1.75762	-17.6722	-0.35287	-4.51951	-30.505	-0.9264	5.84403	22.471	40.37	2.91383
0.75	2.47795	-24.582	-0.92069	-8.60293	-41.735	1.3187	13.56868	20.768	65.956	0.45508
1.00	3.16343	-32.8235	-1.16788	-12.75235	-55.11	-0.3834	18.27023	29.422	85.709	4.49981
1.25	3.74028	-38.7998	-2.36278	-15.28112	-48.5	0.3875	21.25607	39.51	95.711	6.74987
1.50	3.96499	-42.0778	-3.06102	-16.01822	-37.845	-1.1981	19.73221	36.839	108.559	12.15308
1.75	3.69522	-40.042	-3.62485	-13.73364	-31.695	-1.4007	14.95999	39.742	98.388	14.38483
2.00	3.17695	-34.3599	-3.98479	-10.06481	-23.614	5.9537	15.52287	48.181	70.742	1.19056
2.25	2.54203	-27.8488	-4.42917	-5.57026	-15.667	5.2247	6.89102	46.516	50.153	4.09602
2.50	1.93646	-21.3641	-4.22219	-1.90718	-15.994	5.1956	1.34853	47.557	27.611	3.48476
2.75	1.64684	-18.4268	-4.50556	0.22396	-11.197	5.0949	-2.57366	49.259	16.174	4.5606
3.00	1.58403	-17.9608	-4.4421	0.56159	-7.635	6.066	-2.11206	48.03	16.126	2.48878
3.25	1.84422	-21.9484	-4.56541	-0.96411	-2.223	3.7746	-2.00762	49.005	28.617	7.28473
3.50	2.32149	-27.3088	-4.57251	-4.00249	1.528	4.2728	3.31839	48.531	47.316	6.31746
3.75	2.86007	-35.4916	-3.79586	-8.17387	13.02	7.1856	13.85154	44.04	78.901	-1.77626
4.00	3.52241	-43.4275	-3.74323	-12.44714	15.761	2.4115	16.29906	43.702	105.987	7.26694
4.25	4.18491	-52.8584	-3.75654	-16.63881	17.478	2.9322	23.4521	43.277	137.933	6.26266
4.50	4.82121	-61.3184	-4.12935	-20.31864	25.852	-1.7901	24.22835	37.512	171.614	16.6589
4.75	5.33928	-68.7429	-4.41872	-23.34458	24.7	3.8166	33.91458	37.937	196.185	6.67136
5.00	5.80246	-73.1287	-3.19177	-27.51882	20.495	18.7373	56.57897	38.918	210.435	-26.21613
5.25	6.44223	-79.2346	-5.10161	-29.74714	33.649	17.5226	56.18781	50.413	219.306	-17.92498
5.50	7.17335	-86.1747	-5.53665	-33.9533	35.037	6.9886	52.15683	35.086	257.011	3.99827
5.75	6.8188	-87.1346	-4.86936	-32.37045	28.729	-1.12	43.18943	42.423	253.701	17.49413
6.00	6.41896	-87.7396	-0.88027	-33.73032	32.935	1.0047	53.37414	28.725	269.298	0.87602

non-scanning IPs, over the three-dimensional space of fractional tunes (q_x, q_y) and round-beam-equivalent beam-beam parameter (ξ_R) values defined in Table 9. The values of β^* , γ , σ_0 and n in Eq. (13) were chosen such that they correspond to typical settings for pp vdM -calibration sessions during Run 2 of the LHC.

The coefficients p_{0k}, \dots, p_{9k} are listed in Tables 10 and 11 for horizontal and vertical vdM scans, respectively; they are also publicly accessible, in computer-readable form, in Ref. [45]. In the unabridged tables, i.e. for the studies reported in Sects. 4.2.3 and 4.2.4, all parameters were stored with a number of digits corresponding to the precision of the IEEE 754 double-precision floating-point data format. In the tables below, the coefficients are listed with enough precision to reproduce the values of $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ in the unabridged tables to at least five decimal digits. For verification purposes, 50×10^3 points were randomly chosen across the (ξ_R, q_x, q_y) space defined in Table 9, and for each of these, 200 Δ values were randomly generated. The resulting $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}$ values, obtained using the unabridged tables or the ones below, were then compared. The maximum differences found in 10×10^6

evaluations were 5.4×10^{-6} and 9.5×10^{-6} for horizontal and vertical scans, respectively.

The procedure to correct an x - y pair of experimentally measured vdM -scan curves for beam-beam-induced optical distortions is detailed below. It is valid only for on-axis, one-dimensional vdM scans with zero crossing angle; off-axis one-dimensional scans, as well as two-dimensional grid scans in the (δ_x, δ_y) beam-separation plane, require dedicated simulations.

1. For a given colliding-bunch pair, calculate the values of σ_R and ξ_R from the measured bunch-current and uncorrected (Σ_x, Σ_y) values, following the prescriptions in Sect. 4.2.3;
2. for that same bunch pair, calculate the values of the effective fractional tunes (q_x, q_y) , following the prescription in Sect. 4.6.5 with, as input, the nominal tune values (Q_x, Q_y) , the appropriate number N_{NSIP} of non-scanning IPs, and the resulting multi-IP equivalent tune shift ΔQ_{mlP} ;
3. for each scanning plane (x and y) and each tabulated value of the nominal separation Δ_k , calculate the cor-

responding value of $\{[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}\}_k = [\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\xi_R, q_x, q_y, |\Delta_k)$ using Eq. (50), with the polynomial coefficients taken from the relevant row of Table 10 or 11, as appropriate;

4. use all pairs of numbers $(\Delta_k, \{[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}\}_k)$ as input to a suitably chosen numerical-interpolation algorithm²⁰ that can return the value of $[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\Delta/\sigma_R)$ for any nominal separation Δ ;
5. at each scan step, multiply the measured collision rate by $1/[\mathcal{L}/\mathcal{L}^0]_{\text{Opt}}(\Delta/\sigma_R)$, where Δ is the nominal separation for the scan step and in the scanning plane considered. This final step yields, for the bunch pair considered, the horizontal and vertical $v dM$ -scan curves corrected for optical distortions.

References

1. P. Grafström, W. Kozanecki, Luminosity determination at proton colliders. *Prog. Part. Nucl. Phys.* **81**, 97–148 (2015). <https://doi.org/10.1016/j.pnpnp.2014.11.002>
2. S. van der Meer, Calibration of the effective beam height in the ISR. CERN-ISR-PO-68-31 (1968). <http://cds.cern.ch/record/296752>
3. C. Rubbia, Measurement of the luminosity of pp collider with a (generalized) van der Meer method. CERN-pp-Note-38 (1977). <http://cds.cern.ch/record/1025746>
4. ATLAS Collaboration, Precision measurement and interpretation of inclusive W^+ , W^- and Z/γ^* production cross sections with the ATLAS detector. *Eur. Phys. J. C* **77**, 367 (2017). <https://doi.org/10.1140/epjc/s10052-017-4911-9>
5. CMS Collaboration, Measurements of differential Z boson production cross sections in proton–proton collisions at $\sqrt{s} = 13$ TeV. *JHEP* **12**, 61 (2019). [https://doi.org/10.1007/JHEP12\(2019\)061](https://doi.org/10.1007/JHEP12(2019)061)
6. CMS Collaboration, Measurements of the W boson rapidity, helicity, double-differential cross sections, and charge asymmetry in pp collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D* **102**, 092012 (2020). <https://doi.org/10.1103/PhysRevD.102.092012>
7. ATLAS Collaboration, Inclusive and differential cross-sections for dilepton $t\bar{t}$ production measured in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. *JHEP* **7**, 141 (2023). [https://doi.org/10.1007/JHEP07\(2023\)141](https://doi.org/10.1007/JHEP07(2023)141)
8. CMS Collaboration, Measurement of differential $t\bar{t}$ production cross sections in the full kinematic range using lepton+jets events from proton–proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D* **104**, 092013 (2021). <https://doi.org/10.1103/PhysRevD.104.092013>
9. ATLAS Collaboration, Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. CERN-EP-2022-281, *Eur. Phys. J. C* **83** (2023) 982. <https://doi.org/10.1140/epjc/s10052-023-11747-w>. [arXiv:2212.09379](https://arxiv.org/abs/2212.09379)
10. A. Dainese et al., Report on the physics at the HL-LHC, and perspectives for the HE-LHC. Technical report. CERN-2019-007. Geneva, Switzerland (2019). <https://doi.org/10.23731/CYRM-2019-007>
11. CMS Collaboration, The phase-2 upgrade of the CMS beam radiation instrumentation and luminosity detectors. Technical report. CERN, Geneva (2021). <https://cds.cern.ch/record/2759074>
12. ATLAS Collaboration, Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C* **73**, 2518 (2013). <https://doi.org/10.1140/epjc/s10052-013-2518-3>
13. W. Herr, Beam–beam effects and dynamic β^* , in *Proc. LHC Lumi Days* (2012). https://indico.cern.ch/event/162948/contributions/1417430/attachments/191879/269237/S3_WH.pdf
14. M. Bassetti, G.A. Erskine, Closed expression for the electrical field of a two-dimensional Gaussian charge. CERN-ISRTH/80-06 (1980). <https://cds.cern.ch/record/122227>
15. MAD (Methodical Accelerator Design) home page. <https://madx.web.cern.ch/madx/>
16. V. Balagura, van der Meer scan luminosity measurement and beam–beam correction. *Eur. Phys. J. C* **81**, 26 (2021). <https://doi.org/10.1140/epjc/s10052-021-08837-y>
17. COMBI (COherent Multibunch Beam-beam Interactions) home page. <https://twiki.cern.ch/twiki/bin/view/ABPComputing/COMBI>
18. A. Chao et al., *Handbook of Accelerator Physics and Engineering*, 2nd edn. (World Scientific, Singapore, 2013). <https://doi.org/10.1142/8543>
19. V. Balagura, Notes on van der Meer scan for absolute luminosity measurement. *Nucl. Instrum. Methods A* **654**, 634–638 (2011). <https://doi.org/10.1016/j.nima.2011.06.007>
20. Y. Cai, Luminosity of asymmetric e+e– collider with coupling lattices. SLAC-PUB-8479 (2000). [arXiv:hep-ph/0007052](https://arxiv.org/abs/hep-ph/0007052)
21. K.M. Potter, J.C. Sens (eds.), Luminosity and particle detection studies at the CERN intersecting storage rings. Internal Report CERN-NP-71-6 (1971). <https://cds.cern.ch/record/1161781/files/NP-71-6.pdf>
22. ATLAS Collaboration, Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C* **76**, 653 (2016). <https://doi.org/10.1140/epjc/s10052-016-4466-1>
23. W. Herr, B. Muratori, Proceedings of the CAS-CERN Accelerator School: Intermediate Accelerator Physics, Zeuthen, Germany, 15–26 September 2003, edited by D. Brandt, CERN-2006-002 (CERN, Geneva, 2006), pp. 361–378, DOI: e-proceedings: <https://doi.org/10.5170/CERN-2006-002>
24. S. Peggs, Beam–beam collisions and crossing angles in RHIC, in *BNL-66426. Workshop on Beam–Beam Effects in Large Hadron Colliders, CERN, 12–16 April 1999*. https://inis.iaea.org/collection/NCLCollectionStore/_Public/30/046/30046444.pdf?r=1
25. J. Wenninger, Operation and configuration of the LHC in Run 2. CERN-ACC-NOTE-2019-0007 (2019). <https://cds.cern.ch/record/2668326/files/CERN-ACC-NOTE-2019-0007.pdf>
26. P. Bambade et al., Observation of beam–beam deflections at the interaction point of the SLAC Linear Collider. *Phys. Rev. Lett.* **62**, 2949–2952 (1989). <https://doi.org/10.1103/PhysRevLett.62.2949>
27. V. Ziemann, Beyond Bassetti and Erskine: beam–beam deflections for non-Gaussian beams. SLAC-PUB-5582 (1991). <https://www.slac.stanford.edu/pubs/slacpubs/5500/slac-pub-5582.pdf>
28. W. Kozanecki, T. Pieloni, J. Wenninger, Observation of beam–beam deflections with LHC orbit data. CERN-ACC-NOTE-2013-0006 (2013). <https://cds.cern.ch/record/1581723>
29. S.-Y. Lee, *Accelerator Physics*, 4th edn. (World Scientific, Singapore, 2018). <https://doi.org/10.1142/11111>
30. W. Herr, T. Pieloni, in *Proceedings of the CAS-CERN Accelerator School: Advanced Accelerator Physics, Trondheim, Norway, (2013), edited by W. Herr, CERN-2014-009 (CERN, Geneva, 2014)*, pp. 431–456. <http://dx.doi.org/10.5170/CERN-2014-009>
31. A.W. Chao, P. Bambade, W.T. Weng, Nonlinear beam–beam resonances, in *Joint US/CERN School on Particle Accelerators: Topical Course on Nonlinear Dynamics. SLAC-PUB-3545 (1985)*. <http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-pub-3545.pdf>

²⁰ Cubic-spline interpolation is recommended.

32. A. W. Chao, Beam-Beam Instability, SLAC-PUB-3179 (1983). https://inis.iaea.org/collection/NCLCollectionStore/_Public/15/021/15021887.pdf
33. R. Tomás et al., Beam–beam amplitude detuning with forced oscillations. *Phys. Rev. Accel. Beams* **20**, 101002 (2017). <https://doi.org/10.1103/PhysRevAccelBeams.20.101002>
34. T. Pieloni, W. Herr, Coherent beam–beam modes in the CERN Large Hadron Collider (LHC) for multiple bunches, different collision schemes and machine symmetries, in *Proc. of Particle Accelerator Conference (PAC2005)*, Knoxville, TN, USA (2005), pp. 4030–4032. <https://accelconf.web.cern.ch/p05/PAPERS/TPAT078.PDF>
35. T. Pieloni, A study of beam–beam effects in hadron colliders with a large number of bunches. PhD thesis. Ecole Polytechnique Federale de Lausanne (EPFL) (2008). <https://cds.cern.ch/record/1259906/files/CERN-Thesis-2008-Pieloni.pdf>
36. T. Pieloni, W. Herr, Models to study multibunch coupling through head-on and long-range beam–beam interactions. CERN LHC Project Report 937, and Proc. EPAC06, Edinburgh, UK (2006). <https://cds.cern.ch/record/972663>
37. F. Jones, W. Herr, T. Pieloni, Parallel beam–beam simulation incorporating multiple bunches and multiple interaction regions. CERN LHC Project Report 1038, and Proc. 2007 Particle Accelerator Conference, Albuquerque, NM, USA. <https://accelconf.web.cern.ch/p07/PAPERS/THPAN007.PDF>
38. X. Buffat, Transverse beams stability studies at the Large Hadron Collider. PhD thesis. Ecole Polytechnique Federale de Lausanne (EPFL) (2015). <https://doi.org/10.5075/epfl-thesis-6321>
39. Open MPI home page. <https://www.open-mpi.org>
40. K. Hirata, Who is afraid of a crossing angle?, in *Proc. of European Particle Accelerator Conference (EPAC 94)*, London, UK (1994), pp. 190–192. https://accelconf.web.cern.ch/e94/PDF/EPAC1994_0190.PDF
41. L. Barraud, X. Buffat, Mode coupling instability of colliding beams in the HL-LHC. CERN-ACC-NOTE-2019-0032 (2019). <https://cds.cern.ch/record/2684699>
42. W. Herr, M.P. Zorzano, F. Jones, Hybrid fast multipole method applied to beam–beam collisions in the strong–strong regime. *Phys. Rev. ST Accel. Beams* (2001). <https://doi.org/10.1103/PhysRevSTAB.4.054402>
43. CMS Collaboration, CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV. CMS-PAS-18-002 (2019). <https://cds.cern.ch/record/2676164/files/LUM-18-002-pas.pdf>
44. CMS Collaboration, CMS luminosity measurement using nucleus–nucleus collisions at $\sqrt{s_{NN}} = 5.02$ TeV in 2018. CMS-PAS-LUM-18-001 (2022). <http://cds.cern.ch/record/2809613/files/LUM-18-001-pas.pdf>
45. A. Babaev et al., Parameterization of beam–beam-induced optical distortions of van der Meer scans at the LHC. Zenodo (2023). <https://doi.org/10.5281/zenodo.8033453>
46. CMS Collaboration, Precision luminosity measurement in proton–proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 in CMS. *Eur. Phys. J. C* **81**, 800 (2021). <https://doi.org/10.1140/epjc/s10052-021-09538-2>
47. LHCb Collaboration, Precision luminosity measurements at LHCb. *J. Instrum.* **9** (12), P12005 (2014). <https://doi.org/10.1088/1748-0221/9/12/P12005>
48. ALICE Collaboration, ALICE 2016–2017–2018 luminosity determination for pp collisions at $\sqrt{s} = 13$ TeV. ALICE-PUBLIC-2021-005 (2021). <https://cds.cern.ch/record/2776672>
49. H. Bartosik, G. Rumolo, Production of single Gaussian bunches for Van der Meer scans in the LHC injector chain. CERN-ACC-NOTE-2013-0008 (2013). <http://cds.cern.ch/record/1590405>
50. R.L. Workman et al., Review of Particle Physics. *PTEP* **2022**, 083C01 (2022). <https://doi.org/10.1093/ptep/ptac097>
51. A. Babaev, Coherent deflection of elliptic bunches colliding at crossing angle (2021). <https://doi.org/10.48550/ARXIV.2104.02595>
52. K. Yokoya et al., Tune shift of coherent beam–beam oscillations. *Part. Accel.* **27**, 181–186 (1990). <https://cds.cern.ch/record/200504>
53. D. Brandt et al., Is LEP beam–beam limited at its highest energy?, in *Proc. of Particle Accelerator Conference (PAC 1999)*, New York, NY, USA. (1999), pp. 3005–3007. <https://accelconf.web.cern.ch/p99/PAPERS/THP25.pdf>
54. F. Carlier, R. Tomás, Accuracy and feasibility of the β^* measurement for LHC and high-luminosity LHC using k modulation. *Phys. Rev. Accel. Beams* **20**, 011005 (2017). <https://doi.org/10.1103/PhysRevAccelBeams.20.011005>
55. O. Karacheban, Y.C. Cekmecelioglu, Emittance scans for CMS luminosity calibration in run 2. PoS EPS-HEP2019 (2020), p. 193. <https://doi.org/10.22323/1.364.0193>
56. E. Métral et al., Measurement and interpretation of transverse beam instabilities in the CERN Large Hadron Collider (LHC) and extrapolations to HL-LHC, in *57th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams. 2016, TUAM2X01* (2016). <https://doi.org/10.18429/JACoW-HB2016-TUAM2X01>
57. F.S. Carlier et al., MD2723—amplitude detuning studies at 6.5 TeV with various configurations of the crossing scheme. CERN-ACC-NOTE-2018-0022 (2018). <https://cds.cern.ch/record/2306325>
58. J.W. Dilly et al., Report and analysis from LHC MD 3311: amplitude detuning at end-of-squeeze. CERN-ACC-NOTE-2019-0042 (2019). <https://cds.cern.ch/record/2692810>
59. E.H. Maclean et al., New approach to LHC optics commissioning for the nonlinear era. *Phys. Rev. Accel. Beams* **22**, 061004 (2019). <https://doi.org/10.1103/PhysRevAccelBeams.22.061004>