UNIVERSITY OF OXFORD

Improvement of optics control in the Large Hadron Collider with forced Resonance Driving Terms



Sasha Horney

Jesus College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2026

Abstract

This thesis presents the development, validation, and operational implementation of a framework for the routine measurement and correction of non-linear optics in the Large Hadron Collider (LHC), based on forced resonance driving terms (RDTs). Non-linear magnetic fields, arising from both intentional lattice elements and intrinsic imperfections, play a critical role in limiting beam lifetime, dynamic aperture, and ultimately collider luminosity. Standard approaches to the characterisation of these effects have primarily relied on global observables. While these quantities provide valuable information, they do not directly quantify the strength of individual resonances. In contrast, the use of RDTs enables a more complete and localised picture of the relevant beam dynamics.

RDTs have traditionally been studied using free-beam excitation, which is subject to several limitations, most notably decoherence effects that degrade measurement quality. These limitations can be overcome through the use of sustained excitation provided by an AC-dipole, enabling the acquisition of high-quality Turn-by-Turn data and the precise extraction of forced RDTs. This methodology affords direct access to individual resonances and facilitates systematic benchmarking against model predictions. The framework is first applied to sextupolar resonances at injection energy, where an unexpectedly large contribution is identified and subsequently corrected, resulting in measurable improvements in beam performance.

The framework is subsequently extended to higher-order effects, including octupolar and dodecapolar resonances. Novel diagnostic and correction techniques are introduced, such as a non-linear Segment-by-Segment method for localising error sources. A feed-down-based strategy is also created, enabling efficient measurement and compensation of higher-order field errors without directly measuring higher-order RDTs. These developments significantly reduce the complexity and time required for characterisation, making routine operational implementation

feasible.

In addition, this work extends beyond the single-particle framework to incorporate multi-particle effects, in particular long-range beam-beam interactions. For the first time, direct correction strategies for long-range beam-beam-driven non-linearities, using insertion-region magnets, are experimentally demonstrated. Furthermore, the first successful corrections of β-beating induced by long-range beam-beam interactions are presented.

This thesis establishes forced RDT analysis as a robust and versatile tool for beam-based diagnostics and correction in strongly non-linear machines. The results provide both immediate operational benefits for the LHC and essential methodologies for the High-Luminosity LHC, where increased sensitivity to non-linear effects demands advanced control strategies.

Conference Proceedings

The following Conference Proceedings are based on work presented in this thesis:

- S. Horney et al, "Correction of long-range beam-beam driven optics perturbations for the LHC", **THP4005**
- S. Horney et al, "Non-linear resonance feed-down - a new technique for correcting high order errors in the LHC", **THP4006**
- S. Horney et al, “Energy Sensitivity of the High Luminosity LHC Optics at the End of $\boldsymbol{\beta}^*$ squeeze”, **MOPM018**
- S. Horney et al, “Correction of Long-Range Beam-Beam driven normal sextupolar RDTS”, **MOPM019**
- S. Horney et al, “Investigation of octupolar resonances in the LHC”, **MOPM020**
- S. Horney et al, “Sextupole RDTs in the LHC at injection and in the ramp”, **MOPC13**

Dedication and acknowledgements

I am indebted to a great many people whose support, guidance, and encouragement made this work possible. It's hard to do them justice in a few lines, and I hope that I will not be leaving anyone out, but I'll give it a proper go.

First and foremost, my deepest thanks go to my CERN supervisor, Ewen. Despite an alarming number of professional commitments, he still gave me the opportunity to undertake this DPhil, even after previously surviving the experience of supervising me as a technical student, and somehow continued to invest a remarkable amount of time in my development. His patience in explaining complicated ideas, his help preparing for presentations, and his steady guidance across all manner of responsibilities have been invaluable throughout and beyond the studentship.

I am equally grateful to my university supervisor, Phil, for enabling me to pursue this DPhil at the University of Oxford. Despite juggling a long list of leadership roles, he consistently made time for discussions, offering both reassurance and perspective at the moments they were needed.

My sincere thanks also go to my section leader, Rogelio, for his continued support and insight. His guidance and enthusiasm has been instrumental in shaping this work, and has left a clear imprint on the final form of this thesis.

There are also several honourable mentions for people who have helped along the way: Laurent Deniau for his assistance with MAD-NG, and Yannis Papaphilippou and Rhodri Jones for their comments on my various IPAC papers over the years.

This research would not have been possible without the invaluable contributions of colleagues who participated in dedicated machine development periods and commissioning shifts. In particular, I am deeply grateful to the members of the Optics Measurements and Corrections team, some of whose work also formed a fundamental foundation and/or essential complement to my studies, as well as to the Operations team for their continued support. The collaborative spirit,

and collective endurance, fostered during long and often unconventional hours remain among the most memorable aspects of this experience.

I am also grateful to the many friends and colleagues I've met along the way, who made my time here both enjoyable and memorable. Whether during lunches, coffee breaks, other social gatherings, or more competitive pursuits like snooker, skiing, GeoGuessr, and, most recently, poker (where I have very much enjoyed crushing you all), their company has been a constant highlight. To the wine drinkers: I fully expect to see you again at the Cave Ouvertes, and for the others I am sure we will catch up again at some other time. And to those friends from before the DPhil that I've somehow managed to stay in touch with despite distance - thank you for sticking around.

Finally, I would like to thank my partner and my family, whose unwavering support and encouragement have been fundamental to getting this over the line.

List of Figures

2.1 Diagram illustrating the Frenet-Serret coordinate system (orange). The blue arrow indicates the distance s_{dist} travelled by the particle from s_0 to s_1 along a circular orbit of radius R. The green vector $\vec{r}_0$ (green) is the polar vector coordinates of the particle at its initial position. 8

2.2 Schematic of multipoles up to order $n = 6$, where some example magnetic field lines are in black, and the red arrows indicate the direction of force on a positively charged particle moving out of the page. 11

2.3 Schematic of a simple FODO cell where QF denotes focusing quadrupole and QD is a defocusing quadrupole. Green denotes the betatron envelope, while pink denotes a possible particle trajectory within the envelope. 12

2.4 Diagram showing the key coordinates of a trace-space ellipse. 15

2.5 Illustration of the principle of dispersion showing an on-momentum (blue) and an off-momentum (pink) particle trajectory through a bending magnet. 16

2.6 Illustration of a beam-based DA measurement showing the surviving beam intensity following single-kick excitations of increasing amplitude (expressed in units of beam sigma) at injection energy. The predicted (blue) and measured (red) losses are shown [1]. 21

2.7 Diagram illustrating the principle of decoherence by showing the trajectories of three representative particles as functions of s (particle A in blue, particle B in pink, particle C in green) and their combined average trajectory in orange. 22

2.8 Diagram illustrating beam decoherence following a single-kick excitation by showing the beam oscillation amplitude as a function of turn number. 23

2.9 Tune diagram (Q_y vs Q_x) showing the resonance conditions, with the leading resonance order indicated by the following colours: orange for second order, green for third order, cyan for fourth order, pink for fifth order and dashed grey for sixth order. 23

2.10 Distortion of normalised phase-space due to non-linearities. 24

2.11 Diagram showing an example of the transformation of linearly normalised phase-space coordinates into normal form. 32

3.1 The overall layout of the LHC injector complex [2]. 37

3.2 The layout of the LHC with LHCB1 in blue and LHCB2 in red. 38

3.3 A visual representation of the naming conventions for hardware placement in the LHC, where LSS denotes Long Straight Section. 38

3.4 Illustration of an LHC FODO cell. 39

3.5 Cross-section of a LHC dipole [3]. 41

3.6 Cross-section of a LHC quadrupole [4]. 42

3.7 Simulated β-function in the ATLAS IR with *round* beams i.e. β^* = 0.6/0.6 m (horizontal/vertical plane). The schematic above shows the strengths of the magnets in the lattice where: $Q1$ to $Q3$ are the triplet quadrupoles responsible for beam focusing; $D1$, and $D2$ are the separation/recombination dipoles which guide each beam from its individual aperture into the common beam pipe around the IP and back into its own aperture on the far side; and $Q4$ to $Q7$ are the matching quadrupoles. 44

3.8 Similar to Fig. 3.7 but with *flat* beams i.e. β^* = 0.60/0.18 m (horizontal/vertical plane). 45

3.9 The layout of the IR and the triplet corrector package, where $Q1$, $Q2a/b$ and $Q3$ are the triplet quadrupoles and $C0$-$C3$ are the corrector packages with the order to which they can correct listed below. $D1$ is the separation dipole used to split LHCB1 and LHCB2. 45

3.10 Simulated horizontal trajectories x vs s depicting a crossing-angle of 160 μrad for β^* = 0.6/0.18 m (horizontal/vertical) at 6.8 TeV. 46

3.11 Depiction of the LHC cycle for a typical fill in 2025, with the various operational stages highlighted in different colours. The evolution of β^* is shown in pink (left-hand axis), the beam energy, defined to be the energy per nucleon, is indicated by the black dashed line (first right-hand axis), and the current in a Landau octupole circuit is represented in green (second right-hand axis), all as functions of time. 49

3.13 Diagram illustrating the resulting Turn-by-Turn data retrieved using an ACD excitation. 55

3.14 Illustration of the LHC collimator hierarchy. 56

3.15 Example horizontal TbT frequency spectrum (amplitude vs frequency) where green denotes sextupolar RDTs, cyan the octupolar RDTs, and red the ACD frequencies. . . 57

3.16 Illustration of the principle of the 3-BPM method with the chosen BPM highlighted in pink and the nearby BPMs used in the analysis highlighted in green. 58

3.17 Illustration of the principle of the N-BPM method. 59

4.1 The LHCB1 BLM lifetime in red, with the LHCB1 Q_y target in grey vs time during the 2023 ion run. 64

4.2 Vertical TbT frequency spectrum, where green denotes third-order RDTs, and red the ACD frequencies with natural tunes $q_x = 0.28$ and $q_y = 0.31$ for LHCB1. 65

4.3 Tune diagram (Q_y vs Q_x) showing the $2Q_y+Q_x$ resonance line in blue, and the nominal WP as a red star. 66

4.4 Vertical TbT frequency spectrum, where green denotes third-order and cyan the fourth-order RDTs, and red the ACD frequencies with natural tunes $q_x = 0.275$ and $q_y = 0.293$ for LHCB1. The unlabelled line at 0.02 corresponds to another sextupolar RDT. 67

4.5 Measured amplitude of f_{1020} vs s for LHCB1 at e-cloud (blue) and sextupolar resonance measurement (green) tunes for LHCB1. 67

4.6 Measured amplitude of f_{1020} vs s for LHCB1 at sextupolar resonance measurement tunes with new (green) and old (purple) optics for LHCB1. 68

4.7 Measured amplitude of f_{1020} vs s for LHCB1 at sextupolar resonance measurement tunes for 2022 (dark green) and 2023 (purple) for LHCB1. 68

4.8 Comparison of amplitude of f_{1020} vs s, between measurement (dark blue for LHCB1, dark red for LHCB2) and MAD-X simulation (light blue for LHCB1, light red for LHCB2). 69

4.9 f_{1020} amplitude vs s, comparing with (blue) and without (pink) any error models for LHCB1 in MAD-X simulation. 70

4.10 f_{1020} amplitude vs s, comparing an ACD-driven (pink) and a free-kick MAD-X simulation (purple) for LHCB1. 71

4.11 f_{1020} amplitude vs s, comparing a MAD-NG (light green) and a MAD-X free-kick simulation (purple) for LHCB1. 71

4.12 f_{1020} amplitude vs s, comparing with (light green) and without (orange) the MS powered for LHCB1 in MAD-NG simulation. 71

4.13 f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (pink) and without (blue) correction for all normal sextupolar RDTs with no constraints on powerings for LHCB1 in MAD-NG simulation. 74

4.14 Simulated horizontal chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (pink) and without (blue) correction. 74

4.15 f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (red) and without (blue) correction for only f_{1020} RDT with no constraints on powerings for LHCB1 in MAD-NG simulation. 75

4.16 Simulated vertical chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (red) and without (blue) correction. 75

4.17 Final MS circuit powerings, KSF in pink and KSD in green, used for the sextupolar correction, with their initial powerings shown in a paler version of the same colour. . 77

4.19 Measured amplitude of f_{1020} vs s for LHCB1, with (orange) and without (blue) correction. 77

4.18 f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (orange) and without (blue) final correction for LHCB1 in MAD-NG simulation. 78

4.22 Fraction of surviving beam intensity vs action for vertical kicks with (orange) and without (blue) f_{1020} correction. 81

4.23 The effect of f_{1020} correction on lifetime vs time with raw signal in pale red (left-hand axis), and the corresponding sextupolar current in blue (right-hand axis). 82

4.24 Amplitude of f_{1020} RDT vs action for LHCB1 (blue) and LHCB2 (red). 83

4.25 2024 amplitude of f_{1020} RDT vs s for LHCB2, during the initial 35 minutes after arriving at injection, where purple indicates the first minutes after reaching injection, going to yellow which is around 35 minutes after reaching the injection plateau. . . . 85

4.26 2024 average amplitude of f_{1020} RDT vs s for LHCB2, during the initial 35 minutes after arriving at the injection plateau, where the line is simply a visual guide to see the trend. 85

4.27 2025 amplitude of f_{1020} RDT vs s for LHCB2, between 10 and 45 minutes after arriving at the injection plateau where purple indicates 10 minutes after injection, going to yellow which is around 45 minutes after reaching the injection plateau. 86

4.28 2025 average amplitude of f_{1020} RDT vs s for LHCB2, between 10 and 45 minutes after arriving at the injection plateau, where the line is simply a visual guide to see the trend. 86

4.29 2025 amplitude of f_{1020} RDT vs s for LHCB2, between 16 hours (yellow) at injection and 40 minutes (purple) after arriving at injection. 87

4.30 2025 amplitude of f_{0030} RDT vs s for LHCB2, between 16 hours (yellow) at injection and 40 minutes (purple) after arriving at injection. 88

4.31 Units of b_3 expected to be generated in HL-LHC by a new separation dipole LMBRD1 magnet for the inner and outer apertures (AP1 (blue) and AP2 (orange)) [5]. 88

4.32 Units of b_3 expected to be generated in HL-LHC by some of the superconducting quadrupole LMQXFB magnets [6]. 89

4.33 2023 measurement of $\langle|f_{1020}|\rangle$ (mean f_{1020} amplitude) vs time of LHCB1 (light red) and LHCB2 (red), the lines are shown only as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy, in blue. . . . 89

4.34 2022 (green), 2023 (light red), 2024 (pink) and 2025 (purple) measurement of $\langle|f_{1020}|\rangle$ vs time for LHCB1, the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy, in blue. . 90

4.35 2023 measurement (light red) of $\langle|f_{1020}|\rangle$ vs time of LHCB1 compared with simulation (pink), the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy (blue). 90

4.36 2023 measurement of $\langle|f_{0030}|\rangle$ (mean f_{0030} amplitude) vs time of LHCB1 (light red) and LHCB2 (red), the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy (blue). . . 91

5.1 Q_y vs Q_x where the blue lines correspond to the resonances $4Q_x$, $4Q_y$, and $2Q_x - 2Q_y$, and the red shape is the tune spread/footprint caused by the MO at injection with 2023 optics configuration. 95

5.2 BLM lifetime at injection energy vs time for LHCB1 (red, left-hand axis) following a change in MO current (grey, right-hand axis). Measured using pilot bunches during optics commissioning. See Sec 3.8.3 for discussion on lifetime. 96

5.3 Horizontal TbT frequency spectrum (amplitude vs frequency), where green denotes third-order and cyan the fourth-order RDTs, and red the ACD frequencies with natural tunes of $q_x = 0.275$ and $q_y = 0.293$ for LHCB1 where the MO is powered at 5 m^{-4}. The other peaks are the natural tunes, other normal octupolar and sextupolar resonances showing the b_3 and b_4 contributions. 97

5.4 2024 measured (former colour) and simulated (latter colour) amplitude of f_{2002} RDT vs s, with (purple/pink) and without (green/black) MO powered for LHCB1. 98

5.5 2024 measured (former colour) and simulated (latter colour) amplitude of f_{2002} RDT vs s, with (purple/pink) and without (green/black) MO powered for LHCB2. 98

5.8 2024 (dark blue) and 2023 (black) measured shift to imaginary part of f_{2002} RDT vs s with and without MO powered for LHCB1. 100

5.9 Comparison of free-kick MAD-X (dark green) and MAD-NG (orange) simulated amplitude of f_{2002} RDT vs s for LHCB1. 100

5.10 Comparison of simulation (pink) and measurement (dark blue) for the shift to the imaginary part of the f_{2002} RDT vs s caused by the MCOX trim for LHCB1. 101

5.11 Comparison of the measured shift to the imaginary part of the f_{2002} RDT vs s caused by the MCOX trim (dark blue) and its reverse (cyan) for LHCB1. 101

5.12 Comparison of the simulated (pink) and measured (dark blue) shift to the imaginary part of the f_{2002} RDT vs s caused by an alternating MO trim for LHCB1. 102

5.13 Comparison of amplitude of f_{2002} RDT vs s with e-cloud (green) and sextupolar resonance measurement (brown) tunes for LHCB1. 102

5.14 Comparison of the simulated (pink) and measured (dark blue) shift to the imaginary part of the f_{2002} RDT vs s with all the MO being powered at 5 m^{-4} for LHCB1 with the old optics. 103

5.15 Comparison of simulation (pink) and measurement (dark blue) for the difference in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} for LHCB1. 104

5.16 Comparison of simulation (teal) and measurement (dark red) for the difference in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} for LHCB2. 104

5.17 Comparison of measurements for the shift in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} (black), 2 m^{-4} (purple), and -2 m^{-4} (green) for LHCB1. 105

5.18 Shift to the imaginary part of the measured (dark blue) and simulated (pink) f_{2002} vs s caused by the MOF in arc12 being powered at 10 m^{-4} for LHCB1. 105

5.19 Shift to the imaginary part of measured (dark blue) and simulated (pink) f_{2002} vs s caused by the MOF in arc56 being powered at 10 m^{-4} for LHCB1. 105

5.20 Example SbS analysis for the phase error vs s with different optics i.e. β at IP5 = 3 m (green), 1 m (blue), 0.6 m (red) for LHCB1 with the dipole multipoles shown in red above [7]. 106

5.21 Comparison of the difference between simulation and propagated f_{2002} phase vs s for when all the MO were turned off (blue), and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1. 108

5.22 Comparison of the difference between simulation and propagated f_{4000} phase vs s for when all the MO were turned off (blue), and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1. 108

5.23 Comparison of the difference between simulation and propagated f_{4000} real part of the RDT vs s for when all the MO were turned off (blue) and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1. 109

5.24 Comparison of the difference between measurement and propagated f_{2002} imaginary part of the RDT vs s for when all the MO were switched off (blue) and when all the MO uniformly powered at 5 m^{-4} (orange) for LHCB1. 110

5.25 Comparison of the difference between measurement and propagated f_{2002} imaginary part of the RDT vs s for when all the MO were switched off (blue), when all the MO were uniformly powered at 2 m^{-4} (orange), and when all the MO were uniformly powered at -2 m^{-4} (green) for LHCB1. 110

5.26 Comparison of the measured (dark blue) and simulated (pink) shift to the imaginary part of the f_{2002} RDT vs s caused by the MO in arc67 being powered at 10 m^{-4} for LHCB1. 111

6.1 Simulated β-function vs s for LHCB1 (blue) and LHCB2 (red) in the ATLAS IR, with β^* = 0.6/0.18 m (horizontal/vertical plane) (shown also as Fig. 3.8). 113

6.2 Simulated horizontal trajectories x vs s for LHCB1 (blue) and LHCB2 (red) through the ATLAS IR, with the crossing-angle (θ) indicated. Magnet layout as in Fig. 3.10. . 115

6.3 Schematic layout of an LHC experimental insertion - right side of IP1 shown - including the triplet corrector package, as in Fig. 3.9. 115

6.4 Schematic layout of an HL-LHC experimental insertion - right side of IP1 shown - where $Q1a/b$, $Q2a/b$, and $Q3a/b$ are the triplet quadrupoles, and $C0$-$C3$ the corrector packages, with the multipole orders each can correct shown in the lower part of the figure. $D1$ is the separation dipole. 116

6.5 Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural q_x = 0.28, q_y = 0.31. 118

6.6 $\Delta\text{Amp}(f_{0030})$ RDT vs s for crossing-angles of 25 (blue) and 170 (green) μrad from a reference crossing-angle of 0μrad. Lines are included only to guide the eye. 119

6.7 $\langle\Delta\text{Amp}\rangle$ vs vertical IP2 crossing-angles for LHCB1, the line is included only as a visual guide to see the trend. 119

6.8 Demonstration of the matching of the measurements of both beams with the simulated response of specific correctors. This is shown through the depiction of the partial derivatives of the real and imaginary parts of the RDT with respect to the θ vs s for both simulation (black) and measurement (blue for LHCB1, red for LHCB2). Lines are included only to guide the eye. 120

6.9 Workflow for the creation of a RDT feed-down-based correction of higher-order errors using the GUI, where the links are to the relevant figure in this section. 121

6.10 Input tab of the RDTFD GUI, where measurements are entered to evaluate the RDT shift as a function of crossing-angle. 123

6.11 Validation tab of the RDTFD GUI where the measured response to crossing-angle can be displayed. This shows the BPM sub-tab depicting ΔRe and ΔIm vs θ at the BPM level. 125

6.12 Validation tab of the RDTFD GUI where the measured response to crossing-angle can be depicted. This shows the RDT sub-tab depicting the ΔAmp, ΔRe, and ΔIm vs s. . . 126

6.13 Validation tab of the RDTFD GUI where the measured response to crossing-angle can be depicted. This shows the RDT shift sub-tab depicting the ring-averaged ΔAmp vs crossing-angle, as well as $\frac{\partial\text{Re}}{\partial\theta}$ and $\frac{\partial\text{Im}}{\partial\theta}$ vs s. 127

6.14 Correction tab of the RDTFD GUI, with the Graph sub-tab selected, where measurements can be matched with simulated response. 128

6.15 Correction tab of the RDTFD GUI, with the Response sub-tab selected, where the simulated response to crossing-angle can be calculated. 129

6.16 Vertical TbT frequency spectrum, where green denotes 3rd order and cyan the 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.31$. 131

6.18 The ΔIm/ΔRe of the f_{0030} RDT with respect to a horizontal IP1 crossing orbit bump for the BPM.30L2.B1, where the line (orange) is the linear fit of the data. 131

6.17 The ΔAmp/ΔIm/ΔRe vs s for each horizontal IP1 crossing orbit bump each in a separate colour for LHCB1. Lines are included only to guide the eye. 132

6.19 The ⟨ΔAmplitude⟩ vs θ, as well as the real and imaginary partial derivatives vs s, of the f_{0030} RDT with respect to a horizontal IP1 crossing orbit bump for LHCB1. Lines are included only to guide the eye. 133

6.20 The matching of the simulated corrector response (black) to the measured (blue) response, using the partial derivatives of the real and imaginary parts of the f_{0030} RDT, vs s for LHCB1. Lines are included only to guide the eye. 134

6.21 ⟨ΔAmp⟩ of the f_{0030} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB1 with (blue) and without (pink) correction, the line is included only as a visual guide to see the trend. 134

6.22 ⟨ΔAmp⟩ of the f_{0030} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2 with (blue) and without (pink) correction, the line is included only as a visual guide to see the trend. 135

6.23 Measured amplitude of the f_{1012} RDT vs s with (blue) and without (pink) correction. Lines are included only to guide the eye. 135

6.24 Vertical TbT frequency spectrum, where green denotes 3rd order and pink denotes 5th order RDTs, and red the ACD frequencies, with natural q_x = 0.28, q_y = 0.31. . . . 136

6.25 ΔAmp(f_{0140}) RDT vs s for crossing-angles of 60 (purple) and 120 (teal) μrad from a reference crossing-angle of 0 μrad. Lines are included only to guide the eye. 136

6.26 ⟨ΔAmp⟩ of the f_{0140} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2, the line is included only as a visual guide to see the trend. Error bars are reflective of a greater variation in the RDT around the ring. 137

6.27 ⟨ΔAmp⟩ of the f_{0140} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2 for simulation (green) and measurement (pink), the line is included only as a visual guide to see the trend. 138

7.1 Diagram showing bunches from each beam, LHCB1 (blue) and LHCB2 (red), crossing at an experimental insertion and highlighting the difference between head-on (green) and long-range beam-beam (purple). 139

7.2 Beam-beam force (blue) for equal charges in each beam, shown in arbitrary units, vs the transverse distance from the beam centre in units of σ. The head-on interaction region is indicated by the green circle and arrows, and the LR interaction range is shown in purple. The region in which the beam-beam force is approximately linear is highlighted in light pink. 142

7.3 Bunch intensity vs bunch number for LRBB measurements depicting how nominal pilots (blue, right-hand axis) and bunch trains (red, left-hand axis) compare. 144

7.4 Horizontal TbT frequency spectrum, where cyan denotes 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.261$, $q_y = 0.31$. 145

7.5 Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.32$. 145

7.6 Amplitude of f_{0030} RDT vs s for LHCB2 with full (purple) and no LR (orange) interactions. 146

7.7 Amplitude of f_{4000} RDT vs s for LHCB2 with full (purple) and no LR (orange) interactions. 146

7.8 Amplitude of f_{0030} RDT vs s for LHCB2 with full LR, with (green) and without (purple) the LRBB wire. 147

7.9 Amplitude of f_{4000} RDT vs s for LHCB2 with full LR, with (green) and without (purple) the LRBB wire. 147

7.10 Simulated horizontal β-function vs s with the same schematic of magnet strengths as in Fig. 3.7 with the inclusion of the distribution of LRBB kicks in the model. 148

7.11 Amplitude of f_{0030} RDT vs s for LHCB2 with full LR in simulation (light red) and measurement (red). 149

7.12 Amplitude of f_{4000} RDT vs s for LHCB2 with full LR in simulation (light red) and measurement (red). 149

7.15 Measured β-beating vs s with (orange) and without (blue) correction for non-colliding beams in 2024. 151

7.16 Simulated effect of switching to flat-optics (blue) from round (purple) optics for f_{0030} RDT vs s for LHCB2. 152

7.17 Measured effect of switching to flat-optics (blue) from round (purple) optics for f_{0030} RDT vs s for LHCB2. 152

7.21 Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.32$. 154

7.20 Comparison of the diffusion maps for Q_y vs Q_x (top), and vertical $\sigma\sin(\theta)$ vs horizontal $\sigma\cos(\theta)$ initial particle amplitude (bottom) with no correction and with the a_3 correction for MO = 15 m^{-4}. 155

7.22 Measured amplitude of f_{0030} RDT vs s with (orange) and without (blue) correction for LHCB1 with full LR. 156

7.23 MCSSX (skew sextupolar corrector) current (R1 in dark purple and L1 in lighter purple) and beam lifetime (green) as a function of time. 156

7.24 Simulated horizontal amplitude detuning (horizontal tune vs horizontal action) for LHCB1 with full LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange). 158

7.25 Simulated vertical amplitude detuning (vertical tune vs vertical action) for LHCB1 with full LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange). 158

7.26 Simulated amplitude of f_{4000} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange). 159

7.27 Simulated amplitude of f_{0040} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange). 159

7.28 Simulated amplitude of f_{2002} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange). 159

7.29 Comparison of the diffusion maps for Q_y vs Q_x (top), and vertical ($\sigma\sin(\theta)$) vs horizontal ($\sigma\cos(\theta)$) initial particle amplitude (bottom) with no correction and with the b_4 correction for MO = 15 m^{-4}. 160

7.30 Horizontal TbT frequency spectrum, where green denotes 3rd order RDTs, cyan denotes 4th order RDTs, and red the ACD frequencies, with natural q_x = 0.261, q_y = 0.31. 161

7.31 Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, cyan denotes 4th order RDTs, and red the ACD frequencies, with natural q_x = 0.261, q_y = 0.277. . 161

7.32 Amplitude of f_{4000} RDT vs s for LHCB1 with full (purple) and half (blue) LR interactions.161

7.33 Measured amplitude of f_{4000} RDT vs s for LHCB1 with LR, with (orange) and without (blue) correction, q_x = 0.261, q_y = 0.31. 162

7.34 Measured amplitude of f_{0040} RDT vs s for LHCB1 with LR, with (orange) and without (blue) correction, q_x = 0.261, q_y = 0.277. 162

7.35 Histogram of f_{0040} RDT measured for LHCB1 with LR, with (orange) and without (blue) correction, q_x = 0.261, q_y = 0.277. 162

7.36 Horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 160 μrad. 163

7.37 Vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 160 μrad. 163

7.38 Horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 120 μrad. 164

7.39 Vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 120 μrad. 164

7.40 Simulated horizontal β-beating vs s with (orange) and without (blue) the Xsuite correction for LHCB1. 165

7.41 Simulated vertical β-beating vs s with (orange) and without (blue) the Xsuite correction for LHCB1. 165

7.42 Measured horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1 with crossing-angle = 160 μrad. 166

7.43 Measured vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1 with crossing-angle = 160 μrad. 167

7.44 Horizontal β-beating vs s for simulation (pink) and measurement (brown) for LHCB1 with full LR effects. 167

7.45 Vertical β-beating for simulation (pink) and measurement (brown) for LHCB1 with full LR effects. 168

7.46 Measured horizontal β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB1 with full LR effects. 168

7.47 Measured vertical β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB1 with full LR effects. 169

7.48 Horizontal β-beating vs s for simulation (pink) and measurement (brown) for LHCB2 with full LR effects. 169

7.49 Vertical β-beating vs s for simulation (pink) and measurement (brown) for LHCB2 with full LR effects. 170

7.50 Measured horizontal β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB2 with full LR effects. 170

7.51 Measured vertical β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB2 with full LR effects. 171

7.52 Measured horizontal β-beating vs s with the Xsuite (orange) and GUI (green) correction for LHCB2 with full LR effects. 171

7.53 Measured vertical β-beating vs s with the Xsuite (orange) and GUI (green) correction for LHCB2 with full LR effects. 172

8.1 Measured relative optics shift vs s for LHCB1 between 9 and 26 May 2022 for the same accelerator configuration with β^* = 30 cm (re-analysed from historic data). . . . 174

8.2 Simulated vertical β-beating (where the β-function with the quadrupole trim is defined to be $\beta_{y,\mathrm{dE}}$, and the β-function without it is β_y) for HL-LHCB1 vs s, corresponding to a $\Delta p/p$ offset of $\approx 2\times10^{-4}$ for round-optics. 177

8.3 Simulated maximum vertical β-beating (same definition as in Fig. 8.2) for HL-LHCB1 for a $\Delta p/p$ offset of $\approx 2\times10^{-4}$, for various quadrupole types in the HL-LHC, where orange signifies quadrupoles in the matching sections, blue the main arc quadrupoles, green the triplet magnets, and pink is all the quadrupoles. 177

8.4 Simulated vertical β-beating (where the β-function with the quadrupole trim is defined to be $\beta_{y,\mathrm{dE}}$, and the β-function without it is β_y) vs s, for HL-LHCB1 corresponding to a $\Delta p/p$ offset of $\approx 2\times10^{-4}$ for flat-optics. 178

8.5 $\Delta p/p$ vs s, calculated from the LHCB1 RF frequency data over a one-week period in August in 2024, filtered on flat-top periods; the different colours signify separation of shifts. 179

8.6 Maximum change in YASP-predicted $\Delta p/p$ over an hour vs time. Spring (green) and neap (black) tide times taken from [8]. 180

8.7 Distribution of the predicted tidal variation in $\Delta p/p$, predicted on several time intervals relevant to optics commissioning. 180

8.8 Simulated vertical off-momentum β-beating (where the β-function with the RF frequency trim is defined to be $\beta_{y,\mathrm{dp}}$, and the β-function without it is β_y) vs s, for HL-LHCB1 corresponding to a $\Delta p/p$ offset of $\approx 2\times10^{-4}$ for flat-optics. 181

8.9 Predicted maximum vertical off-momentum β-beating (same definition as for Fig. 8.8) HL-LHCB1 as a function of $\Delta p/p$ where the lines are a linear fit of the points. 182

8.10 Frequency distribution of the predicted maximum change in Δchrom β_y/β_y for HL-LHCB1 flat-optics over several different time intervals. 182

8.11 Simulation of the SbS phase advance vs s, starting from BPM.10L5, expected to be generated over a 2 hour period at the worst time during the spring tide with flat-optics. The line is included only to guide the eye. 183

LIST OF TABLES

2.1 Selected resonance driving terms and the corresponding lines in the horizontal and vertical TbT spectra. 34

7.1 Corrector trim strengths for the b_3 correction. The maximum available current is 100 A. RCSX denotes a normal sextupole corrector 151

7.2 Corrector trim strengths for the b_3 (RCSX) and a_3 (RCSSX) correction where the maximum current available is 100 A. 153

7.3 Amplitude detuning values for direct and cross terms for different configurations, i.e. without LR without correction (orange), with LR without correction (blue), and with LR with correction (green). 158

7.4 Corrector trim strengths for the b_4 correction where the maximum current = 100 A. . 159

7.5 Simulated luminosity imbalance for each crossing-angle. 165

7.6 Simulated luminosity imbalance for each crossing-angle and correction. 166

7.7 Measured average reduction in β-beating ($\langle\Delta\beta$-beating$\rangle$) due to the correction. 168

7.8 Measured average change in β-beating ($\langle\Delta\beta$-beating$\rangle$) due to the correction. 170

7.9 Measured average change in β-beating ($\langle\Delta\beta$-beating$\rangle$) due to each correction. 171

TABLE OF CONTENTS

Page

List of Acronyms **1**

1 Introduction **4**

1.1 Thesis Outline 5

2 Linear and Non-linear Optics **8**

2.1 Coordinate System and Reference Trajectory 8

2.2 Beam Rigidity 9

2.3 Magnet Definition 10

2.4 Linear Optics 12

2.4.1 Twiss Parameters 13

2.4.2 Betatron Tune 14

2.4.3 Emittance, Phase and Action 14

2.4.4 Dispersion 15

2.4.5 Chromatic Optics 16

2.4.6 Transfer Maps 17

2.5 Non-linear Optics 20

2.5.1 Dynamic Aperture 20

2.5.2 Amplitude Detuning 21

2.5.3 Decoherence 22

2.5.4 Resonances 23

2.5.5 Hamiltonian Formalism 24

2.5.6 Non-linear Transfer Maps . . . 26
2.5.7 Normal Form . . . 31
2.6 Luminosity . . . 35

3 The LHC **36**
3.1 The LHC Injector Complex . . . 36
3.2 The LHC . . . 37
3.3 An LHC arc cell . . . 39
3.3.1 b_3 decay . . . 42
3.3.2 WISE and FiDeL . . . 42
3.4 The LHC IRs . . . 43
3.5 Radio Frequency (RF) system . . . 46
3.6 Beam injection and extraction . . . 47
3.7 The LHC cycle . . . 47
3.7.1 Beam types . . . 49
3.8 Beam Instrumentation/Hardware . . . 50
3.8.1 Base-Band Tune (BBQ) System . . . 50
3.8.2 Beam Position Monitors (BPMs) . . . 50
3.8.3 Beam Loss Monitors (BLMs) . . . 52
3.8.4 Fast Beam Current Transformers (FBCTs) . . . 53
3.8.5 BWS and BSRT . . . 53
3.8.6 AC-Dipole (ACD) . . . 54
3.8.7 Transverse Damper (ADT) . . . 55
3.8.8 Collimators . . . 56
3.9 Optics Measurements in the LHC . . . 57
3.9.1 Non-linear optics analysis . . . 60

4 Sextupolar Resonance Driving Terms **63**
4.1 Identification of a significant b_3 resonance at LHC injection . . . 65
4.2 Benchmarking . . . 69

4.3 Correction at Injection . . . 72
4.4 Dynamic RDT measurements . . . 82
4.4.1 Decay RDT measurements . . . 83
4.4.2 RDT measurements in the Ramp . . . 88
4.5 Conclusion . . . 91

5 Octupolar RDTs 93
5.1 Landau Octupoles (MO) . . . 94
5.2 MO Response . . . 97
5.3 Verifying the MO response discrepancy . . . 100
5.4 Segment-by-Segment (SbS) Studies . . . 106
5.5 Conclusion . . . 112

6 Novel correction of higher-order RDTs 113
6.1 Feed-down . . . 116
6.2 Creation of correction process using GUI . . . 121
6.3 Beam-based validation of correction method . . . 130
6.4 First measurement of IR dodecapole errors . . . 135
6.5 Conclusion . . . 138

7 Long-Range Beam-Beam Compensation 139
7.1 Beam-beam force . . . 141
7.2 LHC context and motivation . . . 143
7.3 Simulation benchmarking . . . 148
7.4 Sextupolar LRBB-driven RDT correction with round-optics . . . 149
7.5 Sextupolar LRBB-driven RDT correction with flat-optics . . . 152
7.6 Octupolar LRBB-driven RDT correction . . . 157
7.7 LRBB-driven $\boldsymbol{\beta}$-beating correction . . . 163
7.7.1 Creating the correction . . . 164
7.7.2 Testing the correction . . . 166

7.8 Conclusion 172

8 Effects of Energy Shifts **174**

8.1 Impact of energy shifts in HL-LHC 176

8.2 Impact of terrestrial tides on HL-LHC optics 178

8.3 Conclusion 184

9 Conclusion **185**

Bibliography **189**

List of Acronyms

Page numbers indicate the first occurrence and other significant mentions.

ACD AC-dipole. 50

AD Antiproton Decelerator. 36

ADT Transverse Damper. 55

ALICE A Large Ion Collider Experiment. 37

ATLAS A Toroidal LHC ApparatuS. 37

AWAKE Advanced WAKEfield experiment. 36

BBQ Base-Band Tune. 50

BLM Beam Loss Monitor. 52

BPM Beam Position Monitor. 50

BSRT Beam Synchrotron Radiation Telescope. 47, 53

BWS Beam Wire Scanner. 47, 53

CMS Compact Muon Solenoid. 37

DA Dynamic Aperture. 20

ELENA Extra Low ENergy Antiproton ring. 36

ESRF European Synchrotron Radiation Facility. 63

FBCT Fast Beam Current Transformer. 53

FCC Future Circular Collider. 63

FFT Fast Fourier Transform. 57

FiDeL Field Description of the LHC. 43

HERA Hadron-Electron Ring Accelerator. 63

HL-LHC High-Luminosity Large Hadron Collider. 63

IP Interaction Point. 37

IR Insertion Region. 37

LHC Large Hadron Collider. 17

LHCb Large Hadron Collider beauty. 37

LR Long-Range. 140

LRBB *sim.* LR, Long-Range Beam-Beam. 140

LSA LHC Software Architecture. 122

MD Machine Development. 50

n-ToF Neutron Time-of-Flight facility. 36

NA North Area. 36

NL non-linear. 12

OMC Optics Measurements and Correction. 50

OP Operations. 82

ppb particles per bunch. 48

PS Proton Synchrotron. 36

PSB Proton Synchrotron Booster. 36

RDT Resonance Driving Term. 19

RHIC Relativistic Heavy Ion Collider. 63

SbS Segment-by-Segment. 105

SIS-18 Heavy Ion Synchrotron. 63

SPS Super Proton Synchrotron. 36

SuperKEKB Super KEK Electron B-factory. 63

SVD Singular Value Decomposition. 57

TbT Turn-by-Turn. 30

TCLA absorber collimator. 56

TCP primary collimator. 56

TCSG secondary collimator. 56

TCT tertiary collimator. 56

WISE Windows Interface to Simulation Errors. 42

WP Working Point. 24

YASP Yet Another Steering Programme. 180

CHAPTER 1

INTRODUCTION

In particle accelerators, the precise control of beam optics is essential for the optimisation of key performance parameters, including beam lifetime, dynamic aperture, and overall beam stability. These quantities fundamentally determine the achievable performance of an accelerator.

Luminosity is the principal figure of merit for colliders, and it depends directly on these underlying properties. For the Large Hadron Collider (LHC), achieving high luminosity is of central importance for high-energy physics experiments, in particular for the Compact Muon Solenoid (CMS) and A Toroidal LHC ApparatuS (ATLAS) during proton-proton operation.

The LHC exhibits highly non-linear beam dynamics as a consequence of the strongly powered sextupole and octupole magnets required for chromaticity correction and Landau damping, together with intrinsic magnetic field imperfections. The significance of these non-linear effects is reflected in quantities such as amplitude detuning. For example, the amplitude detuning at injection arising from the LHC Landau octupoles is approximately two orders of magnitude larger than that of a detuned SuperKEKB lattice, illustrating the substantially stronger non-linear behaviour of the LHC compared with another collider.

Such non-linearities lead to increasingly complex and potentially chaotic particle motion, which can ultimately result in beam loss. Consequently, significant effort has been devoted

to improving the characterisation of non-linear beam dynamics and to developing advanced measurement techniques. During Run I and Run II, studies of non-linear effects were typically performed either indirectly or in a global manner, through the measurement of observables such as amplitude detuning and non-linear chromaticity. While these approaches provide valuable information, they do not yield a complete description of the underlying beam dynamics.

A more direct and quantitative framework for investigating non-linear resonances is provided by resonance driving terms (RDTs), which characterise the strength of individual resonances. Early attempts to measure RDTs in the LHC using a simple kicker magnet were largely unsuccessful, primarily due to limited excitation duration and insufficient data quality. To overcome these limitations, the technique of forced-RDT measurements was introduced. This method employs Turn-by-Turn acquisition under sustained excitation from an AC-dipole, enabling improved signal quality and resolution.

This thesis presents the first routine application of this technique in the LHC, together with methodological extensions that enable the exploration of previously unstudied regimes, including higher-order non-linear dynamics and developments beyond the single-particle framework considered in earlier studies.

1.1 Thesis Outline

Chapter 2 presents the key concepts of both linear and non-linear optics, which form the foundation for the work carried out. Chapter 3 provides an overview of the LHC, the primary accelerator considered in this study, including a brief description of the injector chain, as well as the principal hardware components relevant to the topics addressed. In addition, the optics measurement procedures employed are described.

Chapter 4 constitutes the first of the results chapters and focuses on the lowest-order non-linear RDTs. It begins with the observation of an unexpectedly large normal sextupole resonance at injection energy. This finding motivated the development and implementation of a correction, performed using the MAD-NG [9] simulation framework, exploiting its capability for RDT-based matching. The ability to measure this resonance without the large excitation amplitudes typically

required for non-linear optics studies enabled further investigations of transient phases of the LHC cycle, including the study of the normal sextupolar decay, and measurements during the energy ramp which were subsequently benchmarked against simulation.

Analysis is then extended to the next order for non-linear effects in Chapter 5, which presents a detailed investigation of the normal octupolar RDT response to the powering of the Landau octupoles at injection. This work was motivated by the implementation of a correction, developed prior to these studies, to mitigate their resonant impact in 2023. As part of these studies, a new technique was developed to improve the localisation of non-linear error sources, building upon the linear Segment-by-Segment approach previously established for linear optics.

Chapter 6 reports the further extension of the study to even higher-order RDTs, up to the dodecapolar level. It introduces a novel methodology for addressing higher-order field errors originating from the interaction regions. This approach was inspired by established correction techniques in linear optics, in particular the use of feed-down from sextupole and octupole errors to linear observables such as tune and coupling, and extends this concept to feed-down from high-order errors such as decapolar and even dodecapolar. This development is of particular relevance for the LHC's upgrade to High-Luminosity LHC (HL-LHC), where increased sensitivity to non-linear effects is expected. In support of this method, a dedicated graphical user interface was developed to facilitate the implementation of the corresponding corrections.

Chapter 7 contains the broadening of the scope from single-particle dynamics to multi-particle effects, with a particular focus on the influence of long-range beam-beam interactions in driving sextupolar and octupolar resonances. A novel measurement procedure, developed in collaboration with the operations team, enabled the first direct benchmarking of the long-range beam-beam-driven RDT response.

Subsequently, a novel correction scheme utilising the available non-linear magnets in the interaction regions was proposed and successfully demonstrated. These studies were further extended to the linear regime following the observation of a luminosity imbalance between the CMS and ATLAS experiments. The corresponding long-range beam-beam-induced β-beating was corrected for the first time using two complementary approaches: one based on simulation and the other employing a response-matrix formalism.

Chapter 8 continues the discussion of linear optics by returning to a single-particle framework and investigates the impact of energy errors in the context of the HL-LHC. In particular, it examines energy variations induced by terrestrial tides, which have previously been disregarded due to their expected negligible effect, and evaluates their potential implications for optics commissioning in both current and future operation.

Finally, the conclusions of this thesis are presented in Chapter 9.

CHAPTER 2

LINEAR AND NON-LINEAR OPTICS

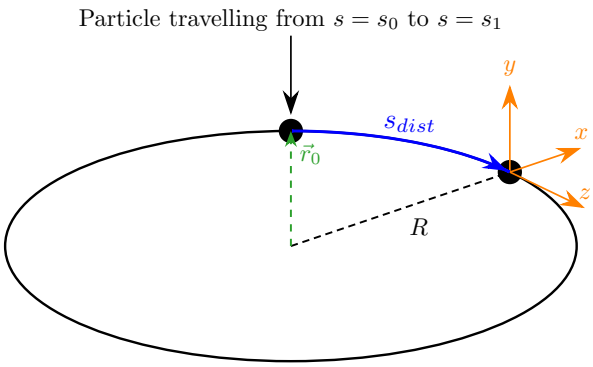


Figure 2.1: Diagram illustrating the Frenet-Serret coordinate system (orange). The blue arrow indicates the distance s_{dist} travelled by the particle from s_0 to s_1 along a circular orbit of radius R. The green vector $\vec{r}_0$ (green) is the polar vector coordinates of the particle at its initial position.

2.1 Coordinate System and Reference Trajectory

In accelerator physics, particle motion is conventionally described using the Frenet-Serret coordinate system, a curvilinear reference frame that follows the trajectory of a reference particle. In circular accelerators such as synchrotrons, this trajectory may correspond either to the *reference orbit*, defined as the path passing through the centres of all magnets under the assumption of perfect alignment, or to the *closed orbit*.

The latter represents the actual trajectory of an ideal reference particle with the nominal energy. It may be displaced from the reference orbit due to dipole field errors or deliberately applied orbit bumps, but remains periodic, returning to the same trajectory after one turn. For a centred bunch, the closed orbit coincides with the path of the beam centroid.

The longitudinal coordinate, s, denotes the arc length along this trajectory, while the local curvature is characterised by the radius $\rho(s)$, which varies according to the magnetic field. The transverse dynamics are described in terms of the trace-space coordinates (x, x', y, y'), where x and y denote horizontal and vertical displacements from the reference orbit, respectively. The primes indicate derivatives with respect to the path length s, i.e. $x' = \frac{\mathrm{d}x}{\mathrm{d}s}$ and $y' = \frac{\mathrm{d}y}{\mathrm{d}s}$.. Alternatively, the dynamics may be expressed in phase-space using the coordinates (x, p_x, y, p_y) where p_x and p_y refer to the conjugate momenta [10] (see pp. 53-54).

2.2 Beam Rigidity

In order to maintain particles on a circular trajectory, the Lorentz force acting on the particles must balance the centrifugal force, leading to the following equation:

$$q\left(\vec{E} + \vec{v} \times \vec{B}\right) = \frac{\gamma_{\mathrm{rel}} m v^2}{\rho}, \tag{2.1}$$

where q is the particle charge, $\vec{E}$ and $\vec{B}$ denote the electric and magnetic field strengths, respectively, $\vec{v}$ is the particle velocity, and m its mass, with γ_{rel} representing the Lorentz factor and ρ the bending radius. The beam rigidity, $B\rho$, quantifies a particle's resistance to magnetic deflection, fixing the bending radius achieved by a given dipole field.

By neglecting the electric field contribution in Eq. (2.1) (appropriate for a purely magnetic bending system), and using the relation $p = \gamma_{\mathrm{rel}} m v$ [11], the following expression can be obtained:

$$B\rho = \frac{p}{q}. \tag{2.2}$$

The beam rigidity (units of $\mathrm{T}\cdot\mathrm{m}$) can also be expressed in terms of the particle momentum. In this case, a commonly used approximation is given by

$$B\ [\mathrm{T}]\,\rho\ [\mathrm{m}] \approx 3.33\, pc\ [\mathrm{GeV/c}] \tag{2.3}$$

2.3 Magnet Definition

As discussed in the preceding section, magnetic fields are employed to maintain particles on the closed orbit. In general, the magnetic field generated by a multipole of order n (where $n = 1$ corresponds to a dipole, $n = 2$ to a quadrupole, etc.) can be described by the multipole expansion:

$$B_y(x,y,s) + iB_x(x,y,s) = \sum_{n=1}^{\infty} [b_n(s) + ia_n(s)](x+iy)^{n-1}, \tag{2.4}$$

where B_x and B_y are the transverse components of the magnetic field, and $b_n(s)$ and $a_n(s)$ are the normal and skew multipole coefficients, respectively. Skew components represent normal multipoles rotated through an angle of $\frac{\pi}{2n}$ rad [12].

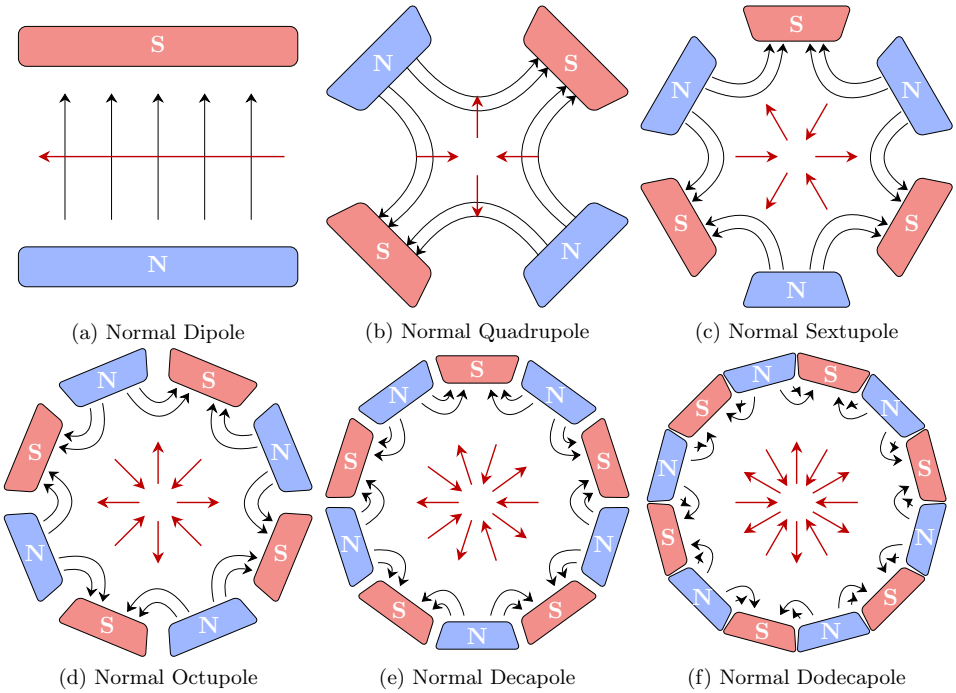
These coefficients are defined in terms of transverse derivatives of the magnetic field evaluated on the reference orbit as:

$$\begin{aligned} b_n(s) &= \frac{1}{(n-1)!} \left.\frac{\partial^{n-1} B_y}{\partial x^{n-1}}\right|_{(0,0,s)}, \\ a_n(s) &= \frac{1}{(n-1)!} \left.\frac{\partial^{n-1} B_x}{\partial x^{n-1}}\right|_{(0,0,s)}. \end{aligned} \tag{2.5}$$

These magnetic field components can be normalised to remove the momentum dependence of the kick. This is achieved by expressing them in terms of the beam rigidity. The normalised normal and skew multipole strengths, K, for a multipole of order n are therefore defined by [13],

$$K_n = \frac{q}{p}(n-1)!\, w_n(s), \qquad w \in \{b, a\} \tag{2.6}$$

where K_n has units of m^{-n}. The quantity K_n is used throughout this thesis to denote the magnetic field strength, unless explicitly stated otherwise. The notation for K_n, b_n, and related quantities can vary between codes and laboratories. For example, in MAD [14], a quadrupole is denoted by K_1, while quadrupole errors are expressed as b_2. In this thesis, the convention adopted is that the index n corresponds to the order of the multipole.



(a) Normal Dipole (b) Normal Quadrupole (c) Normal Sextupole

(d) Normal Octupole (e) Normal Decapole (f) Normal Dodecapole

Figure 2.2: Schematic of multipoles up to order $n = 6$, where some example magnetic field lines are in black, and the red arrows indicate the direction of force on a positively charged particle moving out of the page.

Figure 2.2 illustrates examples of field profiles and corresponding pole arrangements for several multipole orders. The various multipoles serve several purposes within the accelerator, as summarised below:

- **Dipole magnets**: produce a uniform magnetic field used to bend the beam and define its trajectory.
- **Quadrupole magnets**: generate a linear magnetic field gradient which is used to focus and defocus the beam in the transverse planes. They provide focusing in one plane and defocusing in the orthogonal plane; the quadrupole depicted in Fig. 2.2 can be seen to be horizontally focusing.
- **Sextupole magnets**: these are primarily used to correct chromaticity within the accelerator which can arise through energy-dependent focusing in the quadrupoles.
- **Octupole magnets**: provide amplitude-dependent focusing, primarily employed to damp collective instabilities in the beam.

All of these magnets, including decapole and higher-order multipoles, can be employed to correct magnetic field errors arising from various sources, such as imperfections in the design and construction of the main magnets, as well as misalignments.

2.4 Linear Optics

In the linear optics regime, the magnetic elements used to guide and focus the beam are idealised such that the forces they generate are linear functions of the transverse displacements. This approximation is most likely to remain valid for sufficiently small deviations from the reference trajectory, under conditions where non-linear (NL) effects can be neglected.

Transverse oscillations around the reference trajectory arise from the structure of a linear lattice, which consists solely of dipole and quadrupole magnets. In particular, a repeating sequence of a focusing quadrupole, drift space, defocusing quadrupole, and another drift space (a FODO cell, depicted in Fig. 2.3) provides the alternating focusing required to confine particles close to the closed orbit.

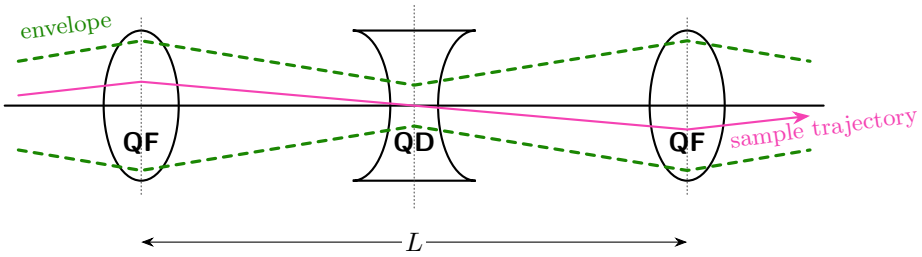


Figure 2.3: Schematic of a simple FODO cell where QF denotes focusing quadrupole and QD is a defocusing quadrupole. Green denotes the betatron envelope, while pink denotes a possible particle trajectory within the envelope.

The transverse dynamics of a particle in either the horizontal or vertical plane are therefore governed by Hill's equation:

$$\frac{\mathrm{d}^2 z}{\mathrm{d}s^2} + K_z(s) z = 0, \tag{2.7}$$

where z denotes the transverse displacement in either transverse plane, and $K(s)$ is the periodic

position-dependent focusing strength determined by the lattice configuration according to:

$$K_z(s) = k_2(s), \qquad \text{(Focusing)} \tag{2.8}$$

$$K_z(s) = -k_2(s), \qquad \text{(Defocusing)} \tag{2.9}$$

where $k_2(s)$ is the quadrupolar field strength. In the horizontal plane there is an additional contribution to $K_z(s)$ arising from the weak focusing of dipole magnets $\equiv \frac{1}{\rho(s)^2}$.

The general solution to Hill's equation describes betatron oscillations:

$$z(s) = \sqrt{\beta_z(s)\varepsilon_z}\cos\left(\phi_z(s) + \phi_{z_0}\right), \tag{2.10}$$

where $\beta_z(s)$ denotes the Courant-Snyder β-function [15], and ε_z is the single-particle emittance. Together, these quantities determine the amplitude of the particle's betatron oscillation. In the absence of non-conservative forces, ε_z is an invariant of the particle's motion, whereas $\beta_z(s)$ characterises the focusing properties of the magnetic lattice and therefore varies periodically around the accelerator. The quantity $\phi_z(s)$ represents the accumulated betatron phase advance. Since both $\beta_z(s)$ and $\phi_z(s)$ inherit the periodicity of the lattice, and by using the identity $w^2 = \beta$ as a substitution into Eq. (2.10) [15] it follows that

$$\phi' = \frac{1}{\beta} \quad \text{or} \quad \phi = \int \frac{ds}{\beta}. \tag{2.11}$$

2.4.1 Twiss Parameters

The evolution of the transverse beam envelope is described by the Twiss parameters, also known as the Courant-Snyder parameters [15]. These are defined as:

$$\alpha_z = -\frac{1}{2}\frac{\mathrm{d}\beta_z}{\mathrm{d}s}, \quad \gamma_z = \frac{1+\alpha_z^2}{\beta_z}. \tag{2.12}$$

Each parameter has a specific physical interpretation:

- $\boldsymbol{\beta_z}$: Describes the local transverse envelope of the beam, reflecting the position amplitude of the oscillations due to the focusing structure.
- $\boldsymbol{\alpha_z}$: Characterises the slope of the betatron envelope, corresponding to the correlation between position and angle.

- γ_z: Complements β_z by characterising the momentum component of the phase-space ellipse.

These parameters can be used to define a normalised coordinate system in which any linear transformation is reduced to a pure rotation. This is achieved through the following linear transformation from the physical coordinate w to the corresponding normalised coordinate $\hat{w}$:

$$\hat{w} = Nw \quad \text{where } w = \begin{pmatrix} z \\ p_z \end{pmatrix}, \tag{2.13}$$

where the matrix N is defined to be

$$N = \frac{1}{\sqrt{\beta}} \begin{pmatrix} 1 & 0 \\ \alpha & \beta \end{pmatrix}. \tag{2.14}$$

2.4.2 Betatron Tune

The betatron tune Q_z (which will be referred to henceforth as simply the tune) is a key parameter quantifying the number of transverse oscillations per turn around the ring:

$$Q_z = \frac{1}{2\pi}\Delta\phi_z = \frac{1}{2\pi}\oint \frac{\mathrm{d}s}{\beta_z}, \tag{2.15}$$

where $\Delta\phi_z$ is the total phase advance over one revolution. The tune plays a central role in resonance conditions and beam stability.

2.4.3 Emittance, Phase and Action

In Hamiltonian systems,Liouville's theorem states that, under the action of conservative forces, the particle density in phase-space remains constant [11]. This invariant quantity can be expressed in terms of the action, J, which characterises the amplitude of a particle's oscillation:

$$J_z = \frac{1}{2}\left(\gamma_z z^2 + 2\alpha_z z' z + \beta_z (z')^2\right). \tag{2.16}$$

The *single-particle emittance* is defined as $\varepsilon_z = 2J_z$ and corresponds to the area enclosed by the trace-space ellipse shown in Fig. 2.4. During acceleration, Liouville's theorem is no longer strictly applicable in terms of the geometric emittance, and ε_z decreases with increasing beam energy through the process of adiabatic damping [10].

However, it is possible to define a normalised emittance, $\varepsilon_{z,\gamma}$, which remains invariant with energy. This quantity is defined as

$$\varepsilon_{z,\gamma} = \beta_{\text{rel}}\gamma_{\text{rel}}\varepsilon_z, \tag{2.17}$$

where β_{rel} and γ_{rel} denote the relativistic β and Lorentz factor, respectively. The particle motion can alternatively be expressed in terms of action-angle variables:

$$z = \sqrt{2J_z\beta_z}\cos(\phi_z), \tag{2.18}$$

$$z' = -\sqrt{\frac{2J_z}{\beta_z}}\left[\sin(\phi_z) + \alpha_z\cos(\phi_z)\right]. \tag{2.19}$$

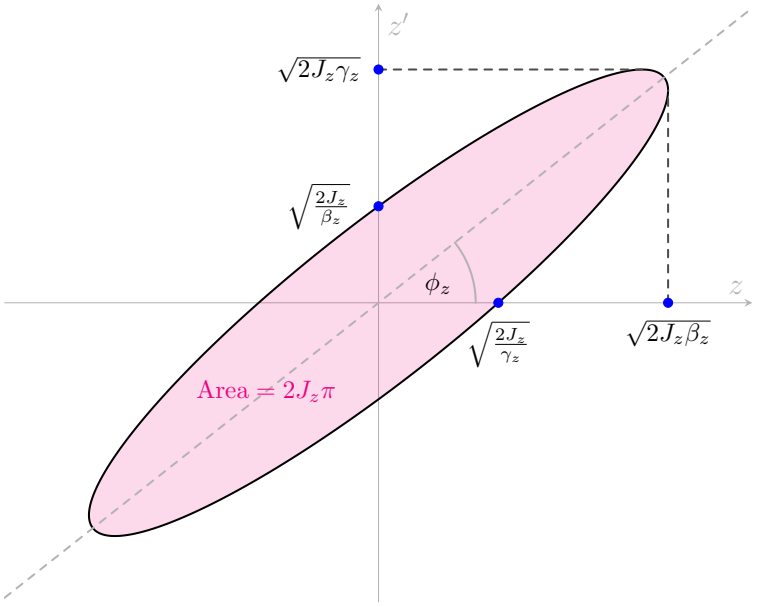


Figure 2.4: Diagram showing the key coordinates of a trace-space ellipse.

2.4.4 Dispersion

If a particle does not have the design momentum, p_0, but instead has a momentum offset δ, arising either from the intrinsic momentum spread of the beam or from a deliberate RF trim (see Section 3.5), this leads to several observable effects. The quantity δ is defined as

$$\delta = \frac{\Delta p}{p_0}, \tag{2.20}$$

where Δp is defined as $p - p_0$ and $p_0 = \beta_{\text{rel}}\gamma_{\text{rel}}mc$. One of the resulting effects is the appearance of dispersion, which arises from the fact that particles with a momentum offset are deflected differently compared to particles at the design momentum. An example of such a deflection is illustrated in Fig. 2.5.

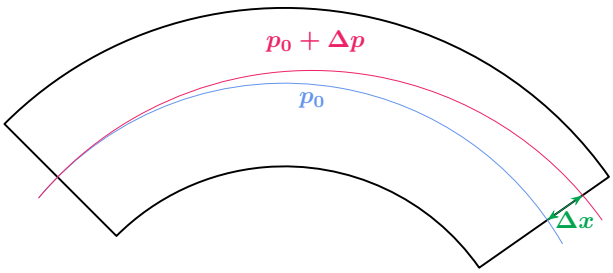


Figure 2.5: Illustration of the principle of dispersion showing an on-momentum (blue) and an off-momentum (pink) particle trajectory through a bending magnet.

As a consequence, the closed orbit remains subject to the magnetic properties of the lattice but follows a displaced trajectory, which is described by the dispersion function in the following equation:

$$D_z(s) = \frac{\Delta z(s)}{\delta}, \tag{2.21}$$

where $\Delta z(s)$ denotes either the horizontal or vertical transverse displacement from the closed orbit.

2.4.5 Chromatic Optics

In circular accelerators, a critical effect arises from the dependence of the betatron tune on the energy deviation of particles. This phenomenon is known as chromaticity. Since a beam has a finite momentum spread, different particles will experience slightly different focusing strengths in the quadrupoles, leading to a tune spread across the bunch. If left uncorrected, this can lead to resonance crossing, beam instability, and degradation of performance.

The chromaticity Q'_z in either transverse plane is defined as the derivative of the betatron tune with respect to the relative momentum deviation $\delta = \Delta p/p$:

$$Q'_z = \frac{dQ_z}{d\delta}. \tag{2.22}$$

In a purely linear lattice composed of ideal quadrupoles, chromaticity arises naturally due to the fact that the focusing strength $K_2 \propto 1/p$ [16]. Thus, off-momentum particles experience weaker or stronger focusing depending on the sign of δ, leading to a shift in tune. In the Large Hadron Collider (LHC), the natural chromaticity is around -80 units at injection. Its negative value is characteristic of a strong-focusing (i.e. use of quadrupoles) ring.

The natural chromaticity can be corrected through the introduction of normal sextupole magnets in regions with non-zero horizontal dispersion. In this configuration, off-momentum particles traverse the sextupoles off-centre, and therefore experience an additional or reduced focusing, depending on the sign of the momentum deviation, compensating the natural chromaticity.

These chromatic concerns also apply to higher-order, and NL errors and magnets in dispersive regions can also introduce or correct the NL chromaticity. This effect corresponds to the higher-order dependence of the tune on the relative momentum deviation, δ. In particular, it is associated with terms of order Q'' and higher in the Taylor expansion of the tune, given by

$$Q_z(\delta) = Q_0 + Q'\delta + \frac{1}{2!}Q''\delta^2 + \frac{1}{3!}Q'''\delta^3 + \dots. \tag{2.23}$$

Chromatic β-beating represents the local manifestation of the differing focusing experienced by on- and off-momentum particles. While this effect is observed globally through the chromaticity, it can be quantified locally by chromatic β-beating ($\Delta\beta_{\mathrm{dp}}$)

$$\Delta\beta_{\mathrm{dp}} = \frac{\beta_{\mathrm{dp}} - \beta}{\beta}, \tag{2.24}$$

where β_{dp} denotes the off-momentum β-function, and β represents the corresponding on-momentum β-function.

2.4.6 Transfer Maps

Particle motion through the various elements of an accelerator can be described using transfer matrices. For a linear element i, the state vector, w, before and after the element is related via:

$$w_{\mathrm{after}} = M_i w_{\mathrm{before}}. \tag{2.25}$$

Standard transfer matrices for linear elements include [11]:

Focusing quadrupole:

$$\begin{pmatrix} x \\ p_x \end{pmatrix}_L = \begin{pmatrix} \cos\left(\sqrt{|k_2|}L\right) & \frac{1}{\sqrt{|k_2|}}\sin\left(\sqrt{|k_2|}L\right) \\ -\sqrt{|k_2|}\sin\left(\sqrt{|k_2|}L\right) & \cos\left(\sqrt{|k_2|}L\right) \end{pmatrix} \begin{pmatrix} x \\ p_x \end{pmatrix}_0, \tag{2.26}$$

$$\begin{pmatrix} y \\ p_y \end{pmatrix}_L = \begin{pmatrix} \cosh\left(\sqrt{|k_2|}L\right) & \frac{1}{\sqrt{|k_2|}}\sinh\left(\sqrt{|k_2|}L\right) \\ \sqrt{|k_2|}\sinh\left(\sqrt{|k_2|}L\right) & \cosh\left(\sqrt{|k_2|}L\right) \end{pmatrix} \begin{pmatrix} y \\ p_y \end{pmatrix}_0, \tag{2.27}$$

where k_2 is the strength of the quadrupole in m^{-2} and L is the path distance. By convention, this is termed a focusing quadrupole because it focuses in the horizontal plane and defocuses in the vertical plane.

Drift space:

$$\begin{pmatrix} z \\ p_z \end{pmatrix}_L = \begin{pmatrix} 1 & L \\ 0 & 1 \end{pmatrix} \begin{pmatrix} z \\ p_z \end{pmatrix}_0, \tag{2.28}$$

Dipole:

$$\begin{pmatrix} x \\ p_x \end{pmatrix}_L = \begin{pmatrix} \cos\left(\frac{L}{\rho}\right) & \rho\sin\left(\frac{L}{\rho}\right) \\ \frac{1}{\rho}\sin\left(\frac{L}{\rho}\right) & \cos\left(\frac{L}{\rho}\right) \end{pmatrix} \begin{pmatrix} x \\ p_x \end{pmatrix}_0. \tag{2.29}$$

These transfer matrices can be combined into a single transfer matrix for all elements along the beam-line (one-turn map, M_{OTM} in the case of a circular collider), using each matrix M_i corresponding to a linear element as follows:

$$M_{OTM} = M_N \ldots M_2 M_1. \tag{2.30}$$

This can simplify to

$$M_{OTM} = \begin{pmatrix} \cos(\phi_z) + \alpha\sin(\phi_z) & \beta\sin(\phi_z) \\ -\gamma\sin(\phi_z) & \cos(\phi_z) - \alpha\sin(\phi_z) \end{pmatrix}, \tag{2.31}$$

where ϕ_z denotes the one-turn phase advance where the tune is given by $Q \equiv \frac{\phi_z}{2\pi}$.

2.4.6.1 Optics simulation, tracking and matching

There exist several approaches to optics simulation in accelerator physics. The most conceptually straightforward method is particle tracking, as implemented in codes such as SixTrack [17] and Xsuite [18], which propagate particle trajectories through a sequence of transfer maps or kicks.

In contrast, optics codes typically operate directly on transfer maps. These codes construct the one-turn map and analyse it to extract optical functions or propagate these functions through the lattice using the underlying maps, for example through matrix multiplication in the linear regime. For NL dynamics, higher-order effects are incorporated using Lie algebraic methods and truncated power series expansions of the one-turn map [19]. Examples of such codes include MAD-X [14] and MAD-NG [9, 20, 21].

In practice, the distinction between these approaches is not strict. Optics codes can also be used for particle tracking, while tracking codes can approximate the one-turn map by comparing trajectories of neighbouring particles to estimate the corresponding matrix elements, therefore enabling limited optics calculations and matching. However, such finite-difference methods can become unreliable in the presence of strong non-linearities or in challenging regimes, such as near resonances or in strongly NL lattices.

MAD-NG is particularly well suited to NL optics analysis, as it can directly construct the NL one-turn map from the functional description of the accelerator lattice and its associated control parameters. This capability enables efficient analysis in normal-form coordinates by factorising the one-turn map into its Lie normal-form decomposition, from which the resonance driving terms (RDTs) are extracted from the generating function of the normalising transformation. Consequently, MAD-NG provides a robust framework for NL optics matching and optimisation.

The NL matching procedure employed in this thesis is based on the Jacobian matrix, which contains the sensitivities of the matching target variables with respect to the optimisation variables, namely the magnet strengths. Mathematically, the Jacobian, $\mathbf{J}$, is defined as

$$\mathbf{J} = \frac{\partial \mathbf{f}}{\partial \mathbf{k}} = \begin{bmatrix} \frac{\partial f_1}{\partial k_1} & \frac{\partial f_1}{\partial k_2} & \cdots & \frac{\partial f_1}{\partial k_n} \\ \frac{\partial f_2}{\partial k_1} & \frac{\partial f_2}{\partial k_2} & \cdots & \frac{\partial f_2}{\partial k_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial k_1} & \frac{\partial f_m}{\partial k_2} & \cdots & \frac{\partial f_m}{\partial k_n} \end{bmatrix}, \tag{2.32}$$

where $\mathbf{f}$ represents the matching target variables, such as the betatron tunes or the NL RDTs, and $\mathbf{k}$ denotes the corresponding circuit strengths used as optimisation variables.

Within MAD-NG, the Jacobian is computed analytically using Generalised Truncated Power Series Algebra (GTPSA) rather than through finite-difference approximations. This approach provides exact derivatives of accelerator observables with respect to the magnet settings, enabling efficient Newton-type optimisation while improving numerical stability. Consequently, both linear and NL optics can be matched accurately and with significantly greater computational efficiency than is achievable using finite-difference methods.

2.5 Non-linear Optics

To introduce the topic of NL dynamics in accelerators, it is first necessary to define a set of observables that provide measurable signatures of NL perturbations, arising from higher-order magnetic field components such as sextupoles, octupoles, decapoles, and higher-order multipoles. These fields may arise either intentionally, as part of the accelerator design, or unintentionally due to imperfections in the magnetic lattice or the presence of fringe fields.

2.5.1 Dynamic Aperture

A key global manifestation of the cumulative effects of non-linearities is the Dynamic Aperture (DA). The DA constitutes a fundamental concept in the study of NL beam dynamics in circular accelerators. In contrast to the physical aperture, it defines the region of phase-space within which particle motion survives in the accelerator for a given time-scale. In the asymptotic limit, the long-term DA approaches the boundary separating stable and unstable (chaotic) motion.

A detailed understanding of the DA is essential for both the design and operation of such accelerators, as it provides insight into the impact of NL effects and informs strategies for optimising beam lifetime. The standard approach for computationally evaluating the DA relies on numerical tracking simulations, in which a large set of initial conditions is distributed throughout phase-space and tracked over a sufficient number of turns to probe time-scales relevant to accelerator operation. In addition, experimental estimates of the DA can be obtained through beam-based measurements, for example by observing beam losses following controlled excitation, as illustrated in Fig. 2.6, or deliberate beam blow-up in the accelerator.

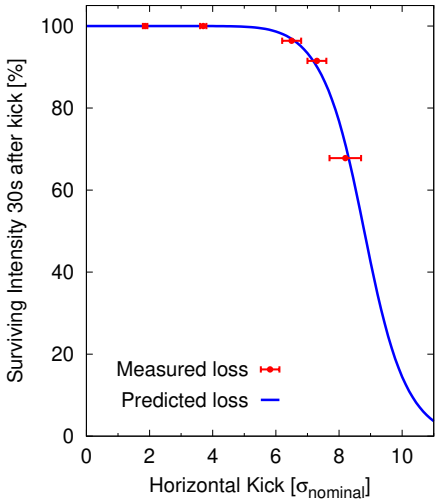


Figure 2.6: Illustration of a beam-based DA measurement showing the surviving beam intensity following single-kick excitations of increasing amplitude (expressed in units of beam sigma) at injection energy. The predicted (blue) and measured (red) losses are shown [1].

This concept can be extended to the notion of a forced DA, which is defined as the DA estimated in the presence of forced beam oscillations. It is generally smaller than the free DA [22].

2.5.2 Amplitude Detuning

The presence of these perturbations means the motion is no longer simple harmonic, and the tune then depends on the amplitude. This dependence is the global observable amplitude detuning, defined as the variation of the natural tune with the transverse action (J_z), or equivalently with the Courant-Snyder invariant (ε_z). This effect can be described by performing a Taylor expansion of the tune around the closed orbit, yielding

$$Q_z(\varepsilon_x,\varepsilon_y) = Q_{z_0} + \left(\frac{\partial Q_z}{\partial \varepsilon_x}\varepsilon_x + \frac{\partial Q_z}{\partial \varepsilon_y}\varepsilon_y\right) + \frac{1}{2}\left(\frac{\partial^2 Q_z}{\partial \varepsilon_x^2}\varepsilon_x^2 + 2\frac{\partial^2 Q_z}{\partial \varepsilon_x \partial \varepsilon_y}\varepsilon_x\varepsilon_y + \frac{\partial^2 Q_z}{\partial \varepsilon_y^2}\varepsilon_y^2\right) + \ldots, \tag{2.33}$$

where $\frac{\partial Q_z}{\partial \varepsilon_z}$ denotes the first-order direct detuning term, while $\frac{\partial Q_z}{\partial \varepsilon_a}$, with $a \neq z$ and $a \in \{x, y\}$, represents the first-order cross-detuning contribution.

2.5.3 Decoherence

Effects such as amplitude detuning and chromaticity can give rise to decoherence, where an initially coherent transverse oscillation of the particle beam progressively loses coherence with time. This phenomenon arises because individual particles oscillate at slightly different betatron frequencies, causing their phases to diverge and the amplitude of the collective beam motion to decrease. This behaviour is illustrated in Fig. 2.7, where the trajectories of three representative particles, labelled A, B, and C, oscillate at different frequencies, resulting in a gradual reduction of the average beam displacement as the individual oscillations become increasingly out of phase. In the presence of chromaticity alone, the frequency spread is reversible, allowing the beam to undergo recoherence after a sufficient number of turns. However, when non-linear effects are also present, the frequency spread becomes effectively irreversible, preventing recoherence and leading to sustained decoherence of the collective signal [23, 24].

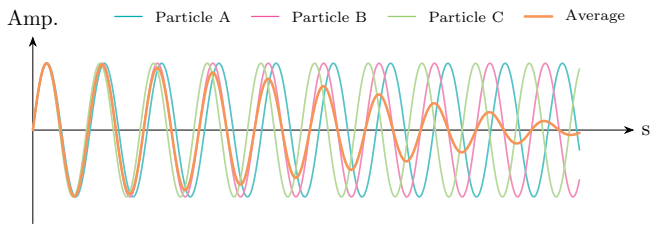


Figure 2.7: Diagram illustrating the principle of decoherence by showing the trajectories of three representative particles as functions of s (particle A in blue, particle B in pink, particle C in green) and their combined average trajectory in orange.

This leads to limitations in measurements based on single-kick excitation, as the number of turns available for spectral analysis is restricted, as illustrated in Fig. 2.8.

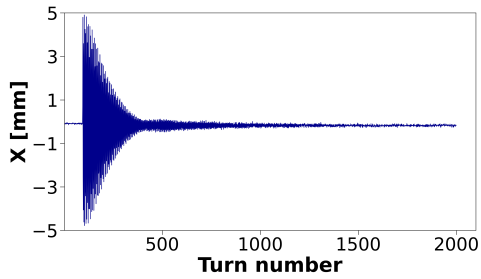


Figure 2.8: Diagram illustrating beam decoherence following a single-kick excitation by showing the beam oscillation amplitude as a function of turn number.

2.5.4 Resonances

Resonances arise when perturbations, either linear or NL, act on a particle in synchronism with its oscillatory motion, causing successive perturbations to add coherently on each turn. The corresponding resonance condition is given by

$$nQ_x + mQ_y = k, \quad n,m,k \in \mathbb{Z}, \tag{2.34}$$

which defines a series of resonance lines in tune space, as illustrated in Fig. 2.9, where the resonance conditions up to sixth order ($|n| + |m|$) are shown.

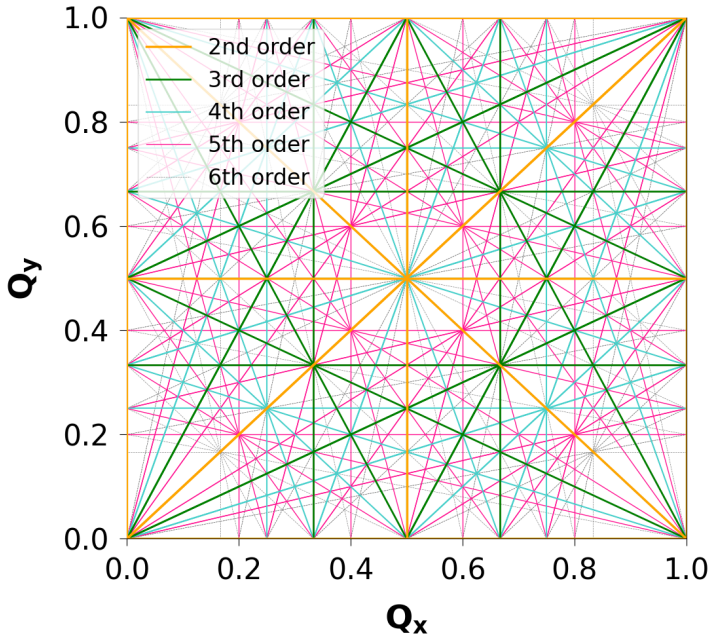
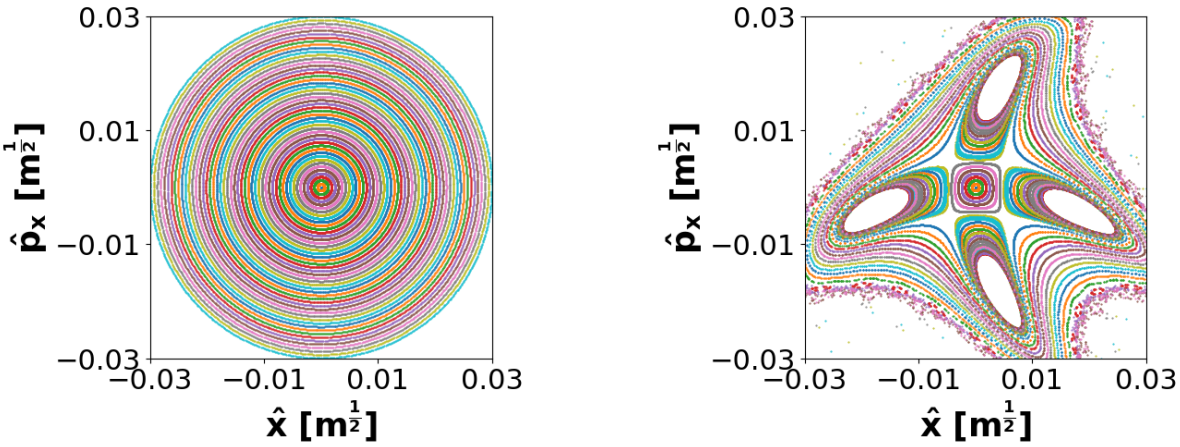


Figure 2.9: Tune diagram (Q_y vs Q_x) showing the resonance conditions, with the leading resonance order indicated by the following colours: orange for second order, green for third order, cyan for fourth order, pink for fifth order and dashed grey for sixth order.

One of the key properties of a circular accelerator is the strength of resonances in the vicinity of the chosen operational accelerator tunes, known as the Working Point (WP). When the tunes approach a resonance condition, the structure of phase-space can become significantly distorted. By transforming to linear normalised coordinates, as in Eq. (2.13), these distortions become clearly visible in the normalised phase-space plots in Fig. 2.10.



(a) Without non-linearities

(b) With non-linearities

Figure 2.10: Distortion of normalised phase-space due to non-linearities.

The presence of NL resonances can induce a transition from stable to chaotic motion, leading to diffusion of particles towards the physical aperture of the accelerator. In addition, resonant excitation can result in rapid particle loss at specific frequencies. Repeated crossing of resonances, for example due to synchrotron motion, can lead to an increase in particle action and consequently contribute to emittance growth. Furthermore, the formation of phase-space islands may occur, as illustrated in Fig. 2.10. All of these effects ultimately influence the achievable DA. In Fig. 2.10, the boundary of the DA may be interpreted as the largest initial amplitude for which the particle follows a bounded, regular trajectory.

2.5.5 Hamiltonian Formalism

In order to describe this NL motion, it is convenient to adopt a formulation based on Hamiltonian mechanics. The Hamiltonian of a relativistic particle in an electromagnetic field can be

written [25]

$$\mathscr{H}_{\text{full}}(z,p_z,t) = c\sqrt{(p - qA(z,t))^2 + m_0^2c^2} + q\Phi(z,t), \tag{2.35}$$

where z denotes the Cartesian coordinates, p the canonical momentum, A the magnetic vector potential, and Φ the electric scalar potential.

Several approximations may be introduced to simplify this expression [25]. First, the independent variable is transformed from time, t, to the longitudinal coordinate, s, by adopting the Frenet-Serret coordinate system, where the reference trajectory is characterised by a curvature $h = 1/\rho$. The contribution from electric fields is neglected, and only static transverse magnetic fields are retained, since the influence of longitudinal magnetic fields is negligible in the LHC. Under these assumptions, the magnetic vector potential is reduced to a single longitudinal component, $A = (0,0,A_s)$, which considerably simplifies the Hamiltonian. The resulting Hamiltonian then takes the form

$$\mathscr{H}_{\text{full}} = \frac{1}{2}(p_x^2 + p_y^2) + \frac{A_s}{B\rho}. \tag{2.36}$$

For an individual multipole of order n, the potential can be expressed in terms of the complex transverse coordinate, as previously introduced in Eq. (2.4):

$$w = x + iy. \tag{2.37}$$

Using the relation $\vec{B} = \nabla \times \vec{A}$, the Hamiltonian describing the motion through a given multipole element can be written:

$$\mathscr{H} = \frac{q}{p}\Re\left[\sum_{n=1}^{\infty}(b_n(s) + ia_n(s))\frac{(x+iy)^n}{n}\right]. \tag{2.38}$$

This expression is commonly referred to as the kick Hamiltonian, as it describes the transverse momentum change experienced by a particle traversing the element. The corresponding momentum kick is obtained from Hamilton's equations [26] (see p. 337)

$$\frac{\mathrm{d}p_z}{\mathrm{d}s} = -\frac{\partial\mathscr{H}}{\partial z}. \tag{2.39}$$

As an example, for a thin sextupole, the normal and skew components of the Hamiltonian are given by

$$\begin{aligned} N_3(x,y) &= \frac{1}{3}\cdot\frac{q}{p}\,b_3(x^3-3xy^2), \\ S_3(x,y) &= \frac{1}{3}\cdot\frac{q}{p}\,a_3(3x^2y-y^3), \end{aligned} \tag{2.40}$$

where N_3 and S_3 denote the normal and skew components, respectively. The corresponding normal component of the transverse momentum kick is then expressed by

$$\begin{aligned} \frac{\mathrm{d}p_x}{\mathrm{d}s} &= -\frac{q}{p}\left[b_3(x^2-y^2)\right], \\ \frac{\mathrm{d}p_y}{\mathrm{d}s} &= -\frac{q}{p}\left[2b_3xy\right]. \end{aligned} \tag{2.41}$$

Therefore, the transfer map of a thin normal sextupole located at s_0 is

$$\begin{pmatrix} x \\ x' \\ y \\ y' \end{pmatrix}_{s_{\text{final}}} = \begin{pmatrix} x \\ x' \\ y \\ y' \end{pmatrix}_{s_0} - \frac{qLb_3}{p}\begin{pmatrix} 0 \\ x^2-y^2 \\ 0 \\ -2xy \end{pmatrix}. \tag{2.42}$$

2.5.6 Non-linear Transfer Maps

Since a NL element cannot be represented by a linear transfer matrix, a more general mapping formalism is required. As described in [25] (see pp. 289-291), for a particle moving in an accelerator, the independent variable is taken to be the longitudinal coordinate s, representing the distance along the reference trajectory. Consider an arbitrary function f of the phase-space variables, defined as

$$f = f(x_i, p_i), \tag{2.43}$$

with no explicit dependence on s. Although f does not depend explicitly on s, its value evolves along the beam-line due to the evolution of the dynamical variables. The rate of change of f with respect to s can be written:

$$\frac{\mathrm{d}f}{\mathrm{d}s} = \sum_{i=1}^{C}\left(\frac{\mathrm{d}x_i}{\mathrm{d}s}\frac{\partial f}{\partial x_i} + \frac{\mathrm{d}p_i}{\mathrm{d}s}\frac{\partial f}{\partial p_i}\right), \tag{2.44}$$

where C denotes the number of degrees of freedom. Using Hamilton's equations, this expression becomes

$$\frac{\mathrm{d}f}{\mathrm{d}s} = \sum_{i=1}^{C} \left(\frac{\partial H}{\partial p_i} \frac{\partial f}{\partial x_i} - \frac{\partial H}{\partial x_i} \frac{\partial f}{\partial p_i} \right). \tag{2.45}$$

The Lie operator, denoted $:g:$, is defined for any function $g(x_i, p_i)$ by

$$:g: := \sum_{i=1}^{C} \left(\frac{\partial g}{\partial x_i} \frac{\partial}{\partial p_i} - \frac{\partial g}{\partial p_i} \frac{\partial}{\partial x_i} \right). \tag{2.46}$$

Assuming that the Hamiltonian, H, has no explicit dependence on s, the corresponding Lie operator $:H:$ can be constructed. In this notation, Eq. (2.45) can be written compactly as

$$\frac{\mathrm{d}f}{\mathrm{d}s} = -:H:f, \tag{2.47}$$

which represents the equation of motion for the function $f(x_i, p_i)$.

The formal solution to this equation can be expressed as

$$f\big|_{s=s_0+\Delta s} = e^{-\Delta s:H:} f\big|_{s=s_0}. \tag{2.48}$$

The operator $e^{-\Delta s:g:}$ is referred to as a Lie transformation generated by g. Such a transformation is governed by the following rules [19]:

$$e^{:f:}g = g + [f,g] + \frac{1}{2}[f,[f,g]] + \ldots, \tag{2.49}$$

where $[f,g]$ denotes the Poisson bracket of the phase-space functions f and g, given by:

$$[f,g] = \frac{\partial f}{\partial x} \cdot \frac{\partial g}{\partial p} - \frac{\partial f}{\partial p} \cdot \frac{\partial g}{\partial x}, \tag{2.50}$$

where the coordinate subscripts have been suppressed for notational simplicity. When the Lie operator acts on a coordinate function x, the result is an expansion

$$x_f = e^{:f:}x = x + [f,x] + \frac{1}{2}[f,[f,x]] + \ldots, \tag{2.51}$$

which may be interpreted as the coordinate value at time $t = 1$, evolved from its initial condition at $t = 0$, for a Hamiltonian system governed by $\mathscr{H} = f$. That is,

$$x_f = e^{:\mathscr{H}:}x_0. \tag{2.52}$$

To construct the one-turn map of a circular accelerator including the effects of multipole kicks, several key properties of Lie algebra must be invoked:

- **Action on polynomial functions:** using Eq. (2.49) it can be shown that:

$$g(x_f) = g(e^{:f:}x) = e^{:f:}g(x). \tag{2.53}$$

- **Composition of exponential Lie operators:** Consider two Lie operators, $e^{:f_1(z_1):}$ mapping $z_1 \to z_2$ and $e^{:f_2(z_2):}$ mapping $z_2 \to z_3$. The composed transformation is

$$z_3 = e^{:f_2(z_2):}e^{:f_1(z_1):}z_1. \tag{2.54}$$

  By applying the relation in Eq. (2.53), this becomes:

$$z_3 = g(z_2) = g\left(e^{:f_1(z_1):}z_1\right) = e^{:f_1(z_1):}g(z_1) = e^{:f_1(z_1):}e^{:f_2(z_1):}z_1. \tag{2.55}$$

  This shows that successive Lie operators may be expressed in terms of the initial coordinates by reversing their order. From this, the following identity can be derived:

$$e^{:g:}e^{:f:}e^{-:g:} = e^{:e^{:g:}f:}, \tag{2.56}$$

  which is analogous to the similarity transformation in linear matrix algebra.

- **The Campbell-Baker-Hausdorff (CBH) theorem:** The concatenation of two exponential Lie operators can itself be written as a single exponential Lie operator:

$$e^{:f_2:}e^{:f_1:} = e^{:f:}, \tag{2.57}$$

  where the generator f is given by a formal power series:

$$f = f_2 + f_1 + \frac{1}{2}[f_2, f_1] + \frac{1}{12}[f_2, [f_2, f_1]] + \ldots. \tag{2.58}$$

  The proof of this theorem can be found in [27].

Following the approach of [28, 29], the Lie operator formalism can be employed to construct the one-turn map by incorporating the kick Hamiltonians of each of the NL thin lenses in the accelerator. This is achieved by representing their effect as successive Lie transformations acting on the initial phase-space coordinates, leading to the following expression:

$$M_{OTM} = \mathcal{M} = M_{N+1}e^{:h_N:}M_N e^{:h_{N-1}:}\cdots e^{:h_2:}M_2 e^{:h_1:}M_1, \tag{2.59}$$

where the M_n denote the maps corresponding to the linear elements, and $h_n = h_n(x_n, y_n)$ represents the aforementioned NL kick Hamiltonians, located at the longitudinal positions $s = s_n$.

By applying the rule that a succession of Lie operators can be rewritten in terms of the initial coordinates by reversing the order of operations, the map $\mathscr{M}$ can be expressed as:

$$\mathscr{M} = M_1 e^{:h_1:} M_2 e^{:h_2:} \cdots e^{:h_{N-1}:} M_N e^{:h_N:} M_{N+1}. \tag{2.60}$$

Taking $\overline{M}_n = M_1 M_2 \ldots M_n$, inserting identities of the form $M_n^{-1} M_n$ and using the similarity relation, this equation transforms into:

$$\mathscr{M} = e^{:\overline{M}_1 h_1:} e^{:\overline{M}_2 h_2:} \cdots e^{:\overline{M}_{N-1} h_{N-1}:} e^{:\overline{M}_N h_N:} \overline{M}_{N+1}. \tag{2.61}$$

The linear one-turn map is denoted by $\overline{M}_{N+1}$. As $\overline{M}_{N+1}$ is a linear, symplectic, and stable operator, there exists a linear transformation of coordinates that reduces it to a pure rotation, as introduced in Eq. 2.13, corresponding to the transformation to linearly normalised coordinates.

This is expressed by the similarity transformation:

$$\overline{M}_{N+1} = NRN^{-1}, \tag{2.62}$$

where N represents the transformation, given in Eq. (2.13), and R the rotation.

By inserting identities of the form NN^{-1} in Eq. (2.61) the one-turn map can be denoted:

$$\hat{\mathscr{M}} = N^{-1} e^{:\overline{NM}_1 h_1:} e^{:\overline{NM}_2 h_2:} \cdots e^{:\overline{NM}_{N-1} h_{N-1}:} e^{:\overline{NM}_N h_N:} RN. \tag{2.63}$$

In the following, the coordinate system considered is the linearly normalised frame. The total one-turn map in this frame is denoted by the symbol $\hat{\mathscr{M}}$. Using the relation $g(x_f) = g(e^{:f:}x) = e^{:f:} g(x)$, the map $\hat{\mathscr{M}}$ can be expressed in the form:

$$\hat{\mathscr{M}} = e^{:h_1:} e^{:h_2:} \cdots e^{:h_N:} R, \tag{2.64}$$

where each h_n is expressed in the coordinates defined in Eq. (2.10) propagated via the normalised linear transfer map between the observation location and the kick. This formulation demonstrates that the one-turn map can be represented, in appropriate coordinates, as a pure rotation

describing the linear motion, combined with a Lie map that accounts for the NL kicks experienced by the particle.

By applying the CBH theorem, Eq. (2.64) can be simplified to:

$$\hat{\mathscr{M}} = e^{:h:} R. \tag{2.65}$$

Assuming that the h_n are small, h can be approximated by

$$h = \sum_{n=1}^{N} h_n + \sum_{n,m<n}^{N} [h_m, h_n] + \ldots. \tag{2.66}$$

In the following, only terms up to first order in h_n are retained, such that the second summation on the right-hand side of Eq. (2.66) is neglected, so h becomes

$$h = \sum_{n=1}^{N} h_n. \tag{2.67}$$

Physically, this corresponds to the approximation that resonances are predominantly driven by magnetic field components of the same order, for example that octupolar resonances are primarily driven by octupole fields. If higher-order terms were included in the CBH, additional contributions to the dynamics would arise, such as sextupoles inducing amplitude detuning, sextupoles driving octupolar resonances (derivation shown in Chapter 4), and combined effects of sextupoles and octupoles generating higher-order resonances, for example at the decapolar level.

The Turn-by-Turn (TbT) motion is often described in terms of the normalised complex coordinates

$$\Gamma_{z,\pm} = \hat{z} \pm i\hat{p}_z = \sqrt{2J_z}\, e^{\mp i(\phi_z + \phi_{z0})}, \tag{2.68}$$

where $\hat{z}$ and $\hat{p}_z$ are the normalised coordinates introduced previously in Eq. (2.13).

To obtain the expressions for h_n in Eq. (2.67), the kick Hamiltonian given in Eq. (2.38) can be expanded in terms of these coordinates. The Hamiltonian, h, can then be written [30] (see p. 172)

$$h = \sum_{jklm} h_{jklm} \Gamma_{x,-}^{j} \Gamma_{x,+}^{k} \Gamma_{y,-}^{l} \Gamma_{y,+}^{m}. \tag{2.69}$$

Upon substitution of the normalised coordinates in Eq. (2.68), this expression becomes

$$h = \sum_{jklm} h_{jklm} (2J_x)^{\frac{j+k}{2}} (2J_y)^{\frac{l+m}{2}} e^{-i[(j-k)(\phi_x+\phi_{x0})+(l-m)(\phi_y+\phi_{y0})]}, \tag{2.70}$$

where the coefficients h_{jklm} represent the Hamiltonian terms associated with contributions from all multipole components of order $n = j+k+l+m$, and are given by

$$h_{jklm} = -\frac{[K_{n-1}\Omega(l+m) + iJ_{n-1}\Omega(l+m+1)]}{j!k!l!m!2^{j+k+l+m}} i^{l+m} (\beta_x)^{\frac{j+k}{2}} (\beta_y)^{\frac{l+m}{2}}, \tag{2.71}$$

where the function $\Omega(i)$ is defined such that $\Omega(i) = 1$ for even i and $\Omega(i) = 0$ for odd i, therefore selecting the normal (B_n) or skew (A_n) multipole components, respectively. A detailed derivation of this expression can be found in [28].

2.5.7 Normal Form

The transfer maps described above remain highly complex; however, a powerful method of analysis is provided by the normal form framework. The objective of this approach is to transform the one-turn map of the accelerator into a simplified representation that depends only on the action variables. Within the normal form framework, the transfer map is reduced to its simplest functional form, typically corresponding to an amplitude-dependent rotation.

This transformation is represented by

$$e^{-:F:} e^{:h:} R e^{:F:}, \tag{2.72}$$

where F denotes the generating function of the canonical transformation, h represents the higher-order component of the Hamiltonian which gives rise to NL dynamics, and R corresponds to the linear rotation operator associated with the unperturbed motion. Figure 2.11 illustrates the effect of this transformation in linearly normalised phase-space, demonstrating that the effects of the non-linear perturbations on the TbT motion are largely encoded in the generating function F.

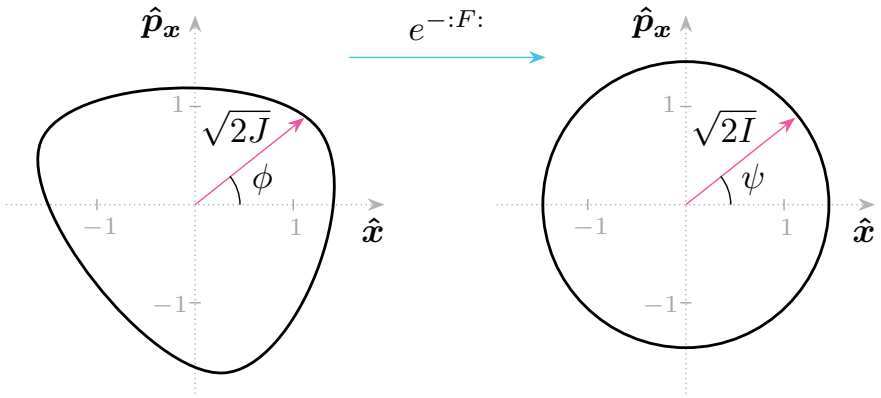


Figure 2.11: Diagram showing an example of the transformation of linearly normalised phase-space coordinates into normal form.

In normal form coordinates, the one-turn map takes the form:

$$e^{:H(I):}R, \tag{2.73}$$

where $H(I)$ is a function that depends only on the action variables in normal form coordinates $I = (I_x, I_y)$ and not on the phase angles. This leads to the following equivalence:

$$e^{-:F:}e^{:h:}Re^{:F:} = e^{:H(I):}R. \tag{2.74}$$

Multiplying both sides on the left by $R^{-1}R$ and applying the identity for the composition of exponential Lie operators yields:

$$e^{:F:}e^{:h:}e^{-:F:}R = e^{:H(I):}R. \tag{2.75}$$

Applying the CBH theorem and truncating at first order in F, h, and $H(I)$, the condition simplifies to:

$$(1-R)F + H(I) = h. \tag{2.76}$$

The formal solution to this equation is given by

$$F = \frac{1}{1-R}(h - H(I)), \tag{2.77}$$

$$H(I) = \bar{h}, \tag{2.78}$$

where $H(I)$ can be approximated by $\bar{h}$ [30] which denotes the average of h over the phase variables ϕ_x and ϕ_y.

Substituting the expansion of h from Eq. (2.70) into Eq. (2.77) [29], the generating function F, which defines the transformation that reduces the one-turn map to a pure rotation, can be expressed as:

$$F = \sum_{jklm} f_{jklm} (2I_x)^{\frac{j+k}{2}} (2I_y)^{\frac{l+m}{2}} e^{-i[(j-k)(\phi_x+\phi_{x0})+(l-m)(\phi_y+\phi_{y0})]}, \tag{2.79}$$

where the generating function coefficients f_{jklm} are related to the Hamiltonian coefficients h_{jklm} by:

$$f_{jklm} = \frac{h_{jklm}}{1 - e^{-i2\pi[(j-k)Q_x+(l-m)Q_y]}}. \tag{2.80}$$

It can be seen that the denominator approaches zero when $(j-k)Q_x + (l-m)Q_y = o$, where $o \in \mathbb{Z}$, giving the resonance condition as detailed earlier in Eq. (2.34), where the motion becomes non-regular.

Each term in the generating function, or equivalently in the transformed Hamiltonian, is associated with a specific resonance condition. These terms are therefore referred to as RDTs, which quantify the extent to which a given resonance distorts the phase-space away from an ideal pure rotation. As such, they provide a direct measure of the strength of individual resonances through their perturbative effect on the particle motion.

In order to relate this formalism to the TbT motion, the normal form coordinates can first be expressed as

$$\zeta_{z,\pm} = \sqrt{2I_z}\, e^{\mp i(\phi_z+\phi_{z0})}. \tag{2.81}$$

Since, in normal form coordinates, the TbT evolution corresponds to a pure rotation, this leads to the following expression for ζ as a function of turn number,

$$\zeta_{z,\pm}(N) = \sqrt{2I_z}\, e^{\mp i(2\pi Q_z N+\phi_{z0})}, \tag{2.82}$$

where N denotes the number of turns.

The TbT motion in normalised coordinates (defined by Eq. (2.13)) can then be reconstructed by transforming back to the normalised Courant-Snyder coordinates, using the generating function,

$$\Gamma_{z,\pm} = e^{:F:} \zeta_{z,\pm}. \tag{2.83}$$

The TbT evolution of the normalised coordinates can be expressed as a function of the turn number as

$$\begin{aligned}
\Gamma_{x,-}(N) &= \hat{x}(N) - i\hat{p}_x(N) = \sqrt{2I_x}\, e^{i(2\pi Q_x N + \phi_{x0})} \\
&\quad - 2i \sum_{jklm} j f_{jklm} (2I_x)^{\frac{j+k-1}{2}} (2I_y)^{\frac{l+m}{2}} e^{i\left[(1-j+k)(2\pi Q_x N + \phi_{x0}) + (m-l)(2\pi Q_y N + \phi_{y0})\right]}, \\
\Gamma_{y,-}(N) &= \hat{y}(N) - i\hat{p}_y(N) = \sqrt{2I_y}\, e^{i(2\pi Q_y N + \phi_{y0})} \\
&\quad - 2i \sum_{jklm} l f_{jklm} (2I_x)^{\frac{j+k}{2}} (2I_y)^{\frac{l+m-1}{2}} e^{i\left[(k-j)(2\pi Q_x N + \phi_{x0}) + (1-l+m)(2\pi Q_y N + \phi_{y0})\right]}.
\end{aligned} \tag{2.84}$$

This expression may be interpreted as a spectral decomposition of the transverse beam motion. The linear contribution, represented by the first term in Eq. (2.84), produces an oscillation at the fundamental betatron tune, Q_x. The remaining terms, which are proportional to the RDTs, give rise to additional spectral components at frequencies corresponding to characteristic combinations of the horizontal and vertical tunes. For the horizontal spectrum, these frequencies are of the form

$$(1-j+k)Q_x + (m-l)Q_y.$$

Consequently, spectral analysis of the transverse TbT beam position at a given observation point provides direct information on the underlying non-linear resonances that perturb the particle motion at that location. Furthermore, Eq. (2.84) shows that only those RDTs satisfying $j \neq 0$ contribute to the horizontal spectrum, whereas only terms with $l \neq 0$ contribute to the vertical spectrum. Resonance driving terms for which $j = 0$ are therefore absent from the horizontal spectrum, while those with $l = 0$ do not appear in the vertical spectrum.

Table 2.1 summarises the relationship between the resonance terms f_{jklm} and the corresponding spectral lines for several RDTs considered in this thesis.

Table 2.1: Selected resonance driving terms and the corresponding lines in the horizontal and vertical TbT spectra.

RDT	Resonance	Horizontal Line	Vertical Line
f_{jklm}	$(j-k)Q_x + (l-m)Q_y$	$(1-j+k)Q_x + (m-l)Q_y$	$(k-j)Q_x + (1+m-l)Q_y$
f_{1020}	$Q_x + 2Q_y$	$-2Q_y$	$-Q_x - Q_y$
f_{0030}	$3Q_y$		$-2Q_y$
f_{4000}	$4Q_x$	$-3Q_x$	
f_{2002}	$2Q_x - 2Q_y$	$-Q_x + 2Q_y$	
f_{0040}	$4Q_y$		$-3Q_y$
f_{1012}	$Q_x - Q_y$	Q_y	$-Q_x + 2Q_y$
f_{0140}	$-Q_x + 4Q_y$		$Q_x - 3Q_y$

2.6 Luminosity

One of the primary objectives of investigating both linear and NL optics is to maximise the achievable integrated luminosity, which constitutes a principal performance metric of a collider such as the LHC. Instantaneous luminosity determines the rate of particle collisions and is therefore of fundamental importance for high-energy physics experiments.

Under the assumption of head-on collisions between two uncorrelated Gaussian beams, the luminosity can be expressed as

$$\mathscr{L} = \frac{N_1 N_2 f_{\text{rev}} n_B}{2\pi\sqrt{\sigma_{1x}^2 + \sigma_{2x}^2}\sqrt{\sigma_{1y}^2 + \sigma_{2y}^2}}, \tag{2.85}$$

where $N_{1,2}$ denote the number of particles per bunch in the two colliding beams, f_{rev} is the revolution frequency, and n_B represents the number of colliding bunch pairs. The quantity σ_{nz} denotes the r.m.s. beam size, where n denotes the beam number, and z denotes either the horizontal (x) or vertical (y) plane [31].

CHAPTER 3

THE LHC

3.1 The LHC Injector Complex

Having established the fundamental performance considerations, the discussion will now turn to the CERN accelerator complex as a whole. It is responsible for preparing the proton and ion beams used in the LHC, as well as supplying beams for several experiments. Examples of such facilities include the Antiproton Decelerator (AD), the Extra Low ENergy Antiproton ring (ELENA), Neutron Time-of-Flight facility (n-ToF), the North Area (NA), and the Advanced WAKEfield experiment (AWAKE), all of which are illustrated in Fig. 3.1.

Since this thesis focuses on proton studies, only the proton injection chain will be described in detail. While ions are accelerated from a different starting point (Linac3), protons begin their journey as H^- ions in Linac4, a linear accelerator that boosts their momentum to 571 MeV/c (the kinetic energy, E_k, being 160 MeV). The beam is then transferred to the Proton Synchrotron Booster (PSB), where the H^- ions are stripped, leaving only the protons, and the momentum is increased to 2.8 GeV/c (the E_k being 2 GeV) for injection into the Proton Synchrotron (PS). The PS is a larger circular accelerator with a circumference of approximately 628 m, which further accelerates the protons to about 26 GeV/c. The protons are then injected into the Super Proton Synchrotron (SPS), the most powerful accelerator in the injector chain, which brings the beam momentum up to 450 GeV/c, the injection momentum required by the LHC [32].

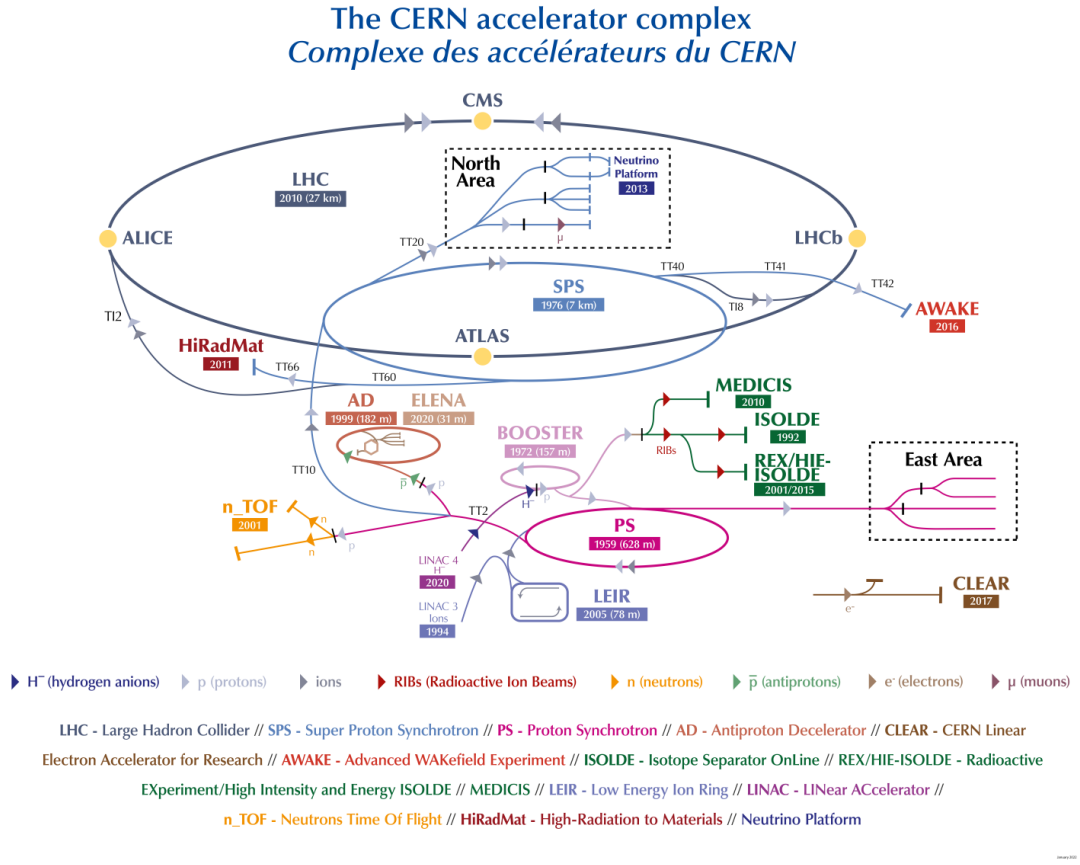


Figure 3.1: The overall layout of the LHC injector complex [2].

3.2 The LHC

The LHC is a twin-ring synchrotron, with a circumference of approximately 27 km. It is composed of eight arcs (see Subsection 3.3) separated by eight straight sections, referred to as Interaction Regions (IRs) (IR1-8). Four of these IRs house the experiments: A Toroidal LHC ApparatuS (ATLAS) in IR1, A Large Ion Collider Experiment (ALICE) in IR2, Compact Muon Solenoid (CMS) in IR5, and Large Hadron Collider beauty (LHCb) in IR8. The remaining IRs are dedicated to accelerator systems, with momentum and betatron collimation located in IR3 and IR7, the RF system and beam instrumentation in IR4, and the beam extraction and dump system situated in IR6. In addition, IR2 and IR8 contain the injection systems for LHCB1 and LHCB2, respectively.

For the experimental IRs, the Interaction Point (IP) is located at the geometric centre of the corresponding insertion region in all cases except IP8. At IP8, the detector is displaced

longitudinally by three half bunch spacings (11.25 m) [33] (see p. 3) to ensure optimal alignment of the colliding bunches. The overall layout of the accelerator is shown in Fig. 3.2, while the associated machine naming conventions are summarised in Fig. 3.3.

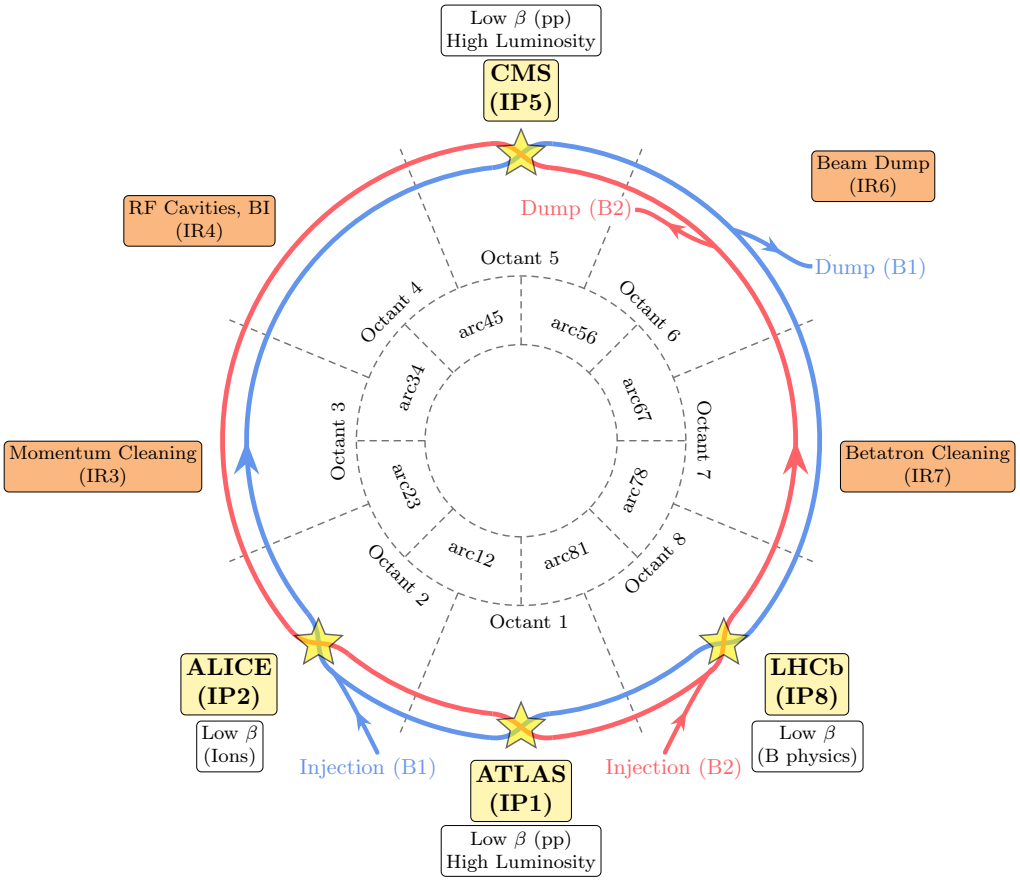
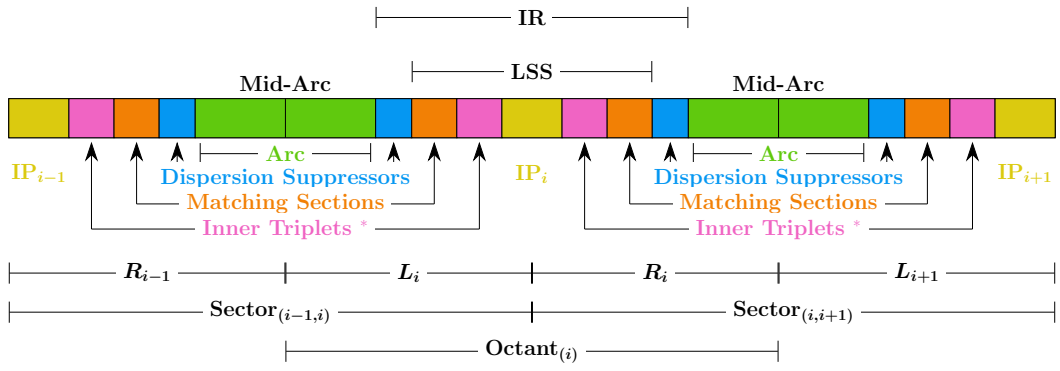


Figure 3.2: The layout of the LHC with LHCB1 in blue and LHCB2 in red.



* Inner Triplets only present in the experimental insertions

Figure 3.3: A visual representation of the naming conventions for hardware placement in the LHC, where LSS denotes Long Straight Section.

Comprehensive technical details of the LHC are available in [34]. As of Run III, which commenced in April 2022, the LHC operates at a centre-of-mass energy of approximately 13.6 TeV, with up to around 2400 proton bunches circulating in each beam.

3.3 An LHC arc cell

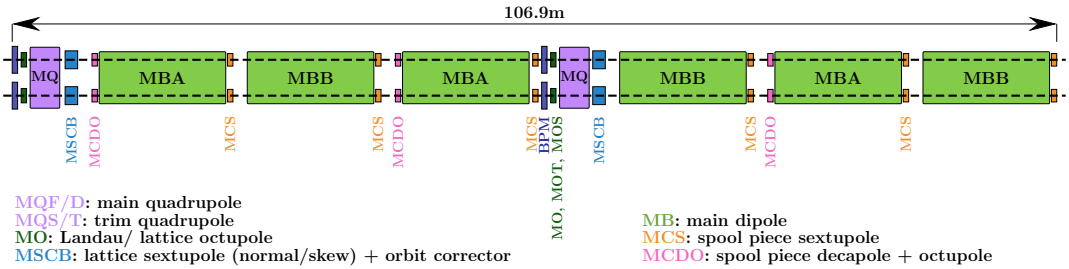


Figure 3.4: Illustration of an LHC FODO cell.

Each arc of the LHC is composed of 23 repeating FODO cells. An example of a single cell is illustrated in Fig. 3.4. The cell comprises several distinct magnetic elements, whose full naming conventions can be found in [34] (see pp. 15-16), described in detail below.

- **Main Dipole** (MBA/B): The main dipole magnets of the LHC are distinctive in two respects. First, they operate in a superconducting state at 1.9 K (apart from some exceptions such as some of the separation and correction dipoles), and second, each magnet contains two beam channels with opposite magnetic field directions within a single cold yoke. The two apertures are mechanically coupled through common collars [35]. These magnets provide the bending field (up to 8.33 T) required to steer the beams along the LHC's circular trajectory.

- **Main Quadrupoles** (MQF/D): Similar to the main dipoles, the main quadrupoles also operate at 1.9 K and feature two apertures, each carrying beams in opposite field directions within the same magnetic yoke. However, unlike the dipoles, the two apertures in the quadrupoles are mechanically independent. These magnets provide the alternating focusing and defocusing fields in the horizontal and vertical planes necessary for transverse beam confinement and stability.

- **Spool piece correctors**: Due to the dual-bore configuration and superconducting nature of the main arc magnets, significant higher-order field errors are present. To correct these, NL correction windings are mounted directly onto the extremities of the main dipoles.

  - **Sextupole spool piece magnets** (MCS): The main dipoles introduce b_3 field integral errors. These errors comprise a dynamic component, arising from persistent current effects within the NbTi superconducting filaments, which is particularly pronounced at injection energy due to the low magnetic field strength [36]. In addition, a static systematic b_3 component originates from the coil and wedge geometry of the dipole design, which prevents self-compensation of the sextupolar field component [37].

    As a result, the presence of a residual sextupolar term is an intrinsic feature of the LHC main dipoles (further details are provided in the following section). To mitigate this effect, each bore of a main dipole is equipped with a small MCS corrector at one end (Type B magnet) [34] (see pp. 31-41). These correctors are connected in series to form two families per sector, one for each beam.

  - **Octupole/decapole spool piece magnets** (MCD/O): The bores of every other dipole are equipped with a combined octupole-decapole spool-piece (MCDO) at the opposite end (Type A magnet). In a similar alternating configuration to the MCS, the MCDO correctors are employed to correct the higher-order b_4 and b_5 multipole components generated by the main dipoles.

- **Lattice correctors**: These magnets are mounted at the main arc quadrupoles, with all correctors powered independently for the two beams. The principal types are described below:

  - **Orbit Correctors** (MCB): The horizontal (MCBH) and vertical (MCBV) arc orbit corrector magnets are installed at each focusing and defocusing quadrupole and are used for orbit correction.

  - **Trim Quadrupoles** (MQT/S): The normal trim quadrupoles (MQT) are primarily used for tune correction, while the skew trim quadrupoles (MQS) serve to correct linear betatron coupling between the horizontal and vertical planes.

– **Lattice Sextupoles** (MS/MSS): Commonly referred to as chromaticity sextupoles, these magnets are mounted at the main quadrupoles of the lattice. The normal sextupoles (MS) are installed at each focusing and defocusing quadrupole, arranged into four families per sector and per ring, MSF and MSD, mounted on the focusing (MQF) and defocusing (MQD) quadrupoles, respectively. Each family is further divided into two interleaved sub-families, designated as 1 and 2. The skew sextupole counterparts (MSS) are installed exclusively at the focusing quadrupoles. The normal sextupoles are responsible for correcting the natural first-order chromaticity, whereas the skew sextupoles compensate for chromatic coupling [38, 39].

– **Lattice (also known as Landau) Octupoles** (MOF/D): Each short straight section without an MQT or an MQS has an MO. These are powered in four families per sector, subdivided into focusing and defocusing magnets, MOF and MOD. Their purpose is to create amplitude detuning to damp the coherent oscillations caused by collective effects.

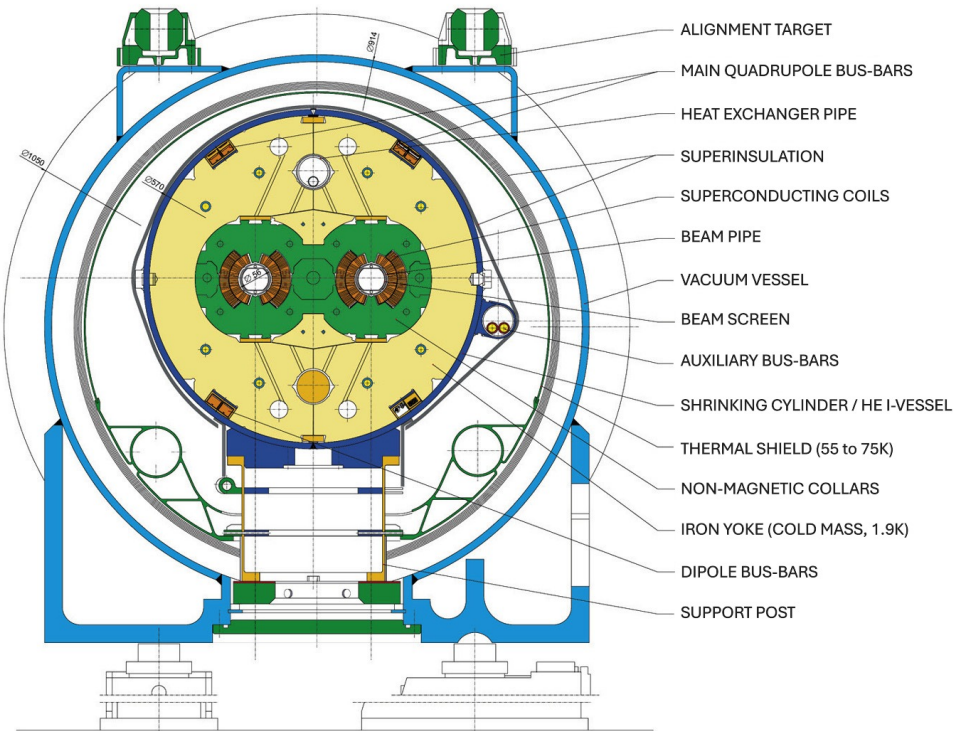


Figure 3.5: Cross-section of a LHC dipole [3].

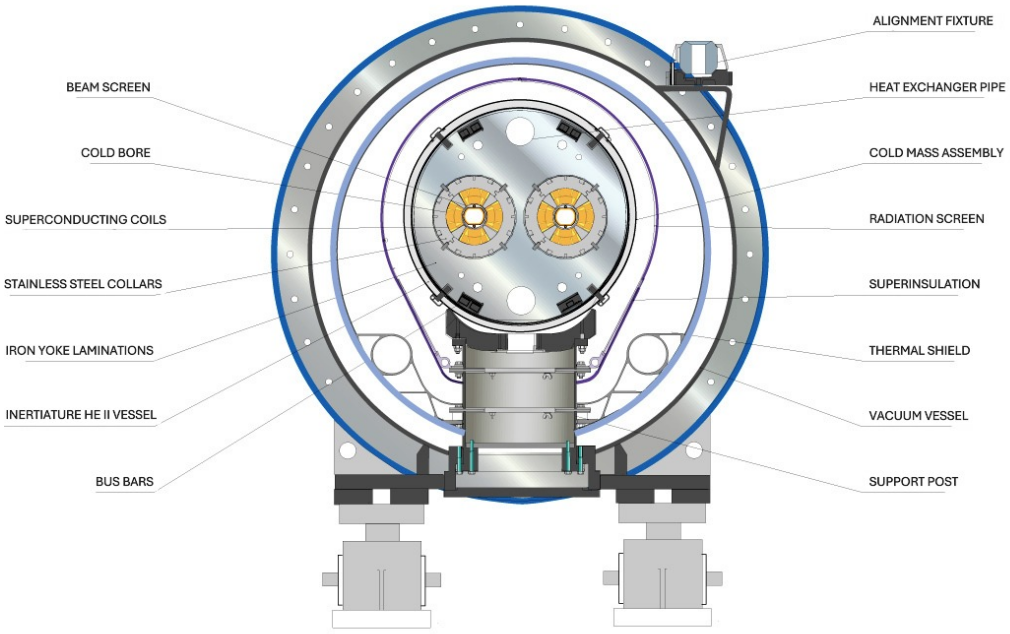


Figure 3.6: Cross-section of a LHC quadrupole [4].

3.3.1 b_3 decay

As discussed previously, dedicated sextupole spool-piece corrector magnets (MCS) are installed at the ends of the LHC main dipoles. Their inclusion is necessary because the dipoles exhibit a significant b_3 field error component, part of which is dynamic [40, 41]. This behaviour originates from persistent currents in the superconducting filaments, which relax gradually after the magnetic field is established. Two distinct mechanisms contribute to this decay. The first is flux creep, a thermally activated process that is independent of the magnet excitation history and can be reduced by lowering the operating temperature. The second mechanism is linked to the internal structure of Rutherford-type cables and occurs only when the magnet current has been cycled [41].

3.3.2 WISE and FiDeL

During the construction of the LHC, measurements were carried out to determine both the individual magnet misalignments within the cryostats and the alignment of the cryostats within the tunnel. The Windows Interface to Simulation Errors (WISE) [42] was developed to process

these data and to generate statistical estimates of the resulting geometric errors in the LHC lattice.

In addition to geometric errors, WISE also provides estimates of magnetic field imperfections based on direct magnetic measurements and their associated uncertainties. Measurements at low current were performed at room temperature for all magnets during industrial production ("warm measurements"), while a subset of magnets underwent "cold measurements" under operational cryogenic conditions after delivery to CERN. For magnets without cold measurements, a warm-to-cold correlation was applied in the field modelling; however, this correlation introduces an additional uncertainty. Typical sources of uncertainty included in the WISE model are those associated with the warm-to-cold correlation, measurement precision, hysteresis effects, and power-supply accuracy.

To capture the possible magnetic configurations of the accelerator, WISE defines sixty distinct magnetic instances, referred to as seeds, each representing a plausible magnetic state of the LHC. Both geometric and magnetic error estimates generated by WISE can be incorporated into lattice models for numerical simulations of beam dynamics.

The Field Description of the LHC (FiDeL) [43] is implemented alongside WISE. In addition, an online version of FiDeL provides snapshots of the LHC magnetic state at specific operating points, for accelerator studies and diagnostics. Its primary objective is to deliver the integral transfer functions, relating integral field to current, in a form that can be inverted to obtain the current corresponding to a desired integral field for each circuit in the LHC.

In addition, for the main ring magnets, FiDeL provides predictions of the magnetic field errors, which can be used to determine the appropriate settings of the corrector circuits.

3.4 The LHC IRs

Each Insertion Region (IR) has configurations specifically optimised for the functional requirements of that insertion [44]. In the experimental IRs (IR1, IR2, IR5, and IR8), which also encompass their respective IPs, the β-function is reduced to small values to maximise luminosity. This effect is most pronounced in IR1 and IR5, where ATLAS and CMS typically operate with the

most tightly squeezed beams, reaching β^* values as low as 18 cm in one plane as of 2025. The value of β^* is generally larger during β^*-levelling, as well as in dedicated high-β runs, Van der Meer optics [45, 46] configurations, low energy operation, and ion and light ion runs.

In contrast, IP2 and IP8 typically operate with less strongly squeezed beams, with β^* values of approximately 10 m for ALICE and 2 m for LHCb during proton-proton (pp, as indicated in Fig. 3.2) collisions. Under specific operating conditions, such as heavy-ion or proton-ion runs, the β^* at ALICE can be reduced in order to enhance collision performance for these dedicated modes of operation.

To achieve small β^* values at the IPs, the beams are focused using a triplet configuration of quadrupole magnets located within the IRs. This arrangement necessitates large β-functions in the triplets and adjacent elements, as illustrated in Figs. 3.7 and 3.8. A schematic of the LHC lattice is also shown, with the various magnet functions listed in the caption of Fig. 3.7.

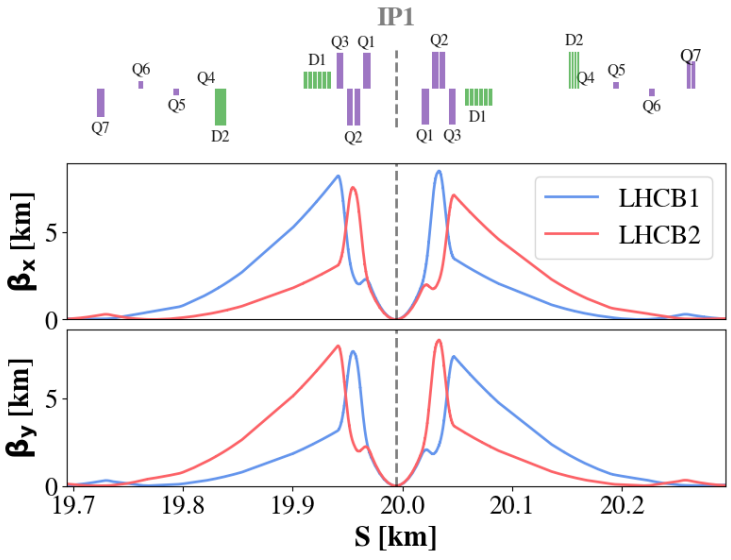


Figure 3.7: Simulated β-function in the ATLAS IR with *round* beams i.e. β^* = 0.6/0.6 m (horizontal/vertical plane). The schematic above shows the strengths of the magnets in the lattice where: $Q1$ to $Q3$ are the triplet quadrupoles responsible for beam focusing; $D1$, and $D2$ are the separation/recombination dipoles which guide each beam from its individual aperture into the common beam pipe around the IP and back into its own aperture on the far side; and $Q4$ to $Q7$ are the matching quadrupoles.

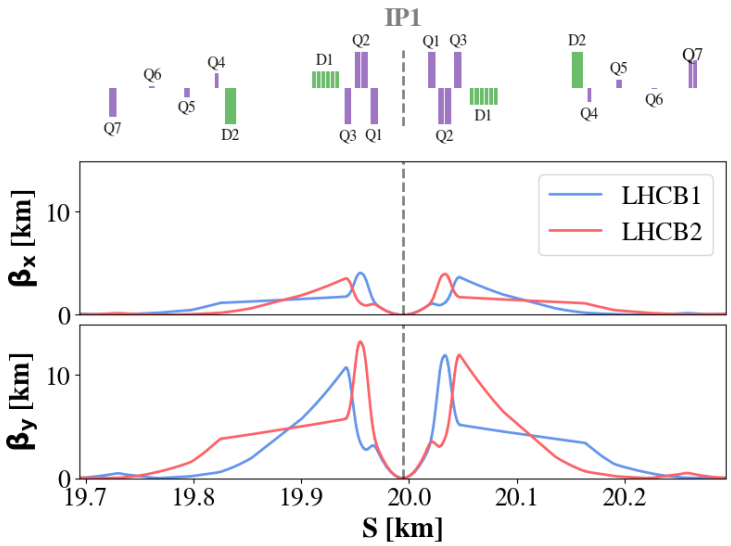


Figure 3.8: Similar to Fig. 3.7 but with *flat* beams i.e. β^* = 0.60/0.18 m (horizontal/vertical plane).

Due to the large β-functions on either side of the IPs, magnetic imperfections in the IRs elements can have a pronounced impact on beam dynamics. To correct for these effects, both linear and NL corrector magnets are strategically installed within the insertions, symmetrically arranged about each experimental IP. Of particular relevance to this work are the b_3 correctors, along with the a_3, b_4, a_4, and b_6 correctors, which are located on the non-interaction-point sides of the Q3 triplet quadrupoles. The layout of these corrector packages within the LHC experimental insertions is shown schematically in Fig. 3.9, which depicts the right side of the IP.

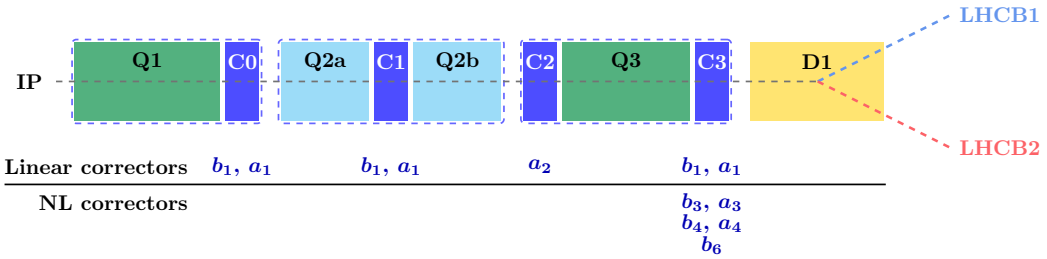


Figure 3.9: The layout of the IR and the triplet corrector package, where $Q1$, $Q2a/b$ and $Q3$ are the triplet quadrupoles and $C0$-$C3$ are the corrector packages with the order to which they can correct listed below. $D1$ is the separation dipole used to split LHCB1 and LHCB2.

During collisions, the LHC operates with many bunches, and there are parasitic crossing points on either side of the IPs. In order to prevent collisions at these locations, closed orbit bumps are introduced in a single plane of each experimental IR to generate a crossing-angle at

the IP. The simulated crossing-angle for IP1 is shown in Fig. 3.10 (for β^* = 0.6/0.18 m at 6.8 TeV). There are differing conventions for the definition of crossing-angle; the convention used in this thesis will be the angle, θ depicted in Fig. 3.10.

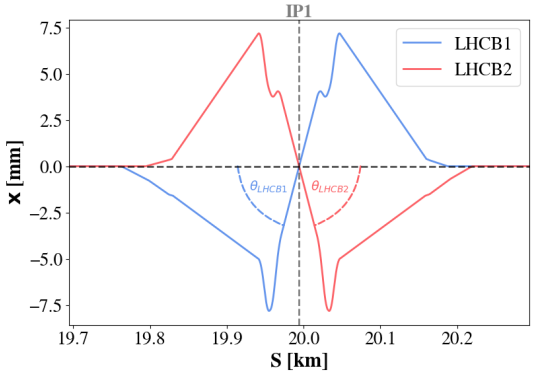


Figure 3.10: Simulated horizontal trajectories x vs s depicting a crossing-angle of 160 μrad for β^* = 0.6/0.18 m (horizontal/vertical) at 6.8 TeV.

The remaining IRs in the LHC serve specific functions, including collimation (see Sec. 3.8.8), RF acceleration, beam injection, and beam extraction.

3.5 Radio Frequency (RF) system

The RF system of the LHC consists of eight RF stations per beam, located in IR4 and operating at a frequency of 400 MHz [34, 47]. The synchronous particle is defined as the particle that remains in phase with the RF system and therefore experiences the nominal accelerating voltage on every revolution. Consequently, it remains in phase with the RF field at the operating frequency of 400 MHz. Particles with small deviations in energy or arrival time execute longitudinal oscillations about the synchronous particle in longitudinal phase space. As a result, the beam is confined into discrete bunches rather than being uniformly distributed around the accelerator circumference. Each bunch is confined within a stable region of longitudinal phase space known as an RF bucket.

The RF system is also responsible for accelerating the beam from the injection energy of 450 GeV to the nominal collision energy. In addition to providing the required energy gain per turn, it compensates for energy losses during operation and provides the longitudinal focusing

necessary to maintain stable beam bunches within the RF buckets. The associated radial feedback loop continuously monitors the beam position and corrects radial deviations from the reference orbit by applying appropriate adjustments to the RF frequency, therefore maintaining the desired beam momentum and orbit.

Beyond its operational role, the RF system is routinely employed in accelerator optics studies to introduce controlled momentum offsets, $\Delta p/p$. Such controlled variations enable the measurement and characterisation of chromatic effects, including chromatic β-beating and dispersion, which are essential for understanding and optimising the machine optics.

3.6 Beam injection and extraction

Beam injection into the LHC occurs approximately 150 m upstream of IP2 for LHCB1 and 160 m downstream of IP8 for LHCB2. In both cases, the beam enters the LHC from outside and below the accelerator plane.

During injection, the beam quality is carefully optimised to ensure favourable conditions for subsequent measurements. In particular, the transverse emittance can be evaluated using diagnostics such as the Beam Synchrotron Radiation Telescope (BSRT) or the Beam Wire Scanner (BWS), as described in Subsection 3.8.5, then corrected for subsequent injections if necessary.

For beam extraction, each beam line is equipped with fifteen beam extraction kicker magnets (MKDs), which are capable of deflecting the entire beam within a single turn. This places stringent requirements on the quality of the optics, particularly at top energy, to ensure that the beam is correctly directed towards the beam dump block.

3.7 The LHC cycle

In the LHC, the period between beam setup and beam dump is referred to as a fill. A standard physics fill begins with a magnet pre-cycle, conducted with no beam present in the accelerator [48]. During this stage, the current in each magnet is first ramped up to a large positive (or, in some cases, negative) value and then ramped down to its required operational setting. The exact current profile depends on the magnet type, for example, the MOs are degaussed. This pre-cycling

procedure aims to ensure that all magnets are aligned on the correct hysteresis branch, so that they exhibit reproducible magnetic behaviour, accounting for effects such as decay, snap-back, and residual magnetisation. This does not occur if the previous fill reached top energy and finished normally, since the previous fill can be considered as the pre-cycle for many magnets.

Following the pre-cycle, the first beam injection takes place from the injector chain at an energy of 450 GeV. This initial injection typically consists of a low-intensity pilot beam, containing a single bunch with an intensity of around 1×10^{10} particles per bunch (ppb), roughly corresponding to the LHC quench limit at injection energy. The pilot beam is used primarily to check that the accelerator is in a safe state, confirming main beam quantities such as orbit, tune, chromaticity, and coupling before injecting an unsafe beam. During this stage, various accelerator subsystems are tested to ensure that a stable circulating beam can be maintained [49].

Once a safe beam condition has been established, an injection setup beam is introduced. While the pilot beam allows only limited measurement precision, the setup beam is designed to closely reproduce the characteristics of the subsequent physics beam, particularly the transverse emittance, but with a reduced number of bunches to remain safely below the machine-protection damage threshold. After optimisation with the setup beam, the physics beam is injected at its nominal intensity with β^* at IP1 and IP5 at 11 m.

The beam is subsequently accelerated to the top energy of 6.8 TeV, while simultaneously undergoing a β^* squeeze, during which the β-functions at the IPs are progressively reduced. This process culminates in the flat-top phase, where the β^* at both CMS and ATLAS is approximately 2 m (in 2025).

The accelerator subsequently enters the adjust phase, during which preparations are made for a further squeeze to a β^* suitable for collisions, typically around 1 m. A sequence of additional manipulations is then performed, including optimisation of the WP, adjustment of the crossing scheme if required, and the controlled approach of the beams into collision, ultimately leading to the stable beams phase.

From the onset of collisions, β^*-levelling is applied, progressively reducing β^* towards its minimum value, while maintaining a constant instantaneous luminosity for the ATLAS and CMS experiments. Once the minimum β^* is reached, operation continues in this configuration until

the beam intensity has decreased to a level at which it becomes more efficient, from a luminosity perspective, to dump the beam and refill the accelerator.

Once the physics fill is completed, the beams are safely dumped, followed by a ramp-down, magnet cycling, and accelerator recovery, after which the next operational cycle can begin.

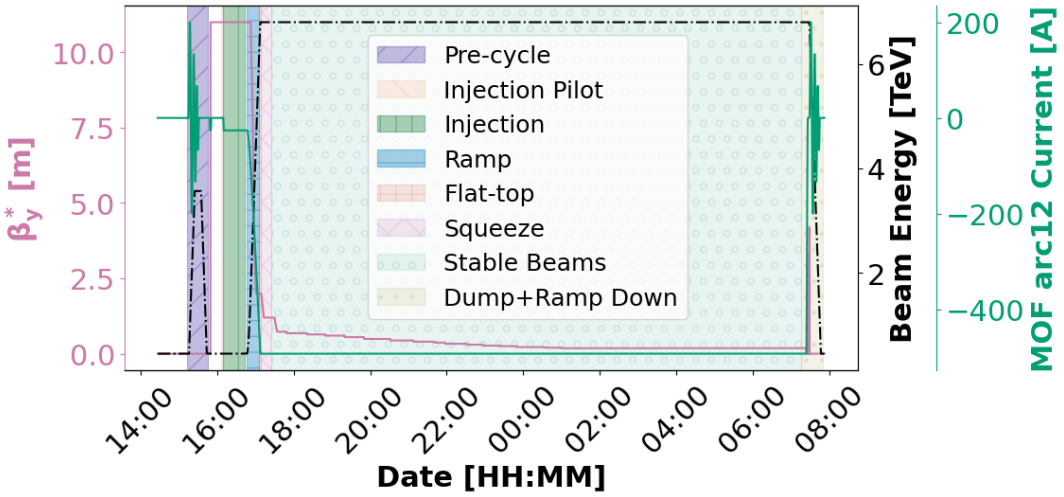


Figure 3.11: Depiction of the LHC cycle for a typical fill in 2025, with the various operational stages highlighted in different colours. The evolution of β^* is shown in pink (left-hand axis), the beam energy, defined to be the energy per nucleon, is indicated by the black dashed line (first right-hand axis), and the current in a Landau octupole circuit is represented in green (second right-hand axis), all as functions of time.

3.7.1 Beam types

Depending on the specific studies being carried out in the LHC, a variety of beam filling schemes can be employed. During standard physics operation, the beam typically consists of multiple bunches organised into trains for each beam, with a nominal bunch population in 2025 of 1.7×10^{11} ppb.

The most commonly utilised operational configuration is the “25 ns physics beam.” In 2025, this scheme consists of 68 PS batches, each containing 36 bunches. These are primarily arranged as four successive injections of 36-bunch trains into the SPS, followed by a final injection of 12 bunches. The complete set is subsequently transferred to the LHC. This injection sequence is repeated multiple times to accumulate approximately 2000 bunches in the collider. The bunch spacing of 25 ns is defined by the PS [50]. Alternative filling schemes have also been implemented.

One notable example is the 8b4e configuration, which comprises trains of eight bunches separated by 25 ns, followed by four empty slots. This pattern is specifically designed to mitigate certain collective effects [51].

Nominal intensity bunches can also be injected one at a time (Individual Bunch Physics Beam - INDIV) for setup or specific accelerator studies.

For dedicated Machine Development (MD) studies or optics measurements, other beam types are used. These include the use of pilot beams, which provide low-intensity, safe conditions that are preferred for machine-protection reasons when applying AC-dipole (ACD) excitations for measurements such as those performed by the Optics Measurements and Correction (OMC) team.

3.8 Beam Instrumentation/Hardware

3.8.1 Base-Band Tune (BBQ) System

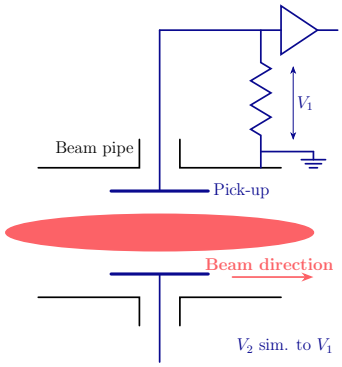
The Base-Band Tune (BBQ) system [52, 53] is an essential diagnostic tool for the non-invasive measurement of the beam tune, Q. The system operates by continuously monitoring the natural oscillations of the beam, which may arise from ambient noise or controlled excitation, using highly sensitive detection electronics. The acquired signal is processed through analogue demodulation to extract the base-band components, which are subsequently digitised and analysed in the frequency domain.

This approach enables a precise determination of the tune without introducing significant perturbations to the beam. As such, the BBQ system plays a key role in establishing and maintaining the correct accelerator configuration for optics studies.

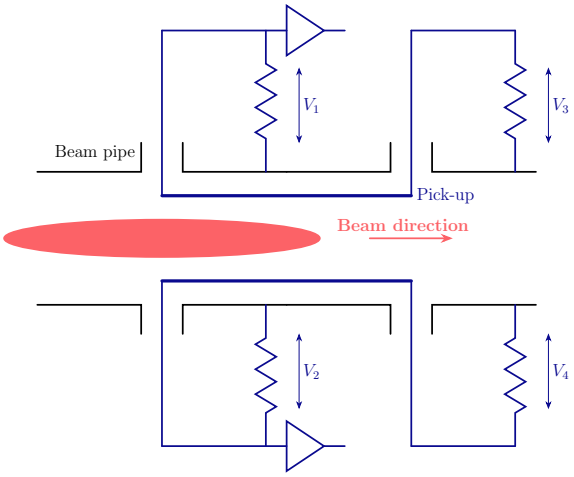
3.8.2 Beam Position Monitors (BPMs)

The LHC is equipped with approximately 500 Beam Position Monitors (BPMs) per beam, distributed throughout the accelerator ring. Two primary types of BPMs are employed: strip-line (inductive) BPMs and button (capacitive) BPMs. Strip-line BPMs (see Fig. 3.12(b)) are located in regions with common apertures, where their ability to distinguish between the two counter-

rotating beams is essential. Button BPMs (see Fig. 3.12(a)), in contrast, are employed throughout the remainder of the ring, particularly in the arcs, and are mounted near the main quadrupole magnets. Some larger versions are also installed in the IRs.



3.12(a) Diagram of button BPM measurement principle.



3.12(b) Diagram of strip-line BPM measurement principle.

All BPMs operate on the principle that a passing charged particle bunch induces an electromagnetic signal on the surrounding electrodes. The beam position is determined using the difference-over-sum method, in which the difference in signal amplitudes between opposing electrodes is normalised by their sum and multiplied by a calibration factor, as expressed in Eq. (3.1):

$$\text{pos} = f \cdot \frac{V_1 - V_2}{V_1 + V_2}, \tag{3.1}$$

where pos represents either the horizontal (x) or vertical (y) position, V_1 and V_2 are the signal amplitudes from the corresponding opposing electrodes (e.g. left/right or up/down), and f is the calibration factor. Each BPM contains two orthogonal pairs of electrodes, enabling independent position measurements in both transverse planes.

The BPMs are spaced to provide a betatron phase advance of approximately 45°, which is advantageous for the reconstruction of linear optics. This configuration also enables phase-space reconstruction through the comparison of TbT data between BPMs separated by approximately 90° in betatron phase.

Before any optics measurements, BPM calibration is required [54, 55], followed by the correction of geometric and electrical non-linearities in the BPMs distributed around the ring [56–58].

3.8.3 Beam Loss Monitors (BLMs)

The primary purpose of the Beam Loss Monitors (BLMs) system is to protect the accelerator against the adverse effects of beam losses. Even the loss of a small fraction of the circulating beam can induce a quench in the superconducting magnets, or cause damage to accelerator components. By detecting the secondary particles produced by lost protons, the BLMs system enables the monitoring of beam losses, and triggers a beam dump when predefined thresholds are exceeded, hence safeguarding the accelerator against quenches and potential damage [59, 60].

The measurement principle is based on detecting the energy deposited by secondary shower particles using ionisation chambers located outside of the magnet cryostats. To observe a representative fraction of the secondary particle flux, detectors are placed at likely loss locations. The calibration of the damage and quench level thresholds with respect to the measured secondary particle energy deposition is simulation-based [34] (see p. 361).

BLMs are employed during optics measurements to monitor beam losses and to ensure that the applied beam excitation does not result in unacceptable degradation of beam quality, therefore preventing unnecessary beam dumps. This is particularly important at top energy, where an unscheduled beam dump results in the loss of several hours of machine availability. In addition, the beam lifetime can be estimated from the secondary radiation measured by the BLMs, as described in [60]. However, these estimates may occasionally indicate an erroneously short lifetime following a large instantaneous loss. Such events are often dominated by losses originating from the beam halo rather than the beam core and can saturate the measurement, therefore reducing the reliability of the inferred beam lifetime. This behaviour can be identified by the apparent beam lifetime gradually increasing following the initial loss event, as the contribution from the transient halo losses diminishes and the measurement returns to representing the losses from the beam core.

3.8.4 Fast Beam Current Transformers (FBCTs)

Fast Beam Current Transformers (FBCTs) are employed in the LHC to provide bunch-by-bunch and TbT intensity measurements [61]. Each FBCT consists of a single high-permeability magnetic core mounted around a dielectric insert located within the vacuum chamber. The core is wound with a secondary wire coil, the ends of which are connected to a sensing resistor [62]. The beam current is measured indirectly by detecting the magnetic field produced by the passage of each bunch, which induces a current in the secondary winding of the transformer surrounding the beam.

The intensity measurements provided by these transformers can be used to ensure that the excitation amplitude during optics measurements remains within acceptable limits. Furthermore, such measurements offer insight into the forced DA by enabling the monitoring of beam losses induced by the excitation of betatron oscillations.

3.8.5 BWS and BSRT

The LHC's Beam Wire Scanners (BWSs) are used to obtain high-precision transverse beam profile measurements. Wire scanners work by moving a thin wire through the particle beam; secondary particles produced by beam-wire interactions are collected and analysed to reconstruct the beam profile in one plane. They are limited to a maximum beam intensity beyond which there is risk of wire breakage or quenching of superconducting magnets due to beam-wire induced losses [63].

The Beam Synchrotron Radiation Telescope (BSRT) provides accurate, continuous, and non-destructive measurements of the transverse beam profile. Operational since 2015, the system has undergone significant consolidation and a series of upgrades, resulting in improved reliability and performance.

The BSRT operates by analysing the synchrotron radiation emitted by protons in the superconducting dipoles located in IR4. This radiation is transported to an optical table positioned below the beam line, where it is processed to reconstruct the transverse beam profile [64, 65].

Both of these instruments can be used to monitor the beam quality, ensuring that it remains suitable for optics measurements. In particular, they provide information on the beam emittance, allowing verification that the transverse beam size is sufficiently small to avoid particle losses

following a transverse excitation. Excessively large emittance would increase the oscillation amplitudes after the beam is kicked, potentially leading to beam loss through aperture limitations. Typical transverse emittances are approximately 3 μm for pilot beams and 2 μm for standard operational beams.

3.8.6 AC-Dipole (ACD)

Many accelerator optics measurements rely on the excitation of transverse betatron oscillations. In the LHC, this capability is provided by a dedicated transverse kicker magnet. This device is a warm, normal-conducting magnet with a very fast field rise and fall time, allowing it to reach its nominal field strength and return to zero within approximately half of an LHC revolution period. In this configuration, it is used for single-turn excitation measurements. The maximum kick strength is sufficient to drive the beam beyond the available physical aperture, and the excitation amplitude must therefore be carefully selected to avoid beam losses while remaining large enough to provide an adequate measurement signal.

Historically, RDTs have been studied using data from free oscillations of the beam, for example following single-turn kicks or injection oscillations. However, this approach can be problematic, particularly in the presence of strong NL errors, where RDT measurements are most relevant. In such cases, the free oscillations can decohere rapidly, therefore limiting the number of turns available for spectral analysis. Furthermore, different RDT spectral lines decohere at different rates [29], which can introduce additional uncertainties and reduce the reliability of the measurement.

Owing to its fast ramping capabilities, the excitation magnet can instead be driven with a sinusoidal waveform to provide continuous AC excitation. In this mode of operation, it is referred to as the ACD [66–68]. This excitation produces a sustained and coherent transverse oscillation of the beam, making it particularly well suited for forced excitation studies, as it significantly increases the number of turns available for analysis.

To ensure adiabatic excitation and therefore avoid significant emittance growth [69], the ACD gradually increases the kick amplitude over approximately 1000 turns until the desired excitation level is reached. The beam is then driven for a specified duration, which in Run I was

typically 1000 turns, referred to as the excitation plateau, after which the amplitude is reduced in a similarly gradual manner, as can be seen in Fig. 3.13.

This capability was further enhanced through hardware upgrades implemented during Run I and Run II, which extended the duration of the excitation plateau. The current hardware configuration allows excitations of up to 6600 turns. As a consequence of this adiabatic excitation scheme, the beam remains available for subsequent measurements, eliminating the need for a beam dump and reinjection. This improves operational efficiency and makes the technique well suited for routine application during commissioning and MD sessions.

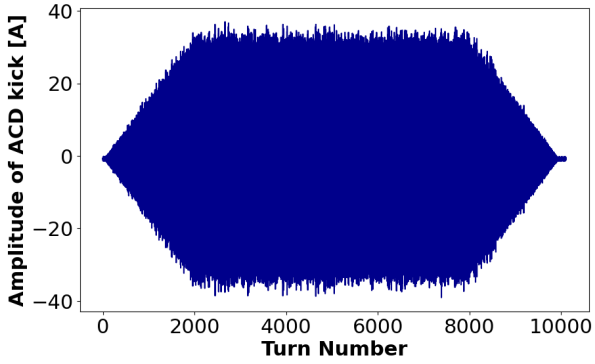


Figure 3.13: Diagram illustrating the resulting Turn-by-Turn data retrieved using an ACD excitation.

An additional advantage of this method is the precise control of the excitation frequency, which improves spectral resolution and facilitates the interpretation of the measured data.

3.8.7 Transverse Damper (ADT)

In the LHC, the Transverse Damper (ADT) is a feedback system that measures bunch-by-bunch oscillations and damps them with electrostatic kickers [70]. Its primary functions include the damping of injection-induced betatron oscillations, the suppression of emittance growth arising from ground motion and magnetic field noise or ripple, and the active stabilisation of the beam against collective effects [71].

Furthermore, the ADT can be operated in an excitation mode analogous to that of the ACD, albeit with a significantly smaller achievable amplitude, owing mainly to machine protection constraints since it can be used on nominal intensity bunches, but for an extended amount of time. This mode is commonly employed for fast linear coupling measurements [72]. However, in

general, the excitation provided by the ADT does not yield a sufficient signal-to-noise ratio for more detailed linear or NL optics measurements.

3.8.8 Collimators

The LHC collimation system is distributed around the ring and incorporates two dedicated warm cleaning insertions: IR3 for off-momentum cleaning and IR7 for betatron cleaning, as shown in Fig. 3.2. The collimation hierarchy begins with the primary collimators (TCPs), which are positioned closest to the beam in order to intercept particles with large transverse amplitudes. Downstream of the primaries, the secondary collimators (TCSGs) capture particles scattered from the TCPs. Further along the hierarchy, the tertiary collimators (TCTs) provide direct protection for the inner triplet magnets located near the IPs, while the absorber collimators (TCLAs) absorb secondary particle showers generated by upstream collimators [73].

Maintaining the integrity of this hierarchy is essential to ensure that the largest energy depositions occur within the designated absorbers rather than in vulnerable accelerator components. However, non-linear beam dynamics can distort the transverse phase space, allowing particles to bypass the primary collimators and therefore compromise the intended cleaning hierarchy. This mechanism has been strongly associated with excitation of the $3Q_y$ resonance, as demonstrated in [74].

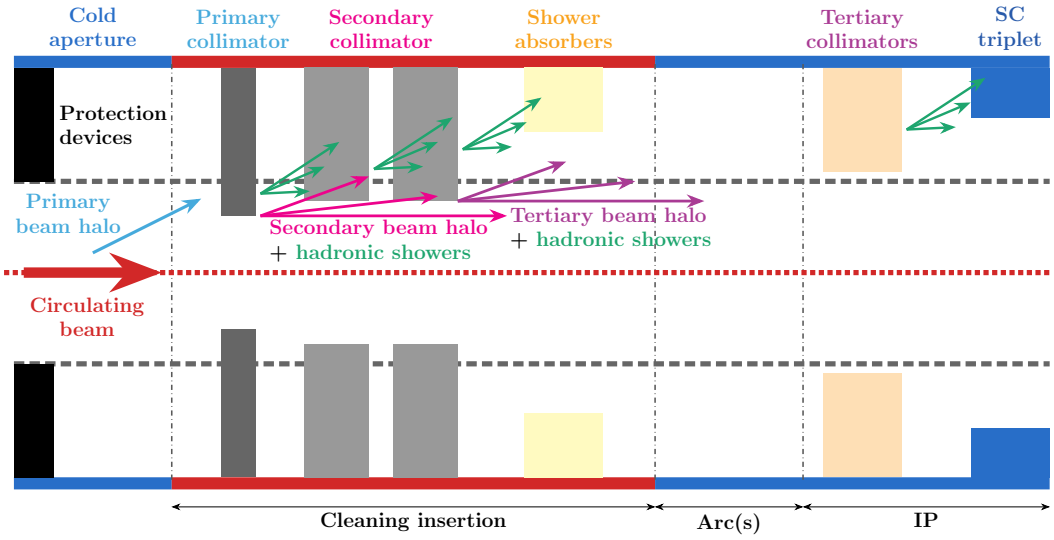


Figure 3.14: Illustration of the LHC collimator hierarchy.

3.9 Optics Measurements in the LHC

Optics measurements are performed by exciting the beam with the ACD for approximately 6600 turns. The excitation frequency is chosen to correspond to a tune offset, ΔQ, from the natural tune that is sufficiently small for the measured RDTs to remain representative of the free-RDTs, while remaining sufficiently large to avoid excessive beam loss, even for small excitation amplitudes. The selection of this offset is based on empirical observations of acceptable beam losses during operation. TbT beam position data are subsequently acquired from the BPMs distributed around the accelerator. Prior to further analysis, the TbT BPM data are processed using Singular Value Decomposition (SVD) [75], retaining the first 12 singular modes to suppress uncorrelated noise between individual BPMs.

The fractional tunes (q) and betatron phases are determined from the spectral analysis of the beam oscillations recorded in the TbT data. This analysis is performed using an interpolated Fast Fourier Transform (FFT) [76]. Within this framework, the natural fractional tunes, corresponding to the transverse oscillations in the absence of ACD excitation, and the fractional tunes associated with the ACD-driven motion are specified.

Frequency windows, typically around ±0.005, are defined around these values, and the dominant spectral peak within each window is required to correspond to the specified tune. An example of such a spectrum is shown in Fig. 3.15, where a peak can be observed immediately to the right of the horizontal ACD tune, labelled in red as $Q_{x,ACD}$, corresponding to the natural tune.

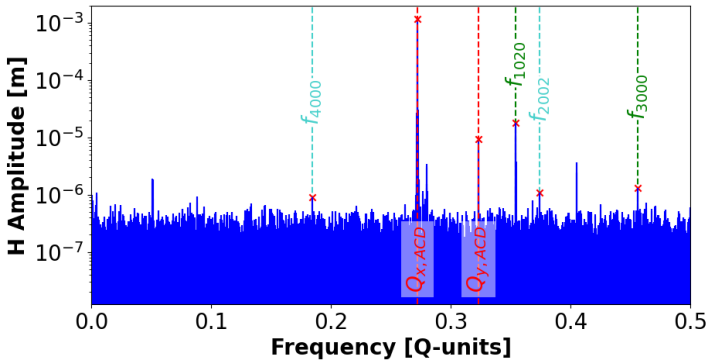


Figure 3.15: Example horizontal TbT frequency spectrum (amplitude vs frequency) where green denotes sextupolar RDTs, cyan the octupolar RDTs, and red the ACD frequencies.

From the spectral analysis of the BPMs, it is possible to determine both the phase advance between BPMs and the total phase advance around the ring. Deviations of these measured phase advances from their design values are extensively used to derive optics corrections, in a manner that is both model-independent and largely insensitive to BPM calibration.

The corresponding corrections are obtained by identifying appropriate quadrupole settings that reproduce the observed phase deviations. This can be achieved either through manual adjustment or by employing a response matrix approach [77].

The β-function can be determined either from phase measurements or from oscillation amplitudes. In the LHC, the amplitude-based approach is generally avoided due to the sensitivity of the result to BPM calibration uncertainties.

Consequently, the β-function is typically derived from measured phase advances. One possible approach is the 3-BPM method [78], which utilises the phase advances between neighbouring BPMs to evaluate the β-function at a selected BPM, according to:

$$\beta_i = \beta_i^{\mathrm{MDL}} \frac{\cot(\phi_{ij}) - \cot(\phi_{ik})}{\cot(\phi_{ij}^{\mathrm{MDL}}) - \cot(\phi_{ik}^{\mathrm{MDL}})}, \tag{3.2}$$

where i, j, and k denote three consecutive BPMs, and ϕ_{ij} represents the phase advance between BPMs i and j. The superscript MDL indicates quantities obtained from the chosen simulation model. This method is illustrated in Fig. 3.16, where i corresponds to BPM 2, while j and k correspond to BPM 1 and BPM 3, respectively.

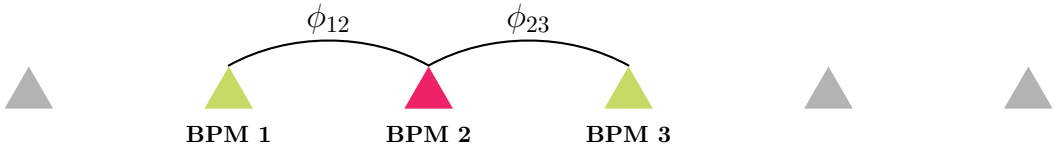


Figure 3.16: Illustration of the principle of the 3-BPM method with the chosen BPM highlighted in pink and the nearby BPMs used in the analysis highlighted in green.

A limitation of this method arises from the uncertainty in the reconstructed β-function, which diverges for phase advances approaching $n\pi/2$ rad [78], where n is an integer. This behaviour leads to a significant amplification of phase measurement errors.

For this reason, the analytical N-BPM formalism is typically employed in the LHC [79]. This approach offers several advantages. In particular, it enables the exclusion of BPMs with unfavourable phase advances from the analysis, such as the grey BPMs shown in Fig. 3.17.



Figure 3.17: Illustration of the principle of the N-BPM method.

Moreover, it provides an analytical estimate of the associated systematic uncertainties, based on the expected RMS quadrupole errors, sextupole misalignments, and longitudinal misalignments between the BPMs. Furthermore, known quadrupolar field errors obtained from WISE can be incorporated into the calculation to further improve the accuracy of the reconstruction, for example, the systematic b_2 component in the main dipoles.

Accordingly, the β-function is computed as follows:

$$\beta(s_1) = \frac{\cot(\phi_{12}) - \cot(\phi_{13})}{\cot(\phi_{12}^{\mathrm{MDL}}) - \cot(\phi_{13}^{\mathrm{MDL}}) + h_{12} - h_{13}} \beta^{\mathrm{MDL}}(s_1) + \mathcal{O}(\delta K_1^2), \tag{3.3}$$

where h_{ij} is given by

$$h_{ij} = \frac{\sum_{i<w<j} \beta_w^{\mathrm{MDL}} \delta K_{w,1} L \sin^2(\phi_{wj}^{\mathrm{MDL}})}{\sin^2(\phi_{ij}^{\mathrm{MDL}})}. \tag{3.4}$$

here $\delta K_{w,1} L$ is the integrated quadrupolar field error of element w.

Another quantity that requires careful monitoring is the transverse coupling. One method to quantify the coupling is based on the closest tune approach ($\Delta Q_{min} = |C^-|$) [80]. Alternatively, the coupling can be characterised through the corresponding resonance driving terms, f_{1001} and f_{1010}, which are determined from the relevant spectral components of the TbT beam position [54].

It should be noted that, for all optics measurements presented in this thesis, the transverse coupling was carefully corrected and maintained at a high level of control. In particular, the coupling RDT was typically kept below $|C^-| = 0.001$, compared to fractional tune separations of at least 0.015.

3.9.1 Non-linear optics analysis

An important quantity for NL analysis is the transverse action (I_z). For an ACD excitation, the action can be quantified through the relation

$$2I_z = \frac{1}{N_{\mathrm{BPM}}} \sum_{\mathrm{BPM}} \frac{A_z^2}{\beta_{z,\mathrm{MDL}}}, \tag{3.5}$$

which is based on an average over the arc BPMs. Here, $\beta_{z,\mathrm{MDL}}$ denotes the model β-function, under the assumption that the β-beating is small and averages to zero over the considered BPMs, and A_z represents the amplitude of the ACD excitation, obtained from the corresponding spectral peak in the analysis of the TbT BPM data.

The calculated action can subsequently be used in the measurement of amplitude detuning, where the variation of the natural tune is evaluated as a function of the transverse action induced by the ACD excitation. To mitigate the impact of slow tune drifts, which could otherwise bias the measurement, the BBQ tune signal is used as a reference. Consequently, the analysis is performed on the difference between the natural tune extracted from the spectral analysis of the excited beam, and that measured by the BBQ system immediately prior to the excitation.

A fit to this dependence is then carried out to extract the corresponding detuning coefficients. This procedure requires an additional correction (the inclusion of extra terms driven by the ACD) to account for the influence of the forced oscillation, as described in [81].

By analysing the spectral lines in the frequency spectrum resulting from an ACD excitation, it is possible to determine the forced-RDTs (which are defined with respect to the ACD frequency rather than the natural tune), therefore quantifying the strength with which a given resonance perturbs the beam. Since individual BPMs do not allow an unambiguous distinction between spectral components in the ranges $[0, 0.5]$ and $[-0.5, 0]$, the complex spectrum is reconstructed by combining data from pairs of BPMs separated by a phase advance of approximately 90°, as expressed by the following relation

$$z_1 - i p_{z,1} = z_1 - i \left(\frac{z_2}{\sin(2\pi\Delta\phi)} - \frac{z_1}{\tan(2\pi\Delta\phi)} \right), \tag{3.6}$$

where z_1 and z_2 denote the TbT position data measured at a given pair of BPMs, and $\Delta\phi$ represents the measured phase advance of the driven oscillation between the two BPMs. This

reconstruction can be performed either directly at the level of the TbT data or, alternatively, by combining the FFTs of the individual BPMs signals to obtain the complex FFT.

Having determined the transverse actions and reconstructed the complex spectra, the forced-RDT can be evaluated on a BPM-by-BPM basis. Since each BPM corresponds to a specific longitudinal position s along the accelerator, this approach allows the spatial evolution of the RDT around the ring to be reconstructed.

RDTs manifest as secondary peaks in the TbT spectra at characteristic frequencies given by linear combinations of the tunes (and, in the case of ACD measurements, of the ACD excitation frequencies). In particular, the forced-RDT f'_{jklm} appears at the frequencies $[(1-j+k)Q_{x,\mathrm{ACD}} + (m-l)Q_{y,\mathrm{ACD}}]$, in the horizontal spectrum, and $[(k-j)Q_{x,\mathrm{ACD}} + (1-l+m)Q_{y,\mathrm{ACD}}]$ in the vertical spectrum.

These spectral lines are identified using a procedure analogous to that employed for the determination of the fractional tunes, but with substantially tighter tolerance windows, with a minimum width of 10^{-4}. Initial estimates of the spectral-line frequencies are obtained from the expected frequency combinations, using the measured ACD fractional tunes as input. For each expected frequency, the corresponding spectral line is subsequently identified as the dominant peak within the associated tolerance window.

The amplitude of the forced-RDT can then be determined either on a kick-by-kick basis or by combining multiple ACD excitations. The corresponding expression for the spectral line amplitude of the forced-RDT is given by

$$H(1-j+k, m-l) = 2j|f'_{jklm}|(2I_x)^{\frac{j+k-1}{2}}(2I_y)^{\frac{l+m}{2}}, \tag{3.7}$$

$$V(k-j, 1-l+m) = 2l|f'_{jklm}|(2I_x)^{\frac{j+k}{2}}(2I_y)^{\frac{l+m-1}{2}}. \tag{3.8}$$

The spectrum is normalised to the amplitude of the main ACD spectral line. This normalisation reduces the sensitivity to linear BPM calibration errors, leading to the following relation:

$$H(1-j+k, m-l) = 2j|f'_{jklm}|(2I_x)^{\frac{j+k-2}{2}}(2I_y)^{\frac{l+m}{2}}, \tag{3.9}$$

$$V(k-j, 1-l+m) = 2l|f'_{jklm}|(2I_x)^{\frac{j+k}{2}}(2I_y)^{\frac{l+m-2}{2}}, \tag{3.10}$$

where the reduction in the exponents of the transverse actions, relative to Eq. (3.7), reflects the additional normalisation applied to the measured spectral-line amplitudes. This relation enables

the forced RDT, f'_{jklm}, to be determined by fitting the measured normalised line amplitudes to their expected dependence on the transverse actions using unweighted least squares. Assuming Gaussian-distributed errors, the corresponding 1σ uncertainties, as shown in the RDT amplitude plots, are estimated from the diagonal elements of the parameter covariance matrix obtained from the fitting procedure. In cases for which the fit does not yield a finite uncertainty estimate, the corresponding uncertainty is set to zero.

The corresponding forced-RDT phases, Φ'_{jklm}, are defined by,

$$\Phi'_{H,jklm} = \theta_{H(1-j+k,m-l)} - (1-k+j)\phi_{x,\mathrm{ACD}} - (m-l)\phi_{y,\mathrm{ACD}} + \frac{\pi}{2}, \tag{3.11}$$

$$\Phi'_{V,jklm} = \theta_{V(k-j,1-l+m)} - (k-j)\phi_{x,\mathrm{ACD}} - (1-l+m)\phi_{y,\mathrm{ACD}} + \frac{\pi}{2}, \tag{3.12}$$

where θ denotes the reconstructed phase of the spectral peak in the complex spectrum, and $\phi_{z,\mathrm{ACD}}$ correspond to the measured phases of the driven oscillations at the BPM. Throughout the remainder of this thesis, the prime symbol used to differentiate forced- from free-RDTs is omitted for simplicity of notation.

This procedure inherently assumes that the considered RDT is the sole contributor to its corresponding spectral line. It further relies on the assumption that the RDTs are dominated by their leading-order Hamiltonian terms. For example, Eq. (3.9) provides an appropriate description for resonance lines such as $4Q_x$ driven by octupolar fields, but does not account for second-order contributions arising from sextupoles [28, 82].

Finally, Eq. (3.9) does not explicitly account for the presence of noise in the spectrum. As a result, a manual cleaning procedure is applied, where, for example, BPMs exhibiting anomalously high noise levels are excluded from the analysis. The formulation also implicitly assumes that the line amplitude vanishes for zero action. Consequently, the analysis must be restricted to excitation conditions for which the RDT spectral lines are well above the noise floor.

For this reason, the amplitude of the ACD excitation is typically chosen to be larger than that used for linear optics measurements. The excitation amplitude is therefore increased progressively in order to avoid triggering a beam dump while ensuring sufficient signal quality.

CHAPTER 4

Sextupolar Resonance Driving Terms

Non-linear optics modelling and correction is a subject of broad relevance within the accelerator community. The optimisation of NL beam dynamics represents a major consideration during the design, construction, and operation of many past, present, and future facilities, including, for example, Hadron-Electron Ring Accelerator (HERA) [83], Tevatron [84], SPS [29, 85], Relativistic Heavy Ion Collider (RHIC) [86, 87], Heavy Ion Synchrotron (SIS-18) [28, 88], Super KEK Electron B-factory (SuperKEKB) [89], European Synchrotron Radiation Facility (ESRF) [82, 90], LHC [91, 92], High-Luminosity Large Hadron Collider (HL-LHC) [93] (see pp. 579-592), [94], and Future Circular Collider (FCC) [95, 96].

In any accelerator, non-linearities can constitute a significant operational challenge. Although they may play a beneficial role in the stabilisation of collective effects, they can also drive emittance growth, reduce beam lifetime, and limit the DA. In the case of colliders such as the LHC, these effects directly constrain the maximum achievable luminosity, making the control and understanding of NL dynamics a critical aspect of accelerator performance. Consequently, considerable effort has been devoted to improving the characterisation of NL beam dynamics in the LHC and to the development of advanced beam-based measurement techniques.

During the first two operational runs of the LHC, initial investigations concentrated mainly on observables that were relatively straightforward to access experimentally, including NL

chromaticity, amplitude detuning, tune feed-down, and dynamic aperture [1, 22, 38, 97, 98]. Although these observables are essential for model validation and for informing optimisation efforts, a central aspect of NL beam dynamics is the strength of the underlying resonances. This importance is reflected in the careful choice of the LHC WP, which must be placed sufficiently far from certain harmful resonances to preserve beam lifetime. This dependence is clearly illustrated in Fig. 4.1, where a change in tune was seen to result in a marked reduction in beam lifetime.

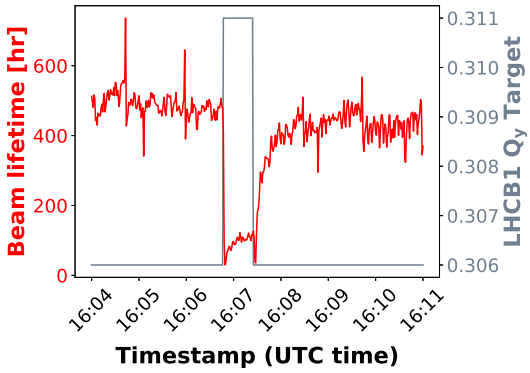


Figure 4.1: The LHCB1 BLM lifetime in red, with the LHCB1 Q_y target in grey vs time during the 2023 ion run.

Therefore, a powerful complementary approach to these methods is the direct quantification of resonance strength through the measurement of RDTs. Historically, RDT measurements have relied on single-kick beam excitations, commonly referred to as free-RDT measurements, as demonstrated, for example, at RHIC [87] and ESRF [82]. Free-RDTs, however, are limited by rapid decoherence of the oscillatory signal, as seen in Fig. 2.8, and by uncertainties in the precise spectral frequencies, which complicate their interpretation [29]. These limitations motivated the development of so-called forced-RDT measurements, in which sustained transverse oscillations are driven using a sinusoidally powered dipole kicker, the ACD, as seen in Fig. 3.13. This technique was first investigated in the SPS, where dedicated experiments demonstrated its practical feasibility [29].

These initial studies primarily aimed at validating the measurement principle, by artificially enhancing resonances, rather than probing residual non-linearities of the accelerator. Early applications in the LHC were also pursued towards the end of Run II, but were limited to top energy and focused on a restricted set of issues in the IRs [98, 99], and no studies at injection energy had been carried out following unsuccessful free-RDT attempts in 2011 [100]. This

lack of success was largely due to concerns regarding the practicality of the measurements, notably potential non-linearities in the BPMs, the shortness of the ACD excitation plateau (only 1000 turns), and the stability of the resonances in the presence of the strongly dynamic conditions characteristic of injection, such as the b_3 decay as discussed in the following section.

This provided strong motivation to investigate whether RDT measurements could be successfully performed at injection using the newer measurement techniques, and whether such measurements could be used to validate the accelerator model. This study addresses these questions and further explores the possibility of correcting a resonance that had previously been considered negligible prior to the present analysis.

4.1 Identification of a significant b_3 resonance at LHC injection

During the 2023 commissioning period, large-amplitude transverse excitations were intentionally applied using the ACD to evaluate the feasibility of measuring a range of RDTs at injection energy. Some of these measurements were carried out at the natural tunes of Q_x = 62.28, Q_y = 60.31 (referred to as the nominal tunes), using the technique described in Sec. 3.9, with the excitation amplitude increased progressively in order to avoid a beam dump.

Figure 4.2 shows the frequency spectrum obtained from one of the earliest beam excitations performed with a small kick amplitude, a regime in which only linear optics would normally be studied. Despite this, a clear spectral line associated with the f_{1020} RDT, corresponding to the normal sextupolar $Q_x + 2Q_y$ resonance, is observed and, notably, it is the only NL resonance line present in the spectrum.

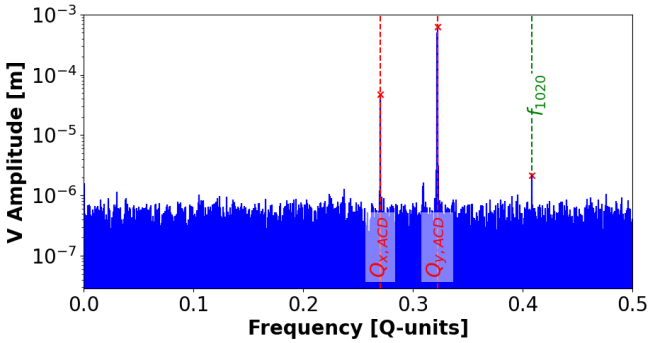


Figure 4.2: Vertical TbT frequency spectrum, where green denotes third-order RDTs, and red the ACD frequencies with natural tunes $q_x = 0.28$ and $q_y = 0.31$ for LHCB1.

This behaviour is unexpected, as the resonances traditionally assumed to be most critical at injection were the skew sextupolar $3Q_y$ resonance, and those driven by the Landau octupoles (MO) since previous studies had shown them to lead to a loss in DA [101]. The prominence of the f_{1020} line is particularly surprising given that, for this accelerator configuration, the WP lies far from the $Q_x + 2Q_y$ resonance, as illustrated in Fig. 4.3.

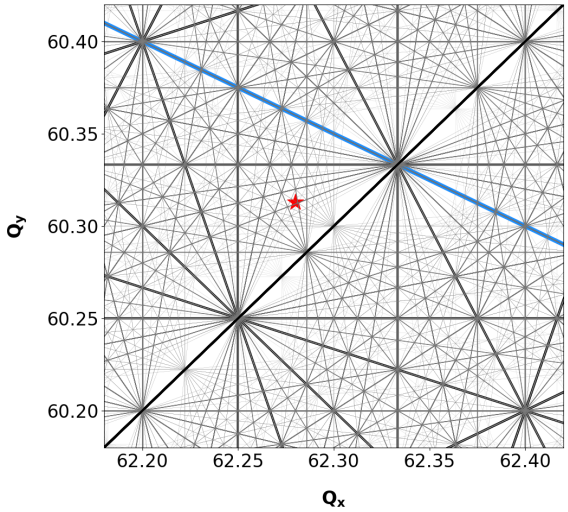


Figure 4.3: Tune diagram (Q_y vs Q_x) showing the $2Q_y + Q_x$ resonance line in blue, and the nominal WP as a red star.

It was further observed that, even when operating with an alternative WP (Q_x = 62.275, Q_y = 60.293), commonly referred to as the electron-cloud or “e-cloud" WP, a strong f_{1020} resonance line remained visible, as shown in Fig. 4.4. This WP was selected to mitigate beam-lifetime limitations associated with electron-cloud effects by avoiding the NL resonances responsible for these issues. As a result, the tunes are even further from the $Q_x + 2Q_y$ resonance.

Despite this increased separation, the spectral line corresponding to the f_{1020} RDT remained nearly as prominent as the f_{2002} line, now present in the spectrum, associated with an octupolar resonance. This is particularly striking because, in this configuration, the accelerator tunes lie much closer to the octupolar resonance conditions. Also, the Landau octupoles were powered with K_4 = -5 m^{-4}, settings that would normally be expected to significantly enhance octupolar resonance effects (in the case of the nominal accelerator optics at injection corresponding to an increase in the RDT amplitude by approximately a factor of ten.). In addition, the amplitude of the f_{1020} RDT was found to be considerably larger than that of the f_{0030} RDT, which corresponds to the skew sextupolar $3Q_y$ resonance discussed previously.

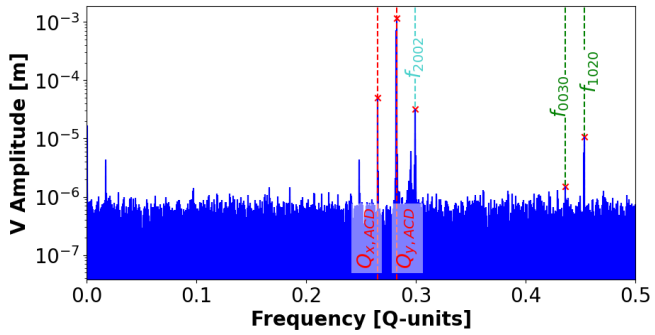


Figure 4.4: Vertical TbT frequency spectrum, where green denotes third-order and cyan the fourth-order RDTs, and red the ACD frequencies with natural tunes $q_x = 0.275$ and $q_y = 0.293$ for LHCB1. The unlabelled line at 0.02 corresponds to another sextupolar RDT.

Given the unexpectedly large amplitude of the sextupolar resonance lines, it was initially suspected that the observation might be an artefact arising from BPM non-linearities. This possibility was investigated using measurements taken during the 2023 commissioning period at different accelerator tune settings: the so-called “e-cloud" WP, and one specifically for sextupolar resonance studies with $Q_x = 62.28$ and $Q_y = 60.313$, hereafter referred to as the “sextupolar" resonance measurement tunes. Spurious contributions originating from BPM non-linearities are not expected to display a dependence on the WP. In contrast, a clear variation in the measured RDT was observed when the tunes were changed. As shown in Fig. 4.5, the amplitude of the f_{1020} RDT as a function of s position around the ring has an overall decrease by approximately a factor of two in magnitude when the WP is shifted further away from the corresponding resonance, consistent with expectations (see Eq. (2.80)).

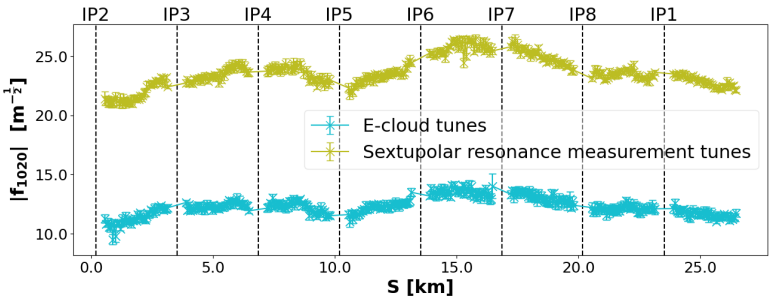


Figure 4.5: Measured amplitude of f_{1020} vs s for LHCB1 at e-cloud (blue) and sextupolar resonance measurement (green) tunes for LHCB1.

It was then conjectured that the unexpected strength of the resonance could also be associated with the introduction of the new injection optics implemented in 2023 [102]. To examine this possibility, dedicated measurements were carried out during the 2023 commissioning period.

The objective was to determine whether the updated optics configuration led to any significant change in any of the measured RDTs. These were again completed at the “sextupolar" resonance measurement tunes. As shown in Fig. 4.6, the comparison between the new optics introduced in 2023 and the previous optics configuration revealed no significant difference, with any variations remaining within the limits of measurement reproducibility. This is expected because the level of linear optics correction is not identical from year to year.

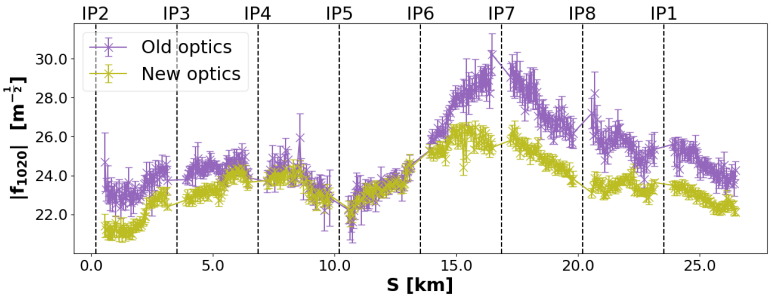


Figure 4.6: Measured amplitude of f_{1020} vs s for LHCB1 at sextupolar resonance measurement tunes with new (green) and old (purple) optics for LHCB1.

Having established that the new optics deployed in 2023 was not enhancing the resonance strength, historical data was revisited in order to assess consistency across different years with an identical accelerator configuration. When measurements from 2022 were re-analysed, the f_{1020} RDT was found to be clearly present, with an amplitude largely consistent with that observed in 2023, within the associated uncertainties, as illustrated in Fig. 4.7.

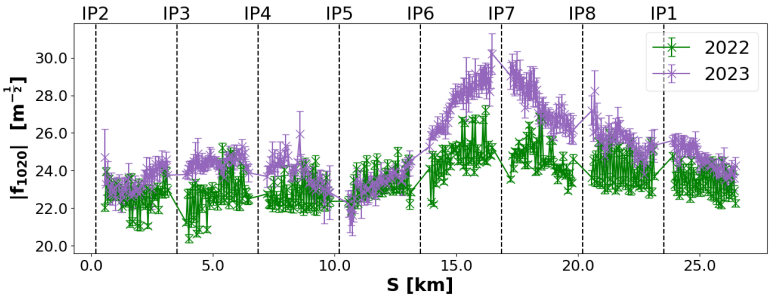


Figure 4.7: Measured amplitude of f_{1020} vs s for LHCB1 at sextupolar resonance measurement tunes for 2022 (dark green) and 2023 (purple) for LHCB1.

Given that this RDT was both real and consistently observed, and that the resonance had not been expected to play a significant role, its prominence provided clear motivation for further investigation. The subsequent step was therefore to examine the effect in simulation, with the objective of determining whether it could be reproduced by the magnetic model and, if so, using

the model to identify its underlying physical origin. Provided this could be established, the study would naturally extend to exploring the feasibility of optimising the corresponding RDT.

4.2 Benchmarking

As mentioned previously in Subsection 3.3.1, the LHC dipole magnets are known to exhibit a significant b_3 field error component, for which dedicated correctors (MCS) are installed at the end of each main dipole magnet. Following the observation of this strong sextupolar resonance, it was suspected that the large amplitude of the f_{1020} RDT might have resulted from an error in this applied MCS correction scheme. To investigate this, simulations were carried out in MAD-X [14], incorporating magnetic field errors based on FiDeL [43] as well as the actual MCS settings used in the accelerator.

These simulations, carried out at injection energy, exhibited good agreement with the measured RDT amplitudes, as shown in Fig. 4.8. Such a close correspondence between model and measurement, to within approximately 15%, is relatively rare; for comparison, benchmarking of skew-sextupole effects shows discrepancies of up to a factor of two, even when using magnetic models that include measured field errors [103]. This level of agreement therefore indicates that the strength of the resonance is unlikely to originate from an issue with the MCS correction scheme. Instead, it is a highly encouraging result, as it suggests that the magnetic model is sufficiently accurate to identify the underlying source of the effect, and to potentially guide an effective correction strategy. It was also noted that the amplitude of this RDT was consistently larger in LHCB1 than in LHCB2, both in measurements and in simulations. For this reason, the subsequent analysis focuses exclusively on LHCB1.

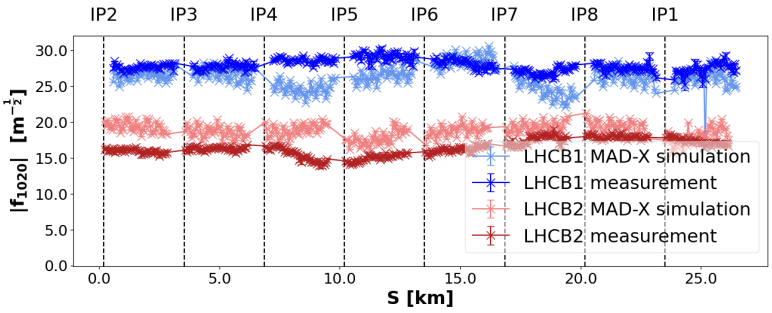


Figure 4.8: Comparison of amplitude of f_{1020} vs s, between measurement (dark blue for LHCB1, dark red for LHCB2) and MAD-X simulation (light blue for LHCB1, light red for LHCB2).

To completely rule out the possibility that the observed behaviour was driven by the MCS correction, the simulations were repeated with all magnetic errors and MCS corrections removed.

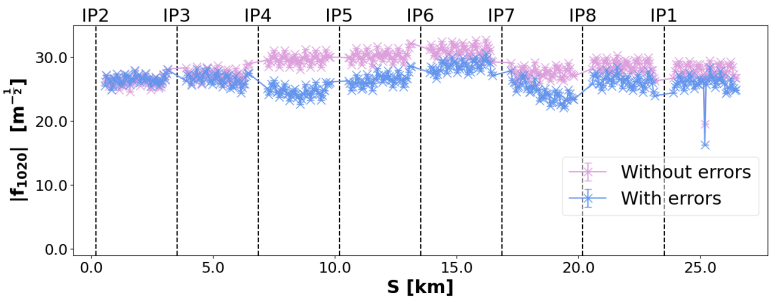


Figure 4.9: f_{1020} amplitude vs s, comparing with (blue) and without (pink) any error models for LHCB1 in MAD-X simulation.

Figure 4.9 shows that eliminating all magnetic errors from the simulation (along with the corresponding MCS correction) results in only a small increase in the RDT amplitude. This leaves the nominal powering of the main lattice sextupoles (MS), used for chromaticity correction, as the remaining dominant candidate.

To verify unambiguously that the MS circuits were responsible, a dedicated simulation study was carried out to exclude the possibility that the observed effect arose from higher-order perturbations or unintended tracking artefacts. This required evaluating the RDT with the MS completely switched off. Rather than attempting this study using tracking-based RDT methods, which could be influenced by the extremely large chromatic effects present when the MS are disabled, the newly developed MAD-NG code [20] computes RDTs directly from the normal form using differential techniques, making it particularly well suited to this type of investigation.

Before carrying out this analysis, it was essential to confirm that the RDT magnitudes obtained with MAD-NG were consistent with those derived from MAD-X, therefore ensuring that methodological differences, and in particular the absence of an ACD element in the MAD-NG normal-form calculation, did not lead to significant discrepancies. As a first validation step, the impact of explicitly modelling the ACD in MAD-X was examined. Figure 4.10 shows that including the ACD produces a modest reduction in the calculated RDT amplitude. This behaviour is expected, since forced-RDTs obtained with an ACD excitation are intrinsically different from free-RDTs generated by a single kick [104].

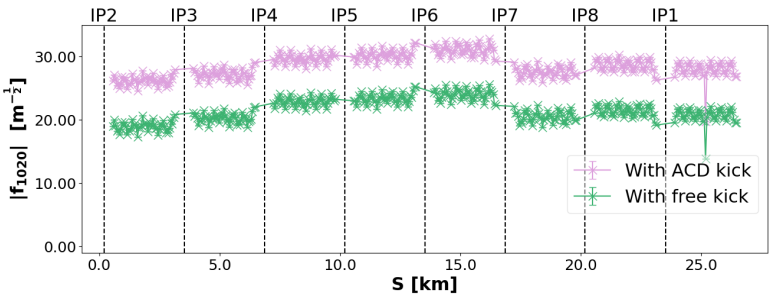


Figure 4.10: f_{1020} amplitude vs s, comparing an ACD-driven (pink) and a free-kick MAD-X simulation (purple) for LHCB1.

The free-kick MAD-X simulation was subsequently adopted as a direct benchmark against which the MAD-NG implementation was validated. As shown in Fig. 4.11, the RDTs computed using the two codes exhibit excellent agreement, with an average relative difference of approximately 4%. This agreement confirmed the suitability of MAD-NG for the present study, with the additional advantage of providing a significantly more efficient computational framework.

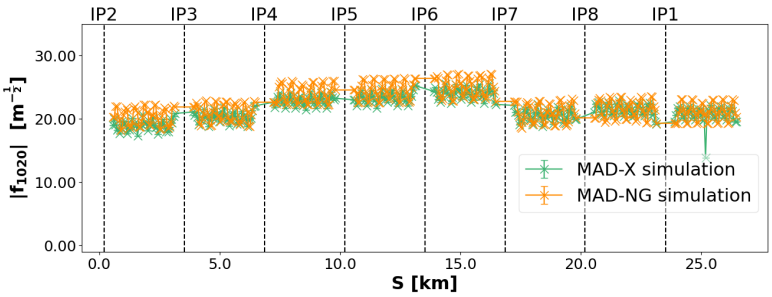


Figure 4.11: f_{1020} amplitude vs s, comparing a MAD-NG (light green) and a MAD-X free-kick simulation (purple) for LHCB1.

As shown in Fig. 4.12, the RDT amplitude was negligible when the MS magnets were switched off. In contrast, when the MS circuits were powered at their nominal operational settings, the simulated amplitude of f_{1020} was found to be in close agreement with the measured values.

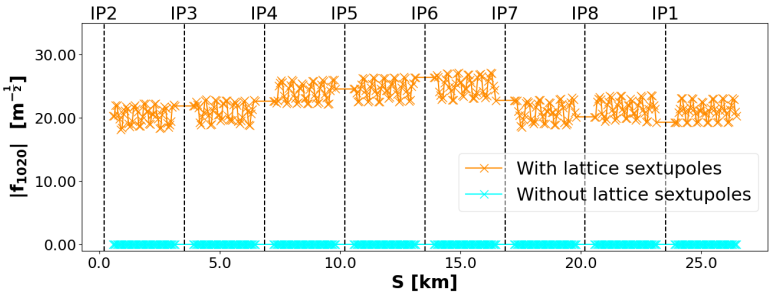


Figure 4.12: f_{1020} amplitude vs s, comparing with (light green) and without (orange) the MS powered for LHCB1 in MAD-NG simulation.

This result demonstrated that the dominant contribution to the observed f_{1020} RDT originated from the nominal powering of the MS circuits. The pronounced magnitude of this effect subsequently motivated the development of a dedicated correction strategy.

4.3 Correction at Injection

Most NL optics corrections implemented in the LHC to date have relied on beam-based methodologies. Examples include the correction of amplitude detuning [105] and NL chromaticity [106]. This approach is necessary due to the poor agreement typically observed between beam-based measurements and magnetic field models. For instance, beam-based measurements of the b_4 multipole component at injection energy have historically differed from magnetic model predictions by up to an order of magnitude [12].

In contrast, the present case exhibits an unusually high level of agreement between MAD-X simulations and beam-based measurements, as demonstrated in Fig. 4.8. This provided a rare opportunity to pursue model-based optimisation strategies with a reasonable degree of confidence.

Two strategies were considered for a model-based correction of the f_{1020} RDT. The first strategy consisted of optimising the linear optics, in particular the betatron phase advances, in order to minimise the driving term. However, a phase-based correction had already been implemented in a previous optics configuration to suppress resonances generated by the Landau octupoles [102].

An alternative approach employed the available normal sextupole magnets in the lattice to compensate the resonance. In order to avoid perturbing the previously optimised betatron phase, this second strategy was adopted. The correction was therefore implemented by using the MS circuits, identified as the primary source of the resonance, as dedicated correctors.

The simulation code MAD-NG was chosen for this study because it provides an analytical calculation of RDTs based on normal form theory, without the need for particle tracking. This results in a substantial gain in computational efficiency, as discussed previously. Although this approach yields free-RDTs rather than forced RDTs, the resulting magnitudes were found to be comparable to those obtained from forced-excitation measurements. In addition, MAD-NG is

currently the only available code that supports NL matching through Jacobian-based optimisation algorithms. These algorithms enable systematic adjustment of model parameters, such as magnet strengths, to match target values of both linear and NL observables, including the real and imaginary components of a given RDT, in direct analogy with standard linear optics matching procedures used for quantities such as tune and chromaticity.

In this study, MAD-NG's RDT matching functionality was employed to determine an optimised configuration of the MS circuits that would minimise the f_{1020} RDT in LHCB1. The optimisation matched both the real and imaginary parts of f_{1020} at IP1 by varying the MS circuit strengths. Simultaneously, the chromaticity was constrained to remain at its nominal value, ensuring that the correction scheme did not adversely affect other operational parameters.

The first optimisation strategy sought to drive the amplitudes of all normal sextupolar RDTs to zero. By applying the Jacobian-based matching algorithm available in MAD-NG, a solution was identified rapidly and appeared, at first glance, to be effective (see Fig. 4.13).

However, this solution introduced excessively large chromatic β-beating, with peak values reaching nearly 60% (see Fig. 4.14), making it unacceptable for operational deployment. It should be noted that, for both the simulation and measurement studies, the chromatic β-beating was evaluated at a relative momentum deviation typically used for LHC off-momentum loss-maps [107], namely $\Delta p/p_0 = 1.43 \times 10^{-3}$.

Since the other normal sextupolar resonances appeared to be only weakly affected by the correction, a subsequent strategy was to optimise exclusively the f_{1020} RDT, in order to assess whether this would significantly perturb the other sextupolar terms. As shown in Fig. 4.15, this selective optimisation did indeed reduce the target resonance, while only slightly degrading the other sextupole resonances, making it the chosen strategy moving forward. In addition, the associated chromatic β-beating was lowered to a peak value of approximately 30% (see Fig. 4.16). Although this represented a clear improvement, the level of chromatic β-beating remained above the target optics quality of around 20%, a threshold relevant for maintaining acceptable conditions in loss map studies, indicating that further refinement of the correction strategy was required.

Subsequent analysis revealed that the excessive chromatic β-beating was driven by a limited

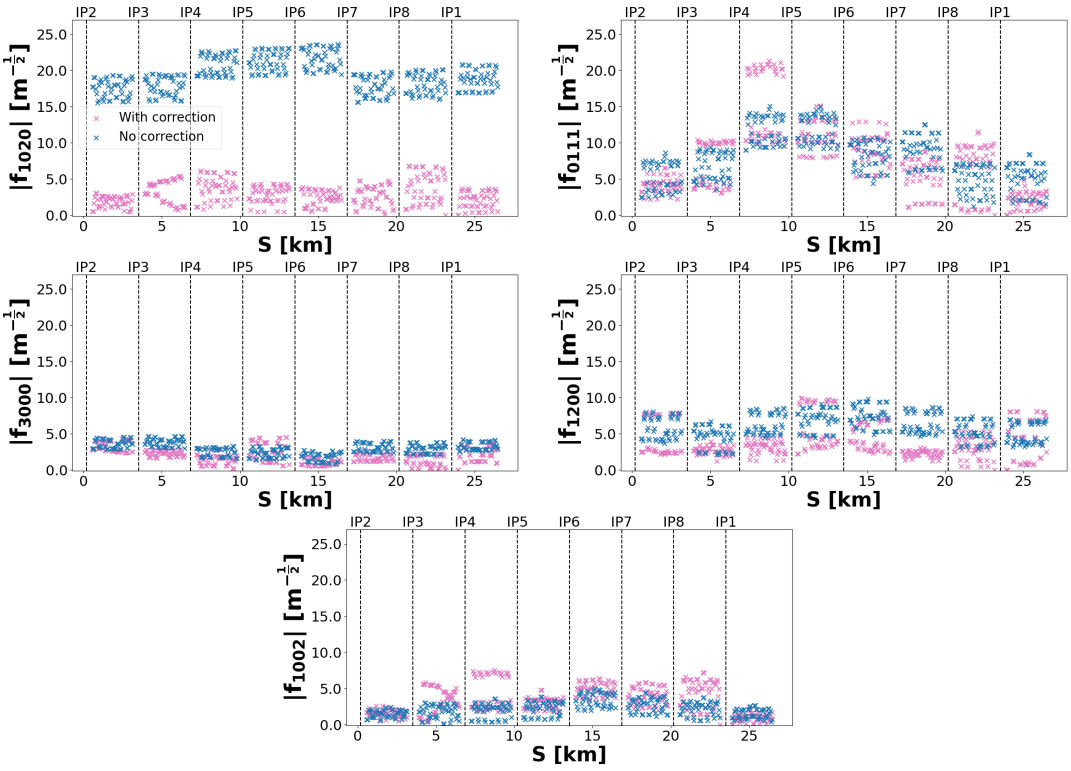


Figure 4.13: f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (pink) and without (blue) correction for all normal sextupolar RDTs with no constraints on powerings for LHCB1 in MAD-NG simulation.

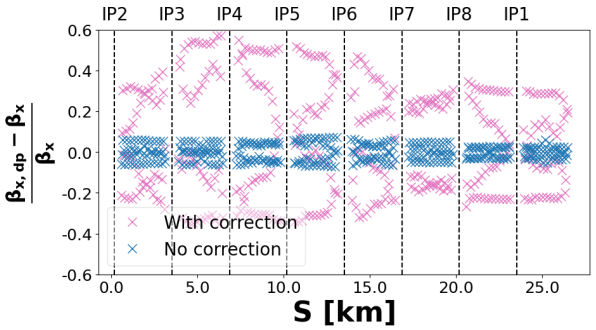


Figure 4.14: Simulated horizontal chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (pink) and without (blue) correction.

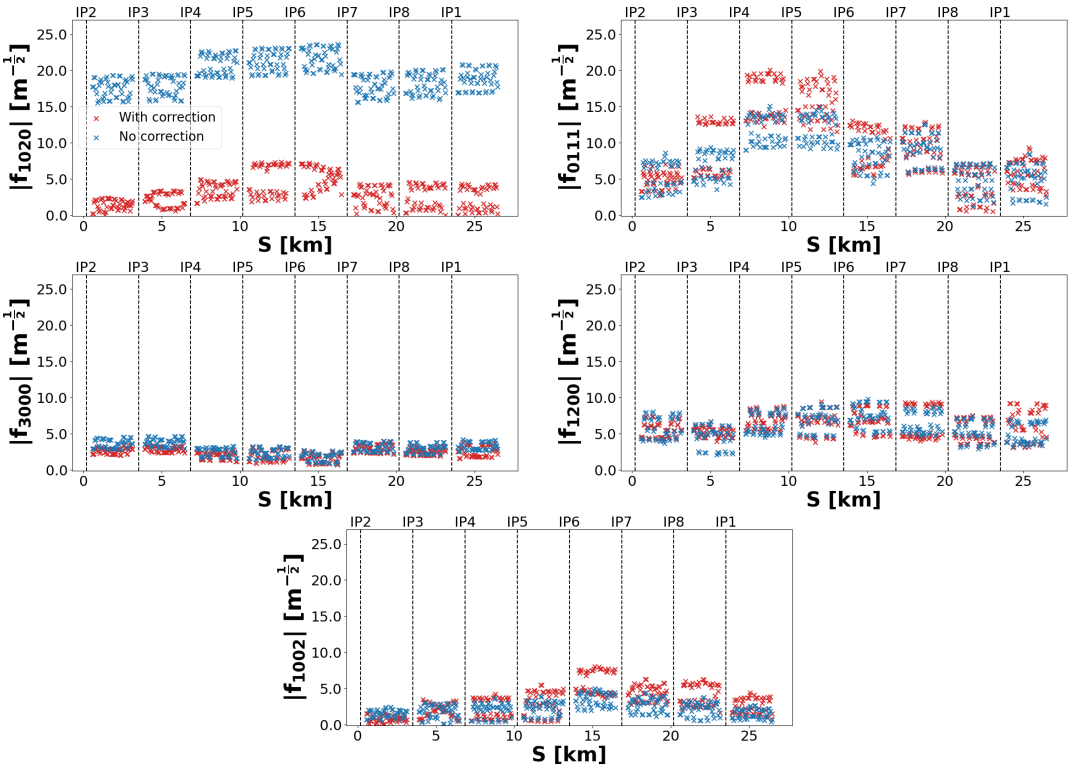


Figure 4.15: f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (red) and without (blue) correction for only f_{1020} RDT with no constraints on powerings for LHCB1 in MAD-NG simulation.

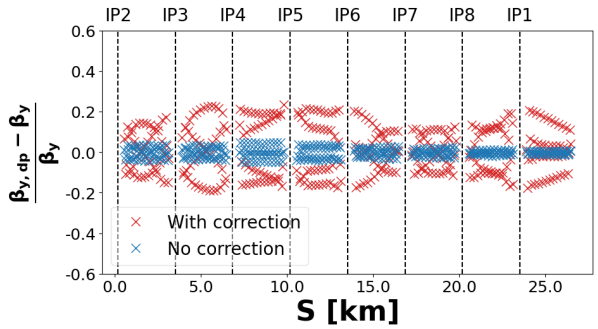


Figure 4.16: Simulated vertical chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (red) and without (blue) correction.

number of MS circuits that were pushed to relatively large strengths while contributing only weakly to the overall suppression of the target RDT. The origin of this effect was identified as an imbalance between the KSD1 and KSD2 (the “D" denoting the sextupoles close to the defocusing quadrupoles) families (and their focusing counterparts). These magnet families are separated by approximately 90° of betatron phase and, when powered symmetrically, their chromatic contributions largely cancel, resulting in minimal chromatic β-beating. Effective resonance correction, however, requires breaking this symmetry by introducing differential powering between the two families 1 and 2, which inevitably generates chromatic β-beating. This effect was further amplified by the fact that the resonance was significantly more sensitive to the KSD families than to the KSF (the “F" denoting the sextupoles close to the focusing quadrupoles) families, leading the optimisation to rely predominantly on KSD trims, and therefore increasing the imbalance.

To mitigate this effect, explicit constraints were imposed on the allowable variation of the MS circuit strengths. The correction trims were limited to lie between 10% and 200% of the nominal MS powering at injection. These bounds ensured that no circuit was driven beyond acceptable operational limits, that the current polarity remained unchanged, and that all currents stayed above the minimum levels required for stable regulation, in accordance with LHC operational constraints.

With these restrictions in place, however, the matching algorithm failed to converge, indicating that the constraint set was overly restrictive. To address this, the RDT matching targets were relaxed away from zero, providing additional freedom in the optimisation and enabling the possibility of solutions that satisfied the requirement of keeping the chromatic β-beating below the aforementioned threshold of 20%.

By combining the constrained Jacobian-based matching procedure with targeted manual refinement, a satisfactory compromise solution was obtained. This approach achieved a substantial reduction in the amplitude of the f_{1020} RDT, while maintaining chromatic β-beating within acceptable limits and preserving overall operational feasibility. The resulting MS trim settings, expressed relative to the nominal configuration, are presented in Fig. 4.17. The corresponding simulated reduction in the f_{1020} RDT amplitude is shown in Fig. 4.18, together with the predicted

negligible effect on the remaining normal sextupolar RDTs.

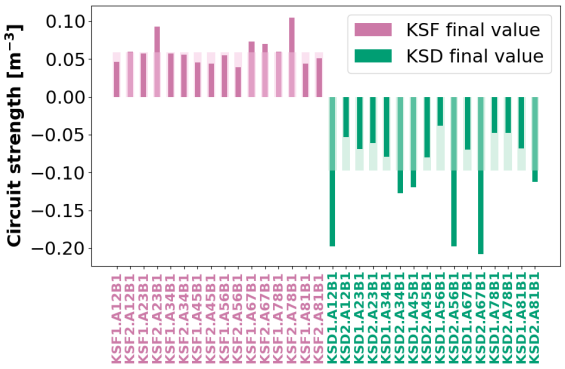


Figure 4.17: Final MS circuit powerings, KSF in pink and KSD in green, used for the sextupolar correction, with their initial powerings shown in a paler version of the same colour.

Figures 4.18(a) and 4.18(b) depict the predicted horizontal and vertical chromatic β-beating resulting from the application of the correction. These results demonstrate that the induced β-beating remains within acceptable operational limits, confirming the suitability of the correction for use in the accelerator. Consequently, the next step was to validate the correction through beam-based measurements.

The proposed correction scheme was experimentally evaluated during the 2024 LHC commissioning period. The measurements showed a reduction of approximately 50% in the magnitude of the f_{1020} resonance driving term, as illustrated in Fig. 4.19. This result agrees with the MAD-NG prediction to within approximately 2%, therefore validating the effectiveness of the model-based correction approach and demonstrating that an improvement of the free RDT directly translates into a corresponding improvement of the forced-RDT.

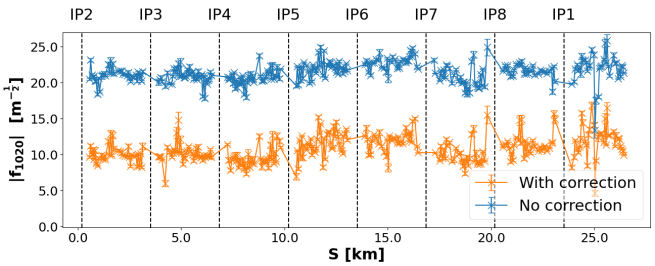


Figure 4.19: Measured amplitude of f_{1020} vs s for LHCB1, with (orange) and without (blue) correction.

As shown in Figs. 4.20(a) and 4.20(b), both the horizontal and vertical chromatic β-beat

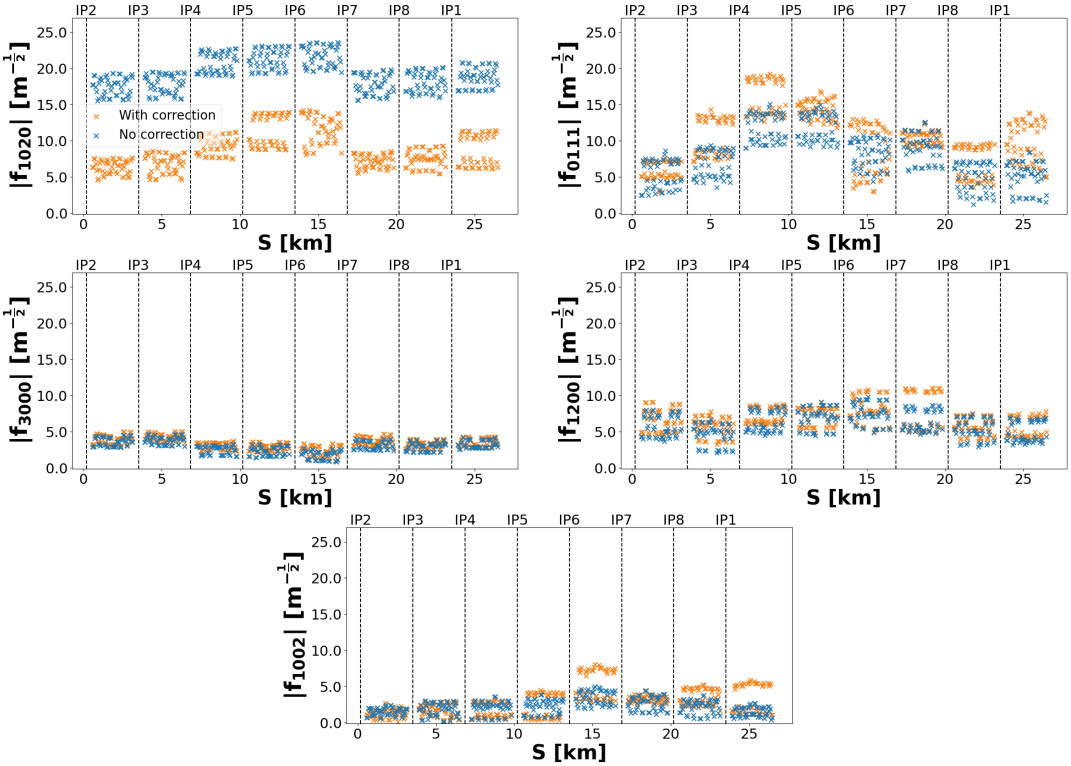
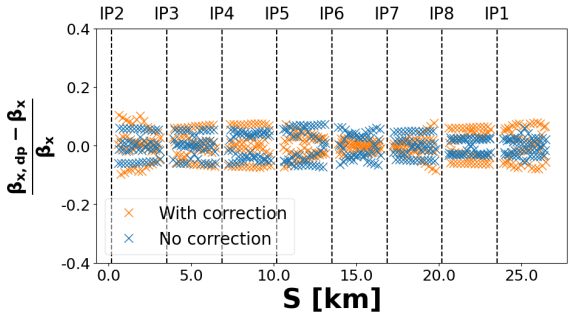
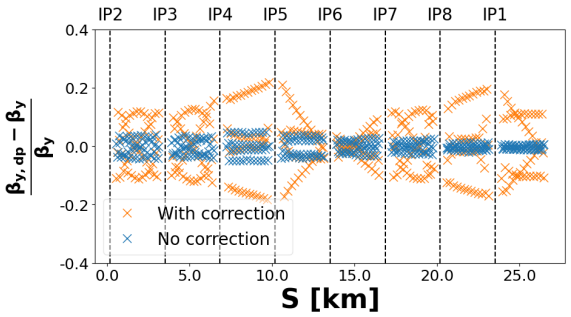


Figure 4.18: f_{1020}, f_{1011}, f_{3000}, f_{1200}, and f_{1002} amplitude vs s, comparing with (orange) and without (blue) final correction for LHCB1 in MAD-NG simulation.

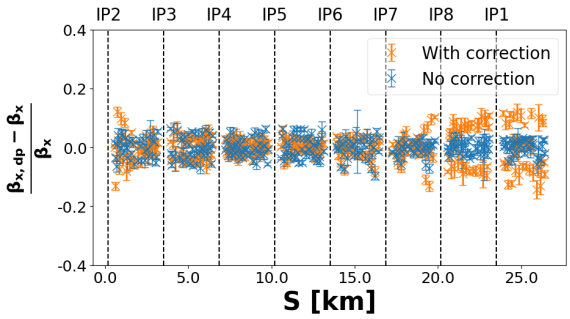


4.18(a) Simulated horizontal chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (orange) and without (blue) correction.

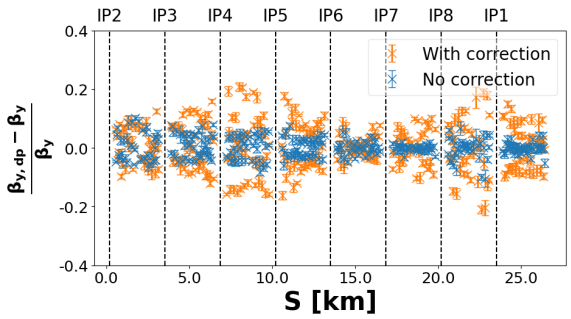


4.18(b) Simulated vertical chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (orange) and without (blue) correction.

remained below 22% using the same value of $\Delta p/p$ as simulation, with and without the f_{1020} correction applied. The observed increase in chromatic β-beating following the correction was consistent with predictions from the MAD-NG model (see Figs. 4.18(a) and 4.18(b)) and was considered acceptable for routine operation.



4.20(a) Measured horizontal chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (orange) and without (blue) correction.



4.20(b) Measured vertical chromatic β-beating vs s generated by $\Delta p/p_0 = 1.43\times10^{-3}$ with (orange) and without (blue) correction.

Tune and chromaticity measurements showed no appreciable changes following the application of the f_{1020} correction, remaining stable within the reproducibility limits of the measurements.

As mentioned previously, sextupoles can generate octupole-like effects through second-order contributions. A simple example of such behaviour is provided by two sextupole magnets separated by a drift space. As introduced previously, the Hamiltonian describing the normal component of a sextupole is given by

$$F_s = \frac{1}{3!}K_3(x^3 - 3xy^2). \tag{4.1}$$

The Hamiltonian corresponding to a drift of length L is

$$F_d = \frac{1}{2}(p_x^2 + p_y^2). \tag{4.2}$$

After the drift,

$$X = x + Lp_x, \quad Y = y + Lp_y \tag{4.3}$$

If the first sextupole has coordinates x, y, then the second sextupole would have coordinates X, Y. The resulting transfer map for the system can therefore be expressed as

$$M = e^{:F_s(X):} e^{:F_s(x):}. \tag{4.4}$$

Expanding the map using the Lie operator expansion given in Eq. (2.51) yields a series of higher-order terms, one of which is

$$\frac{1}{2}\{F_s(X), F_s(x)\}. \tag{4.5}$$

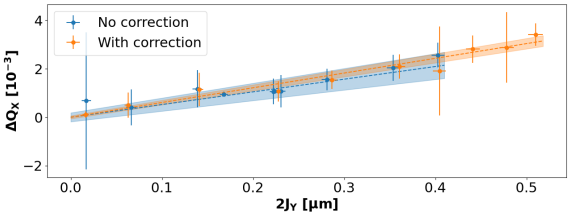
Evaluating the Poisson bracket gives

$$\{F_s(X), F_s(x)\} = \frac{K_A K_B L}{2}\left[\frac{1}{4}(X^2 - Y^2)(x^2 - y^2) + XYxy\right], \tag{4.6}$$

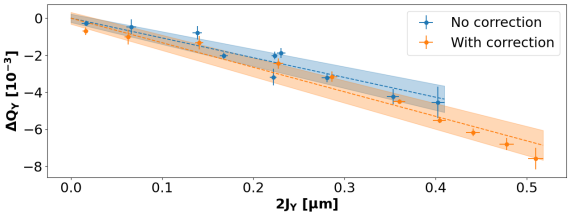
where K_A and K_B are the strengths of the two sextupoles.

The resulting quartic dependence on the transverse coordinates has the same functional form as the Hamiltonian of an octupole magnet. This demonstrates that the combined action of two sextupoles separated by a drift generates an effective octupolar term, illustrating how lower-order magnetic fields can produce higher-order non-linearities through their interaction.

It was therefore essential to verify that the modification of the lattice sextupoles did not degrade the octupolar behaviour of the accelerator. To this end, Figs. 4.21(a) and 4.21(b) compare the amplitude detuning before and after the correction. Within the experimental uncertainties, no measurable variation in detuning was observed, confirming that the correction does not have a detrimental impact on the LHC detuning characteristics.



4.21(a) Measured change in Q_x as a function of $2J_y$, thus showing the amplitude detuning, given by the fit lines, with (orange) and without (blue) the LHCB1 f_{1020} correction at injection energy. The shaded regions indicate the error on the fit.



4.21(b) Measured change in Q_y as a function of $2J_y$, thus showing the amplitude detuning, given by the fit lines, with (orange) and without (blue) the LHCB1 f_{1020} correction at injection energy. The shaded regions indicate the error on the fit.

Notably, as shown in Figs. 4.21(a) and 4.21(b), the application of the correction allowed for higher maximum ACD excitation amplitudes. In the LHC, the achievable kick amplitude is constrained by the beam loss thresholds of the BLMs distributed around the ring, as mentioned previously. After the correction was applied, beam losses caused by each ACD kick were reduced, and these thresholds were reached only at larger excitation amplitudes, indicating an improvement in the forced DA. This conclusion is reinforced by Fig. 4.22, which displays the surviving beam intensity as a function of kick amplitude. For comparable excitation levels, the corrected optics systematically exhibited reduced beam losses relative to the uncorrected configuration.

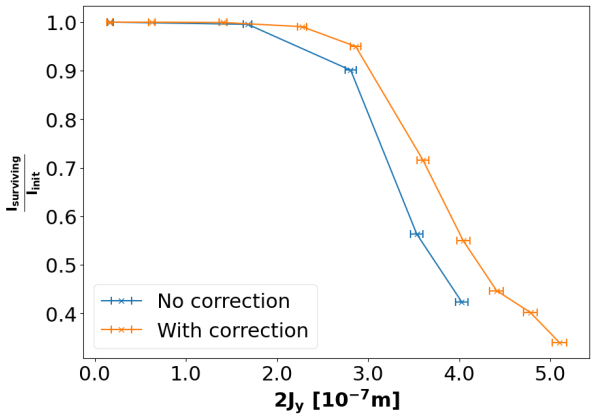


Figure 4.22: Fraction of surviving beam intensity vs action for vertical kicks with (orange) and without (blue) f_{1020} correction.

In addition to the observed improvements in beam dynamics, the application of the correction was found to produce a measurable improvement in beam lifetime under operational conditions using a pilot bunch with the Landau octupoles powered at 40 A. As shown in Fig. 4.23, the initial beam lifetime was measured to be 12.2 ± 0.3 hours. Disabling the correction resulted in a reduction of the beam lifetime to 11.1 ± 0.4 hours. Reversing the correction by applying the MS trim currents with opposite polarity produced a further decrease in beam lifetime to 9.7 ± 0.6 hours. These results provide additional experimental evidence that the f_{1020} correction has a beneficial operational impact.

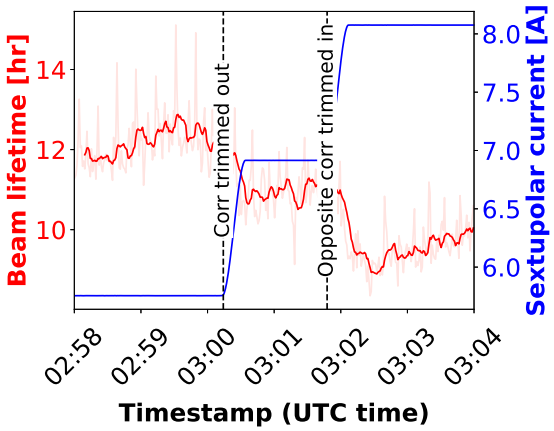


Figure 4.23: The effect of f_{1020} correction on lifetime vs time with raw signal in pale red (left-hand axis), and the corresponding sextupolar current in blue (right-hand axis).

Following consultation with the LHC Operations (OP) group, the correction strategy developed to suppress the f_{1020} resonance was deployed in routine LHC operation at injection throughout 2024, with the correction being trimmed out early in the ramp. Its successful implementation provides strong evidence for the effectiveness of this model-based compensation approach, and not only validates the predictive capability of the MAD-NG simulation framework, despite being based on free- rather than forced-RDTs, but also underlines the operational significance of this previously overlooked resonance, located on the opposite side of the $3Q_y$. By reducing the dominant perturbation observed in the TbT data, namely the f_{1020} RDT, the adverse impact on DA is mitigated, resulting in an improvement in beam lifetime.

4.4 Dynamic RDT measurements

Given the pronounced impact of the f_{1020} resonance at injection, it was of particular interest to study its behaviour throughout the full LHC cycle. Of specific relevance are the dynamically evolving phases of the cycle, namely the decay occurring on the injection plateau, and the energy ramp to top energy. NL optics measurements during these periods had not previously been performed, and therefore constituted the focus of the following investigation.

Conventionally, optics measurements are performed under well-defined and stable accelerator conditions, such as after several hours at injection, during squeeze plateaus, or at other phases of constant beam energy. Under these circumstances, there is sufficient time to correct essential

accelerator parameters, including tune, coupling, and chromaticity, before gradually increasing the excitation amplitudes to obtain reliable RDT measurements while allowing beam losses to be monitored carefully to avoid triggering a dump on the BLM thresholds. In contrast, during periods of dynamic accelerator evolution, the time available for such preparatory corrections is severely constrained. Moreover, beam excitations must be applied rapidly in order to achieve the temporal resolution required to capture the evolving effects. As a result, the excitation amplitudes must be kept relatively small to minimise the risk of beam dumps, particularly during the energy ramp, where a beam dump would imply a loss of around six hours of accelerator time.

To assess the robustness of the RDT measurement at small excitation amplitudes, a dedicated MD session was performed. As illustrated in Fig. 4.24, the amplitude of the f_{1020} RDT was found to remain constant over a wide range of excitation strengths.

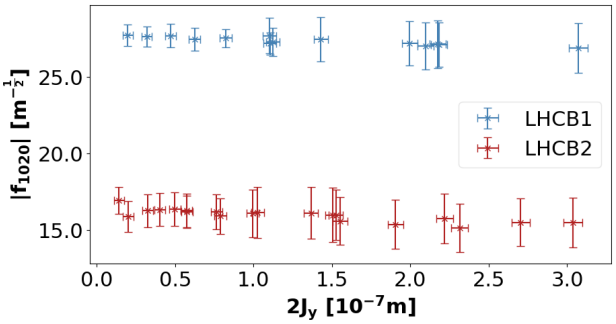


Figure 4.24: Amplitude of f_{1020} RDT vs action for LHCB1 (blue) and LHCB2 (red).

This observed independence from action validates the reliability of the RDT measurement technique, demonstrating that the f_{1020} resonance can be accurately resolved even for relatively small excitation amplitudes. These results provided the basis for a series of subsequent studies targeting the different previously unexplored dynamic phases of the LHC operational cycle.

4.4.1 Decay RDT measurements

At injection energy in the LHC, the main dipole magnets exhibit a significant decay of the b_3 field component as explained in Subsection 3.3.1.

In future accelerators such as FCC-hh, the magnitude of the b_3 decay is expected to be substantially larger [108], and large b_3 errors have already been shown to adversely affect

the DA even when chromaticity is well corrected [109]. A detailed understanding and direct measurement of resonances induced by the decay of the main dipole fields is therefore of clear interest for accelerator benchmarking. Indeed, recent studies suggested that decay effects in the LHC main dipoles could account for large discrepancies observed in the third-order chromaticity, Q''', at injection [110]. This provides strong motivation to investigate the feasibility of directly measuring the decay of NL effects through the evolution of the corresponding RDTs.

Furthermore, since RDT measurements and the associated corrections are generally carried out a considerable time after the injection plateau (defined to be the time at which a beam energy of 450 GeV is reached, or when the main dipoles reach their nominal powering at injection) is reached, it is necessary to evaluate the influence of this time-dependent decay on their validity. The evolving b_3 component is normally compensated using the dedicated MCS, operated with time-dependent settings derived from magnetic field models [111], and complemented by adjustments based on global chromaticity measurements. Assessing the reliability of this compensation therefore requires validation of the underlying models. The decay has historically been characterised through measurements of the global chromaticity, Q'. It is also interesting to evaluate the effectiveness of the compensation through its impact on resonance excitation, which can be obtained by beam-based measurements of the relevant RDTs. Consequently, during Run III several attempts were made to demonstrate measurements of this RDT decay.

A special measurement procedure was prepared with the OP team. This required that the accelerator be prepared for optics measurements prior to reaching the injection plateau, including, for example, the configuration of hardware interlocks to permit ACD excitations and the adjustment of the collimator settings for optics studies. In addition, the upstream injector SPS had to be fully prepared to deliver beam on demand. Within the first few minutes following beam injection, rapid corrections of tune, chromaticity, and coupling were computed by the optics team, together with adjustments made by the OP team to the injection-line steering based on the initial beam. This beam was then dumped and replaced with a fresh injection for the actual measurements. Moderate-amplitude excitation kicks were subsequently applied as rapidly as possible in order to maximise the temporal resolution of the decay effect while remaining safely below beam-loss thresholds to avoid a beam dump.

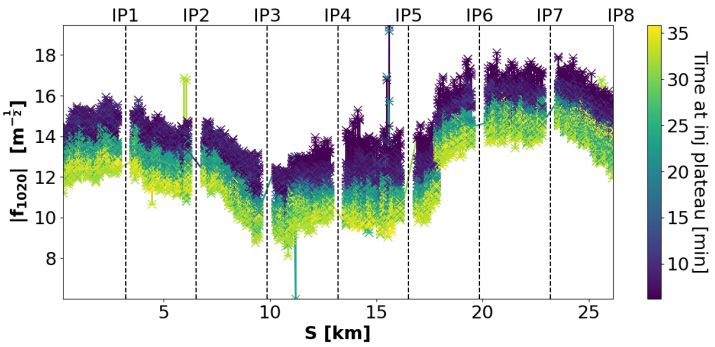


Figure 4.25: 2024 amplitude of f_{1020} RDT vs s for LHCB2, during the initial 35 minutes after arriving at injection, where purple indicates the first minutes after reaching injection, going to yellow which is around 35 minutes after reaching the injection plateau.

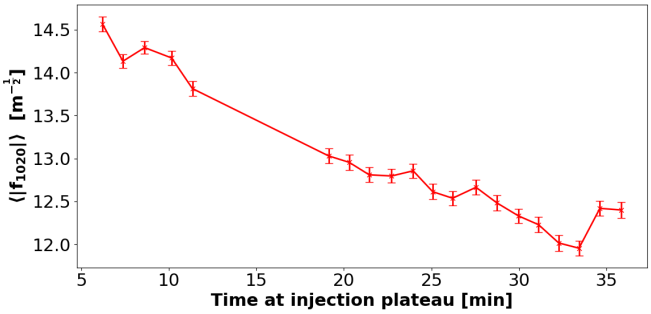


Figure 4.26: 2024 average amplitude of f_{1020} RDT vs s for LHCB2, during the initial 35 minutes after arriving at the injection plateau, where the line is simply a visual guide to see the trend.

With the accelerator preparation coordinated in advance with the OP team, the first optics measurements could be carried out within only a few minutes of ramping the main dipoles to 450 GeV. Figure 4.25 shows a sequence of f_{1020} RDT measurements obtained on the second day of 2024 commissioning, spanning roughly 5 to 35 minutes after reaching injection energy, while the b_3 decay was still in progress. During this interval, the RDT amplitude exhibits a gradual and systematic reduction, suggesting that the time-dependent MCS correction relying exclusively on the FiDeL settings does not fully compensate the evolving decay. This behaviour is illustrated more clearly in Fig. 4.26, which presents the average RDT amplitude around the ring as a function of time.

This study represents the first demonstration that accelerator optics measurements can be successfully performed within such a short time after the beam reaches the injection plateau. Conventionally, approximately 20 minutes are required to establish the appropriate machine

configuration for optics measurements, including verification and adjustment of the tune and chromaticity. The results further indicate that, at this stage of the 2024 commissioning programme, additional optimisation of the MCS settings was still necessary to achieve the desired accelerator performance.

This observed decay behaviour was revisited during 2025, later in the commissioning period. In this later configuration, the time dependence appears reduced, as shown in Fig. 4.27 and Fig. 4.28. This more oscillatory behaviour may be attributed to the presence of unreliable measurements acquired during the middle of the measurement session.

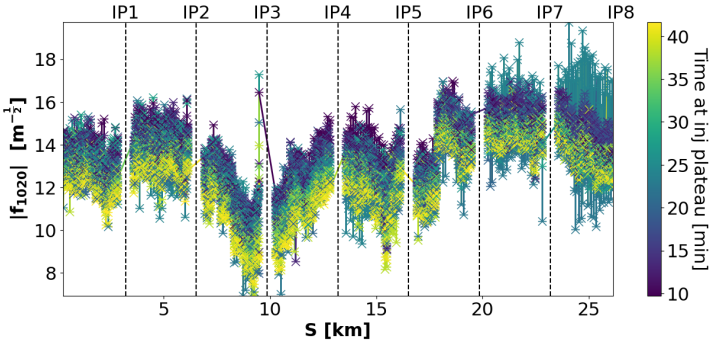


Figure 4.27: 2025 amplitude of f_{1020} RDT vs s for LHCB2, between 10 and 45 minutes after arriving at the injection plateau where purple indicates 10 minutes after injection, going to yellow which is around 45 minutes after reaching the injection plateau.

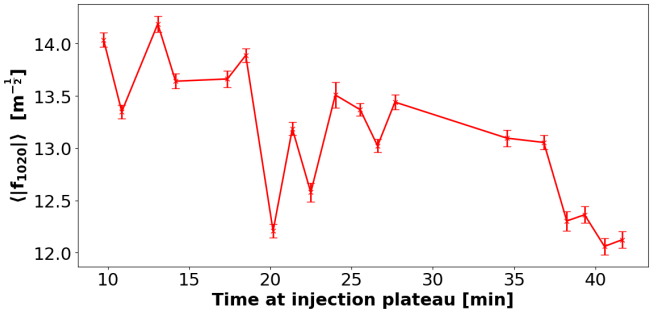


Figure 4.28: 2025 average amplitude of f_{1020} RDT vs s for LHCB2, between 10 and 45 minutes after arriving at the injection plateau, where the line is simply a visual guide to see the trend.

Over longer timescales, the b_3 decay is expected to saturate. It was therefore interesting to compare the RDT stability over a longer period. Figure 4.29 shows measurements taken 40 minutes and 16 hours after reaching the injection plateau, for which the observed difference was negligible. This demonstrates that no significant long-term decay remains after the correction and, therefore, that the f_{1020} RDT can be measured in a reliable and reproducible manner over

longer time periods.

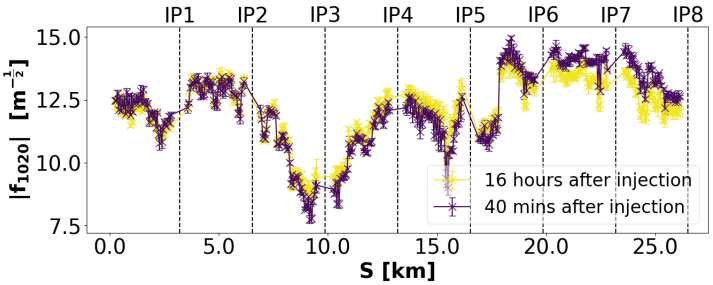


Figure 4.29: 2025 amplitude of f_{1020} RDT vs s for LHCB2, between 16 hours (yellow) at injection and 40 minutes (purple) after arriving at injection.

Although the decay of the f_{1020} resonance did not require further correction, a noteworthy behaviour was observed for the skew-sextupolar $3Q_y$ resonance. This is relevant, as correction of the $3Q_y$ resonance is also of interest for LHC operation [103], given its potential to degrade the lifetime of high-intensity beams. During the 2025 commissioning period, the measured strength of the $3Q_y$ resonance exhibited large variations between optics commissioning shifts. Repeated measurements indicated that this variability correlated with the time spent on the injection plateau (since in some cases optics measurements were scheduled following a long time spent at injection, or immediately following a ramp-down).

The same dataset used to compare the f_{1020} RDT at 40 minutes and 16 hours after injection was used to investigate the time-dependent behaviour of the f_{0030} RDT. The linear optics and coupling were controlled to the same level in each case. The results, shown in Fig. 4.30, reveal an apparent reduction in the RDT magnitude with time.

It should be noted that measurements on comparable time scales over multiple LHC fills were mutually consistent, whereas measurements separated by approximately ten hours exhibited substantial differences. This may suggest the presence of a very slow long-term decay affecting the $3Q_y$ resonance. However, given the limited number of observations, no definitive conclusion can be drawn. A dedicated MD study would be required to fully identify the origin of this apparent decay. Nevertheless, the fact that the $3Q_y$ resonance strength can differ significantly between days implies that this effect should be taken into account during future LHC commissioning activities, and as a consequence, no correction for the $3Q_y$ resonance was implemented in 2025.

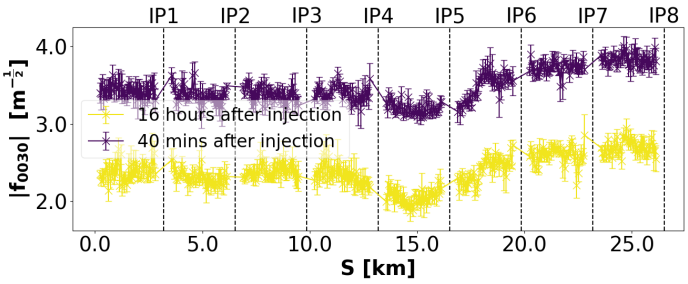


Figure 4.30: 2025 amplitude of f_{0030} RDT vs s for LHCB2, between 16 hours (yellow) at injection and 40 minutes (purple) after arriving at injection.

4.4.2 RDT measurements in the Ramp

As part of the HL-LHC upgrade programme, new triplet magnets as well as separation and recombination dipoles will be installed in the ATLAS and CMS IRs. Recent studies indicate that these magnets are expected to exhibit significant variations in their NL field errors throughout the energy ramp [5, 6], as illustrated in Figs. 4.31 and 4.32, which depict the changing b_3 strength with current in magnetic units. It should be noted that the units presented in these figures correspond to 10^{-4} of the main dipole field evaluated at a reference radius of 17 mm. This convention differs from that adopted in Eq. (2.6), where the multipole strengths are defined using the normalised coefficients described previously. This provides strong motivation to investigate whether NL RDT measurements can be reliably performed during the ramp. If feasible, such measurements would offer a powerful beam-based method for benchmarking and potentially correcting dynamic NL errors originating from the IR magnets during the ramp.

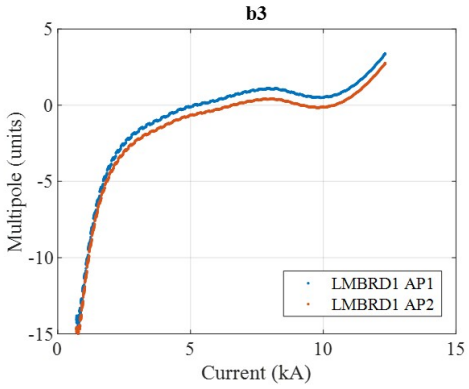


Figure 4.31: Units of b_3 expected to be generated in HL-LHC by a new separation dipole LMBRD1 magnet for the inner and outer apertures (AP1 (blue) and AP2 (orange)) [5].

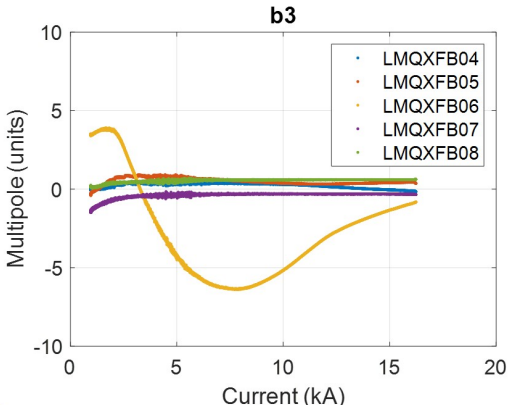


Figure 4.32: Units of b_3 expected to be generated in HL-LHC by some of the superconducting quadrupole LMQXFB magnets [6].

During LHC commissioning, it is standard practice to verify the quality of the linear optics throughout the energy ramp from injection (450 GeV) to top energy (6.8 TeV) [112]. Historically, NL optics measurements were not performed during the ramp, as the studies were constrained to small-amplitude excitations, usually suitable only for linear optics verification.

However, due to the prominence of the f_{1020} RDT even at low excitation amplitudes, a NL analysis of this sextupolar RDT was carried out for the first time during the 2023 ramp. The results, shown in Fig. 4.33, revealed a general increase in the amplitude of f_{1020} as the energy increased. This trend was observed in both beams, with LHCB1 consistently exhibiting a higher RDT amplitude than LHCB2, similar to at injection energy.

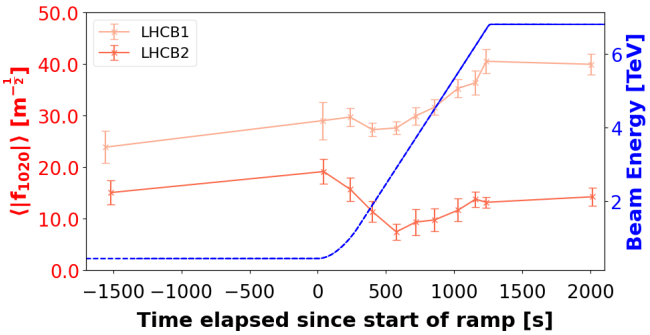


Figure 4.33: 2023 measurement of $\langle|f_{1020}|\rangle$ (mean f_{1020} amplitude) vs time of LHCB1 (light red) and LHCB2 (red), the lines are shown only as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy, in blue.

A re-analysis of 2022 data, along with additional measurements performed during the 2024 commissioning period, confirmed the same behaviour, as shown in Fig. 4.34. The effect of the

correction can be seen in the 2024 measurement, with the first two points displaying the RDT amplitude with it in place.

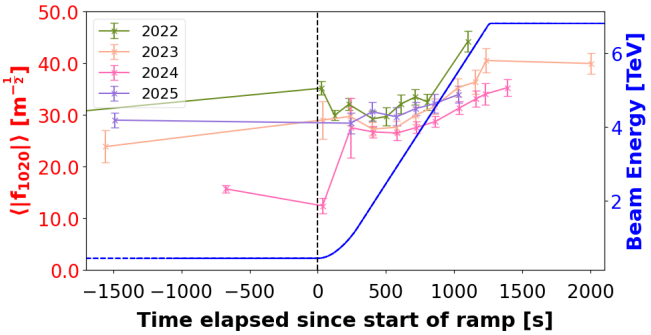


Figure 4.34: 2022 (green), 2023 (light red), 2024 (pink) and 2025 (purple) measurement of $\langle|f_{1020}|\rangle$ vs time for LHCB1, the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy, in blue.

As with the static measurements at injection energy, the observations were benchmarked against MAD-X simulations, which were used since, at the time, an ACD element was available only in this code, and not yet implemented in MAD-NG. This comparison showed that the measured increase in the f_{1020} amplitude during the energy ramp (Fig. 4.35) was approximately a factor of two larger than expected based on the known evolution of the magnetic fields and optics parameters. This discrepancy indicates that an important contribution is likely missing from the model. A plausible explanation is that sextupolar field components in the IRs become increasingly significant as β^* is squeezed towards the end of the ramp. At present, however, there is no reliable model describing these IR b_3 contributions.

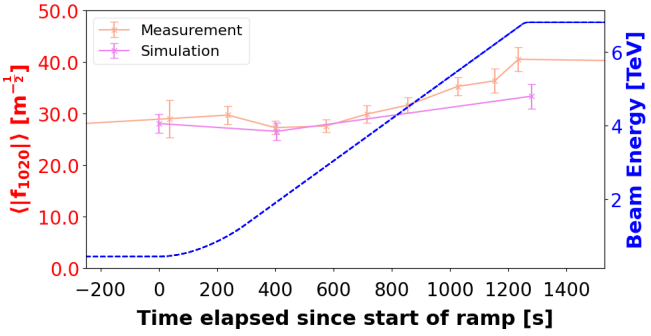


Figure 4.35: 2023 measurement (light red) of $\langle|f_{1020}|\rangle$ vs time of LHCB1 compared with simulation (pink), the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy (blue).

This approach has since been extended to other RDTs, including the one associated with the $3Q_y$ resonance, as illustrated in Fig. 4.36. Such measurements enable a more comprehensive validation of the NL optics model and provide a direct means of verifying the correct application of optics corrections throughout the energy ramp. In contrast, conventional NL optics observables, such as amplitude detuning or NL chromaticity, are not well suited to measurements during a continuous ramp. Implementing these would require a stepwise ramp with extensive operational overhead, rendering such an approach impractical for routine studies.

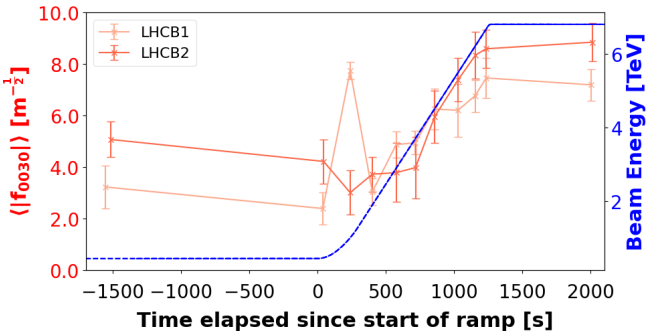


Figure 4.36: 2023 measurement of $\langle|f_{0030}|\rangle$ (mean f_{0030} amplitude) vs time of LHCB1 (light red) and LHCB2 (red), the lines are shown simply as a visual guide to the trend, during the ramp shown as the energy per nucleon, also known as the beam energy (blue).

4.5 Conclusion

This work demonstrated that NL forced-RDT measurements could be reliably exploited at injection in the LHC to study and correct the naturally driven sextupolar resonances of the accelerator. These studies revealed that phase-space distortions associated with the $Q_x + 2Q_y$ resonance are already observable even during nominal linear optics measurements, an unexpected result that was subsequently confirmed to be a genuine accelerator effect. The pronounced nature of this resonance motivated dedicated investigations aimed at its mitigation. Using MAD-NG, the strengths of the MS circuits were optimised to minimise the magnitude of the f_{1020} RDT in LHCB1. The resulting correction was successfully validated through beam-based measurements, demonstrating clear improvements in beam lifetime and forced DA. Based on these results, the correction scheme was deployed in routine LHC operation throughout 2024. The transition to a new reversed-polarity optics in 2025 required the optics commissioning to be repeated.

Unfortunately, there was no time during this commissioning to revalidate the b_3 RDT correction. However, offline analysis showed that the same scheme would also have been valid. This work demonstrated clearly that a simulation-driven correction strategy based on free-RDTs can be effectively translated into a reduction of the corresponding forced-RDT, establishing a robust framework for future NL optics optimisation.

Future work could focus on extending the matching procedure by introducing an explicit constraint on the chromatic β-beating in addition to the RDTs. Implementing such a constrained optimisation would require modifications to the MAD-NG code.

Furthermore, the ability to measure RDTs in a robust and reproducible manner even at small excitation amplitudes motivated the development of novel measurement procedures. These techniques enabled, for the first time, successful NL RDT measurements during dynamically evolving accelerator states. This constitutes a significant extension of earlier NL optics studies in the LHC, which had been limited to static accelerator conditions.

More generally, throughout Run III, as a result of this work, the use of forced NL RDTs has become a standard and reliable tool for accelerator commissioning and beam dynamics studies. These techniques are now routinely employed during LHC operation, including for applications such as the correction of the $3Q_y$ resonance [103]. In addition, investigations of higher-order field errors [110] have relied on improvements in the forced DA made possible by this work. The f_{1020} RDT measurements were also instrumental in extending injection studies to higher-order RDTs, as described later in this thesis.

CHAPTER 5

Octupolar RDTs

In the previous chapter, it was demonstrated that an accurate magnetic model enabled the successful design and implementation of an effective normal sextupolar RDT correction scheme. However, such agreement between model and accelerator cannot always be assumed. This chapter examines a situation in which substantial discrepancies were observed between model predictions and beam-based measurements.

The response of the Landau octupoles (MO) had previously been assessed prior to this work, for example through measurements of amplitude detuning and NL chromaticity at the end of Run I, where an approximate agreement with simulation was reported [1, 113]. A new injection optics configuration was introduced in 2023 [102]. This optics configuration was designed through simulation to reduce the excitation of the normal octupolar resonances through a phase-optimisation strategy. The success of this simulation-based correction relies critically on the accurate modelling of the RDT response to changes in the octupole circuits.

While agreement between measurement and simulation in global observables such as amplitude detuning can be reassuring, it does not necessarily guarantee that the RDT behaviour is correctly captured. Furthermore, these checks were done in previous runs with much weaker excitation currents than are used today. For this reason, a key objective during the Run III commissioning was to perform the first direct beam-based measurements of the RDT response

driven by the MO, thus providing the principal motivation for the work presented in this thesis. Measurements of the targeted octupolar RDTs revealed that the new injection optics did not deliver the performance expected from simulation, prompting an extensive investigation in which both modelling limitations and potential hardware-related effects had to be carefully examined.

5.1 Landau Octupoles (MO)

In circular accelerators, octupoles are routinely employed as a controlled source of non-linearity. Octupolar fields play an important role; they allow the deliberate manipulation of amplitude-dependent tune shifts, and therefore the NL stability properties of the beam through the use of Landau damping. Landau damping occurs as, when particles oscillate at different tunes, coherent motion is suppressed, preventing the growth of collective instabilities [114]. Their relevance has been demonstrated in a wide range of accelerators, such as the SPS [115], Tevatron [116], LHC [117], and, in the future, FCC-hh [118] where collective effects impose stringent stability requirements.

In the LHC, there are three sources of octupolar fields: intrinsic magnetic field errors, dedicated octupole correctors (MCO in the arcs and MCOX in the low-β IRs), and the MO. While sextupoles can also give rise to second-order octupolar-like effects through geometric non-linearities that distort phase space and modify tunes as a function of amplitude, these contributions are negligible at injection energy. In practice, octupolar field errors are locally compensated using the appropriate corrector magnets, whereas the MO are deliberately powered to generate the tune spread required for Landau damping. An example of the tune footprint produced by the LHC Landau octupoles powered at 18 m^{-4} at injection generating detuning of the order 100 km^{-1}, is shown in Fig. 5.1, where the resulting tune spread is highlighted in red.

While octupoles are beneficial for mitigating collective instabilities, they also strongly excite higher-order resonances, which can adversely affect single-particle stability, DA, and beam lifetime. Since the beams remain at injection for a significant period, these effects can ultimately reduce the achievable luminosity by increasing beam losses prior to collisions. Figure 5.1 illustrates the principal resonances driven by normal octupolar field components, shown in blue.

Particular attention should be paid to the diagonal resonance corresponding to the f_{2002} RDT, as it lies closest to the nominal WP. The same amplitude-dependent tune spread that enhances Landau damping can therefore become detrimental to beam stability when the tune footprint overlaps strongly with resonance lines. As shown in Fig. 5.1, such an overlap occurs with the f_{2002} resonance, which would lead to enhanced resonant excitation, emittance growth, a reduction in DA, and ultimately beam loss. Consequently, the optimisation of octupole settings represents an inherently delicate compromise between maximising collective stability and preserving favourable single-particle dynamics.

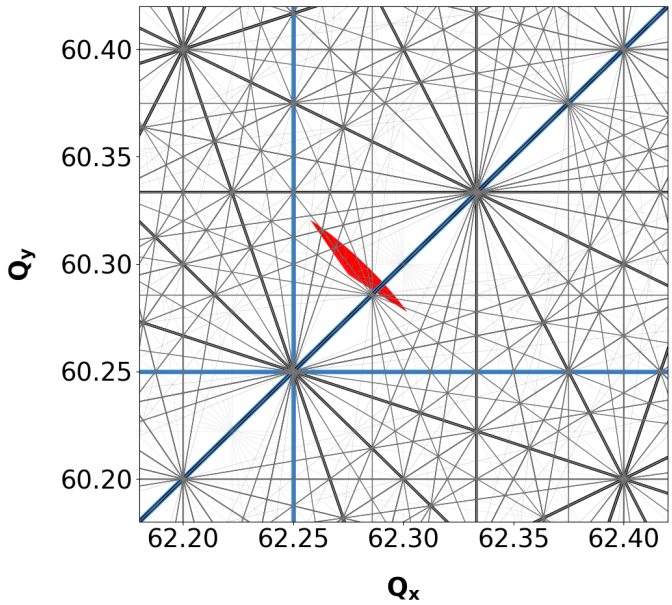


Figure 5.1: Q_y vs Q_x where the blue lines correspond to the resonances $4Q_x$, $4Q_y$, and $2Q_x - 2Q_y$, and the red shape is the tune spread/footprint caused by the MO at injection with 2023 optics configuration.

In the LHC, the effect of the MO has been studied in both the design, for example in [117], and in beam-based studies in the accelerator itself. Experience from operation has shown that strong octupole currents are often required for high intensity operation to mitigate effects such as electron cloud and impedance-driven instabilities. At the same time, these settings were seen to significantly reduce DA and enhance resonance excitation, with direct consequences for beam lifetime, as seen in Fig. 5.2 where the powering of the MO at 40 A resulted in a clear decrease in beam lifetime.

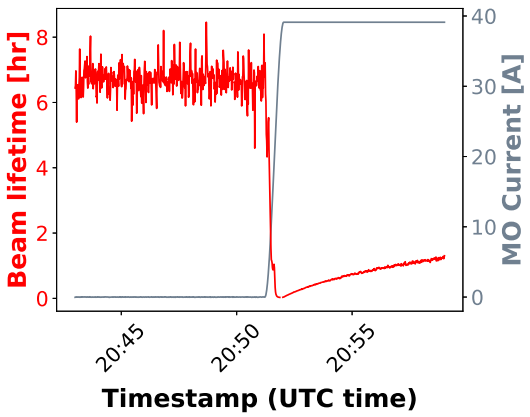


Figure 5.2: BLM lifetime at injection energy vs time for LHCB1 (red, left-hand axis) following a change in MO current (grey, right-hand axis). Measured using pilot bunches during optics commissioning. See Sec 3.8.3 for discussion on lifetime.

At injection energy, the MO are operated at relatively high strength, with $K_4 \approx 18$ m^{-4}. This corresponds to direct horizontal and vertical amplitude-dependent detuning coefficients of approximately 1.5×10^5 m^{-1}, while the cross-term is approximately -1.1×10^5 m^{-1}. In addition to their significant influence on beam lifetime, the MO, in combination with chromaticity effects, were identified as a principal source of collimation hierarchy breakage during the 2022 LHC run [119]. This behaviour has direct implications for machine protection [120], as MO-driven resonances can induce phase-space distortions that allow particles to bypass specific collimators, therefore compromising the intended collimation hierarchy. Similar effects were observed again during the 2025 operational period [121].

In parallel, electron-cloud simulations indicated that the MO also contribute to transverse emittance growth, an effect that could potentially be mitigated by reducing the excitation of specific RDTs [122]. These observations motivated the development of a phase-compensation scheme [102] designed to minimise the resonances driven by the MO, in particular the f_{2002} RDT, while preserving the required amplitude-dependent tune spread [102]. This development motivated the work presented in this thesis, which includes the first direct experimental measurements of MO-driven RDTs performed during the subsequent Run III commissioning period.

5.2 MO Response

Measurements of the MO-driven RDT response at injection were carried out following the standard optics measurement procedure and were performed with low-intensity pilot beams. The principal modifications relative to routine optics measurements were, first, the adjustment of the working point to the so-called “e-cloud” tunes ($Q_x = 62.275$, $Q_y = 60.293$) to enhance the visibility of the relevant spectral lines, as this choice brings the tunes closer to the relevant resonances. The ACD excitation frequencies were set accordingly (the shift from the natural tune being $\Delta Q_x = -0.01$, $\Delta Q_y = -0.011$). An MO strength of 5 m^{-4} was applied, corresponding to a current of approximately 10.8 A. This value was chosen as a compromise: it was sufficiently large to produce a measurable change in the RDT, while remaining low enough to avoid excessive degradation of the DA, which would otherwise prevent excitation with the ACD due to beam losses.

Fig. 5.3 shows an example horizontal frequency spectrum for measurements with the MO powered at 5 m^{-4} where the f_{2002} resonance is shown to be clearly visible above the BPM noise.

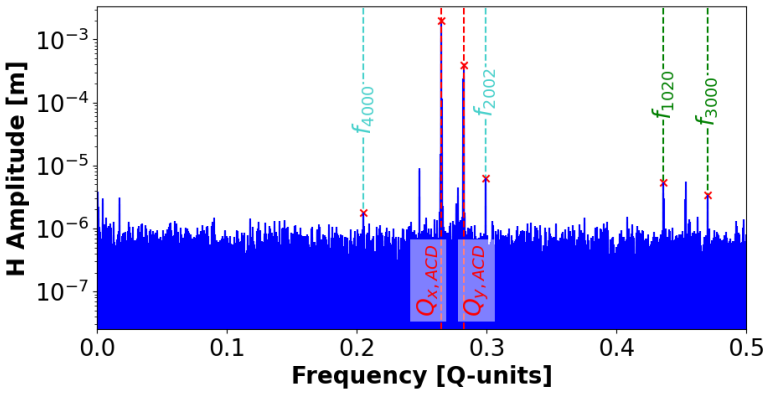


Figure 5.3: Horizontal TbT frequency spectrum (amplitude vs frequency), where green denotes third-order and cyan the fourth-order RDTs, and red the ACD frequencies with natural tunes of $q_x = 0.275$ and $q_y = 0.293$ for LHCB1 where the MO is powered at 5 m^{-4}. The other peaks are the natural tunes, other normal octupolar and sextupolar resonances showing the b_3 and b_4 contributions.

Figures 5.4 and 5.5 show a comparison between simulation and measurements completed in 2024 commissioning with and without the MO powered. It can be seen that only a small change in the RDT is expected when the MO are powered in simulation.

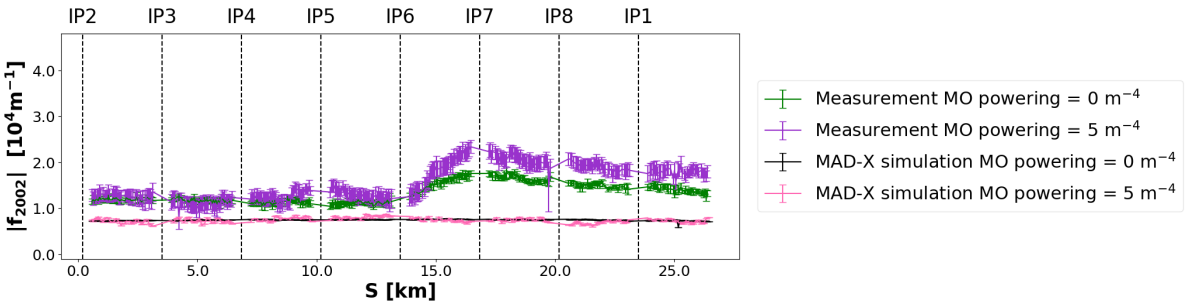


Figure 5.4: 2024 measured (former colour) and simulated (latter colour) amplitude of f_{2002} RDT vs s, with (purple/pink) and without (green/black) MO powered for LHCB1.

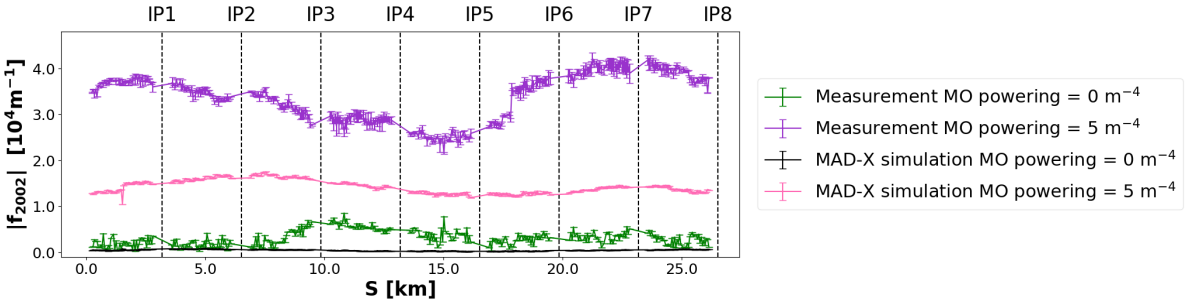
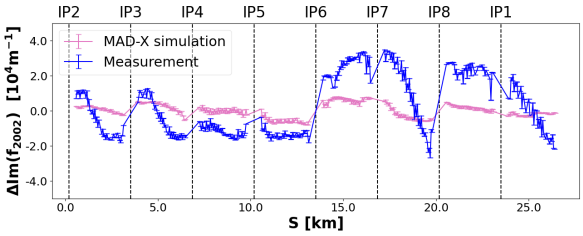


Figure 5.5: 2024 measured (former colour) and simulated (latter colour) amplitude of f_{2002} RDT vs s, with (purple/pink) and without (green/black) MO powered for LHCB2.

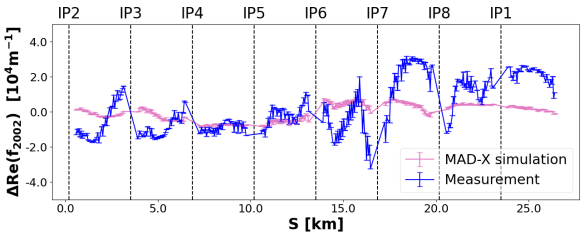
To assess the influence of the powering of the MO on the RDTs, the initial analysis focused on the change in RDT magnitude between the cases with the MO powered on and off. As shown in Fig. 5.4, the shift in the amplitude of the f_{2002} RDT between the MO powered on and off for LHCB1 appeared comparable between measurement and simulation. In contrast, for LHCB2 (Fig. 5.5), the measured change in amplitude was approximately three times larger than that predicted by the simulation. From these initial observations, it appeared that the agreement between measurement and model was poorer for LHCB2 than for LHCB1.

To obtain a more complete understanding of this discrepancy, the shifts in the real and imaginary components of the RDT were analysed, as shown in Figs. 5.6(a), 5.6(b), 5.7(a), and 5.7(b). In this representation, the disagreement between measurement and simulation is more pronounced for LHCB1 than for LHCB2, in contrast to the conclusion suggested by the amplitude comparison alone. This behaviour arises since looking only at the RDT amplitude does not capture phase variations of the RDT around the ring, effectively ignoring the sign of the RDT which could therefore conceal discrepancies. In the present case, when the LHCB1 MO were powered,

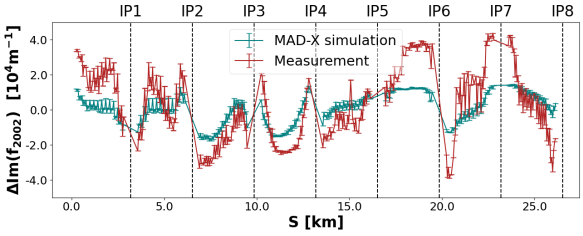
the RDT did not remain unchanged; rather, it crossed zero and subsequently increased again, ultimately attaining a magnitude comparable to its initial value but with an opposite phase.



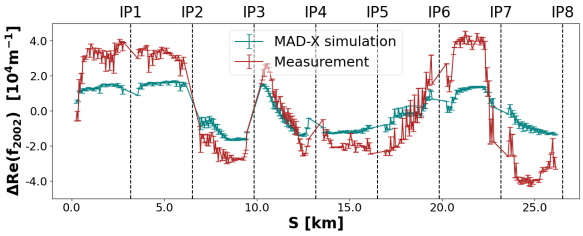
5.6(a) Shift to the imaginary part of f_{2002} vs s caused by all the MO being powered at 5 m^{-4} for LHCB1 in simulation (pink) and measurement (dark blue).



5.6(b) Shift to the real part of f_{2002} vs s caused by all the MO being powered at 5 m^{-4} for LHCB1 in simulation (pink) and measurement (dark blue).



5.7(a) Shift to the imaginary part of f_{2002} vs s caused by all the MO being powered at 5 m^{-4} for LHCB2 in simulation (teal) and measurement (dark red).



5.7(b) Shift to the real part of f_{2002} vs s caused by all the MO being powered at 5 m^{-4} for LHCB2 in simulation (teal) and measurement (dark red).

The observation of a substantial discrepancy between measurement and simulation was unexpected. Consequently, further investigation was required, and the subsequent analysis was based only on the MO-induced shift to the real and imaginary components of the RDT, rather than on the amplitude alone. Discussion in the remainder of this chapter focuses on LHCB1, where the deviation is most significant, although the conclusions will apply to both beams. As no qualitative distinction was observed between the behaviour of the real and imaginary parts, the imaginary component is usually the one shown, without loss of generality, as the representative quantity for the remainder of this study.

The 2024 measurements were presented here, in preference to those from 2023, owing to their improved data quality. Nevertheless, Fig. 5.8 demonstrates the reproducibility of the observed behaviour between the 2023 and 2024 measurement campaigns.

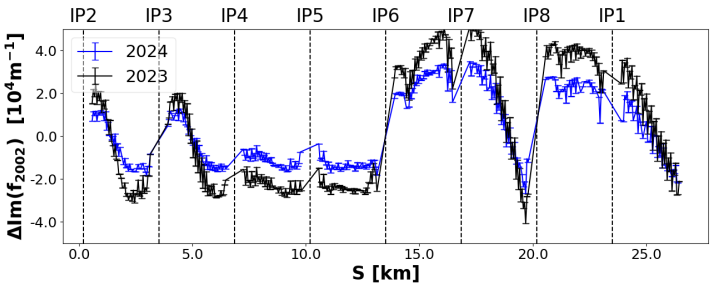


Figure 5.8: 2024 (dark blue) and 2023 (black) measured shift to imaginary part of f_{2002} RDT vs s with and without MO powered for LHCB1.

5.3 Verifying the MO response discrepancy

Given the critical role of the MO in LHC operation, it was essential to establish that the observed discrepancy was not an artefact arising from the measurement or analysis procedures. The next stage of the investigation assessed whether the discrepancy could be attributed to systematic effects, including amplitude detuning and BPM non-linearities. The investigation then examined whether the newly implemented optics optimisation was responsible for the observed behaviour. Finally, the possible influence of magnetic hysteresis was evaluated to determine whether it could account for the discrepancy, particularly at injection energy.

To investigate the measurement and analysis tools, the RDT strengths computed directly from normal form by MAD-NG were benchmarked against tracking simulations performed in MAD-X, with the resulting TbT data from that tracking being processed using the same analysis framework applied to beam measurements. In this comparison, tracking was performed using an initial free oscillation rather than an ACD excitation, owing to the absence of an ACD element in MAD-NG. The two methods agreed to within approximately 10% level, confirming there was no fundamental issue with the analysis tools.

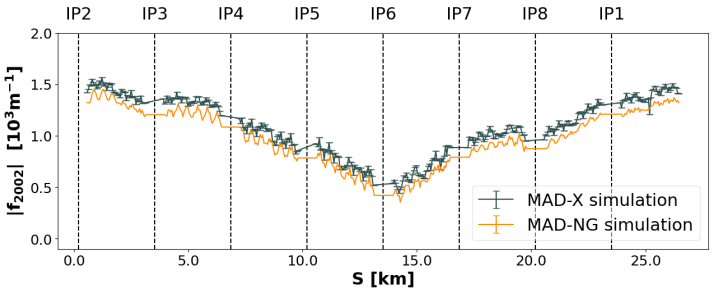


Figure 5.9: Comparison of free-kick MAD-X (dark green) and MAD-NG (orange) simulated amplitude of f_{2002} RDT vs s for LHCB1.

Further validation of the measurement procedure was carried out by testing the analysis tools on an independent magnetic circuit at top energy. Specifically, the RDT response of the octupolar IR corrector circuits (MCOX) in IR1 and IR5 were examined.

Figure 5.10 illustrates the change in the imaginary component of the f_{2002} resonance driving term resulting from the application of a MCOX trim (MCOX left of IP5 = -1.71 m^{-4}, MCOX right of IP1 = 1.54 m^{-4}). A strong agreement was observed between the measured data and the simulation results further confirming that the source was not an issue with the analysis tools.

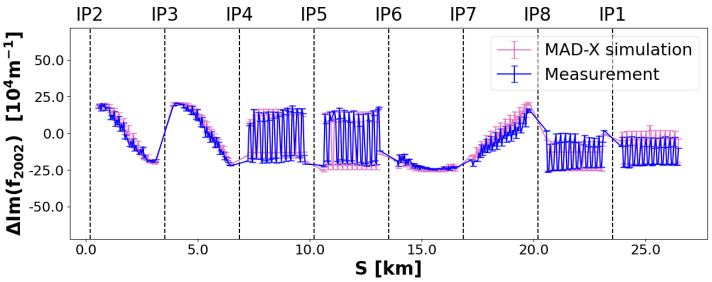


Figure 5.10: Comparison of simulation (pink) and measurement (dark blue) for the shift to the imaginary part of the f_{2002} RDT vs s caused by the MCOX trim for LHCB1.

To verify the sensitivity of the response, measurements were also performed using the inverse of the applied trim. As shown in Fig. 5.11, the reversed MCOX trim produced an equal and opposite shift in the RDT, fully consistent with expectations, and further validating the accuracy of the measurement and analysis procedure.

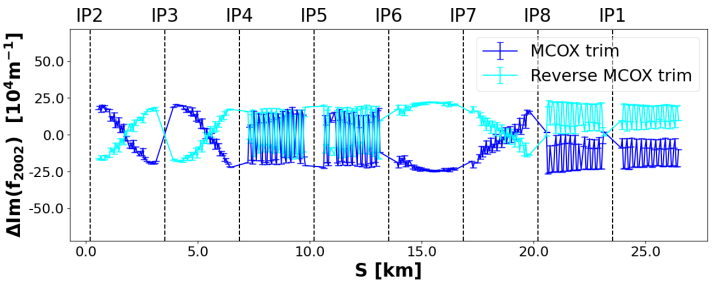


Figure 5.11: Comparison of the measured shift to the imaginary part of the f_{2002} RDT vs s caused by the MCOX trim (dark blue) and its reverse (cyan) for LHCB1.

Another potential explanation for the observed discrepancy is amplitude detuning generated by the MO. The forced-RDT can be influenced by the difference between the ACD excitation frequency and the natural tune [104]; therefore, in principle, substantial detuning over multiple

high-amplitude kicks could bias the measured RDT. To examine this possibility, an alternative MO powering configuration was implemented in which the net detuning effects are self-compensated. In this "alternating" scheme, the polarity of the MO trims was reversed after arcs 1 to 4. As shown in Fig. 5.12, despite this adjustment, a significant discrepancy between measurement and simulation persisted, indicating that amplitude detuning is unlikely to be the primary origin of the observed disagreement.

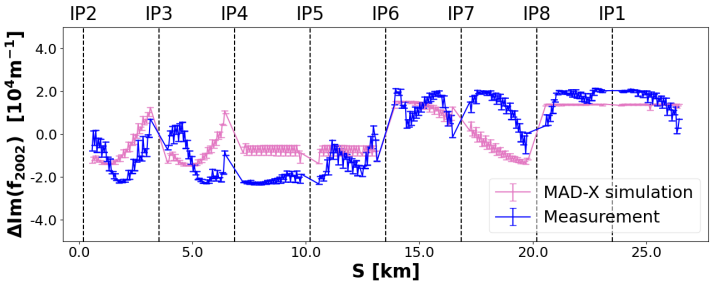


Figure 5.12: Comparison of the simulated (pink) and measured (dark blue) shift to the imaginary part of the f_{2002} RDT vs s caused by an alternating MO trim for LHCB1.

Since no shortcomings were identified in the analysis framework, an additional hypothesis was that the RDT measurement might be affected by an unknown BPM non-linearity. To test this possibility, the accelerator tunes were shifted from the "e-cloud" working point to $Q_x = 62.28$ and $Q_y = 60.313$, with corresponding ACD excitation offsets of $\Delta Q_x = -0.008$ and $\Delta Q_y = 0.01$, matching the configuration typically used for sextupolar resonance measurements. The measured RDT magnitude changed consistently with this tune variation, as shown in Fig. 5.13. This clear dependence on the working point is consistent with the expected behaviour of a genuine RDT. It was therefore unlikely to be a BPM-related artefact.

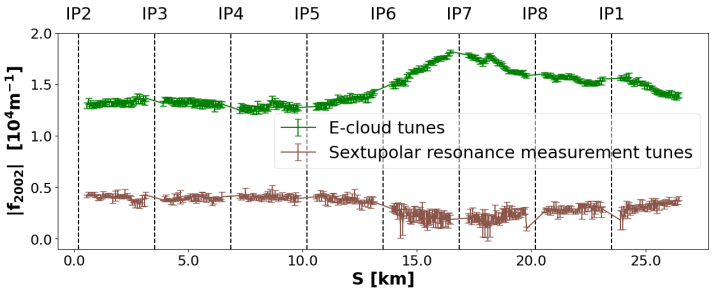


Figure 5.13: Comparison of amplitude of f_{2002} RDT vs s with e-cloud (green) and sextupolar resonance measurement (brown) tunes for LHCB1.

An alternative hypothesis was that the newly implemented optics configuration was not performing as designed. To investigate this possibility, dedicated measurements were carried out in which the accelerator optics were reverted to the 2022 configuration. A significant discrepancy between measurement and simulation was observed even with the 2022 optics, as shown in Fig. 5.14. This demonstrates that an incorrect implementation or flawed design of the new injection optics cannot account for the observed discrepancy in the measured MO response.

Closer examination indicates that, upon reverting to the 2022 optics, the measured shift in the RDT increased. This observation confirms that the new optics configuration does provide partial mitigation of the MO-driven resonances. However, the reduction achieved is smaller than anticipated from simulation.

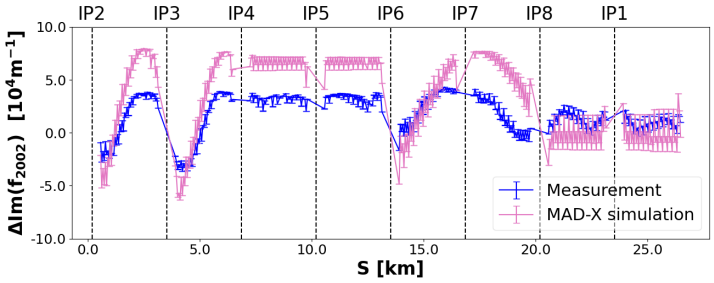


Figure 5.14: Comparison of the simulated (pink) and measured (dark blue) shift to the imaginary part of the f_{2002} RDT vs s with all the MO being powered at 5 m^{-4} for LHCB1 with the old optics.

Since issues related to the analysis framework, BPM non-linearities, and amplitude detuning had been excluded, the remaining evidence pointed towards a real property of the magnets themselves as the source of the discrepancy. Given the operational importance of the MO system and its role in beam-dynamics modelling, this inconsistency was of particular concern, especially if it indicated a hardware-related issue affecting one or more MO circuits.

It was therefore necessary to determine whether the effect was specific to injection energy. To this end, a dedicated test was performed during a Proton Physics reference (PP Ref) run, in which all MO circuits were powered uniformly at the same strength used for the injection measurements. During PP Ref operation, the beam energy is approximately 2.6 TeV, significantly higher than injection energy. As shown in Figs. 5.15 and 5.16 (the latter is included because the agreement is marginally poorer than for LHCB1), a good agreement between measurement and simulation was recovered at this intermediate energy.

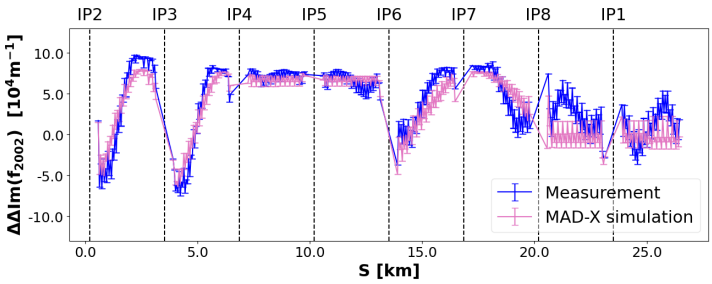


Figure 5.15: Comparison of simulation (pink) and measurement (dark blue) for the difference in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} for LHCB1.

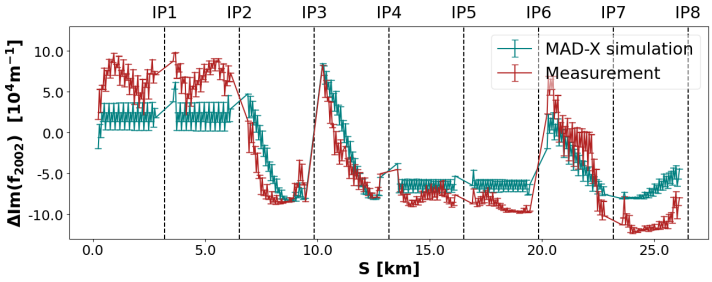


Figure 5.16: Comparison of simulation (teal) and measurement (dark red) for the difference in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} for LHCB2.

Given that the discrepancy in the MO response appears to manifest only at injection energy, the current interpretation is that it is most likely related to an un-modelled hysteresis component in the MO circuits. This is not an anticipated effect, as the MO circuits are degaussed during the pre-cycle. Moreover, significant residual hysteresis contributions in the MO circuits had previously been excluded on the basis of amplitude detuning and Q'' measurements performed during Run I [12, 113]. It should also be noted that this hysteresis must be non-uniform across the circuits: if it affected all circuits identically, the phase-compensation scheme would be expected to cancel the effect globally.

To further investigate this behaviour, additional RDT measurements were performed for several uniform MO strengths, varying from 5 m^{-4} to 2 m^{-4} and subsequently to -2 m^{-4}, as shown in Fig. 5.17. In the absence of hysteresis, the measured response would be expected to exhibit an approximately linear dependence on the octupole strength (Eq. (2.80)). However, such linearity is not observed in the data. This NL behaviour provides further evidence that the RDT response at injection energy is influenced by hysteresis effects within the MO system.

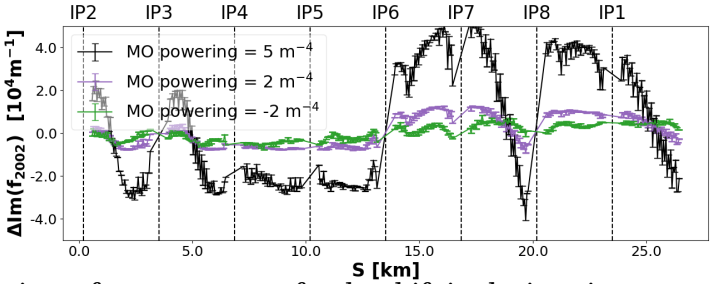


Figure 5.17: Comparison of measurements for the shift in the imaginary part of the f_{2002} RDT vs s caused by all MO being powered at 5 m^{-4} (black), 2 m^{-4} (purple), and -2 m^{-4} (green) for LHCB1.

The next step was to determine which of the MO circuits were exhibiting this hysteresis effect. To investigate this, dedicated measurements were performed in which different arc octupole circuits were powered independently, allowing direct comparison with MAD-X simulation results.

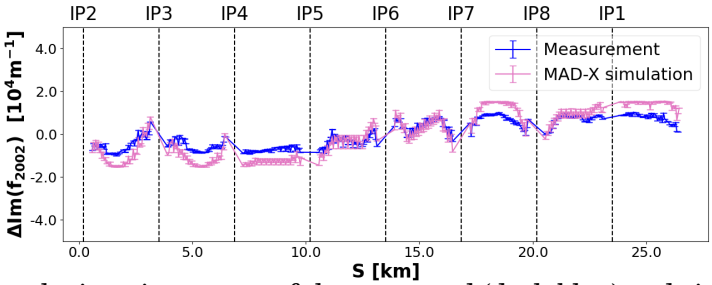


Figure 5.18: Shift to the imaginary part of the measured (dark blue) and simulated (pink) f_{2002} vs s caused by the MOF in arc12 being powered at 10 m^{-4} for LHCB1.

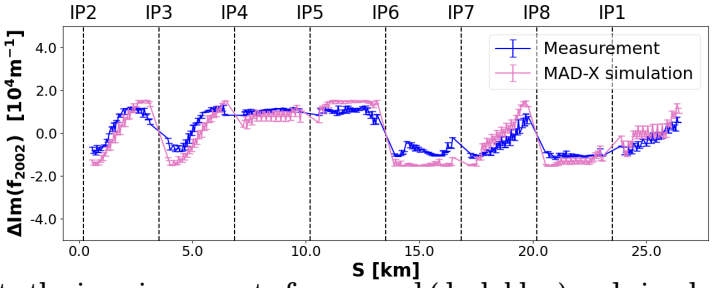


Figure 5.19: Shift to the imaginary part of measured (dark blue) and simulated (pink) f_{2002} vs s caused by the MOF in arc56 being powered at 10 m^{-4} for LHCB1.

As shown in Figs. 5.18 and 5.19, when only the MO circuits in arc12 or arc56 were individually powered, the resulting RDT shifts exhibited significantly better agreement between simulation and measurement than observed for uniform MO powering. The next logical step would be to extend this study to the remaining arcs in order to identify whether any specific circuit deviates from the expected behaviour. However, to individually check every circuit in this manner would require a significant amount of dedicated LHC beam time. Therefore, a new Segment-by-Segment (SbS) method was developed and used to identify the potentially problematic arcs instead.

5.4 Segment-by-Segment (SbS) Studies

One of the key challenges highlighted by this octupolar study is the difficulty of localising the sources of NL errors. Such investigations can be highly time-consuming, often requiring a systematic, arc-by-arc analysis and, in certain cases, an examination at the level of individual circuits. This level of granularity renders the approach impractical within the time constraints of standard commissioning periods.

A well-established and highly effective technique for localising errors in the linear optics regime is the SbS approach [77, 123]. In this method, initial optics functions are propagated in simulation along a segment (from one chosen BPM to another) of the ring, using the designed accelerator model, and optical functions extracted from measured TbT data as initial conditions. The propagated optical functions in successive BPM locations are then compared with the corresponding measurements. An error in the accelerator, differing from the design model, shows up as a discrepancy between the propagated and measured optics functions downstream of the error. The resulting discrepancies can be used to identify local errors [77, 124] and derive local corrections [98, 125, 126]. An example of such a measurement is shown in Fig. 5.20, which shows a difference between measured and propagated phase advance, used to determine local optics corrections in IR5 during LHC commissioning in 2012.

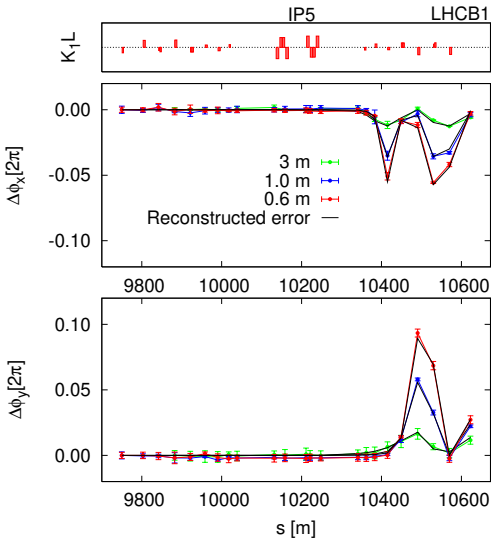


Figure 5.20: Example SbS analysis for the phase error vs s with different optics i.e. β at IP5 = 3 m (green), 1 m (blue), 0.6 m (red) for LHCB1 with the dipole multipoles shown in red above [7].

A preliminary study was conducted to evaluate whether a comparable methodology could be applied to the localisation of NL errors within a simple simulation framework. Tracking simulations were performed in MAD-X including the ACD, and the resulting TbT data were analysed using the OMC team's standard analysis tools. The phase advance of the RDT, determined from an initial BPM, was subsequently compared with its analytical propagation along the accelerator.

The analytical propagation of the RDT phase was derived from the relation given in the following equation [127].

$$f^{(2)}_{jklm} = e^{-i[(k-j)\Delta\phi_x+(m-l)\Delta\phi_y]}\left[f^{(1)}_{jklm} - \sum_{q=1}^{n} e^{i[(k-j)\phi_{xq}+(m-l)\phi_{yq}]}h_{q,jklm}\right], \tag{5.1}$$

where $f^{(2)}_{jklm}$ denotes the RDT evaluated at the second location, whereas $f^{(1)}_{jklm}$ corresponds to the same quantity at the first location. The terms $\Delta\phi_x$ and $\Delta\phi_y$ represent the horizontal and vertical phase advances between these two positions. The summation is taken over the n multipoles located between the two points of interest. The phase advances ϕ_{xq} and ϕ_{yq} are defined between the first location and the qth multipole. The coefficients $h_{q,jklm}$ are real quantities proportional to the strength of the qth multipole and to the product $\beta_{xq}^{(j+k)/2}\beta_{yq}^{(l+m)/2}$. A more comprehensive definition is provided in [128]. For the propagation procedure, the linear optics functions and phase advances were taken from the TbT ACD measurements.

Figure 5.21 shows the difference between the analytically propagated phase and the simulated phase of the f_{2002} RDT, both for the reference configuration in which all MO circuits are switched off and for the case where the MOF magnets in arc56 are intentionally mis-powered. In the absence of a significant localised NL source, the deviation between analytical propagation and simulation is expected to remain close to zero (some deviation can remain since the analytical propagation neglects higher-order contributions to the RDT).

A clear deviation from zero indicates the presence of an additional strong NL contribution, leading to a discrepancy between the simulated phase and its analytically propagated counterpart. A distinct phase shift is observed in the region corresponding to the MOF magnets in arc56, precisely at the location where the artificial error is introduced. This behaviour is even more pronounced in Fig. 5.22, which presents the corresponding phase comparison for the f_{4000} RDT.

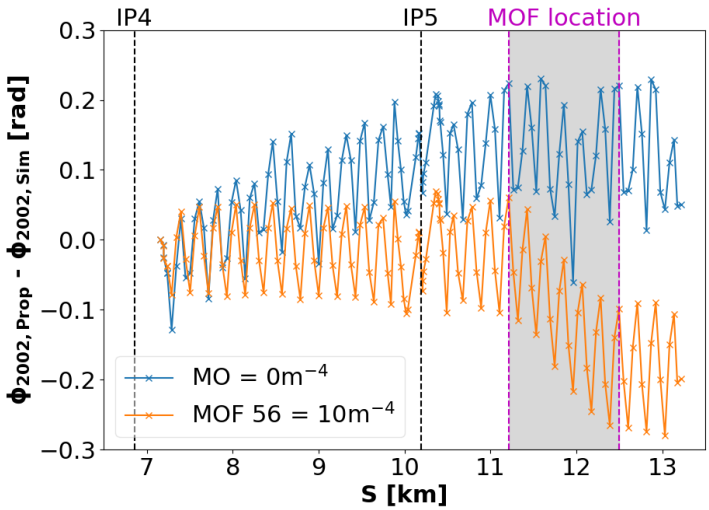


Figure 5.21: Comparison of the difference between simulation and propagated f_{2002} phase vs s for when all the MO were turned off (blue), and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1.

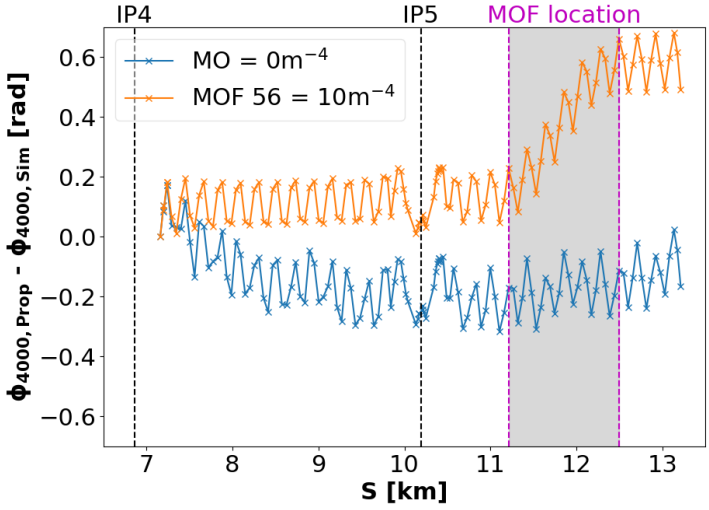


Figure 5.22: Comparison of the difference between simulation and propagated f_{4000} phase vs s for when all the MO were turned off (blue), and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1.

A similar behaviour is evident in Fig. 5.23, which compares the difference between the propagated and calculated real components of the RDT with the same configurations.

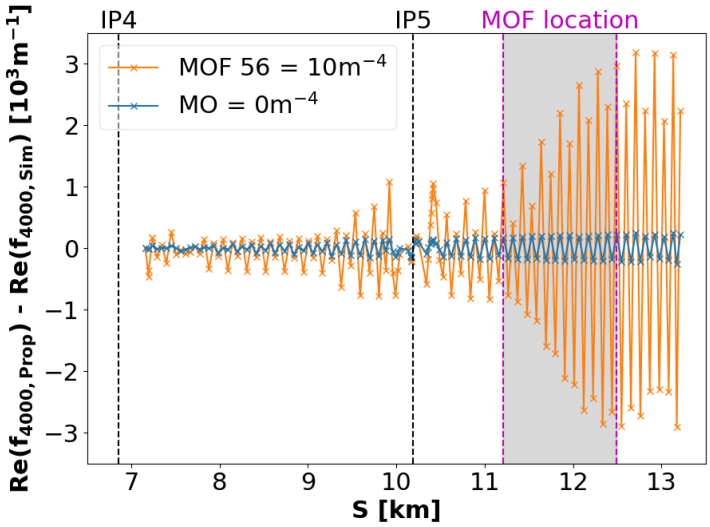


Figure 5.23: Comparison of the difference between simulation and propagated f_{4000} real part of the RDT vs s for when all the MO were turned off (blue) and when the MOF in arc56 were powered at 10 m^{-4} (orange) for LHCB1.

Although these represent simplified test scenarios, the results demonstrate the potential of a NL SbS approach for identifying the location of local NL errors within the accelerator. Therefore, the method was extended to previously acquired experimental data in order to investigate measurement scenarios without requiring additional dedicated beam time. As an initial test case, the configuration in which all MO circuits were uniformly powered to 5 m^{-4} was compared with the reference case in which all MO circuits were switched off.

An example of the results from this analysis is presented in Fig. 5.24. Whereas in most arcs little deviation of the RDT from the ideal model propagation was observed, a pronounced RDT deviation is introduced in arcs 67 and 78 even with the MO off.

The reproducibility of this behaviour was examined for different levels of MO excitation. In one such case, all MO circuits were uniformly powered to 2 m^{-4}, as shown in Fig. 5.25. The same qualitative enhancement of the RDT deviation between the powered and unpowered configurations was observed, albeit with reduced amplitude. Once again, arcs 67 and 78 appear to be of particular significance.

Consistent with the findings from conventional RDT studies, the NL SbS response in these

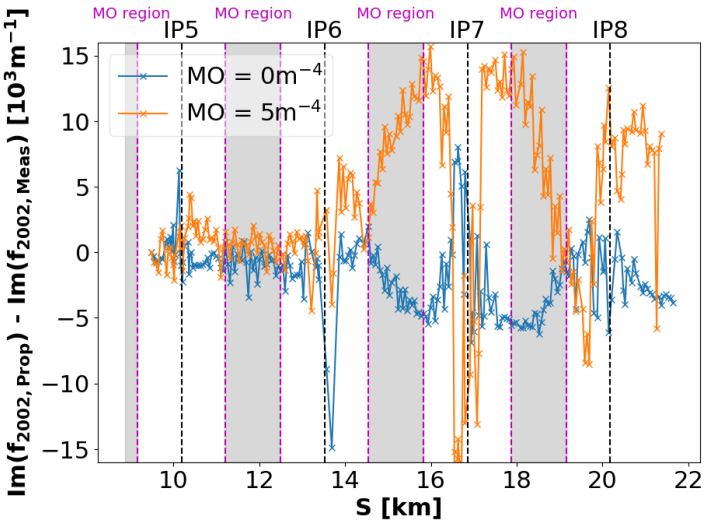


Figure 5.24: Comparison of the difference between measurement and propagated f_{2002} imaginary part of the RDT vs s for when all the MO were switched off (blue) and when all the MO uniformly powered at 5 m^{-4} (orange) for LHCB1.

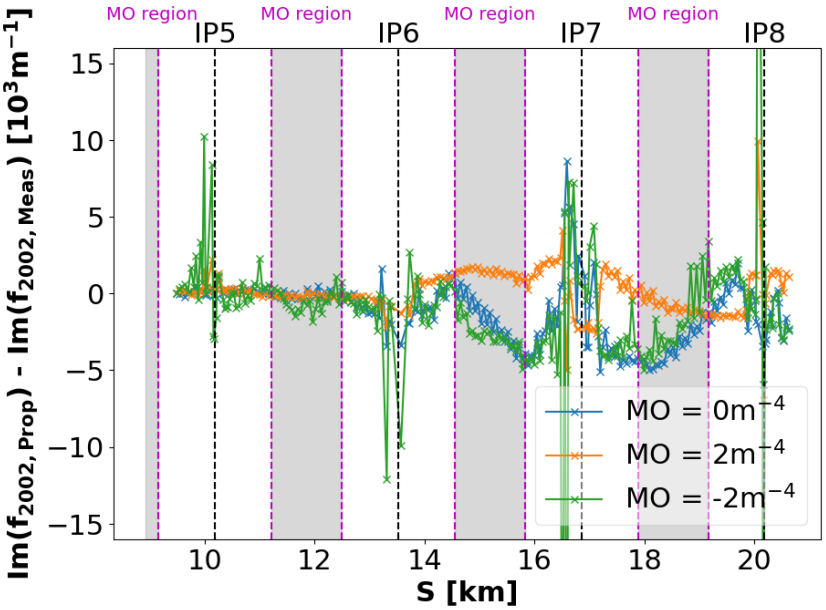


Figure 5.25: Comparison of the difference between measurement and propagated f_{2002} imaginary part of the RDT vs s for when all the MO were switched off (blue), when all the MO were uniformly powered at 2 m^{-4} (orange), and when all the MO were uniformly powered at -2 m^{-4} (green) for LHCB1.

arcs was found not to scale linearly with the applied MO trim when comparing with all MO circuits uniformly powered to -2 m^{-4}, also shown in Fig. 5.25. Such behaviour is indicative of a possible hysteresis effect.

On the basis of the preceding SbS analyses, which suggested that arcs 67 and 78 were plausible sources of the discrepancy, a dedicated study was carried out in which the MO circuits in arc67 were powered individually to 10 m^{-4}.

For this configuration, a clear disagreement between measurement and simulation was observed, as shown in Fig. 5.26. This stands in contrast to the much better level of agreement obtained for arc12 and arc56, illustrated in Fig. 5.18. Arc67 was therefore identified as a strong candidate for the origin of the observed discrepancy.

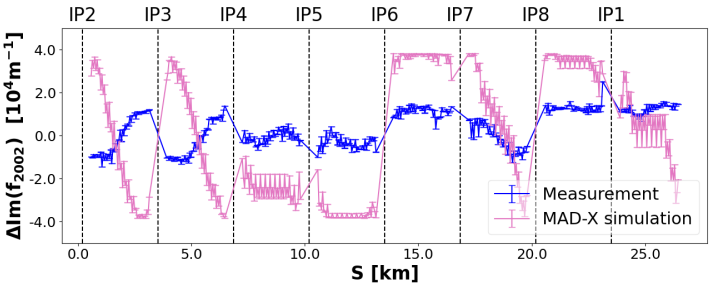


Figure 5.26: Comparison of the measured (dark blue) and simulated (pink) shift to the imaginary part of the f_{2002} RDT vs s caused by the MO in arc67 being powered at 10 m^{-4} for LHCB1.

Having confirmed that one of the potentially problematic circuits identified via NL SbS did indeed show a problem in dedicated conventional RDT measurements, this demonstrates that the NL SbS method can provide a promising approach for identifying unexpected local high-order errors, improving optics diagnosis.

Future studies should also include an investigation of arc78, together with a more systematic characterisation of the responses of the individual octupole circuits in these arcs. Where feasible, repeating the magnetic measurements of these circuits would provide an independent means of validating the observed behaviour and further assessing the origin of the identified discrepancies.

5.5 Conclusion

The response of a key MO-driven resonance f_{2002} was investigated for the first time at injection energy, motivated by the implementation of a new compensation scheme targeting these resonances. This study revealed a substantial discrepancy between measurement and the magnetic model. A comprehensive series of cross-checks of the measurement and analysis procedure established that the inconsistency was genuine and could not be attributed to deficiencies in the analysis tools, inaccuracies in the optics model, or shortcomings in the optimisation procedure. Instead, the evidence pointed towards an underlying hardware-related issue in the MO circuits.

The observed behaviour is consistent with a hysteresis effect that manifests predominantly at injection energy and appears to affect only a subset of circuits. A major challenge for optics commissioning was therefore to localise the specific circuits responsible, while avoiding excessive consumption of dedicated LHC beam time. To address this, a novel NL SbS approach was explored as a means of identifying localised NL error sources.

Preliminary simulation studies demonstrated that the proposed method is capable of identifying anomalous behaviour in individual MO circuits. The technique was subsequently applied to existing LHC measurement data, identifying the MO circuits in arcs 67 and 78 as potential sources of discrepancy. Dedicated follow-up measurements based on conventional RDT analysis confirmed that the arc67 circuit exhibits significantly poorer agreement with the magnetic model than expected. Future work should therefore include a systematic investigation of the arc78 circuit, together with a comprehensive assessment of the response of each individual octupole circuit within these arcs.

CHAPTER 6

NOVEL CORRECTION OF HIGHER-ORDER RDTS

The LHC is particularly sensitive to optics errors originating in the IRs, since the triplet quadrupole magnets operate with very large β-functions in order to achieve the required β^* at the IPs. This behaviour is clearly illustrated in Fig. 6.1, where the vertical plane, corresponding to the smaller β^* in this flat-optics configuration, exhibits substantially larger β-functions in the triplet quadrupoles ($Q1$-$Q3$), and the separation dipoles ($D1$ and $D2$).

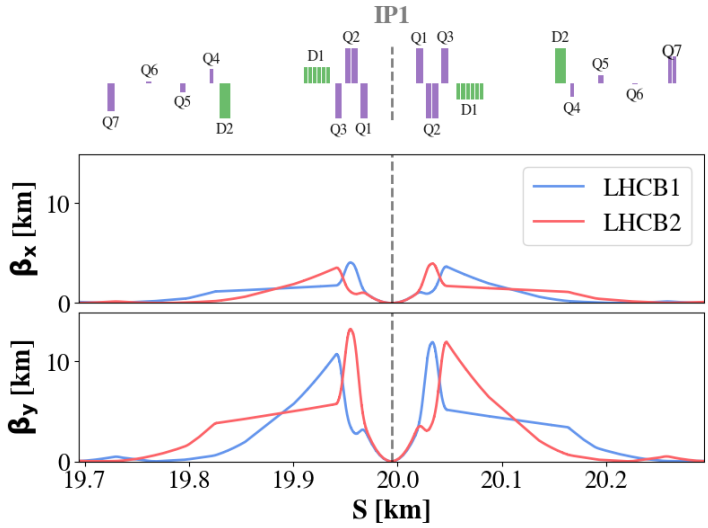


Figure 6.1: Simulated β-function vs s for LHCB1 (blue) and LHCB2 (red) in the ATLAS IR, with β^* = 0.6/0.18 m (horizontal/vertical plane) (shown also as Fig. 3.8).

This sensitivity will be further amplified in the HL-LHC, where β^* will be reduced even further, resulting in correspondingly larger β-functions in the IRs. Consequently, the HL-LHC upgrade will target the correction of optics errors up to the dodecapolar order [94, 129].

Control of NL errors in experimental insertions has long been recognised as an important aspect of collider design and operation. During the design phase of the Tevatron, potential issues arising from such non-linearities were already identified and studied [84]. At RHIC, efforts were made to mitigate these effects through local correction schemes in the IRs [130, 131]. Similarly, SuperKEKB has incorporated dedicated IR corrector magnets to address NL field errors [132–134].

In the LHC, commissioning procedures already include corrections for low-order non-linearities, notably sextupole and normal octupole components [98]. Looking towards future accelerators, the FCC-hh design anticipates the need for dedicated correction strategies to manage such effects [135]. For the HL-LHC upgrade, studies have predominantly been carried out in simulation, focusing on optimisation of the DA in the presence of NL errors [136, 137].

Given the successful application of NL RDT measurements at both injection and top energy, attention has subsequently turned to investigating whether these techniques can be extended to the use of higher-order RDTs. However, the direct measurement of such terms, particularly those of decapolar order and above, is extremely challenging.

This limitation motivates the exploration of alternative correction strategies beyond the approach presented in Chapter 4, which relied on simulating the response of a given RDT to one or more corrector magnets and subsequently validating the predicted response through beam-based measurements.

In this chapter, a new method is introduced for correcting higher-order magnetic errors through their feed-down contributions to lower-order NL RDTs. The approach exploits controlled variations of the crossing-angle (see Fig. 6.2, where θ denotes the crossing-angle formed between the reference orbit and the applied crossing scheme) in order to induce feed-down effects. The resulting variation of the RDT as a function of crossing-angle is then quantified using beam-based measurements.

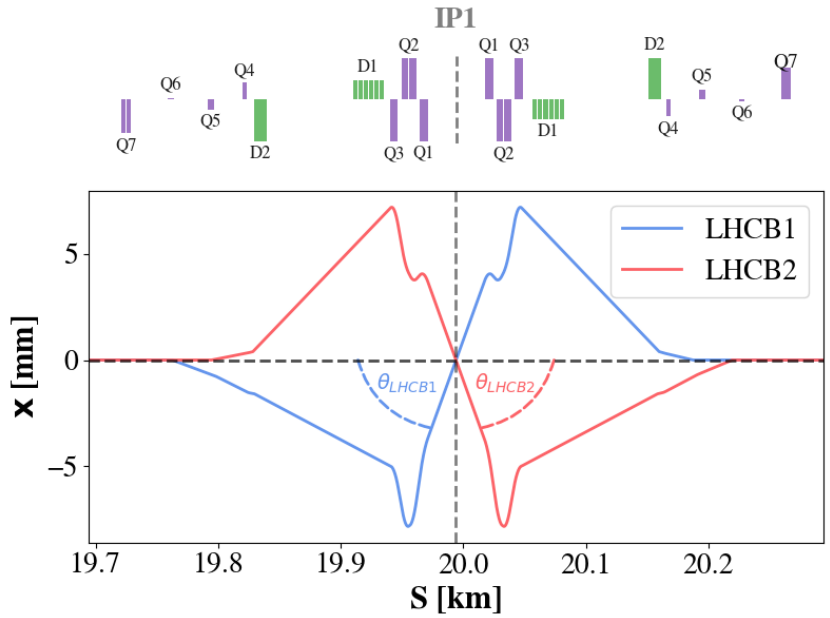


Figure 6.2: Simulated horizontal trajectories x vs s for LHCB1 (blue) and LHCB2 (red) through the ATLAS IR, with the crossing-angle (θ) indicated. Magnet layout as in Fig. 3.10.

The measured response can subsequently be compared with simulation. To determine a correction, simulations are performed where the relevant corrector magnets from the IR corrector package, shown for the LHC in Fig. 6.3 and for the HL-LHC in Fig. 6.4, are powered with the same strength for two different crossing-angle configurations. The resulting change to the RDT as a function of the crossing-angle and corrector powering can therefore be determined in simulation. By matching the simulated corrector response to the measured RDT dependence on crossing-angle, this comparison allows the identification of the corrector strengths required to compensate the higher-order errors.

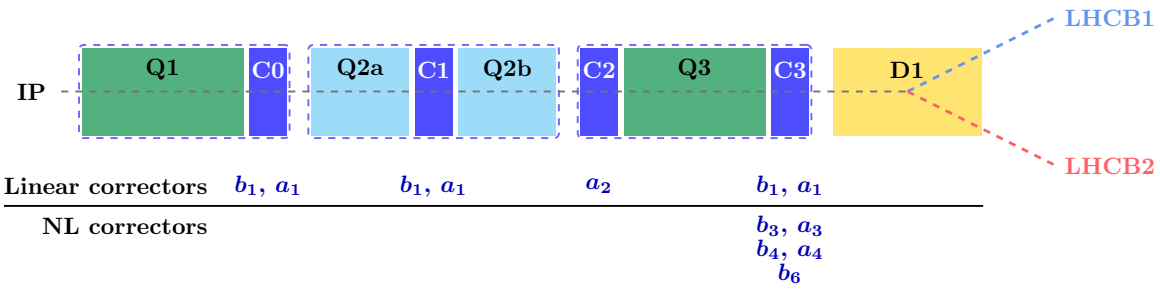


Figure 6.3: Schematic layout of an LHC experimental insertion - right side of IP1 shown - including the triplet corrector package, as in Fig. 3.9.

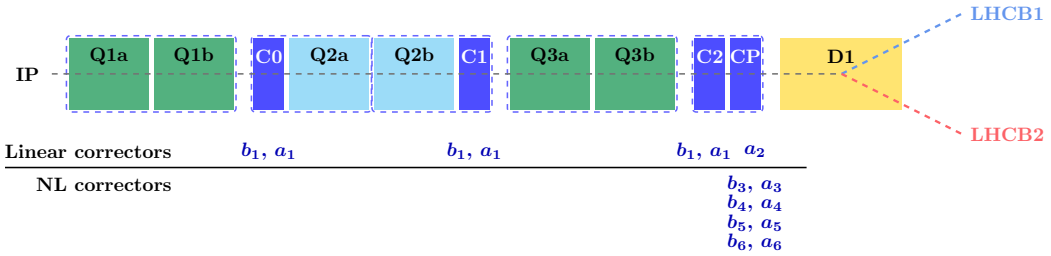


Figure 6.4: Schematic layout of an HL-LHC experimental insertion - right side of IP1 shown - where $Q1a/b$, $Q2a/b$, and $Q3a/b$ are the triplet quadrupoles, and $C0$-$C3$ the corrector packages, with the multipole orders each can correct shown in the lower part of the figure. $D1$ is the separation dipole.

6.1 Feed-down

In order to understand the mechanism underlying this correction scheme, it is first necessary to introduce the concept of feed-down. Feed-down occurs when the beam traverses a magnet with a transverse offset, which may arise either from magnet misalignments or from a displaced closed orbit. Under these conditions, the higher-order magnetic field components experienced by the beam give rise to effective lower-order field components [138].

To provide a mathematical description, the multipole Hamiltonian may be decomposed into its normal (N_n) and skew (S_n) contributions using Eq. (2.6) and Eq. (2.38). Denoting the normal multipole strength by K_n and the skew component by J_n, the Hamiltonian can be written as

$$\begin{aligned} N_n(x,y) &= \Re\left[\frac{1}{n!}K_n(x+iy)^n\right], \\ S_n(x,y) &= \Im\left[-\frac{1}{n!}J_n(x+iy)^n\right]. \end{aligned} \tag{6.1}$$

For the specific case of a normal dodecapole ($n=6$), this expression becomes

$$N_6(x,y) = \Re\left[\frac{1}{6!}K_6(x+iy)^6\right]. \tag{6.2}$$

If a transverse horizontal displacement Δx is introduced, the Hamiltonian is evaluated at the shifted coordinates such that $N_6(x,y)$ becomes $N_6(x+\Delta x, y)$. Expanding this expression in a Taylor series about $\Delta x = 0$ yields

$$N_6(x+\Delta x, y) = N_6(x,y) + \Delta x\, N_6'(x) + \frac{(\Delta x)^2}{2} N_6''(x) + \ldots \quad . \tag{6.3}$$

Substituting the explicit form of N_6 gives

$$\begin{aligned} N_6(x+\Delta x, y) = N_6(x,y) &+ \Delta x \left(\Re \left[\frac{1}{5!} K_6 (x+iy)^5 \right] \right) \\ &+ \frac{(\Delta x)^2}{2} \left(\Re \left[\frac{1}{4!} K_6 (x+iy)^4 \right] \right) + \ldots, \end{aligned} \tag{6.4}$$

from which it can be seen that the offset generates lower-order contributions. In particular, the term lowlighted in grey in Eq. (6.4) corresponds to an effective normal decapolar component, while the subsequent term, shown in lighter grey, corresponds to an octupolar-like contribution.

A similar mechanism occurs for skew multipoles. For example, in the case of a skew-octupole, the introduction of a horizontal offset Δx leads to

$$\begin{aligned} S_4(x+\Delta x, y) = S_4(x,y) &+ \Delta x \left[\Im \left[-\frac{1}{3!} J_4 (x+iy)^3 \right] \right] \\ &+ \frac{(\Delta x)^2}{2} \left[\Im \left[-\frac{1}{2!} J_4 (x+iy)^2 \right] \right] + \ldots, \end{aligned} \tag{6.5}$$

where the first additional term corresponds to an effective skew sextupolar component, followed by a skew quadrupolar-like contribution arising from the horizontal offset.

This feed-down effect can therefore be exploited deliberately, through controlled orbit distortions, i.e. the introduction of crossing-angles, as a means of measuring higher-order errors. The use of feed-down effects on the linear optics as a correction strategy is well established. At RHIC, tune shifts induced by octupole and sextupole errors were mitigated [86]. In the LHC, an analogous approach is routinely applied at top energy, particularly at the end-of-squeeze, where sextupole and normal octupole errors are compensated through their feed-down to the tune [97, 98]. There have also been attempts to apply corrections using feed-down to coupling [97].

This methodology has also been extended to higher-order errors by compensating the corresponding feed-down terms in the amplitude detuning [139, 140]. While effective, the implementation of amplitude detuning corrections is operationally demanding and requires significant accelerator time. Furthermore, it only works for certain multipoles.

Since feed-down from higher-order field errors generates effective lower-order field components, it also enhances the corresponding lower-order RDTs. Since sextupolar RDTs, as mentioned

in a previous chapter, are comparatively easy to measure, this effect could be exploited as a practical way to diagnose and correct higher-order NL errors when direct observation of that multipole term is not possible.

To test this hypothesis, historical beam data from the 2024 ion commissioning period was re-examined. This data had originally been acquired to study the variation of the coupling as a function of the crossing-angle. However, the study was not pursued extensively at the time because it was not sufficiently effective.

During these measurements, the vertical crossing-angles in IP2 were varied, while conventional ACD excitations were applied for different crossing configurations. A subsequent reanalysis of the data revealed that both the f_{1020} and f_{0030} RDTs were clearly visible in the frequency spectrum, as illustrated in Fig. 6.5. This observation enabled the evaluation of the variation of the RDTs as a function of the crossing-angle.

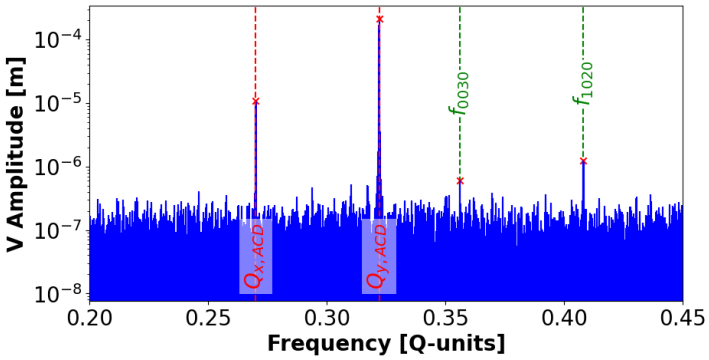


Figure 6.5: Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.31$.

In the following discussion, when referring to the shifts in the real (ΔRe), imaginary (ΔIm), and amplitude (ΔAmp) components of the RDT, the quantities are defined according to

$$\begin{aligned} \Delta\mathrm{Re}(f_{jklm}) &= \Re(f_{jklm}(\theta)) - \Re(f_{jklm}(\theta_{\mathrm{ref}})), \\ \Delta\mathrm{Im}(f_{jklm}) &= \Im(f_{jklm}(\theta)) - \Im(f_{jklm}(\theta_{\mathrm{ref}})), \\ \Delta\mathrm{Amp}(f_{jklm}) &= \sqrt{(\Delta\mathrm{Re}(f_{jklm}))^2 + (\Delta\mathrm{Im}(f_{jklm}))^2}, \end{aligned} \tag{6.6}$$

where $f_{jklm}(\theta)$ denotes the f_{jklm} RDT evaluated for the crossing-angle, θ, while $f_{jklm}(\theta_{\mathrm{ref}})$ corresponds to the value of the RDT at the reference crossing-angle, θ_{ref}.

Examples of the resulting amplitude variations are presented in Fig. 6.6, which shows the magnitude of the shift in the amplitude of the f_{0030} RDT, as defined in Eq. (6.6), for crossing-angles of 25 μrad and 170 μrad relative to a reference crossing-angle of 0 μrad.

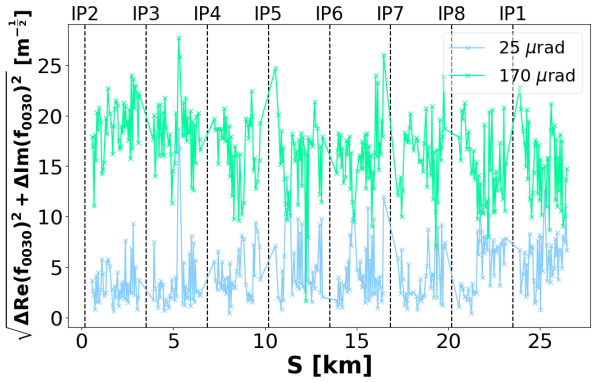


Figure 6.6: $\Delta\mathrm{Amp}(f_{0030})$ RDT vs s for crossing-angles of 25 (blue) and 170 (green) μrad from a reference crossing-angle of 0μrad. Lines are included only to guide the eye.

To enhance the clarity of the observation, and to assess whether the dependence was linear or of higher order, the analysis focused on the ring-averaged RDT amplitude shift (ΔAmp), where

$$\langle\Delta\mathrm{Amp}\rangle = \frac{1}{N_{bpm}} \sum_{n=1}^{N_{bpm}} \Delta\mathrm{Amp}. \tag{6.7}$$

The overall analysis using this average shift to the amplitude is displayed in Fig. 6.7.

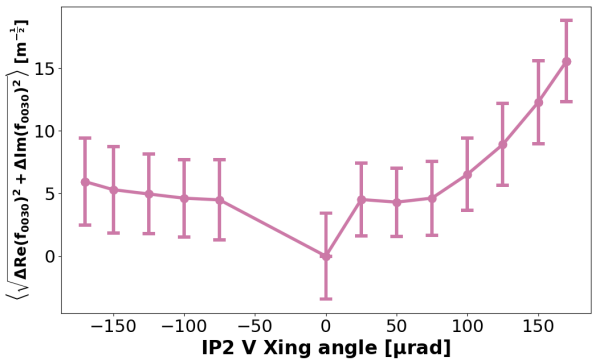


Figure 6.7: $\langle\Delta\mathrm{Amp}\rangle$ vs vertical IP2 crossing-angles for LHCB1, the line is included only as a visual guide to see the trend.

In light of the observed variation with crossing-angle, the implementation of a correction scheme was considered feasible. The principle of this approach is to use the quantities ΔRe and ΔIm to evaluate the following partial derivatives:

$$\frac{\partial \mathrm{Re}(f_{jklm})}{\partial \theta}, \quad \frac{\partial \mathrm{Im}(f_{jklm})}{\partial \theta}. \tag{6.8}$$

These derivatives quantify the sensitivity of the RDT components to variations in the crossing-angle. An analogous calculation is then performed for the simulated response. In this case, the two configurations correspond to simulations carried out at different crossing-angles, while a dedicated corrector magnet is powered, at the same strength in both, to reproduce the effect of a higher-order magnetic field error.

Figure 6.8 illustrates the principle of this procedure. The simulated response, shown in black, has been matched to the measured data for LHCB1 (blue) and LHCB2 (red) by adjusting the strengths of the corresponding corrector circuits. Note that the measurement data used in this example were obtained with a vertical crossing-angle in IP5, and the analysis focuses on the f_{1020} RDT.

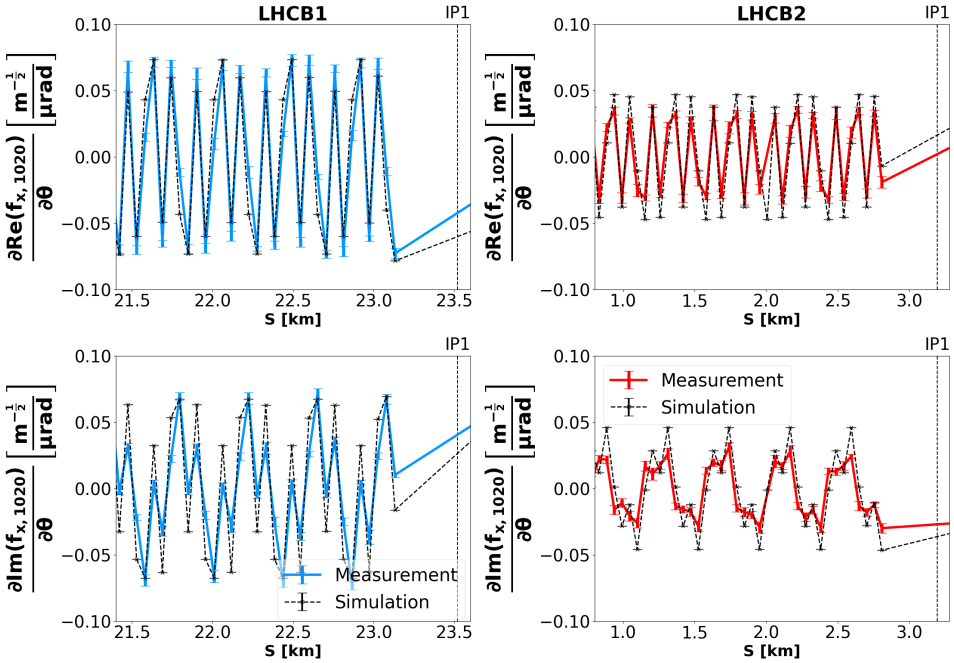


Figure 6.8: Demonstration of the matching of the measurements of both beams with the simulated response of specific correctors. This is shown through the depiction of the partial derivatives of the real and imaginary parts of the RDT with respect to the θ vs s for both simulation (black) and measurement (blue for LHCB1, red for LHCB2). Lines are included only to guide the eye.

6.2 Creation of correction process using GUI

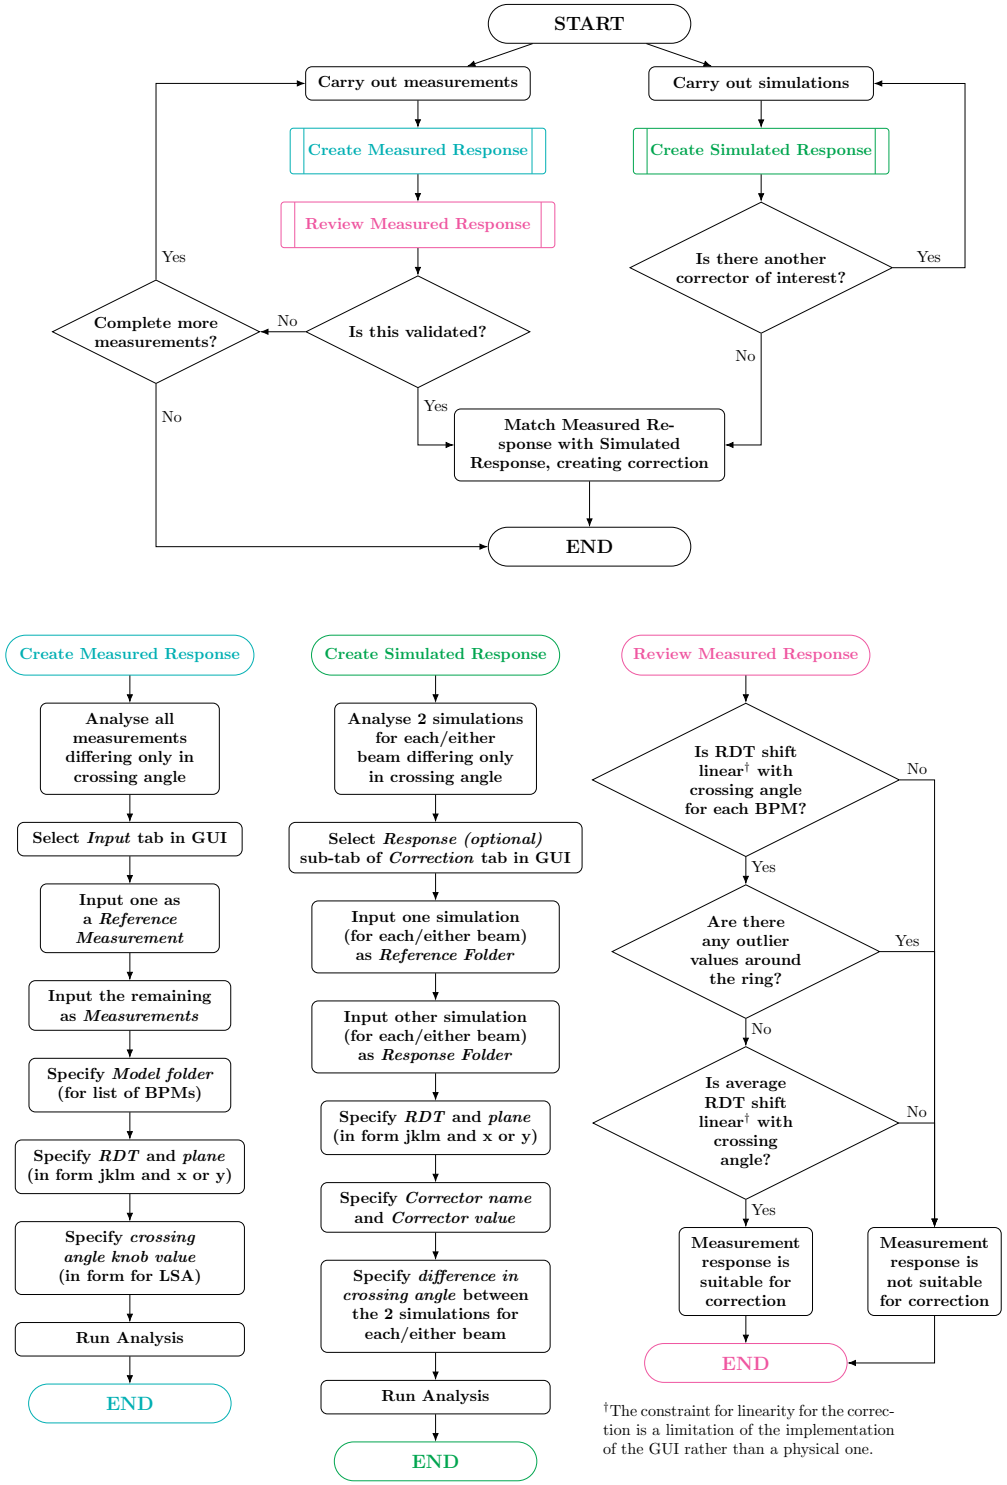


Figure 6.9: Workflow for the creation of a RDT feed-down-based correction of higher-order errors using the GUI, where the links are to the relevant figure in this section.

To support the practical application of this correction strategy during commissioning, a dedicated graphical user interface (GUI) was developed. Its purpose was to enable corrections to be implemented and evaluated rapidly and directly from the control room, therefore avoiding the need for extensive post-measurement analysis or the continuous presence of an expert during beam operation.

Initially, RDT measurements are performed at several different crossing-angles using the ACD, as described in the preceding section. During these measurements, the tune was kept constant while the crossing-angle was varied, in order to avoid modifications to the RDT arising from changes in the proximity to resonances. For each crossing-angle configuration, measurements were carried out on multiple bunches and, in some cases, with multiple kick excitations. The analysis of the RDTs for each individual crossing-angle was then performed using the same analysis framework employed in the preceding chapters [141, 142].

The RDT feed-down (RDTFD) GUI was developed to perform the required analysis directly on the results produced by the measurement procedure described above. One measurement is designated as the reference dataset, while the remaining measurements are used for comparison. The RDTFD GUI provides two distinct input modes, allowing it to operate either with measured data or with simulated measurements. The latter option facilitates the testing and validation of the analysis functionality in the absence of beam.

Within the *Input* tab (the initial page, see Fig. 6.10), the RDTFD GUI evaluates the shift in the real and imaginary components of a selected RDT, as defined in Eq. (6.6), for each specified crossing-angle. For measurement data, the corresponding crossing-angles are automatically retrieved from the LHC Software Architecture (LSA) database [143] using the timestamp contained in the listed TbT file name. For simulation studies, the crossing-angles are instead obtained from a user-defined properties file.

The results may be examined in the *Validation* tab (see Fig. 6.11), which is organised into three sub-tabs. The first sub-tab displays the data at the level of individual BPM (see Fig. 6.11), where the user can select the order of the fit to identify the dominant order of the feed-down contribution.

The second sub-tab illustrates the evolution of the RDT shift as a function of s (see Fig. 6.12).

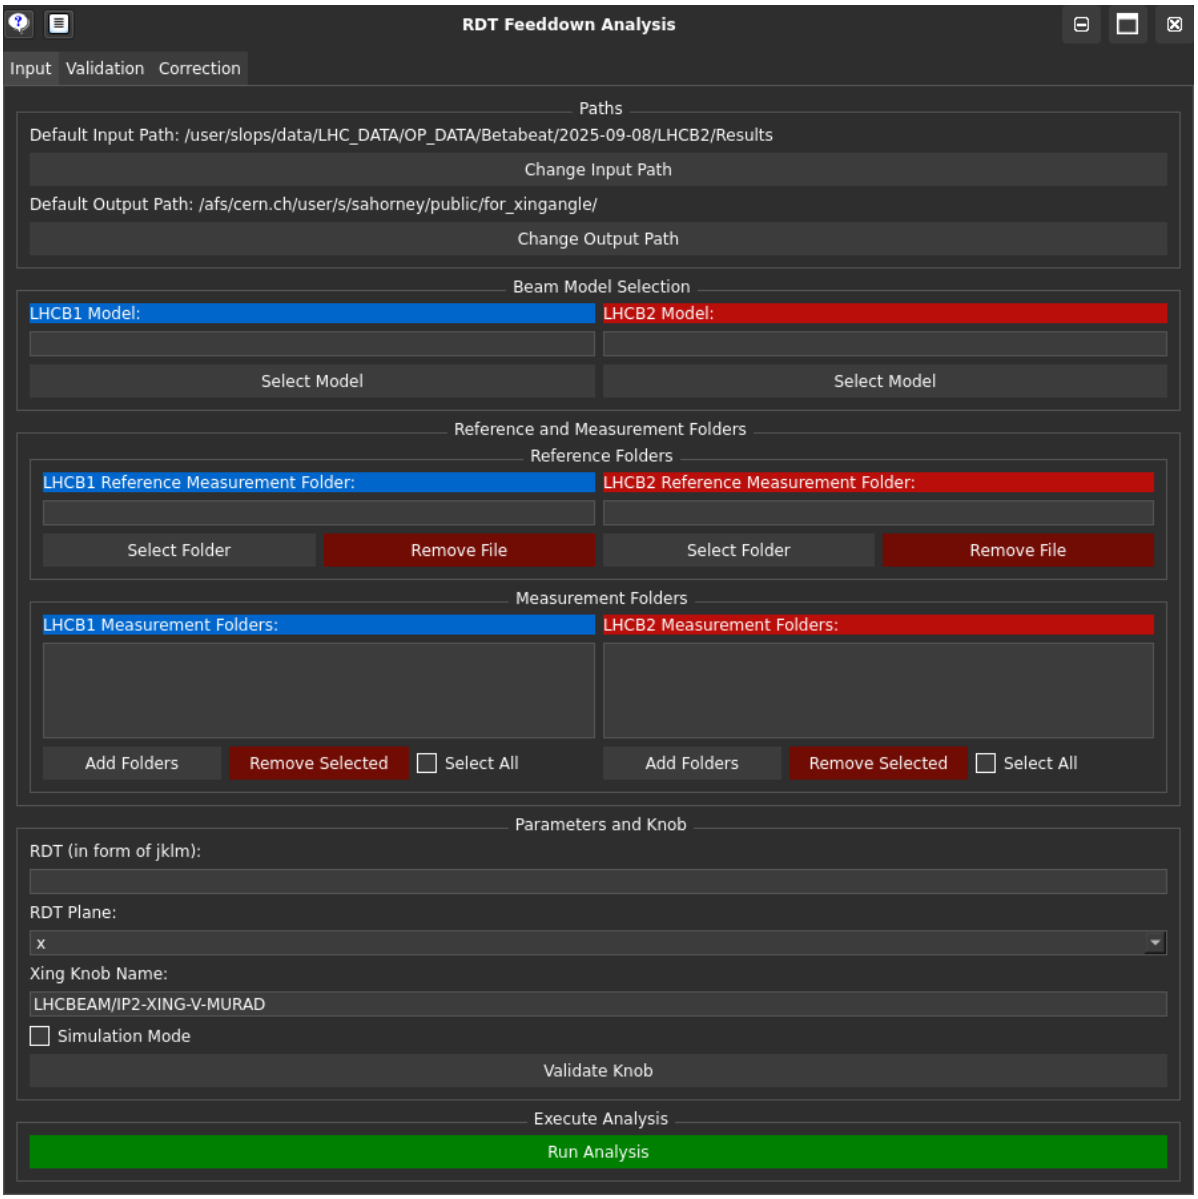


Figure 6.10: Input tab of the RDTFD GUI, where measurements are entered to evaluate the RDT shift as a function of crossing-angle.

This information can be visualised either in terms of the amplitude, calculated from the magnitude of the real and imaginary components, or directly through the real or imaginary component itself.

The third sub-tab presents the ring-averaged RDT shift for each crossing-angle (see Fig. 6.13), referenced to a selected crossing-angle. In addition, this view provides the fitted partial derivatives of the ΔRe and ΔIm RDT components with respect to crossing-angle (Eq. (6.8)) as a function of longitudinal position.

In the *Graph* sub-tab of the *Correction* tab (see Fig. 6.14), the measured partial derivatives defined in Eq. (6.8) can be directly compared with the corresponding derivatives obtained from simulation. To determine these simulated derivatives, which are evaluated separately for each corrector circuit, simulations are performed by applying the same chosen corrector strength with two distinct crossing-angle configurations. This procedure allows the effect of the crossing-angle variation to be isolated from the RDT shift arising directly from the powering of the corrector magnets.

The resulting simulation outputs are subsequently analysed using the same analysis framework employed for the beam-based measurements. One configuration is treated as the reference case, while the second is considered as the response. The corresponding crossing-angle difference, together with the applied corrector strength, must be specified in the *Response* sub-tab of the *Correction* tab (see Fig. 6.15). Through this procedure, the sensitivity of each corrector circuit to variations in the RDT as a function of the crossing-angle can be quantitatively determined.

The magnet circuit settings required for the correction can subsequently be determined within the *Graph* sub-tab by matching the measured and simulated partial derivatives. The correction corresponds to applying the matched corrector strengths with the opposite sign.

It should be noted that this procedure relies on the assumption that the response of the RDT to the applied corrector strength is linear, and that the measured response exhibits the same linear behaviour. In principle, the method could be extended to account for quadratic or higher-order feed-down contributions by evaluating and matching higher-order derivatives with respect to crossing-angle. However, no such functionality has yet been implemented in the GUI.

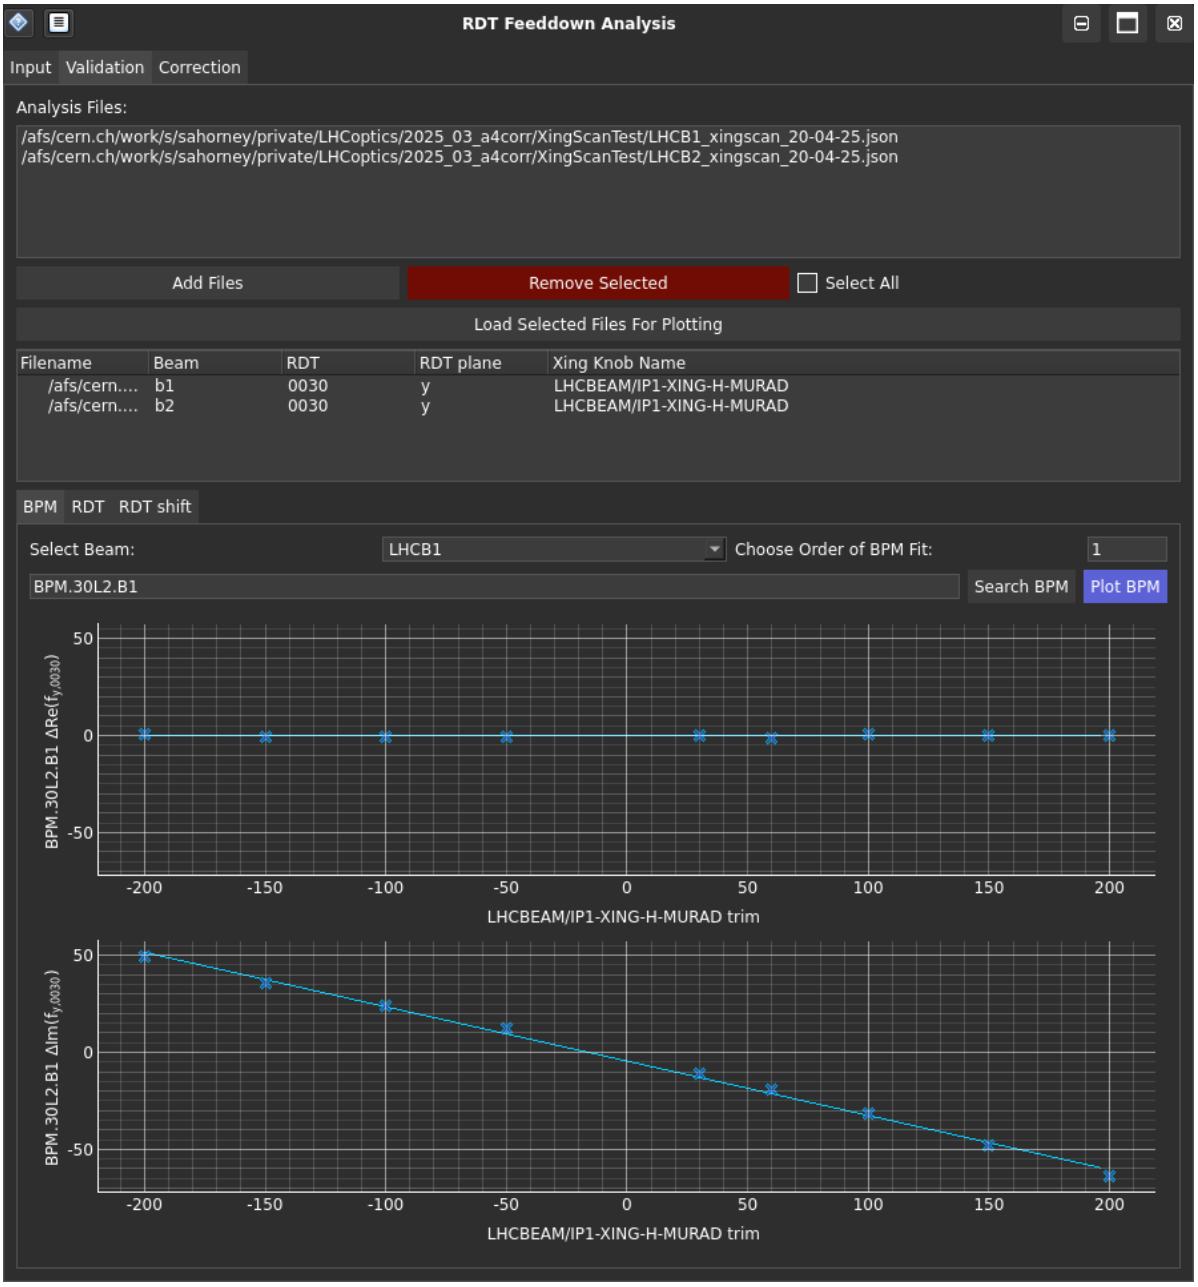


Figure 6.11: Validation tab of the RDTFD GUI where the measured response to crossing-angle can be displayed. This shows the BPM sub-tab depicting ΔRe and ΔIm vs θ at the BPM level.

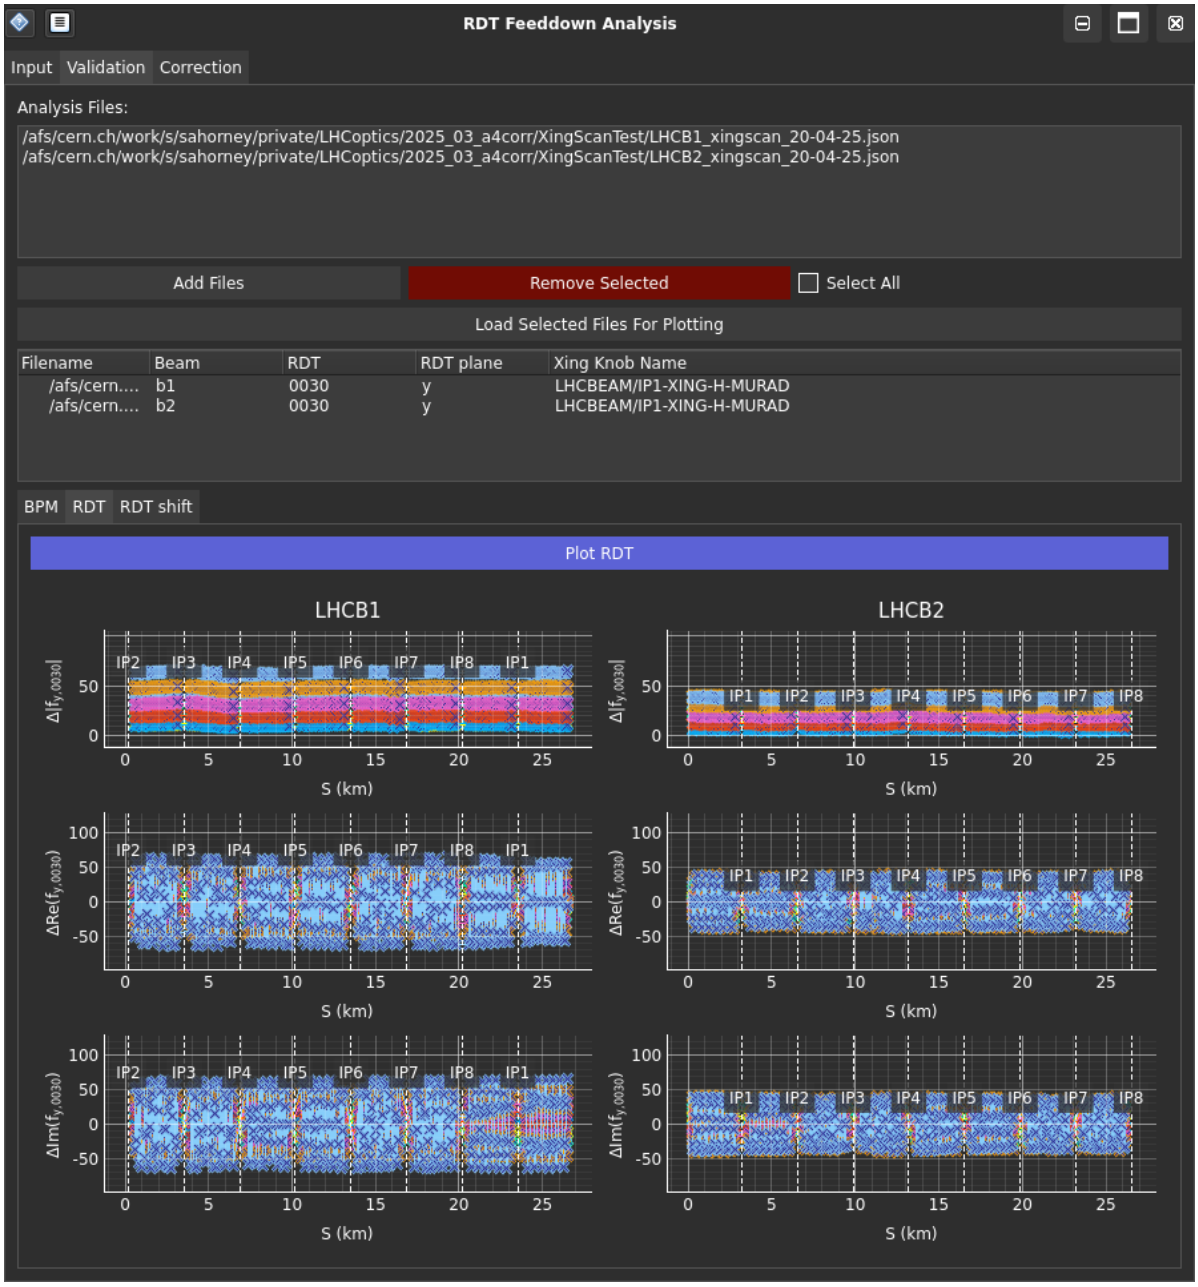


Figure 6.12: Validation tab of the RDTFD GUI where the measured response to crossing-angle can be depicted. This shows the RDT sub-tab depicting the ΔAmp, ΔRe, and ΔIm vs s.



Figure 6.13: Validation tab of the RDTFD GUI where the measured response to crossing-angle can be depicted. This shows the RDT shift sub-tab depicting the ring-averaged ΔAmp vs crossing-angle, as well as $\frac{\partial \mathrm{Re}}{\partial \theta}$ and $\frac{\partial \mathrm{Im}}{\partial \theta}$ vs s.



Figure 6.14: Correction tab of the RDTFD GUI, with the Graph sub-tab selected, where measurements can be matched with simulated response.

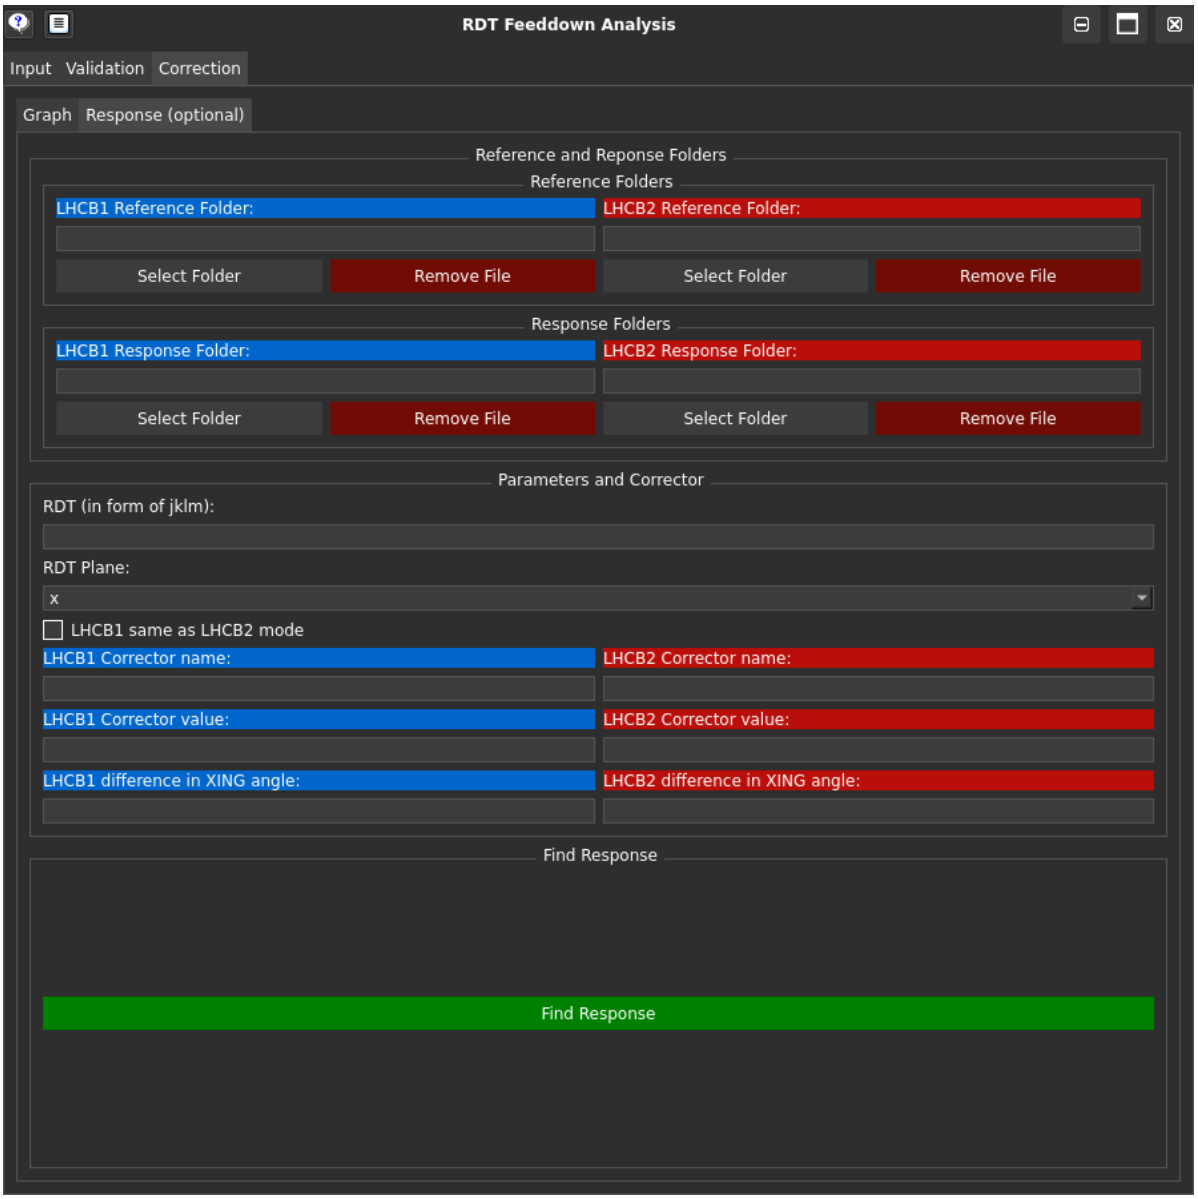


Figure 6.15: Correction tab of the RDTFD GUI, with the Response sub-tab selected, where the simulated response to crossing-angle can be calculated.

6.3 Beam-based validation of correction method

This method was successfully demonstrated during the 2025 commissioning period. Corrections of the skew-octupolar order are of particular relevance for LHC commissioning, as these components are known to influence the DA, and may limit the ability to excite the beam using the ACD [110]. However, skew-octupole RDTs are extremely challenging to measure directly, and a reliable measurement procedure has not yet been consistently achieved.

Previous attempts have explored the use of second-order feed-down to coupling as an indirect diagnostic [97]. These approaches, however, proved to be unreliable and inherently ambiguous, as they could not uniquely distinguish whether the correction should be applied on the left or right side of the IR. Direct corrections based on RDT measurements have also been attempted and have occasionally yielded successful results [99, 110]. Nevertheless, skew-octupole RDTs have not been measured with sufficient reliability at low β^*.

For this reason, the skew-octupole component represents a suitable candidate for the present method, since sextupolar-order RDTs, to which the skew-octupole errors can feed-down, can be measured much more reliably.

Furthermore, during the 2024 studies it was observed that the collimator hierarchy was affected by the presence of the $3Q_y$ resonance [74]. As indicated by Eq. (6.5), a horizontal orbit offset in the presence of a skew-octupole error generates a feed-down contribution corresponding to a skew sextupole component, which would include the $3Q_y$ resonance. Consequently, for HL-LHC operation, the direct correction of this feed-down contribution is of particular interest.

To investigate this effect, measurements of the $3Q_y$ resonance were carried out during the 2025 commissioning campaign at the end-of-squeeze, as a function of both the vertical crossing-angle in IP5 and the horizontal crossing-angle in IP1. As the measurements obtained with the IP1 crossing-angle variation yielded clearer results, only these are presented in the following discussion. An example of the resulting TbT frequency spectrum is shown in Fig. 6.16, in which the f_{0030} spectral line is clearly observed.

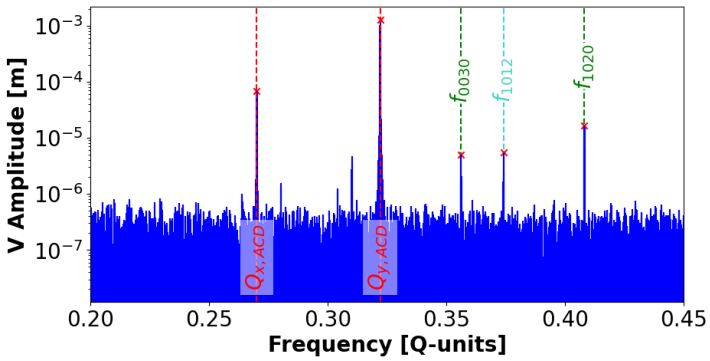


Figure 6.16: Vertical TbT frequency spectrum, where green denotes 3rd order and cyan the 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.31$.

The data was subsequently analysed using the RDTFD GUI described in the previous section. First, Fig. 6.17 was used as a consistency check to verify that the RDT data had been properly cleaned of any problematic BPMs that could otherwise bias the averaged quantities.

Figures 6.18 and 6.19 confirm a linear dependence of the RDT shift on the crossing-angle for LHCB1. This behaviour was observed both at the level of individual BPM measurements, and in the ring-averaged response, indicating that the effect is dominated by the expected skew-octupole feed-down contribution.

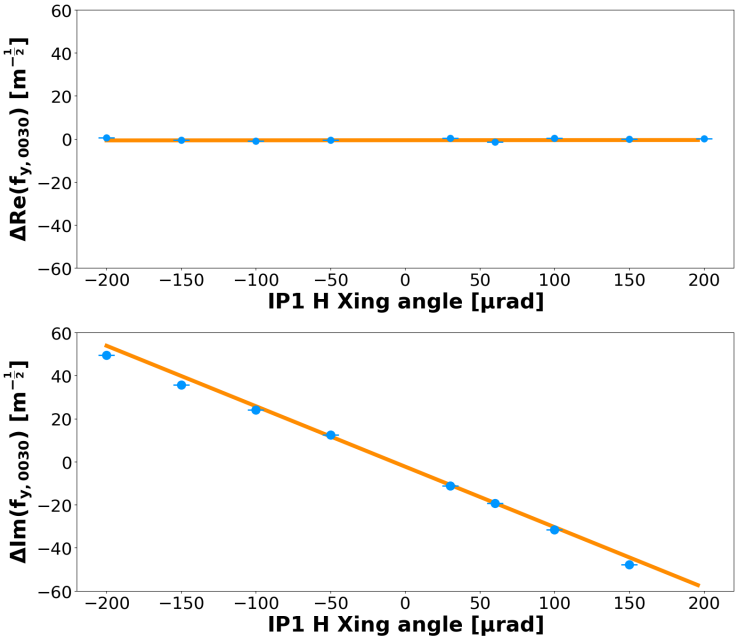


Figure 6.18: The ΔIm/ΔRe of the f_{0030} RDT with respect to a horizontal IP1 crossing orbit bump for the BPM.30L2.B1, where the line (orange) is the linear fit of the data.

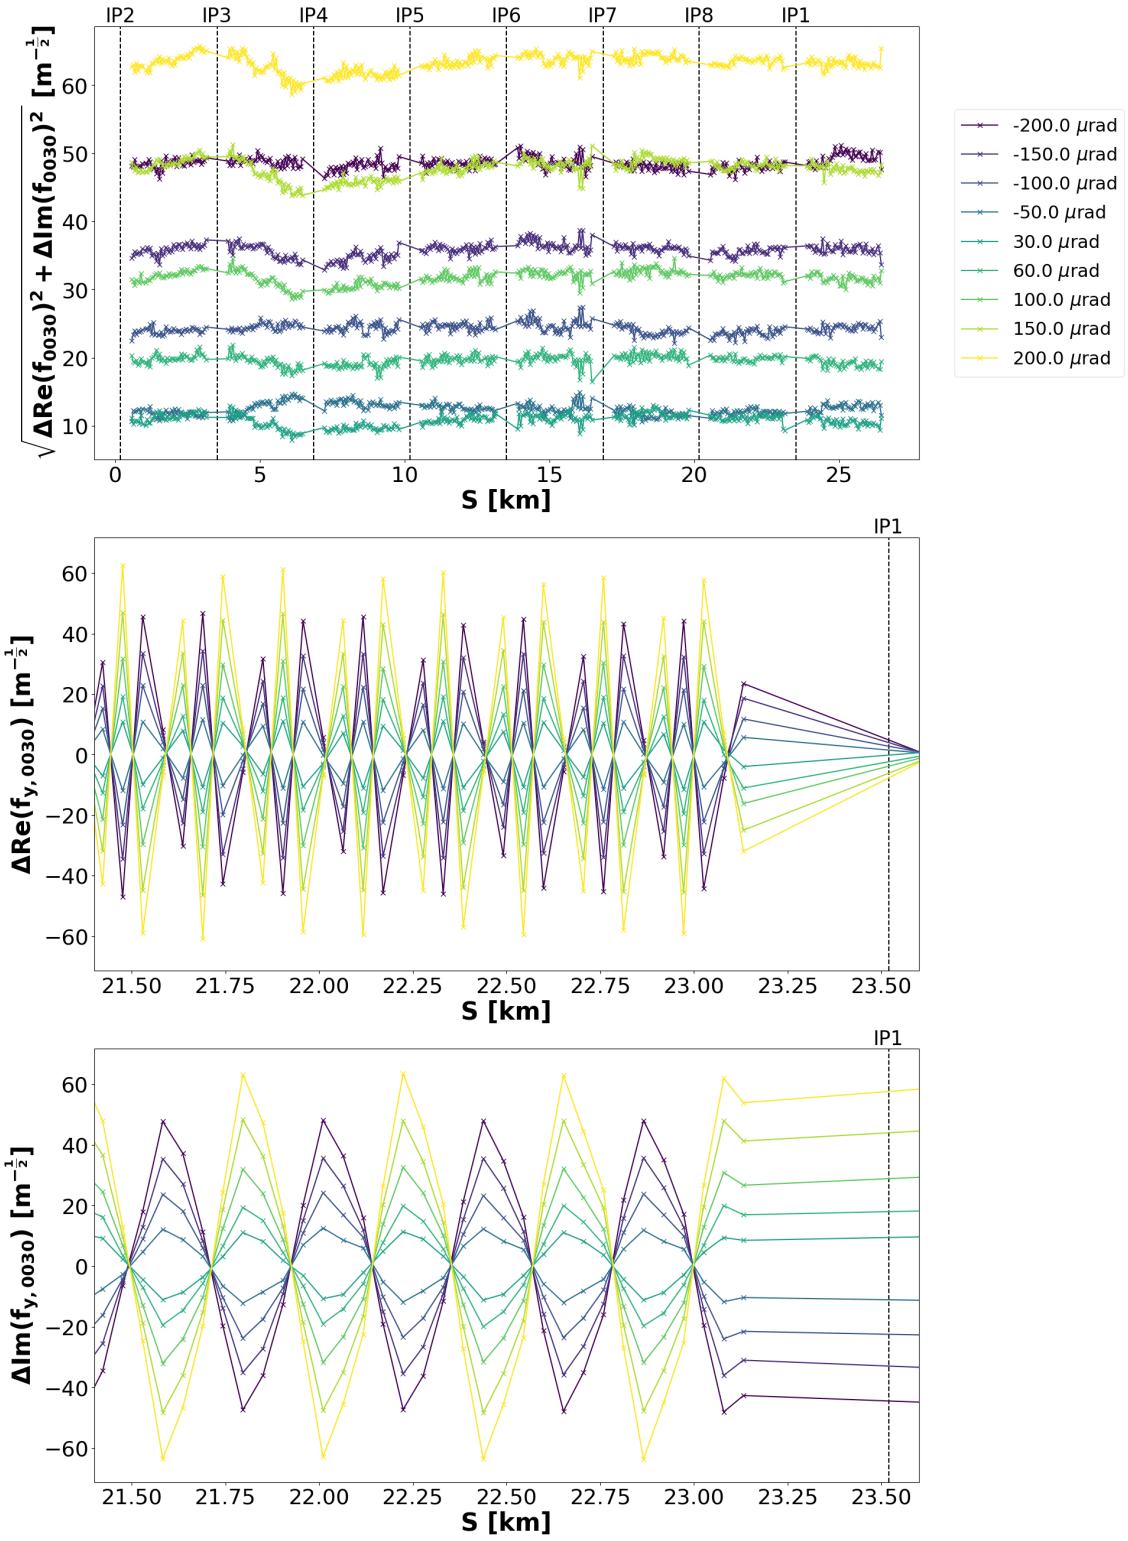


Figure 6.17: The ΔAmp/ΔIm/ΔRe vs s for each horizontal IP1 crossing orbit bump each in a separate colour for LHCB1. Lines are included only to guide the eye.

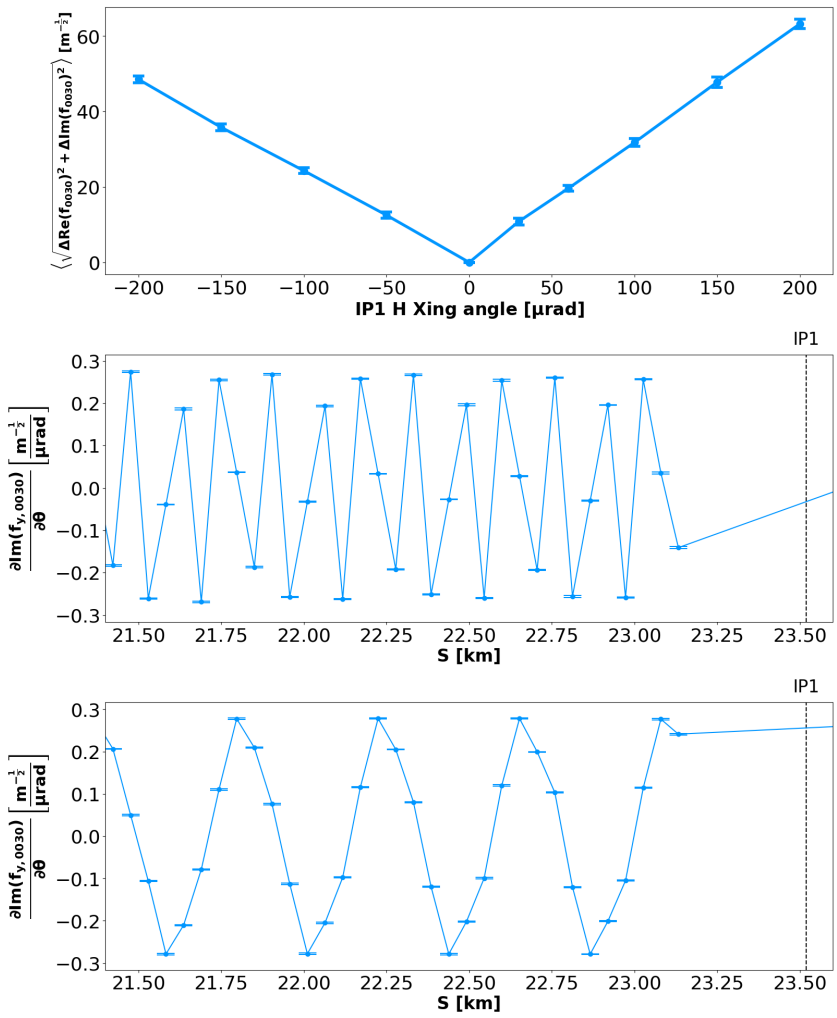


Figure 6.19: The ⟨ΔAmplitude⟩ vs θ, as well as the real and imaginary partial derivatives vs s, of the f_{0030} RDT with respect to a horizontal IP1 crossing orbit bump for LHCB1. Lines are included only to guide the eye.

Since only one operational IP1 skew-octupole corrector was available in IP1 (the other has been broken since the start of the LHC Run I), located on the right side of the IP, it is well-known that a simultaneous correction of the a_4 component in both beams is not possible under these conditions [99]. In such a configuration, one must either prioritise the correction of a single beam, or adopt a compromise solution that partially corrects both.

For the present test, the strategy was to optimise the correction for LHCB1. It should be noted that when correctors located on both sides of the IP are available, a simultaneous correction of both beams becomes feasible, as illustrated in Fig. 6.8.

The correction settings were determined using the skew-octupole corrector MCOSX.3R1 (see the third corrector package in Fig. 6.3) in IR1. The matched measured response is shown in Fig. 6.20, from which the appropriate setting of the corrector MCOSX.3R1 was determined to be $K_4 = 1.2\ \mathrm{m}^{-4}$.

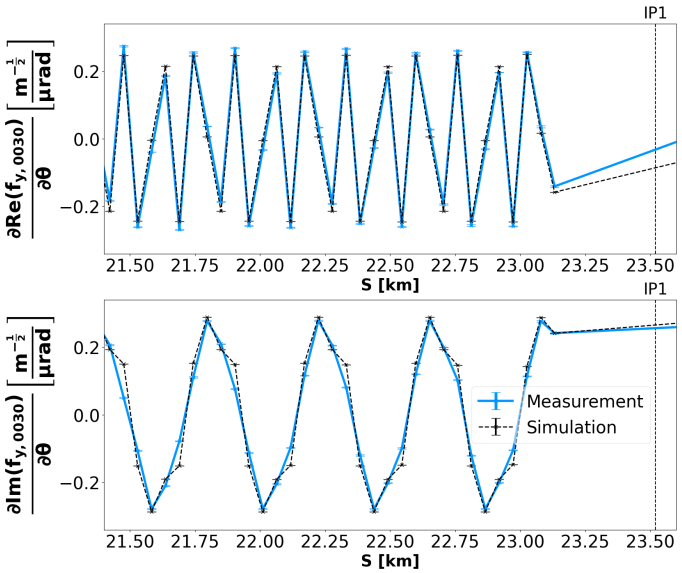


Figure 6.20: The matching of the simulated corrector response (black) to the measured (blue) response, using the partial derivatives of the real and imaginary parts of the f_{0030} RDT, vs s for LHCB1. Lines are included only to guide the eye.

Figure 6.21 shows the measured behaviour before and after the correction was applied, illustrating a clear reduction of the RDT shift as a function of crossing-angle for LHCB1. The applied correction did not significantly degrade the behaviour observed in LHCB2 as seen in Fig. 6.22.

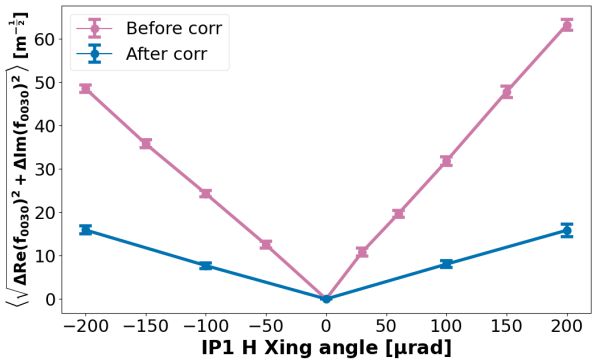


Figure 6.21: $\langle\Delta\mathrm{Amp}\rangle$ of the f_{0030} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB1 with (blue) and without (pink) correction, the line is included only as a visual guide to see the trend.

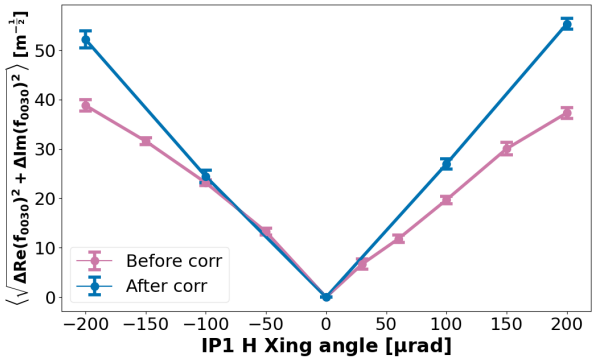


Figure 6.22: ⟨ΔAmp⟩ of the f_{0030} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2 with (blue) and without (pink) correction, the line is included only as a visual guide to see the trend.

Fortunately, the $Q_x - Q_y$ resonance, i.e. the f_{1012} RDT was also, albeit unusually, clearly visible in the measurements, meaning it was possible to assess the direct impact of the correction. Figure 6.23 compares the situation before and after the correction was applied and shows that the correction resulted in a clear reduction of the amplitude of the skew-octupolar f_{1012} RDT.

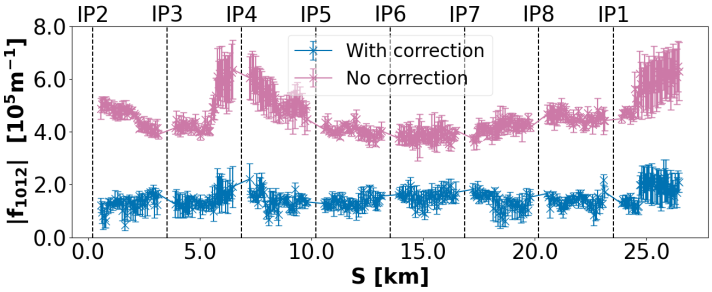


Figure 6.23: Measured amplitude of the f_{1012} RDT vs s with (blue) and without (pink) correction. Lines are included only to guide the eye.

6.4 First measurement of IR dodecapole errors

Having demonstrated the validity of the method, it was of interest to investigate its applicability to more challenging higher-order errors. In particular, the correction of dodecapolar components is of considerable importance, but remains highly difficult in practice. To date, the normal dodecapolar RDT has not been directly measured at the end-of-squeeze. The only approach considered previously relied on the feed-down of the dodecapolar errors to amplitude detuning; however, this method is operationally slow and requires substantial accelerator time.

Dedicated MD studies were therefore carried out in 2025 with the objective of measuring the dodecapolar error feed-down contribution to a decapolar RDT. For this purpose, the working point was selected to lie very close to the decapole resonance. The procedure followed was analogous to that previously employed for the a_4 feed-down to a_3, although larger kick amplitudes were applied in order to enhance the visibility of the relevant spectral lines.

Analysis of these measurements revealed that the normal decapolar resonance $-Q_x + 4Q_y$, associated with the RDT f_{0140}, was clearly visible in the TbT frequency spectra for the different crossing-angle configurations. An example of the corresponding spectrum is shown in Fig. 6.24.

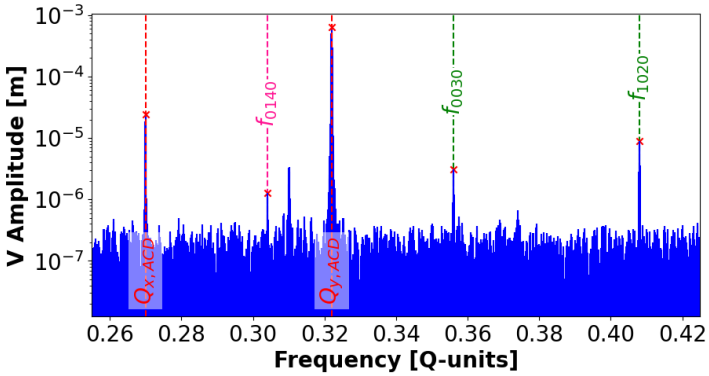


Figure 6.24: Vertical TbT frequency spectrum, where green denotes 3rd order and pink denotes 5th order RDTs, and red the ACD frequencies, with natural q_x = 0.28, q_y = 0.31.

The f_{0140} RDT was observed to vary as a function of the horizontal crossing-angle scan at IP1. This can be seen through the example crossing-angles of 60 μrad and 120 μrad shown in Fig. 6.25.

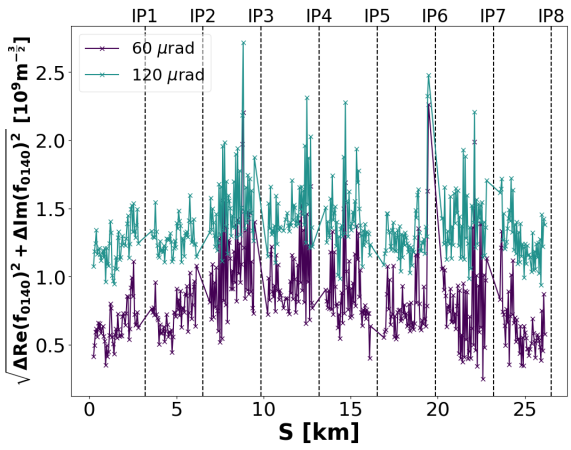


Figure 6.25: ΔAmp(f_{0140}) RDT vs s for crossing-angles of 60 (purple) and 120 (teal) μrad from a reference crossing-angle of 0 μrad. Lines are included only to guide the eye.

The approximately linear dependence, shown in Fig. 6.26, suggests that this behaviour arises from feed-down effects associated with dodecapolar field errors in the IR. This result is particularly encouraging, as it indicates that a proof-of-principle for a correction scheme relevant to the HL-LHC can already be demonstrated in the present LHC.

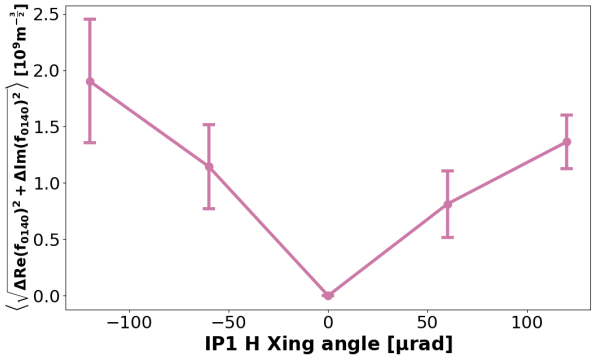


Figure 6.26: $\langle\Delta\mathrm{Amp}\rangle$ of the f_{0140} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2, the line is included only as a visual guide to see the trend. Error bars are reflective of a greater variation in the RDT around the ring.

A dedicated simulation study was performed to benchmark the measurements and to evaluate whether the inclusion of the FiDeL [43] b_6 and b_7 field errors alone could account for the observed behaviour. The results demonstrated that the b_6 component is the dominant contributor, as evidenced by the predominantly linear dependence of the RDT shift on the crossing-angle. Furthermore, the magnitude of the RDT variation was found to be highly reproducible for positive crossing-angle values, whereas the response for negative crossing-angles suggested the presence of additional higher-order effects that are not captured by the current model. An alternative explanation is that the increasing measurement uncertainty observed over the course of the measurement campaign reflects a progressive reduction in beam intensity, leading to a deterioration in measurement quality rather than a purely physical effect.

Taken together, these results indicate that the method is well suited for the correction of even higher-order feed-down effects, and shows particular promise in the context of the HL-LHC. Notably, previous attempts to characterise the b_6 component relied on measurements of amplitude detuning and required multiple dedicated MD shifts [140]. By contrast, the present approach achieved an equivalent objective using only approximately 20 minutes of beam time. This substantial reduction in operational overhead suggests that the implementation of dodecapolar

corrections could realistically become part of routine optics commissioning procedures.

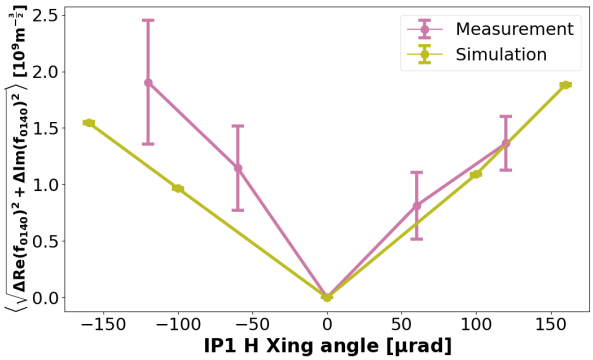


Figure 6.27: $\langle \Delta \text{Amp} \rangle$ of the f_{0140} RDT vs θ, i.e. the horizontal IP1 crossing-angle, for LHCB2 for simulation (green) and measurement (pink), the line is included only as a visual guide to see the trend.

6.5 Conclusion

A novel correction methodology targeting higher-order field errors originating from the LHC IRs was developed and subsequently validated through beam-based measurements in the accelerator in this work. The results demonstrate that this approach enables the identification and compensation of NL errors through feed-down mechanisms, without requiring the direct measurement of higher-order RDTs. Furthermore, the method is compatible with the operational time constraints of routine accelerator operation.

In particular, the skew-octupolar errors in the ATLAS experimental insertion were successfully corrected in LHCB1 through the mitigation of the f_{0030} RDT. This correction is expected to contribute positively to operational issues such as the breaking of the collimator hierarchy, which has been partly attributed to the presence of the $3Q_y$ resonance. In addition, clear evidence of a dodecapolar feed-down effect was observed through the dependence of f_{0140} on the crossing-angle. These measurements were also used to successfully benchmark the FiDeL dodecapole magnetic model of IR1.

The method therefore shows strong potential for application in the HL-LHC era, where the increased sensitivity of the accelerator to NL optics errors will require robust correction strategies extending to the dodecapolar order.

CHAPTER 7

LONG-RANGE BEAM-BEAM COMPENSATION

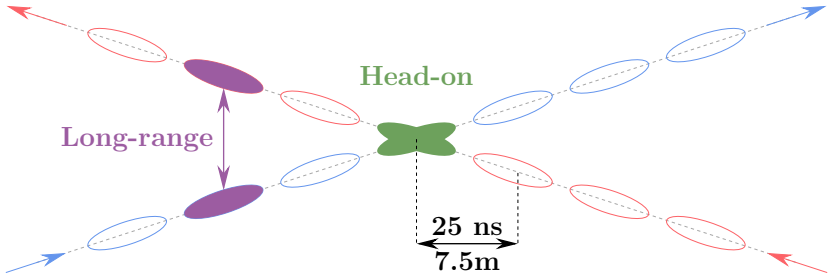


Figure 7.1: Diagram showing bunches from each beam, LHCB1 (blue) and LHCB2 (red), crossing at an experimental insertion and highlighting the difference between head-on (green) and long-range beam-beam (purple).

An effect, specific to colliders, is the beam-beam interaction, which arises from the electromagnetic forces between two beams as they approach and collide in the common aperture regions around the experiments, a simplified schematic of which is shown in Fig. 7.1. Its impact on quantities such as the DA, beam lifetime, and luminosity has been extensively studied using beam-based measurements in many accelerators including the Tevatron [144], RHIC [145], SuperKEKB [146], HERA [147], and the LHC [148].

In contrast, a systematic characterisation of beam-beam effects through beam-based optics measurements has not yet been realised, despite sustained interest from the accelerator community in understanding the impact of beam-beam interactions on luminosity. Previous

investigations have instead been conducted through simulation [149, 150], comparison between measured and simulated DA [151], measurements of dynamic beam-beam effects carried out in SuperKEKB [152], and attempts to measure head-on-generated optics perturbations in a non-operational configuration of the LHC at injection [153, 154]. The lack of direct optics measurements including the effect of beam-beam is largely attributable to the intrinsic difficulty of performing measurements under such conditions. Historically, LHC optics commissioning and measurements have been conducted using low-intensity, non-colliding pilot bunches. While this methodology is well suited to establishing a reference optics configuration, it does not fully reproduce the beam dynamics encountered during high-intensity physics operation.

Under physics operating conditions, the beams are subject to two distinct classes of beam-beam interaction: head-on and Long-Range (LR), or equivalently Long-Range Beam-Beam (LRBB), encounters, as illustrated in Fig. 7.1. LR interactions arise from the finite transverse separation of the two beams in the vicinity of the IPs, imposed by the crossing-angle required to prevent parasitic collisions outside the detector regions. These interactions represent a significant source of NL perturbations. As higher luminosities are pursued, the impact of LR interactions increases, necessitating larger crossing-angles and driving the accelerator further into a strongly NL regime. This effect is expected to become even more pronounced in the HL-LHC era.

The aim of this study was to exploit a novel optics measurement technique to directly measure and correct the optics perturbations induced by the LRBB. Measurements were performed in a weak-strong configuration, in which a low-intensity pilot bunch interacts with high-intensity bunch trains. In this configuration, the beams are separated by approximately 5σ to suppress head-on effects, allowing the impact of LRBB interactions to be isolated and corrected independently. This measurement approach was then used as the basis for a correction scheme targeting LRBB-driven optics distortions, in both the linear and NL regimes, using the existing corrector magnet scheme.

7.1 Beam-beam force

The overall beam-beam interaction originates from the fact that each bunch consists of a distributed charge density [155]. The electromagnetic fields associated with this distribution are obtained by integrating over the bunch profile. For a particle interacting with an opposing bunch, these fields give rise to beam-beam forces that can be either focusing or defocusing, depending on the relative charge and transverse position of the particle with respect to the field-generating beam.

The transverse particle density of a beam can, in the simplest two-dimensional approximation, be described by a bi-Gaussian distribution,

$$\rho(x,y) = \rho_x(x)\rho_y(y), \tag{7.1}$$

where the density in each transverse plane, ρ_z with $z \in \{x,y\}$, is characterised by an RMS beam size of σ_x and σ_y

$$\rho_z(z) = \frac{1}{\sigma_z\sqrt{2\pi}} \exp\left(-\frac{z^2}{2\sigma_z^2}\right). \quad z \in \{x,y\} \tag{7.2}$$

This charge distribution gives rise to an electromagnetic potential $U(x,y,\sigma_x,\sigma_y)$, which can be written [155]

$$U(x,y,\sigma_x,\sigma_y) = -\frac{ne}{4\pi\varepsilon_0}\int_0^\infty \frac{1-\exp\left(-\frac{x^2}{2\sigma_x^2+q}-\frac{y^2}{2\sigma_y^2+q}\right)}{\sqrt{(2\sigma_x^2+q)(2\sigma_y^2+q)}}\,dq, \tag{7.3}$$

where n denotes the line density of particles in the beam and ε_0 is the permittivity of free space. The transverse electromagnetic fields are obtained by taking the gradient of this potential.

For the special case of a round beam, where $\sigma_x = \sigma_y = \sigma$, the transverse Lorentz force can be expressed by

$$\vec{F} = q\left(E_r + \beta c B_\phi\right)\hat{r}. \tag{7.4}$$

Introducing $r^2 = x^2 + y^2$ and using Eq. (7.3), the radial electric and azimuthal magnetic fields entering Eq. (7.4) can be written as

$$E_r = -\frac{ne}{4\pi\varepsilon_0}\frac{\partial}{\partial r}\int_0^\infty \frac{\exp\left(-\dfrac{r^2}{2\sigma^2+q}\right)}{2\sigma^2+q}\,\mathrm{d}q, \tag{7.5}$$

and

$$B_\phi = -\frac{ne\beta c\mu_0}{4\pi}\frac{\partial}{\partial r}\int_0^\infty \frac{\exp\left(-\dfrac{r^2}{2\sigma^2+q}\right)}{2\sigma^2+q}\,\mathrm{d}q. \tag{7.6}$$

From Eqs. (7.5) and (7.6) it follows that the force given by Eq. (7.4) is purely radial. Evaluating the derivative first and using $1/(2\sigma^2+q)$ as the integration variable allows these expressions to be simplified. The resulting radial force can then be written in closed form, making use of $\varepsilon_0\mu_0 = c^{-2}$:

$$F_r(r) = \frac{ne^2(1+\beta^2)}{2\pi\varepsilon_0}\frac{1}{r}\left[1-\exp\left(-\frac{r^2}{2\sigma^2}\right)\right]. \tag{7.7}$$

The impact of this force is illustrated in Fig. 7.2. For small transverse amplitudes, the force is approximately linear, resulting in a tune shift analogous to that produced by a quadrupole. At larger amplitudes, however, the force becomes strongly NL, such that the opposing beam acts as an intense NL electromagnetic lens. In this regime, the resulting tune shift depends sensitively on the particle amplitude.

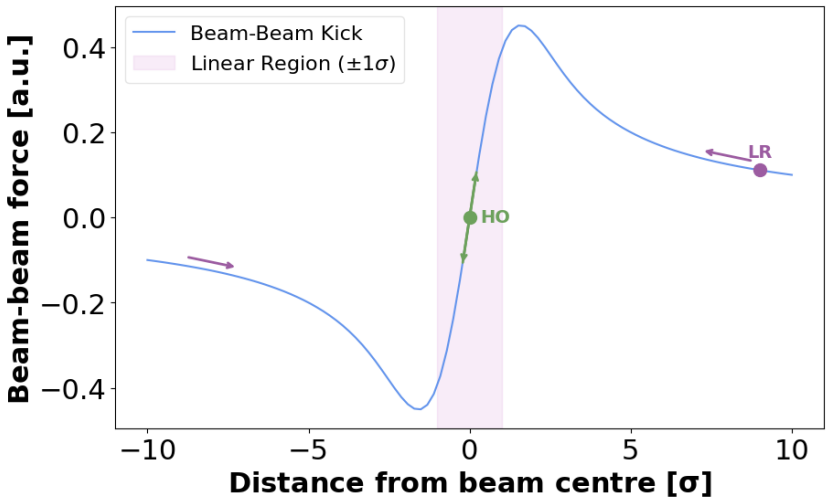


Figure 7.2: Beam-beam force (blue) for equal charges in each beam, shown in arbitrary units, vs the transverse distance from the beam centre in units of σ. The head-on interaction region is indicated by the green circle and arrows, and the LR interaction range is shown in purple. The region in which the beam-beam force is approximately linear is highlighted in light pink.

7.2 LHC context and motivation

As discussed previously, LRBB interactions introduce significant perturbations to the beam optics and are known to strongly excite a range of NL resonances. These effects can reduce the DA and degrade beam lifetime [156], therefore limiting the achievable integrated luminosity. They also constrain the minimum crossing-schemes [157] which then affect luminosity as well. During 2024 operation, notable difficulties were encountered due to a breakdown of the collimation hierarchy [158], with peak losses occurring at the secondary collimators rather than at the primaries, raising concerns for machine-protection. As a result, operation at a β^* of 30 cm could not be sustained, leading to a loss in luminosity.

While new optics corrections eventually allowed 30 cm operation to resume, this behaviour was suspected to be driven, at least in part, by the excitation of specific NL resonances associated with LRBB effects. In particular, it was observed that increasing the vertical tune further exacerbated the breakdown of the collimation hierarchy. Moreover, the hierarchy violation was alleviated when operating with non-colliding bunches, indicating that beam-beam effects were playing a significant role. It was also found that correcting the lattice-driven contribution to the $3Q_y$ resonance led to an improvement in the hierarchy, while LRBB interactions are expected to excite the $3Q_y$ resonance with a comparable strength. For these reasons, the vertical third-order resonance, $3Q_y$, driven by LRBB was identified as a key candidate contributing to the observed performance limitations [74, 159]. This provided strong motivation to perform dedicated beam-based optics measurements to characterise the influence of beam-beam effects, and in particular LRBB interactions, under realistic physics operating conditions.

Therefore, a new procedure for carrying out optics measurements in the presence of LRBB interactions was developed in 2024, in collaboration with T. Persson and the LHC OP team [160]. This configuration consisted of injecting three low-intensity pilot bunches in one beam, which were brought into collision with nominal high-intensity bunch trains circulating in the opposite beam.

For example, one pilot bunch (pilot 1 in Fig. 7.3) was subjected to the full set of LR encounters, namely both the incoming and outgoing interactions, at the two primary IPs, IP1 and IP5. A

second bunch (pilot 2) experienced only the outgoing LR interactions at these locations. The third bunch (pilot 3) served as a reference and did not encounter any LR interactions. This scheme allowed for a direct, bunch-by-bunch evaluation of the impact of LRBB interactions on the optical functions.

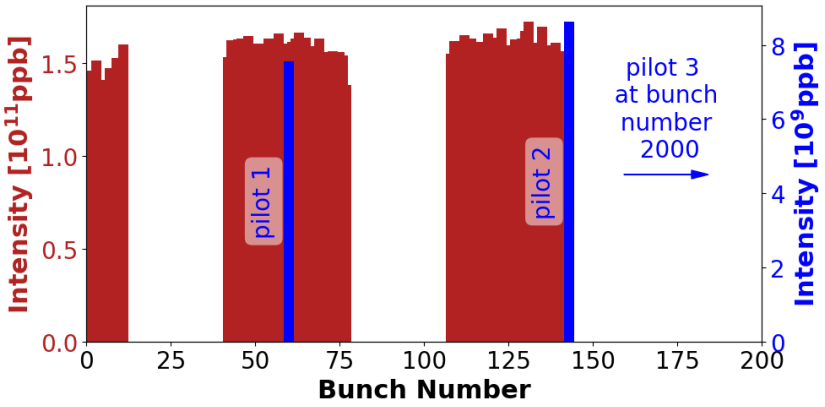


Figure 7.3: Bunch intensity vs bunch number for LRBB measurements depicting how nominal pilots (blue, right-hand axis) and bunch trains (red, left-hand axis) compare.

To realise this configuration, a series of special operational procedures had to be developed. Owing to the asymmetric beam intensities, different BPM gain settings were applied in the two rings, and common BPMs were removed from the orbit feedback. Asymmetric collimator settings were also required: tight operational collimation was maintained in the strong beam to preserve machine-protection, while in the weak beam the collimators were retracted to coarser settings to provide sufficient aperture for the transverse excitation. The ACD was enabled exclusively for the weak beam to ensure that only this beam was excited.

Implementation of these conditions required dedicated operational sequences capable of modifying safe beam flags and machine-protection interlocks independently for the two beams. In addition, the chromaticity and Landau octupole strengths in the weak beam were reduced to enhance the forced DA, therefore permitting ACD excitation while minimising beam losses. Initially, these settings were lowered for the entire accelerator cycle of the weak beam. However, experience during the following optics measurements, detailed below, showed that these lowered settings led to the instability of the pilot bunches early in the cycle, prior to collisions. The procedure was therefore revised such that chromaticity and octupole strengths were reduced only

at the end-of-squeeze, which resolved the instability issue.

This approach was first deployed during a dedicated MD session in 2024, during which reliable measurements of the f_{0030} and f_{4000} RDTs were obtained. Representative frequency spectra are shown in Figs. 7.4 and 7.5, where the distinct spectral lines corresponding to the f_{4000} and f_{0030} RDTs, respectively, are clearly identifiable.

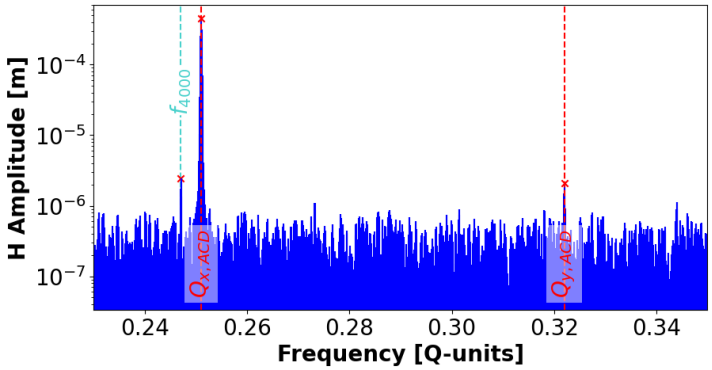


Figure 7.4: Horizontal TbT frequency spectrum, where cyan denotes 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.261$, $q_y = 0.31$.

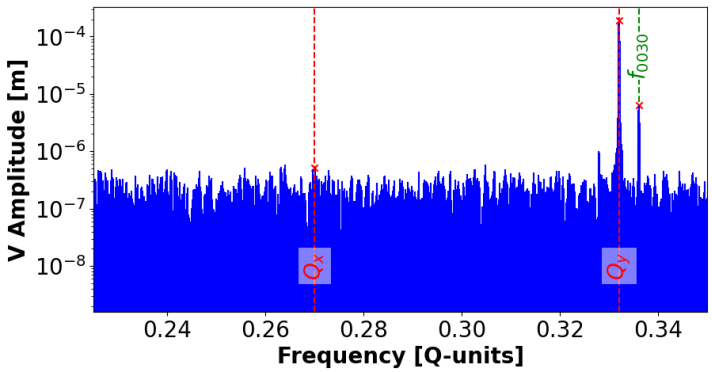


Figure 7.5: Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.32$.

Figure 7.6 presents a comparison of the amplitude of the f_{0030} RDT between configurations including the full set of LR interactions and those without any LR effects. The presence of LR interactions leads to a clear amplification of the $3Q_y$ resonance, a result of particular concern given that this resonance is known to pose significant operational challenges for both lifetime and collimator hierarchy.

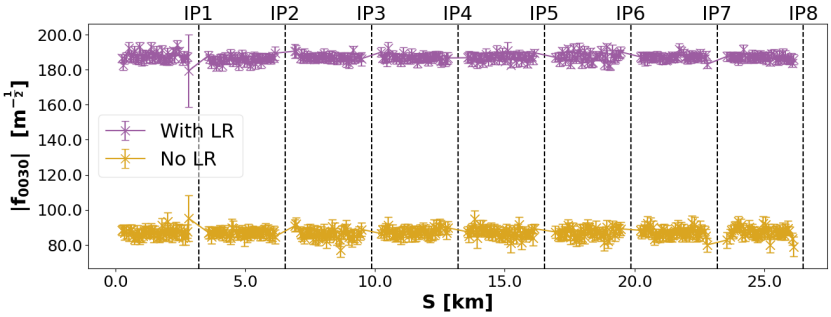


Figure 7.6: Amplitude of f_{0030} RDT vs s for LHCB2 with full (purple) and no LR (orange) interactions.

The $3Q_y$ resonance is not the only resonance significantly driven by LRBB interactions; normal octupolar resonances are also notably enhanced. One representative example is shown in Fig. 7.7, which presents the amplitude of the f_{4000} RDT for configurations with and without LR interactions.

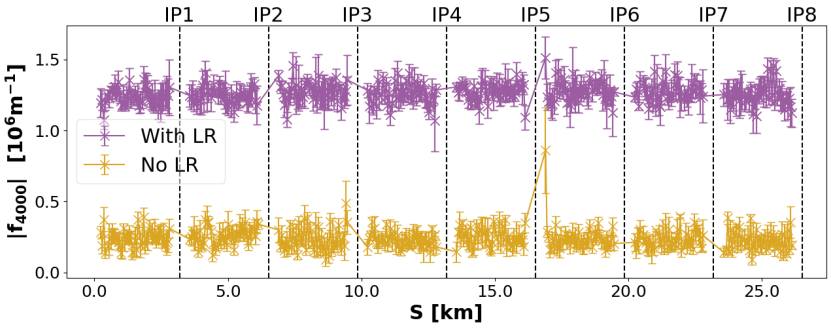


Figure 7.7: Amplitude of f_{4000} RDT vs s for LHCB2 with full (purple) and no LR (orange) interactions.

As with the $3Q_y$ resonance, the clear amplification of the normal octupolar terms under the influence of LRBB interactions highlights their impact on the accelerator's NL dynamics. Given the consequences of these effects for both resonance excitation and key performance indicators, such as luminosity and beam lifetime, the development of effective compensation strategies is of considerable operational importance.

There already exists a correction intended to mitigate LRBB-induced octupolar resonances using current-carrying wires, commonly referred to as LRBB wires [161–164]. These wires are installed within the tungsten jaws of the tertiary collimators in IRs 1 and 5 and are intended to counteract the NL beam-beam forces.

Previously, the effectiveness of this scheme has primarily been evaluated using observables such as measured beam lifetime, or DA in simulation. However, reliance on these global indicators does not allow for a quantitative characterisation of individual resonances, nor does it enable the explicit identification of which specific resonances are being compensated.

Therefore, when these optics measurements in the weak-strong configuration were performed by the optics team during the dedicated MD in 2024, the effect on the RDTs of the LRBB wire was investigated. The results analysed as part of this work indicated that, while this correction produced a clear improvement in the f_{4000} RDT with a 30% decrease in magnitude, it concurrently degraded the f_{0030} RDT at the level of 5%, as shown in Figs. 7.9 and 7.8. It should be noted that this outcome was not unexpected, as the LHC LRBB wires were designed specifically to compensate normal octupolar resonances. They were not intended to address other resonant terms.

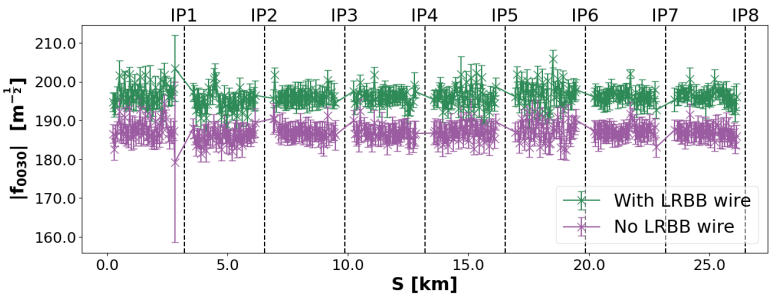


Figure 7.8: Amplitude of f_{0030} RDT vs s for LHCB2 with full LR, with (green) and without (purple) the LRBB wire.

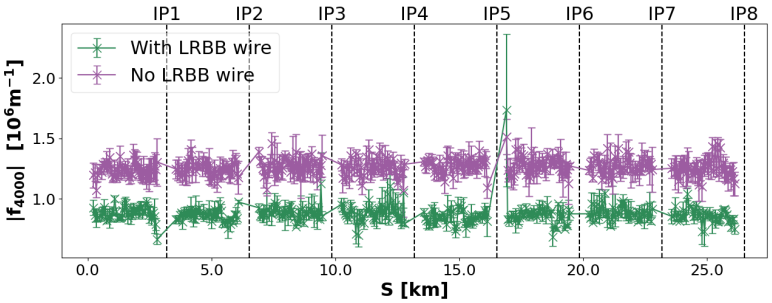


Figure 7.9: Amplitude of f_{4000} RDT vs s for LHCB2 with full LR, with (green) and without (purple) the LRBB wire.

These observations motivated a study into the feasibility of actively correcting the problematic resonances using existing magnet infrastructure, either as an alternative or as a complement to

the LRBB wire scheme. To this end, the NL corrector magnets located in the inner triplet regions of the IRs were selected. While these correctors are primarily designed to compensate magnetic field imperfections in the triplet magnets, they were repurposed for this investigation, despite not having been originally intended for LRBB-driven resonance correction.

7.3 Simulation benchmarking

To evaluate the practicality of a simulation-driven correction strategy, an initial step was to benchmark the model against beam-based measurements. In simulation, LRBB effects can be represented by modelling the LR interaction as a series of effective multipole kicks distributed according to the relevant bunch filling scheme, as illustrated in Fig. 7.10. For this purpose, simulations were performed using CPyMAD [165]. Although Xsuite [18, 166] provides a similarly well-established and extensively validated simulation framework in which the required accelerator settings are already reliably configured that is considerably faster in execution, CPyMAD offers the key advantage of including a model of the ACD, which is essential for a direct comparison with beam-based measurements. MAD-X was not selected because no equivalent framework was available.

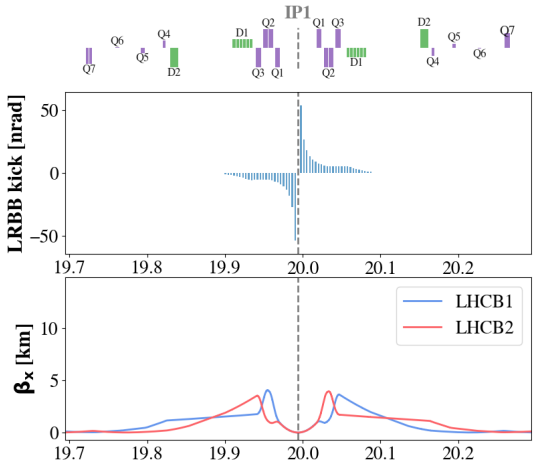


Figure 7.10: Simulated horizontal β-function vs s with the same schematic of magnet strengths as in Fig. 3.7 with the inclusion of the distribution of LRBB kicks in the model.

As discussed earlier, beam-based data was successfully collected during the 2024 MD period. These measurements exhibited very good agreement with simulations for the amplitudes of both

the $3Q_y$ and $4Q_x$ resonances, as shown in Figs. 7.11 and 7.12. It is important to note that, for the f_{0030} term, the quantity shown is not the absolute amplitude of the RDT itself, but the difference between configurations with and without LR interactions. This choice was motivated by the observation that, in the real accelerator, lattice-related effects were by an order of magnitude more pronounced than in the simulations. The resulting level of agreement nevertheless validated the fidelity of the modelling of the LR effect, and justified the use of simulation-driven optimisation studies to develop targeted corrections for these NL beam-beam effects.

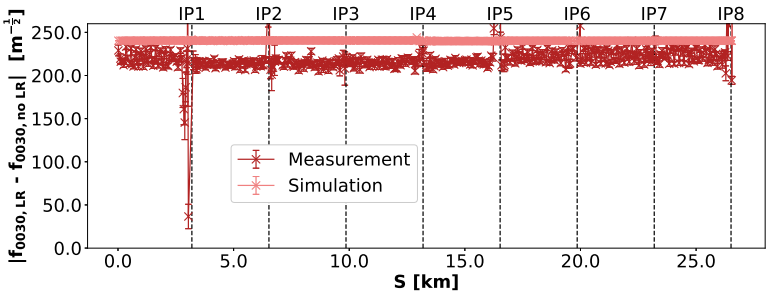


Figure 7.11: Amplitude of f_{0030} RDT vs s for LHCB2 with full LR in simulation (light red) and measurement (red).

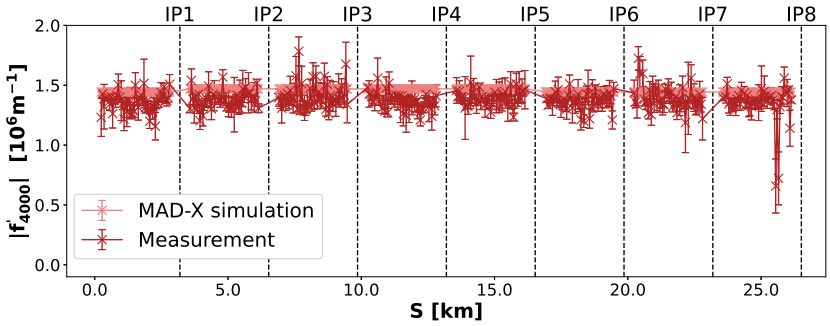


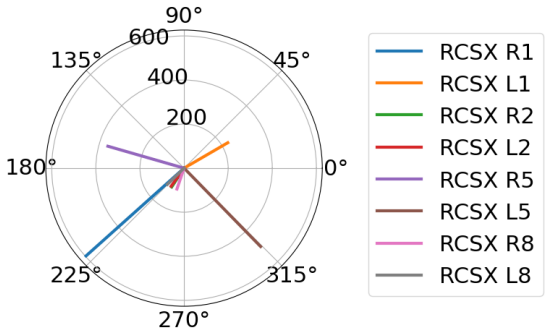
Figure 7.12: Amplitude of f_{4000} RDT vs s for LHCB2 with full LR in simulation (light red) and measurement (red).

7.4 Sextupolar LRBB-driven RDT correction with round-optics

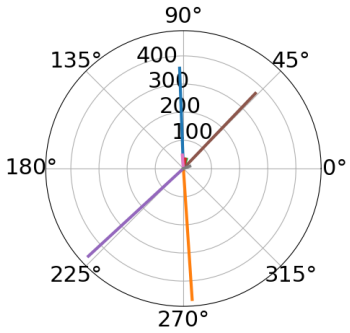
To address the sextupolar resonances $Q_x + 2Q_y$ and $3Q_y$, a correction strategy based on the measured RDT response was implemented. In contrast to the approach adopted at injection energy, this procedure was carried out entirely manually, without recourse to automated matching in MAD-NG, owing to the lack of fully validated beam-beam modelling within that framework.

As an initial step, the sensitivity of the relevant RDTs to the NL corrector magnets installed in the inner triplets was assessed through MAD-NG simulations. These correctors were selected on account of their proximity to the IP, which enables a quasi-local compensation of the LRBB-driven resonances, as opposed to relying on magnets situated further from the IR.

The magnitude of the shift in both the amplitude and phase of the f_{1020} RDT, comparing the baseline configuration with no correctors powered, to a configuration in which the normal sextupolar correctors were set to their maximum operational strength in both beams, is shown in Figs. 7.13(a) and 7.13(b). Because the inner triplet correctors are shared by the two beams, an additional constraint is introduced: the correction must be arranged such that the LRBB-induced RDT shifts are compensated simultaneously in both beams. In this context, the well-defined directionality of the response in the complex plane proved instrumental in identifying coherent combinations of correctors (for example, RCSX R1 and RCSX L1) that could be applied together, therefore achieving a concurrent improvement in both beams. The largest contributions came from the correctors on either side of IP1 and 5.



7.13(a) Polar plot for LHCB1 showing the f_{1020} amplitude and phase of RDT response to maximum powerings of all the normal sextupolar correctors.



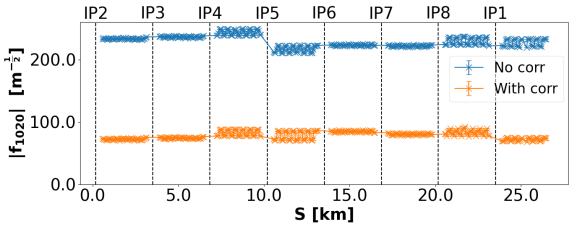
7.13(b) Polar plot for LHCB2 showing the f_{1020} amplitude and phase of RDT response to maximum powerings of all the normal sextupolar correctors.

In parallel, the perturbative impact of LRBB interactions on those RDTs was quantified using Xsuite simulations. It should be noted that, at the time of this study, an ACD model was not available within this framework. As a result, the optimisation was carried out with free- rather than forced-RDTs. Based on previous studies in Chapter 3 demonstrating the validity of using the correction of free-RDTs for forced-RDTs, it was considered an acceptable approximation. A combination of corrector strengths was identified manually, which matched the simulated LRBB-driven RDTs in a manner that preserved effective corrections for both beams.

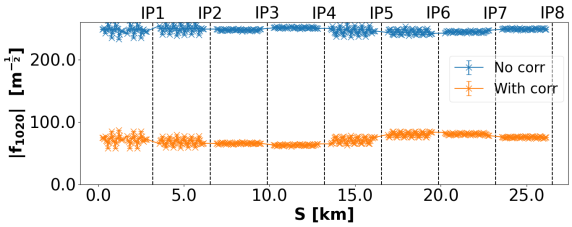
A viable solution was identified for the dominant sextupolar resonances near the WP, applicable to both LHCB1 and LHCB2, and requiring corrector strengths within the available operational limits, shown in Table 7.1. Figures 7.14(a) and 7.14(b) depict the predicted reduction in RDT magnitude. This was based on bunches experiencing full LR interactions, which would be the majority of bunches in the accelerator. The non-colliding bunches may consequently be degraded; however, they are not expected to affect the beam lifetime or collimation hierarchy.

Table 7.1: Corrector trim strengths for the b_3 correction. The maximum available current is 100 A. RCSX denotes a normal sextupole corrector

Corrector name	K value [m^{-3}]	Current [A]
RCSX R1	0.0039	43.8
RCSX L1	0.0003	1.7



7.14(a) Simulated amplitude of f_{1020} RDT vs s with (orange) and without (blue) correction for LHCB1 with full LR.



7.14(b) Simulated amplitude of f_{1020} RDT vs s with (orange) and without (blue) correction for LHCB2 with full LR.

Prior to implementation in weak-strong beam conditions, a dedicated accelerator test with non-colliding beams was conducted in 2024 to ensure that the effects would be within machine-protection constraints. These tests confirmed that the correction introduced negligible β-beating (see Fig. 7.15) and chromatic β-beating.

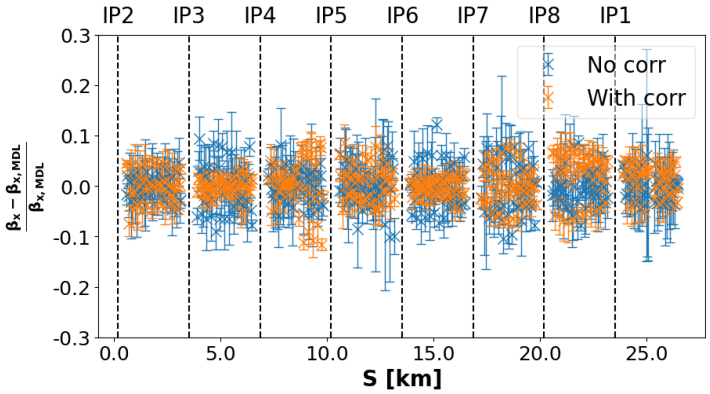


Figure 7.15: Measured β-beating vs s with (orange) and without (blue) correction for non-colliding beams in 2024.

7.5 Sextupolar LRBB-driven RDT correction with flat-optics

Unfortunately, due to beam-time limitations, the correction could not be validated in the full weak-strong configuration during the 2024 run. Furthermore, the operational optics were revised for 2025 to a flat-optics configuration. In this scheme, the β^* values are squeezed asymmetrically in the two planes, with $\beta^* = 60/18$ cm (horizontal/vertical) at IP1 and $\beta^* = 18/60$ cm at IP5.

A reassessment of the correction strategy was therefore required, prompted by the observation of a noteworthy feature associated with flat-optics. It was found that a flat-optics configuration is intrinsically beneficial with respect to the problematic $3Q_y$ resonance. In this case, the LRBB contribution is dominated by the IR employing a vertical crossing scheme, which corresponds to the plane with the weaker β_y functions. This configuration naturally suppresses the strength of the resonance.

This behaviour was initially identified in simulation, as illustrated in Fig. 7.16. It was subsequently validated through beam-based measurements conducted during an MD in 2025, in the context of a study of an LRBB-driven β-beating compensation scheme, discussed in a later section. The corresponding measurement results are presented in Fig. 7.17.

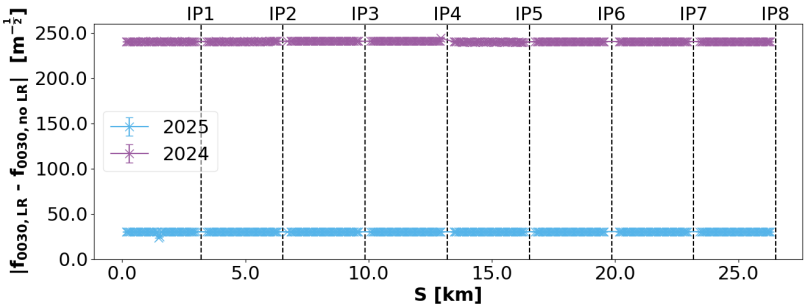


Figure 7.16: Simulated effect of switching to flat-optics (blue) from round (purple) optics for f_{0030} RDT vs s for LHCB2.

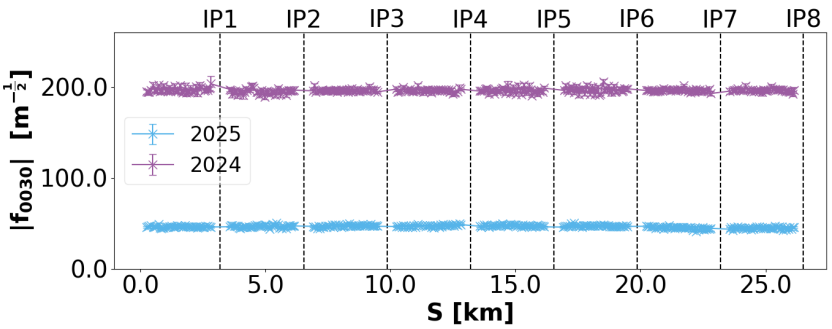


Figure 7.17: Measured effect of switching to flat-optics (blue) from round (purple) optics for f_{0030} RDT vs s for LHCB2.

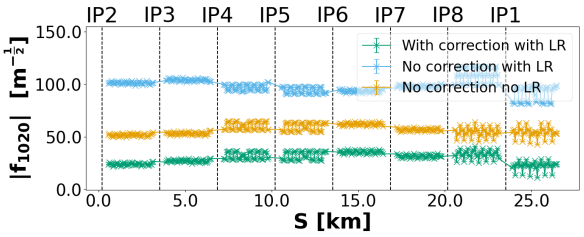
While the flat-optics configuration led to a natural reduction of the $3Q_y$ resonance, it did not provide complete suppression. In particular, for LHCB1, a residual value of approximately $f_{0030} = 50\ \mathrm{m}^{-\frac{1}{2}}$ was still observed in 2025. Consequently, the application of a dedicated correction using the IR correctors remained warranted.

The updated correction settings were derived following the same methodology previously employed for the round-optics configuration. In this instance, the LRBB-induced contributions were found to be primarily driven by IP1 for the b_3 component and by IP5 for the a_3 component. The resulting corrector trim strengths are summarised in Table 7.2.

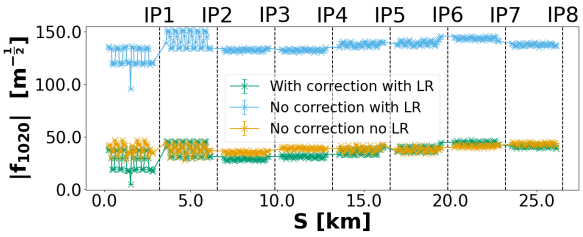
Table 7.2: Corrector trim strengths for the b_3 (RCSX) and a_3 (RCSSX) correction where the maximum current available is 100 A.

Corrector name	K value [m^{-3}]	Current [A]
RCSX R1	0.0022	21.7
RCSX L1	-0.0017	-17.1
RCSSX R5	0.0085	25.6
RCSSX L5	-0.0085	-25.6

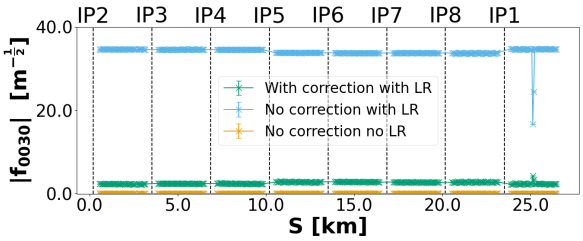
The simulated correction was found to substantially suppress both the f_{1020} and f_{0030} RDTs, simultaneously in both beams, as shown in Figs. 7.18(a), 7.18(b), 7.19(a) and 7.19(b). In these figures, the transition from blue to green indicates the effect of the applied correction, while the orange curve represents the reference configuration without LR interactions and without correction. For the f_{1020} RDT in LHCB1, the correction reduced the RDT amplitude to a level below that obtained in the absence of LR interactions, demonstrating that lattice-driven contributions were also partially compensated. The correction of the f_{0030} resonance was less pronounced; nevertheless, it restored the RDT magnitude to a value close to its nominal level without LR effects.



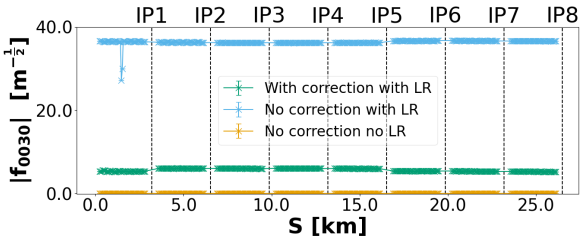
7.18(a) Simulated amplitude of f_{1020} RDT vs s for LHCB1 with full LR with correction (green), with full LR with no correction (blue) and no correction and no LR (orange).



7.18(b) Simulated amplitude of f_{1020} RDT vs s for LHCB2 with full LR with correction (green), with full LR with no correction (blue) and no correction and no LR (orange).



7.19(a) Simulated amplitude of f_{0030} RDT vs s for LHCB1 with full LR with correction (green), with full LR with no correction (blue) and no correction and no LR (orange).



7.19(b) Simulated amplitude of f_{0030} RDT vs s for LHCB2 with full LR with correction (green), with full LR with no correction (blue) and no correction and no LR (orange).

The impact of the corrections was studied in Frequency Map Analysis (FMA) tracking simulations. The particles were tracked for 20,000 turns to examine the single-particle stability. It was then demonstrated that the diffusion, defined as the magnitude of the difference between the tunes calculated over turns 1,000–2,000 and those calculated over turns 19,000–20,000, was reduced when the LRBB corrections were applied. This behaviour is illustrated by the representative comparison shown in the upper panels of Fig. 7.20 as a function of the initial tunes, and by the corresponding dependence on the initial particle amplitude, expressed in units of σ, shown in the lower panels of Fig. 7.20. The observed reduction in diffusion with the inclusion of the corrections indicates an improvement in the underlying single-particle stability.

During the dedicated MD session in 2025 aimed at testing the LRBB-driven RDTs correction, with LHCB1 as the measured beam, the f_{0030} RDT was observed to be clearly visible and dominant in the frequency spectrum. An illustrative example is provided in Fig. 7.21.

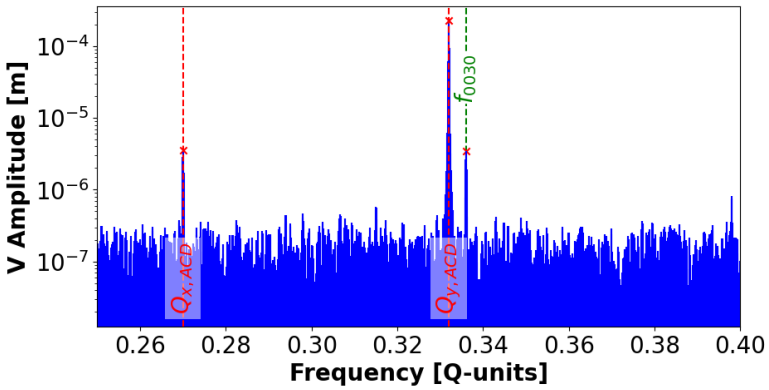
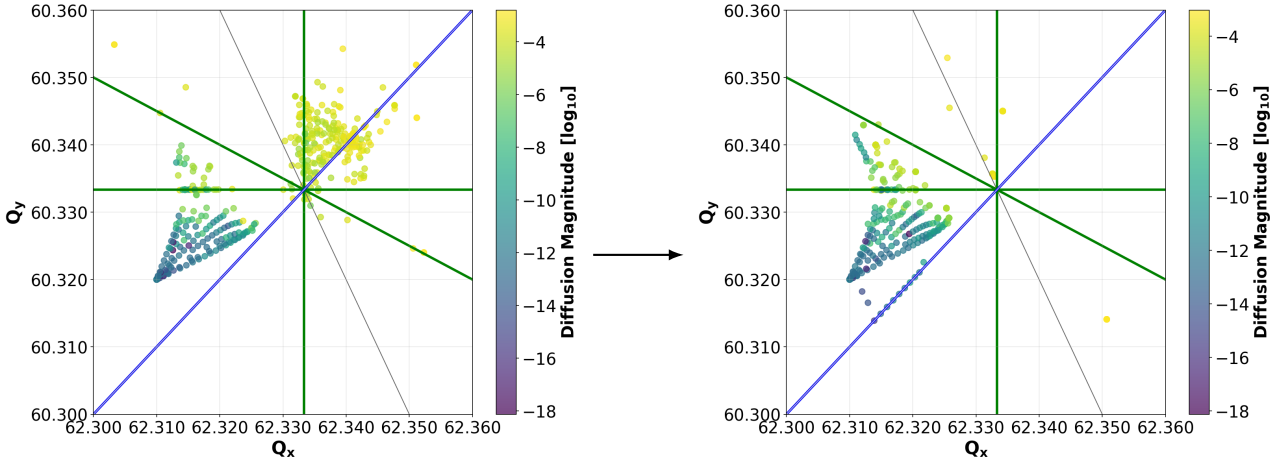
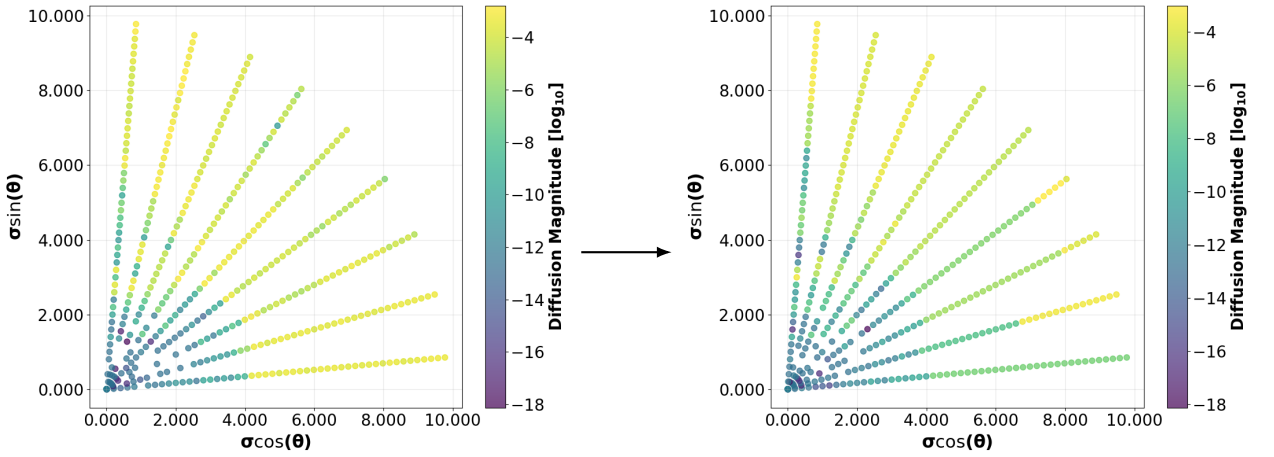


Figure 7.21: Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, and red the ACD frequencies, with natural $q_x = 0.28$, $q_y = 0.32$.



(a) With no correction

(b) With a_3 correction



(c) With no correction

(d) With a_3 correction

Figure 7.20: Comparison of the diffusion maps for Q_y vs Q_x (top), and vertical $\sigma \sin(\theta)$ vs horizontal $\sigma \cos(\theta)$ initial particle amplitude (bottom) with no correction and with the a_3 correction for MO = 15 m^{-4}.

An initial test of the correction demonstrated a substantial reduction of the $3Q_y$ resonance, as seen in Fig. 7.22. However, there was insufficient time during the MD to perform a corresponding evaluation of the f_{1020} correction.

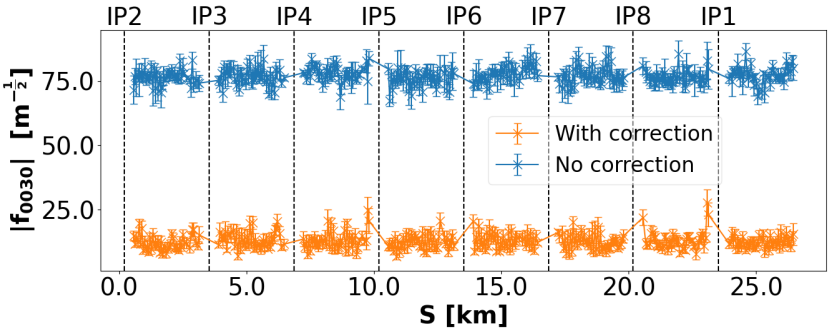


Figure 7.22: Measured amplitude of f_{0030} RDT vs s with (orange) and without (blue) correction for LHCB1 with full LR.

There was also insufficient time to perform a full lifetime validation during that particular session. However, a test conducted at the end of the fill, with collisions, examined the impact of trimming the skew sextupole corrector circuits by an amount comparable to the correction. This test indicated an improvement in beam lifetime of around 2 hours as shown in Fig. 7.23, showing the clear potential value of this type of LRBB corrections to future operation of the LHC or HL-LHC.

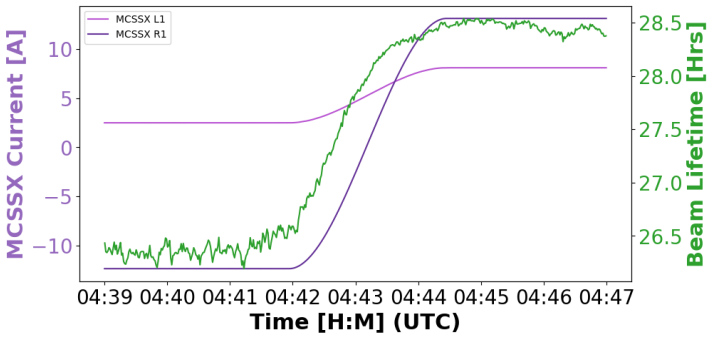


Figure 7.23: MCSSX (skew sextupolar corrector) current (R1 in dark purple and L1 in lighter purple) and beam lifetime (green) as a function of time.

7.6 Octupolar LRBB-driven RDT correction

The wire-based compensation of f_{4000}, introduced previously, could not be implemented in the flat-optics configuration due to the spatial positioning of the wire installations. Following the successful correction of the sextupolar RDTs, attention was therefore shifted to identifying an appropriate correction for the b_4 field error.

As an initial attempt, the same strategy applied to the sextupolar corrections was adopted, namely targeting the octupolar RDT f_{2002}, selected because of its direct proximity to the WP. However, this approach did not lead to a satisfactory solution. No circuit configuration was identified that achieved a substantial reduction of f_{2002} simultaneously for both beams.

As a result, an alternative method was explored, one based on compensating LRBB-induced detuning. This detuning arises because LRBB interactions effectively act as distributed NL lenses, generating amplitude-dependent tune shifts that were quantified in simulations.

The detuning-based correction strategy became feasible after flat-optics were introduced in 2025. The flat-optics configuration had itself been adopted to mitigate potential radiation damage in the low-β regions. In this configuration, the β-functions at the IPs are strongly asymmetric between the horizontal and vertical planes. This pronounced asymmetry leads to a clear separation of the detuning contributions arising from each IR. In particular, at IP1 the much smaller vertical β^* results in significantly larger vertical β-functions in the LR regime, causing the vertical direct detuning term to dominate. Conversely, at IP5 the horizontal plane becomes dominant, reflecting the opposite asymmetry in the optics.

Xsuite simulations were performed for several initial transverse actions, including the LRBB forces, to quantify the contributions from IR1 and IR5 to the amplitude detuning. The detuning response of the accelerator in the absence of LR interactions was then analysed using MAD-NG, based on the anharmonicities produced by powering the corrector magnets. In this analysis, the circuits chosen were the octupolar IR-correctors. This approach demonstrated that a targeted correction could be applied to the direct detuning terms, namely $\partial Q_x/\partial J_x$ and $\partial Q_y/\partial J_y$. Since the LRBB effects at IP1 were dominated by the vertical detuning component, correctors in this region were employed to address this term, while correctors in IP5 were used to target the

horizontal direct detuning. Suppression of these direct detuning terms resulted in a corresponding reduction of the $4Q_x$ and $4Q_y$ RDTs. As the cross-detuning terms were not modified, the $2Q_x - 2Q_y$ resonance remained unchanged; importantly, no degradation of this resonance was observed. Table 7.3 summarises the simulated detuning values before and after the correction to illustrate the resulting changes, while the corresponding horizontal and vertical direct detuning for LHCB1 are shown in Figs. 7.24 and 7.25, respectively.

Table 7.3: Amplitude detuning values for direct and cross terms for different configurations, i.e. without LR without correction (orange), with LR without correction (blue), and with LR with correction (green).

	$[10^3\ \mathrm{m}^{-1}]$	$\frac{\partial Q_x}{\partial 2J_x}$	$\frac{\partial Q_x}{\partial 2J_y}$	$\frac{\partial Q_y}{\partial 2J_y}$
LHCB1	Without LR	8	-13	7
	With LR, without correction	590	-529	554
	With LR, with correction	53	-545	208
LHCB2	Without LR	2	-9	7
	With LR, without correction	577	-512	579
	With LR, with correction	109	-525	167

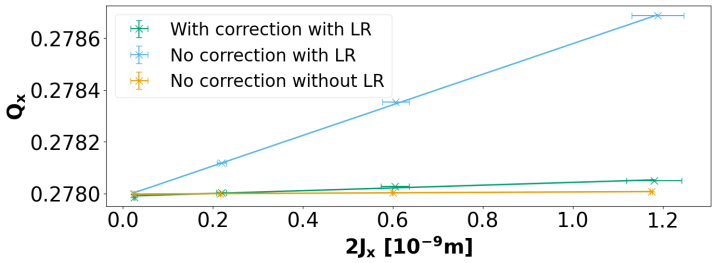


Figure 7.24: Simulated horizontal amplitude detuning (horizontal tune vs horizontal action) for LHCB1 with full LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange).

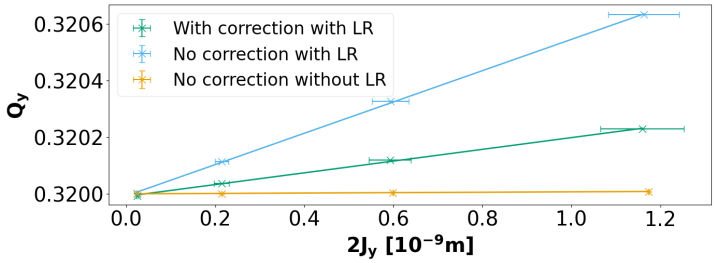


Figure 7.25: Simulated vertical amplitude detuning (vertical tune vs vertical action) for LHCB1 with full LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange).

The settings for the corrector trim strengths are given in Table 7.4.

Table 7.4: Corrector trim strengths for the b_4 correction where the maximum current = 100 A.

Corrector name	K value [m^{-4}]	Current [A]
RCOX R1	1.6	66
RCOX L1	-1.6	-66
RCOX R5	1.6	66
RCOX L5	-1.6	-66

This resulted in a reduction of the simulated amplitudes of the f_{4000} and the f_{0040} RDTs as displayed in Figs. 7.26 and 7.27 to levels close to those obtained without LR. Since the modification acted only on the direct detuning terms, the corresponding change in the f_{2002} RDT was found, as expected, to be minimal, as shown in Fig. 7.28.

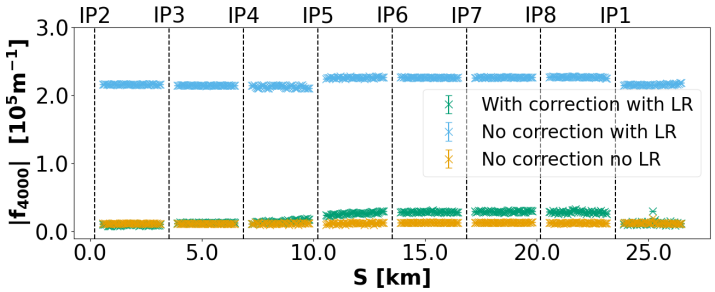


Figure 7.26: Simulated amplitude of f_{4000} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange).

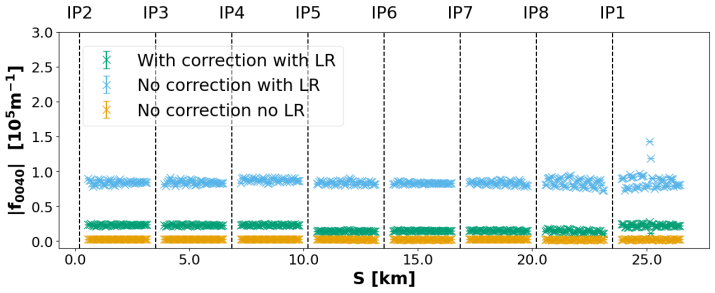


Figure 7.27: Simulated amplitude of f_{0040} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange).

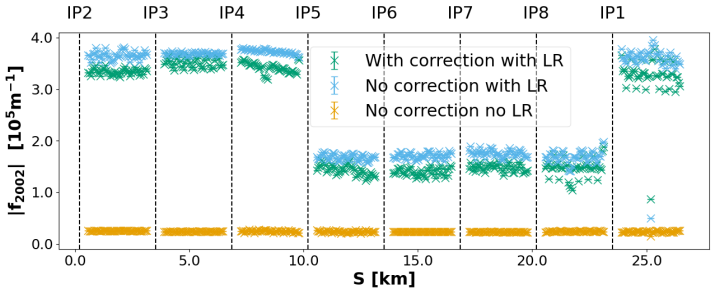
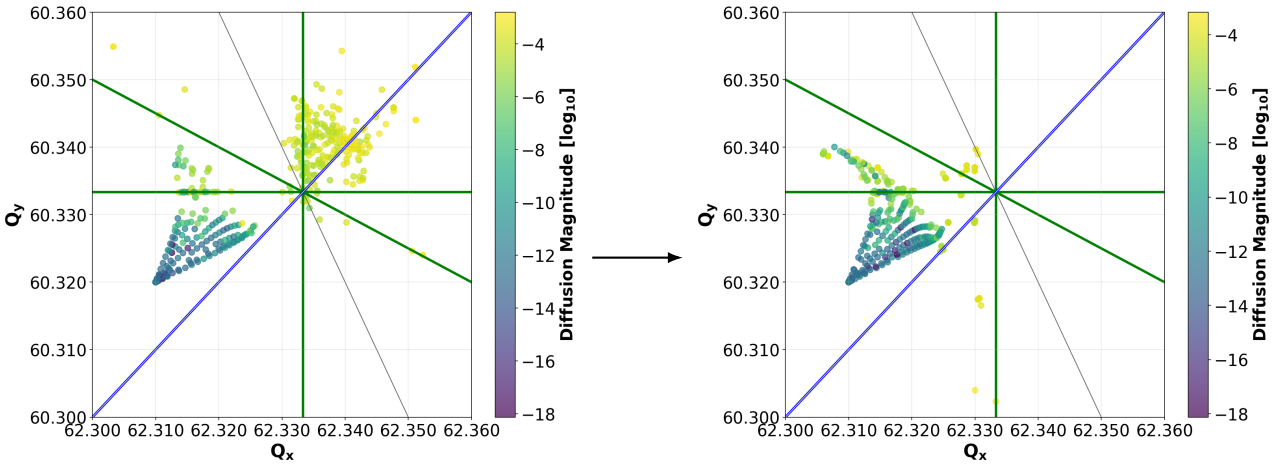


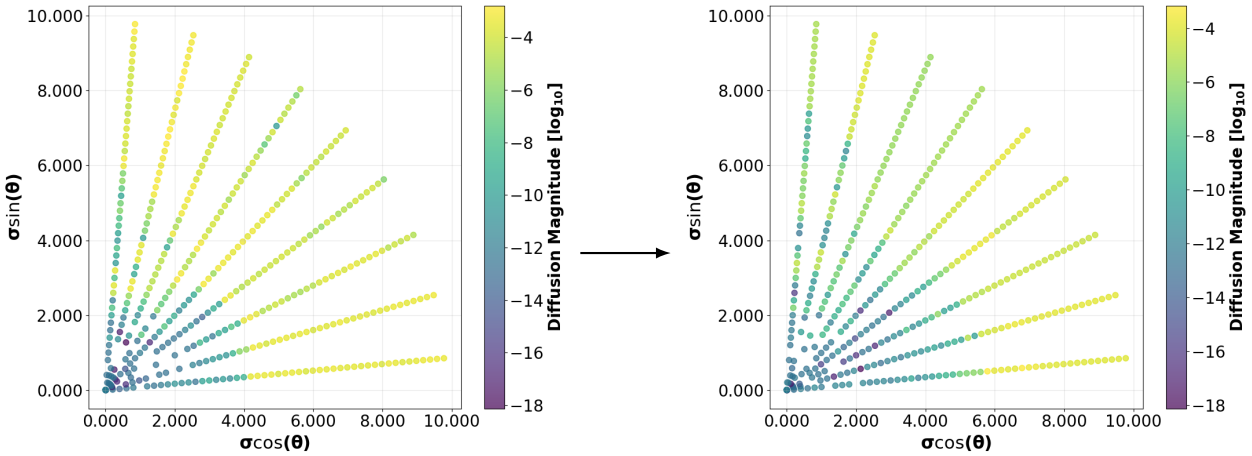
Figure 7.28: Simulated amplitude of f_{2002} RDT vs s for LHCB1 with LR with correction (green), with full LR with no correction (blue), and no correction and no LR (orange).

It was further observed that the application of the correction led to a reduction in diffusion in the FMA simulations. This is illustrated by the example comparisons in Fig. 7.29. Some improvements were observed for both MO polarities and with the MOs powered off in the tracking study. As in the case of the sextupolar correction, this reduction in diffusion is indicative of an improvement in the single-particle stability.



(a) With no correction

(b) With b_4 correction



(c) With no correction

(d) With b_4 correction

Figure 7.29: Comparison of the diffusion maps for Q_y vs Q_x (top), and vertical ($\sigma \sin(\theta)$) vs horizontal ($\sigma \cos(\theta)$) initial particle amplitude (bottom) with no correction and with the b_4 correction for MO = 15 m^{-4}.

Measurements conducted during the dedicated 2025 MD revealed that the f_{4000} and f_{0040} RDTs were visible at the selected WPs, as shown in Figs. 7.30 and 7.31. The f_{2002} RDT, however, was not visible at either of the chosen tunes.

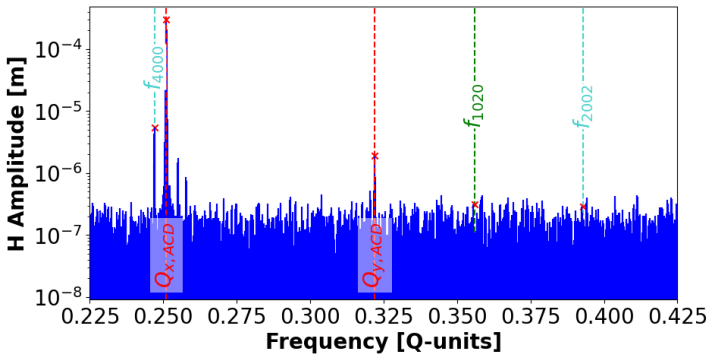


Figure 7.30: Horizontal TbT frequency spectrum, where green denotes 3rd order RDTs, cyan denotes 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.261$, $q_y = 0.31$.

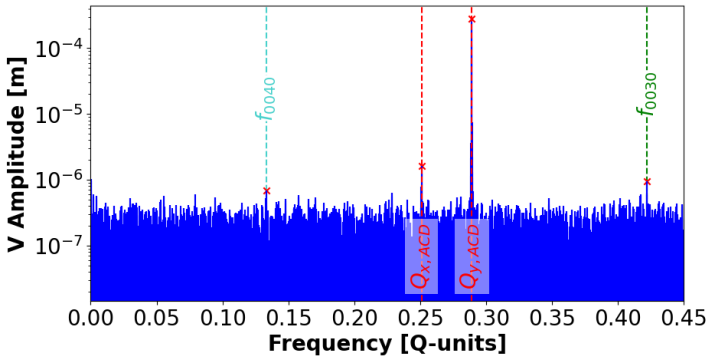


Figure 7.31: Vertical TbT frequency spectrum, where green denotes 3rd order RDTs, cyan denotes 4th order RDTs, and red the ACD frequencies, with natural $q_x = 0.261$, $q_y = 0.277$.

In these measurements, the octupole RDT generated by the lattice was too small to be measured. It was found however, that there was a clear difference between configurations with full and half LR interactions, as demonstrated in Fig. 7.32, where it appeared that the LR interactions dominated the driving of the resonance since half the number of interactions resulted in approximately half the amplitude.

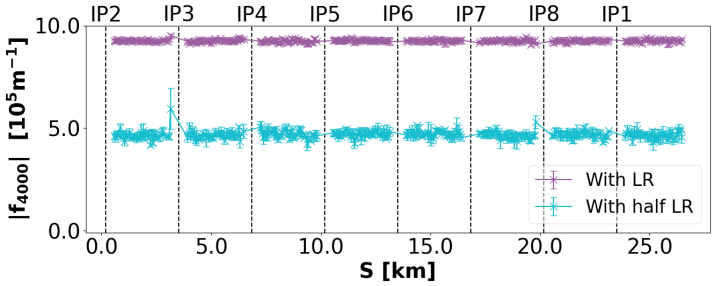


Figure 7.32: Amplitude of f_{4000} RDT vs s for LHCB1 with full (purple) and half (blue) LR interactions.

The measured RDTs, shown in Fig. 7.32 and in the subsequent figures, exhibit larger amplitudes than those predicted by the simulation studies. This difference arises because, during the

measurements, the WP was intentionally moved closer to the corresponding resonance in order to enhance the observable response.

The applied correction was found to be highly effective in reducing the amplitude of the f_{4000} RDT, as illustrated in Fig. 7.33. In contrast, the reduction of the f_{0040} RDT could not be quantified with the same level of precision, since the measured tunes were not positioned sufficiently close to the $4Q_y$ resonance. Nevertheless, a reduction in its amplitude was observed, as shown in Figs. 7.34 and 7.35. The histogram indicates that the peak amplitude was reduced by approximately 0.4×10^5 m^{-1} when the correction was applied.

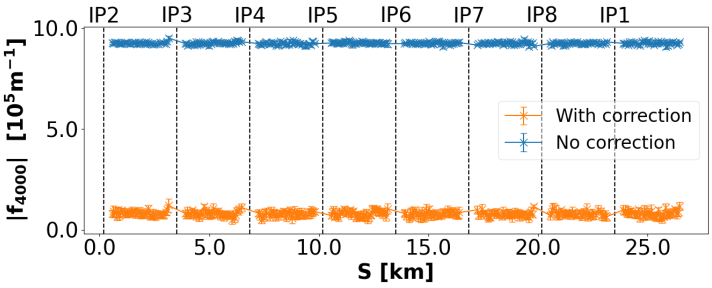


Figure 7.33: Measured amplitude of f_{4000} RDT vs s for LHCB1 with LR, with (orange) and without (blue) correction, $q_x = 0.261$, $q_y = 0.31$.

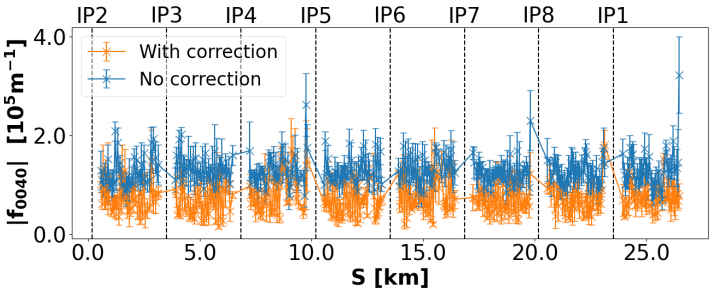


Figure 7.34: Measured amplitude of f_{0040} RDT vs s for LHCB1 with LR, with (orange) and without (blue) correction, $q_x = 0.261$, $q_y = 0.277$.

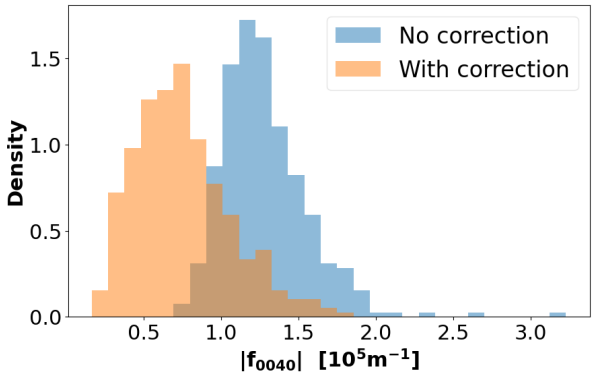


Figure 7.35: Histogram of f_{0040} RDT measured for LHCB1 with LR, with (orange) and without (blue) correction, $q_x = 0.261$, $q_y = 0.277$.

Overall, these results provide a proof-of-principle for a novel correction strategy that can serve as an effective alternative or complement to the LRBB wire scheme. The demonstrated method offers greater flexibility in mitigating specific resonances, while retaining the ability to adapt to different operational configurations. As such, it represents a promising avenue for future resonance control, both for present LHC operation and in the context of the HL-LHC.

7.7 LRBB-driven β-beating correction

During 2025, a luminosity imbalance between CMS and ATLAS was reported [167], prompting investigations into its origin. Since linear optics measurements are always performed as part of the NL optics studies, the MD data were examined to assess the level of β-beating. During the LRBB-driven RDT correction MD, the peak β-beating was approximately 20%, already at the machine-protection limit. This result, corresponding to a crossing-angle of 160 μrad, and an intensity (for the strong beam) of 1.5×10^{11} ppb, is displayed in Figs. 7.36 and 7.37.

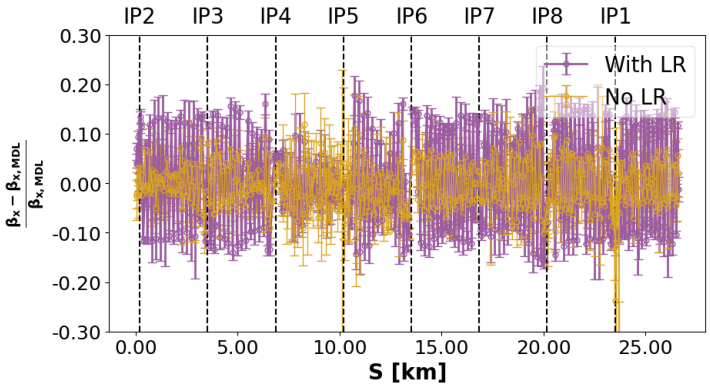


Figure 7.36: Horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 160 μrad.

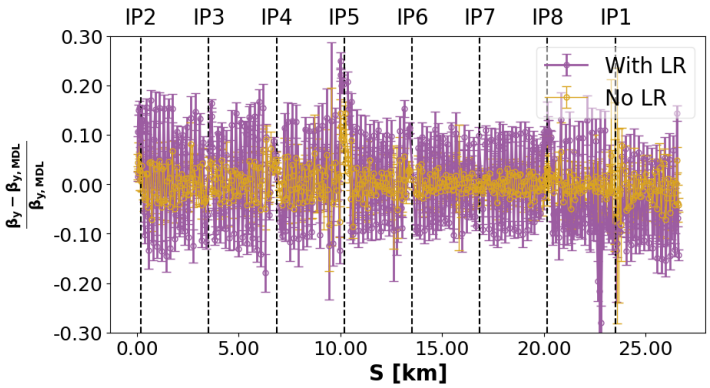


Figure 7.37: Vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 160 μrad.

When the crossing-angle was reduced to 120 μrad, the peak β-beating was found to increase even further, as depicted in Figs. 7.38 and 7.39. Figure 7.39 shows a very poor quality measurement due to the low intensity left at the end of the MD fill. This was of a similar magnitude to the calculated reduction in luminosity between the two crossing-angle configurations.

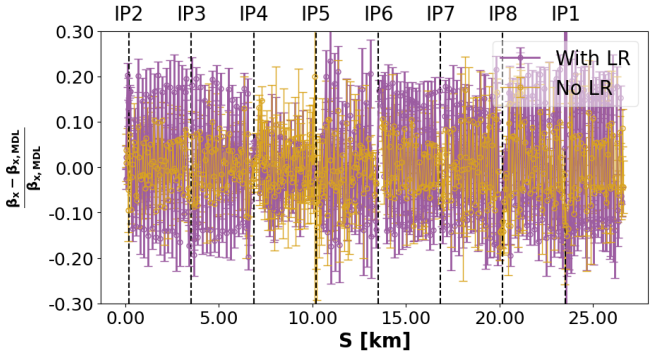


Figure 7.38: Horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 120 μrad.

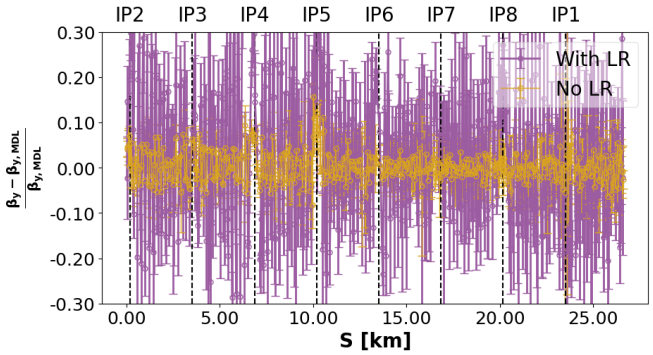


Figure 7.39: Vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1, crossing-angle = 120 μrad.

Maintaining the β-beating at a low level is essential, both to preserve operational margin for optics drift and to enable the application of β-waist knobs, as well as being in the interest of machine-protection. It was therefore critical to identify a suitable correction strategy.

7.7.1 Creating the correction

Before applying a simulation-based correction, it was first necessary to verify whether the observed luminosity imbalance could be reproduced in simulation. To this end, the same Xsuite simulation set-up used for the LRBB-driven RDT correction studies was employed. The resulting luminosity imbalance, summarised in Table 7.5, when scaled to the typical operational intensity of 1.7×10^{11} ppb, was found to be consistent with that reported by the experiments.

Table 7.5: Simulated luminosity imbalance for each crossing-angle.

Crossing angle [μrad]	Luminosity imbalance in favour of CMS [%]
160	6.1
120	10.1

Two complementary approaches to correcting LRBB-driven β-beating were investigated. The first used the matching functionality in Xsuite. The LRBB interactions induce both a tune shift and β-beating. In this study, the tune was re-matched to its initial value following the installation of LRBB elements, firstly in IR1 and then in IR5. This was done using the innermost individually powered quadrupoles for each beam and IR (the Q4 quadrupoles in IR1 and the Q5 quadrupoles in IR5). The simulation was configured with a bunch intensity of 1.5×10^{11} ppb and a crossing-angle of 160 μrad, which corresponds to the typical MD configuration. In practice, however, this can be easily re-scaled or recalculated for different beam intensities and crossing-schemes. This method achieved a significant reduction of the LRBB-induced β-beating in simulation, as can be seen in Figs. 7.40 and 7.41.

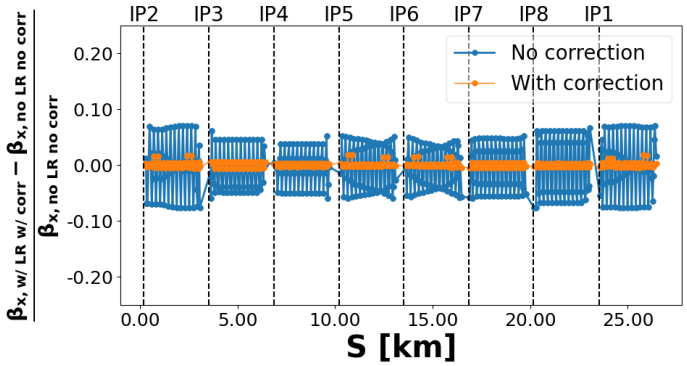


Figure 7.40: Simulated horizontal β-beating vs s with (orange) and without (blue) the Xsuite correction for LHCB1.

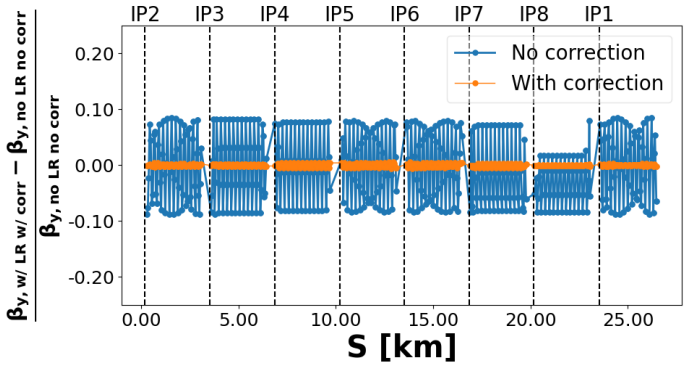


Figure 7.41: Simulated vertical β-beating vs s with (orange) and without (blue) the Xsuite correction for LHCB1.

These simulation studies also indicated that the Xsuite-based correction would reduce the luminosity imbalance in favour of CMS, as summarised in Table 7.6. This result suggests that the proposed correction could provide a viable means of addressing the observed imbalance.

Table 7.6: Simulated luminosity imbalance for each crossing-angle and correction.

Crossing angle [μrad]	Correction	Luminosity imbalance in favour of CMS [%]
160	No correction	6.1
160	With correction	0.1
120	No correction	10.1
120	With correction	4.5

The second approach used the standard β-function correction tools normally used for conventional LHC optics commissioning. In this case, a response matrix of selected quadrupoles in each IR was used to minimise errors on the phase advance between all BPMs around the LHC ring [77, 125]. The optics measurement from the bunch with all LR collisions was then used as input for the correction, rather than the non-colliding bunch, as would be usual during conventional optics commissioning.

7.7.2 Testing the correction

A dedicated MD, mentioned previously, provided the opportunity to test the correction methods. These measurements showed that the results were reproducible from the LRBB-driven RDT correction MD when comparing Figs. 7.36 and 7.42, as well as Figs. 7.37 and 7.43.

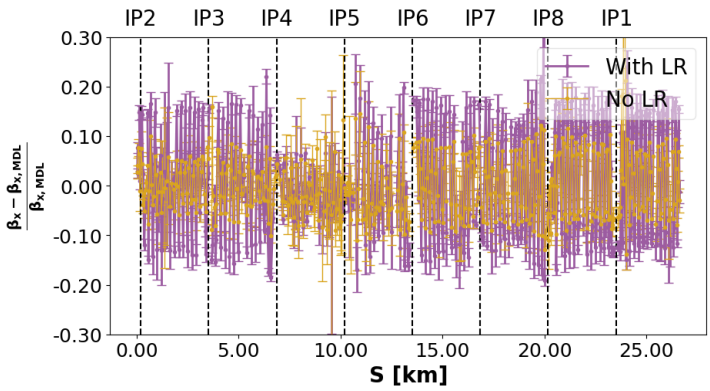


Figure 7.42: Measured horizontal β-beating vs s with (purple) and without (yellow) LR effects for LHCB1 with crossing-angle = 160 μrad.

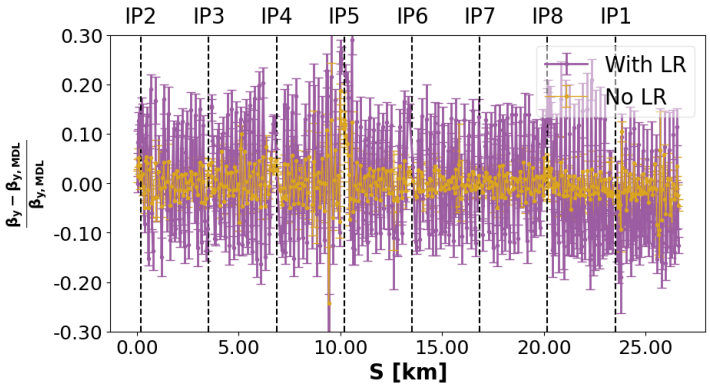


Figure 7.43: Measured vertical β-beating vs s with (purple) and without (yellow) LR effects for LHCB1 with crossing-angle = 160 μrad.

To permit a direct comparison between simulation and measurement, and to quantify the β-beating induced by the applied correction, a modified definition of β-beating was introduced. In this formulation, the β-function measured without LR effects and without correction was subtracted from the β-function measured in the presence of LR effects with the correction applied. The resulting difference was then normalised to the model β-function.

Using this definition, an approximate agreement in both amplitude and phase of the β-beating between simulation and measurement was observed for LHCB1, as shown in Figs. 7.44 and 7.45. The measurement uncertainties were, however, relatively large, likely due to the limited excitation amplitudes achievable.

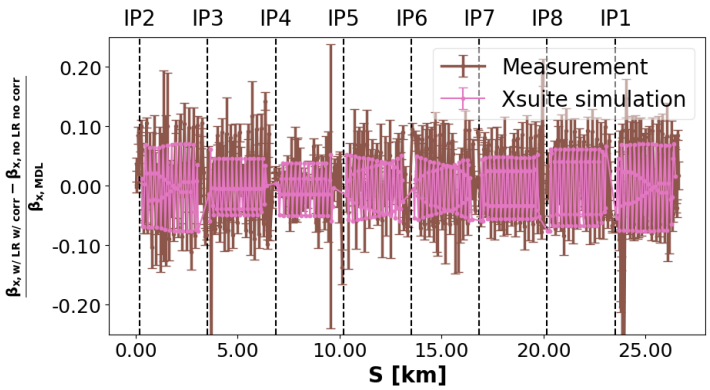


Figure 7.44: Horizontal β-beating vs s for simulation (pink) and measurement (brown) for LHCB1 with full LR effects.

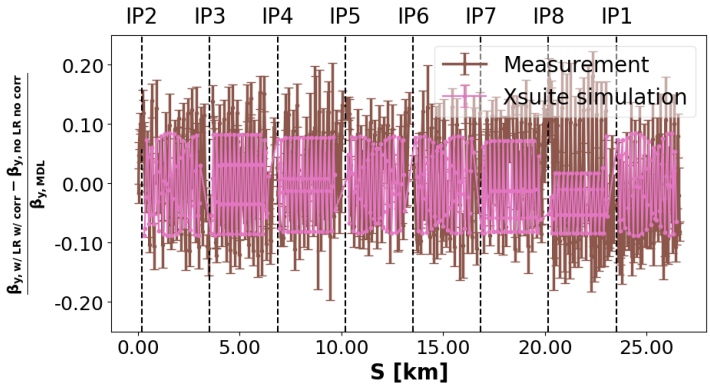


Figure 7.45: Vertical β-beating for simulation (pink) and measurement (brown) for LHCB1 with full LR effects.

The application of the model-based correction resulted in a substantial reduction of the conventional β-beating for the LRBB bunches, as summarised in Table 7.7 and illustrated in Figs. 7.46 and 7.47.

Table 7.7: Measured average reduction in β-beating ($\langle\Delta\beta$-beating$\rangle$) due to the correction.

Plane	$\langle\Delta\beta$-beating$\rangle$ (%)
Horizontal	4.2
Vertical	5.2

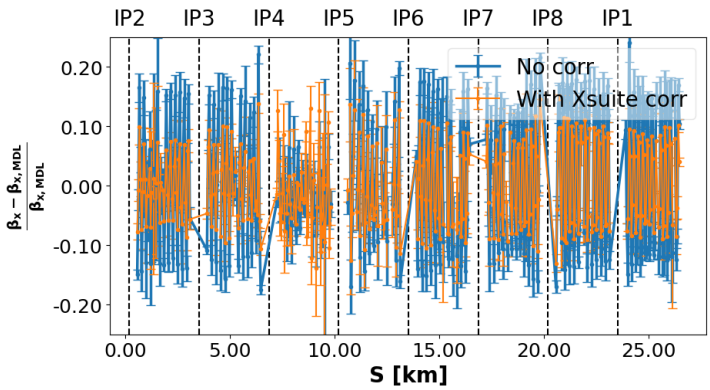


Figure 7.46: Measured horizontal β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB1 with full LR effects.

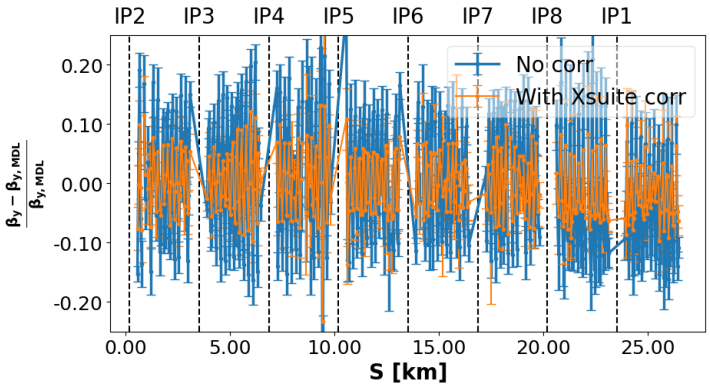


Figure 7.47: Measured vertical β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB1 with full LR effects.

For LHCB2, the benchmarking between simulation and measurement did not exhibit the same level of agreement, as illustrated in Figs. 7.48 and 7.49. Notable asymmetries were observed between the horizontal and vertical planes. For the vertical plane, the measured β-beating was larger than anticipated, and an unexpected jump in magnitude was observed downstream of IP1. While this behaviour may be attributable to residual linear optics errors remaining in the accelerator following conventional commissioning, its precise origin is not yet fully understood. It should also be noted that the overall quality of the measurements was limited due to the reduction in intensity because the measurements were performed at the end of the MD. In contrast, in the horizontal plane, the measured β-beating amplitude was generally smaller than predicted by simulation across several arcs.

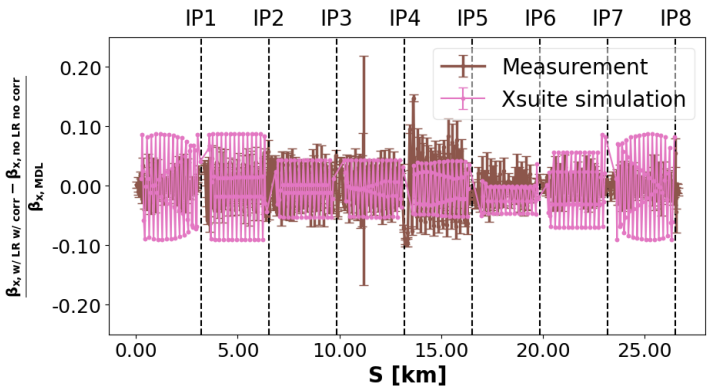


Figure 7.48: Horizontal β-beating vs s for simulation (pink) and measurement (brown) for LHCB2 with full LR effects.

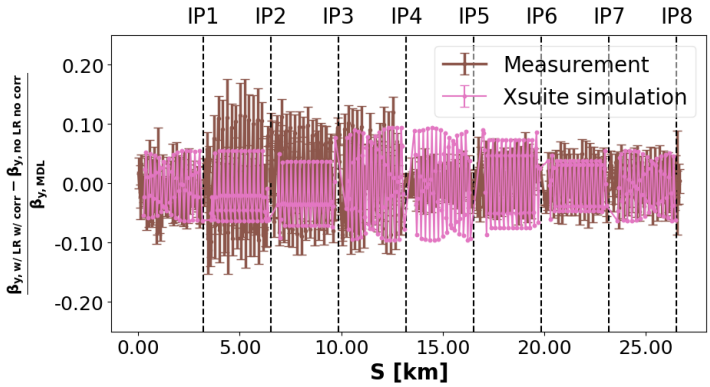


Figure 7.49: Vertical β-beating vs s for simulation (pink) and measurement (brown) for LHCB2 with full LR effects.

The LRBB model-based correction remained effective in reducing the peak β-beat in LHCB2, as summarised in Table 7.8 and shown in Figs. 7.50 and 7.51.

Table 7.8: Measured average change in β-beating ($\langle\Delta\beta$-beating$\rangle$) due to the correction.

Plane	$\langle\Delta\beta$-beating$\rangle$ (%)
Horizontal	4.1
Vertical	4.5

This reduction is particularly significant from the perspective of machine protection, where limiting the peak β-beat constitutes the primary concern. Although the correction resulted in a slight degradation of the horizontal plane, it led to a more balanced distribution between the two transverse planes, which nevertheless represents a substantially improved operational scenario.

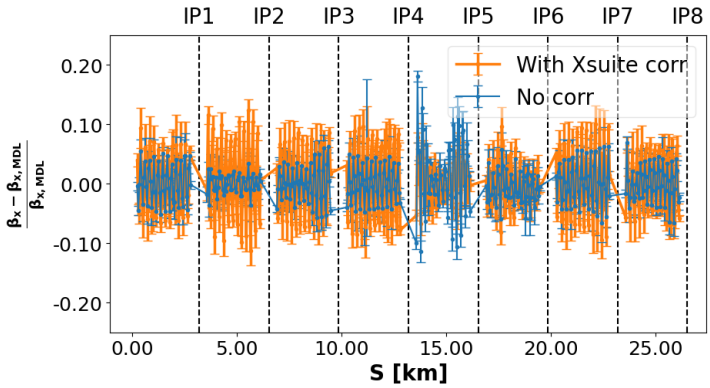


Figure 7.50: Measured horizontal β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB2 with full LR effects.

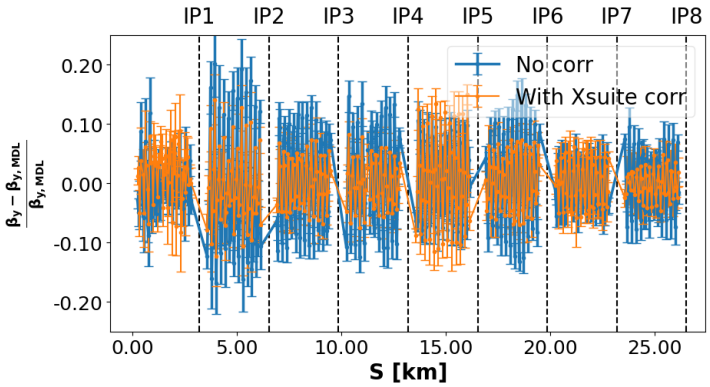


Figure 7.51: Measured vertical β-beating vs s with (orange) and without (blue) Xsuite correction for LHCB2 with full LR effects.

To determine whether the model-based correction for LHCB2 could be improved, a correction generated with the standard optics-correction GUI was also applied. This provided a practical alternative that reduced the β-beating in both planes. Table 7.9 gives a summary of the change to the β-beating with the correction.

Table 7.9: Measured average change in β-beating ($\langle\Delta\beta\text{-beating}\rangle$) due to each correction.

Plane	Correction	$\langle\Delta\beta\text{-beating}\rangle$ (%)
Horizontal	GUI Correction	2.5
Vertical	GUI Correction	3.3

Figure 7.52 shows that the GUI correction was more effective in the reduction of the horizontal β-beating than the Xsuite correction. In the vertical plane, the two methods produced comparable reductions, as shown in Fig. 7.53.

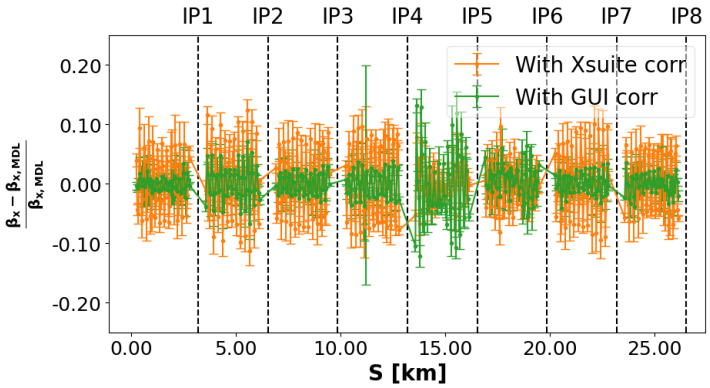


Figure 7.52: Measured horizontal β-beating vs s with the Xsuite (orange) and GUI (green) correction for LHCB2 with full LR effects.

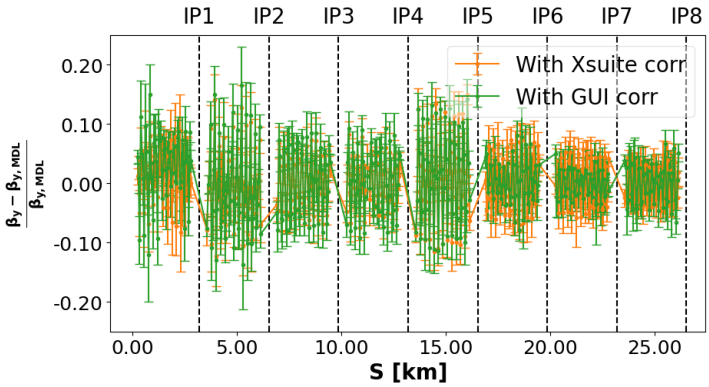


Figure 7.53: Measured vertical β-beating vs s with the Xsuite (orange) and GUI (green) correction for LHCB2 with full LR effects.

Importantly, these results demonstrate that a simultaneous correction for both beams is, in principle, achievable.

7.8 Conclusion

This chapter has presented the first dedicated measurement campaign aimed at the experimental characterisation and direct correction of LRBB-driven RDTs using existing corrector magnets in the LHC. The results demonstrate, as a proof-of-principle, that LRBB-induced RDTs can be robustly extracted from beam data, quantitatively compared with optics-model predictions through their measured amplitudes, and corrected using existing non-linear correctors in the inner triplets. This establishes the proposed methodology as viable and potentially applicable to future LHC and HL-LHC commissioning. The resulting RDT correction represents a credible experimental alternative to the use of LRBB wires for compensating LRBB effects. It may also be employed as a complement to the wire-based scheme, providing greater operational flexibility.

In addition, a reported luminosity imbalance between CMS and ATLAS motivated an investigation of β-beating under the same experimental conditions. Dedicated tests provided the first demonstration that linear optics perturbations induced by LRBB interactions could be measured and corrected. This opens the possibility of extending optics commissioning beyond conventional low-intensity, non-colliding conditions and of considering corrections of LRBB-induced β-beating during future LHC and HL-LHC commissioning. These results further validate the robustness

of the measurement framework and highlight the need for reliable optics measurements in weak-strong beam configurations.

Overall, this campaign constitutes the first experimental validation in the LHC of the direct correction of LRBB-driven RDTs and β-beating using existing magnetic circuits. It demonstrates that both non-linear and linear beam-beam effects can be diagnosed and mitigated using beam-based optics measurement techniques, and provides a concrete foundation for future beam-beam studies and correction strategies in the HL-LHC era.

CHAPTER 8

Effects of Energy Shifts

As introduced in the previous chapter, precise control of the linear optics is essential to complement the control of NL effects. In 2022, a significant optics variation was identified during LHC commissioning at low β^*. In particular, a change in the measured β-functions around the ring was observed at the end-of-squeeze ($\beta^* = 0.3$ m), when the optics were remeasured under nominally identical accelerator conditions after a period of 2.5 weeks. The magnitude of the observed deviation was substantial, reaching approximately 10% [168]. An example of the corresponding relative optics variation is shown in Fig. 8.1.

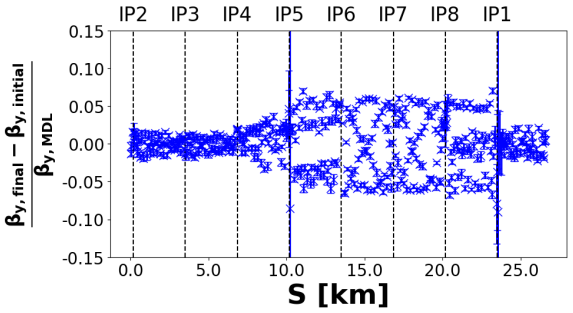


Figure 8.1: Measured relative optics shift vs s for LHCB1 between 9 and 26 May 2022 for the same accelerator configuration with $\beta^* = 30$ cm (re-analysed from historic data).

The beam time required to diagnose the origin of this shift and re-establish the desired optics configuration resulted in a considerable extension of the 2022 optics commissioning period. The observed optics deviation was ultimately traced to a relatively small relative momentum offset of approximately $\Delta p/p \approx 1 \times 10^{-4}$ [168], originating from the procedure used to establish

the nominal closed orbit at top energy. The final closed-orbit configuration is established using nominal-intensity bunches, whereas the preceding stages of optics commissioning are performed with pilot bunches. During the transition from the pilot-beam configuration to the nominal closed orbit, a systematic change in the average strength of the horizontal orbit corrector magnets introduced an unintended shift in the beam momentum.

A momentum shift associated with changes in the closed orbit can generally be interpreted in terms of orbit-length variations. For sufficiently asymmetric orbit-corrector patterns, the resulting change in orbit length, together with the constraint imposed by the RF system, produces a shift in the mean radial beam position and hence in the equilibrium momentum [169]. In the present case, however, this interpretation is not directly applicable because the corrector pattern does not exhibit the required degree of asymmetry. Instead, the observed momentum shift can be understood as arising from an effective modification of the beam rigidity, as defined in Eq. (2.2), resulting from a change in the net integrated dipole field experienced by the beam. Since the beam rigidity is directly proportional to the particle momentum for fixed charge, this modification corresponds to a relative momentum offset and, consequently, to a shift in the beam energy.

This change in the effective ρ and consequently alter the quadrupole focusing strength (see Eq. (2.6)), leading to measurable distortions of the optical functions. This effect becomes particularly pronounced at low β^*, where the optics exhibit increased sensitivity due to the large β-functions in the triplet focusing magnets, required for the β-squeeze. Consequently, the impact of such perturbations scales with the local β-function. Since 2022, similar energy-induced optics shifts have been repeatedly observed during LHC commissioning and have been sufficiently significant to require frequent additional iterations of optics correction.

The sensitivity to energy offsets is expected to increase further for the significantly reduced β^* values foreseen for HL-LHC operation. In addition, slow energy drifts, such as those induced by terrestrial tides, have traditionally been neglected during LHC optics commissioning, and no compensation has been applied over the course of individual measurement campaigns. This may have detrimental implications for optics commissioning in both the LHC and the HL-LHC.

This chapter therefore examines the projected sensitivity of HL-LHC optics corrections to beam energy errors at the end of the β^* squeeze at top energy.

8.1 Impact of energy shifts in HL-LHC

Following the completion of the LHC's third operational run in 2026, the accelerator will be upgraded to the HL-LHC [170]. A central objective of this upgrade is a substantial reduction of the minimum achievable β^* at the end-of-squeeze, targeting $\beta^* = 0.15$ m for round-optics and potentially as low as $\beta^* = 0.075$ m in a flat-optics configuration [171].

To investigate representative scenarios, both a round-optics configuration with $\beta^* = 15$ cm and a flat-optics configuration with a minimum β^* of 7.5 cm, corresponding to the lowest value targeted for operation, were considered. Such a reduction in β^* will further increase the sensitivity of the accelerator to energy-induced optics distortions. To assess the implications of these effects for HL-LHC operation and commissioning, energy shifts representative of those routinely observed during Run III commissioning were simulated using HL-LHC optics models.

The effect of an energy offset was modelled by introducing a systematic perturbation to the quadrupole strengths in the HL-LHC lattice, hence accounting for the change in the effective K values arising from the variation in beam rigidity. This was implemented without modifying the RF settings, such that the reference orbit remained unchanged, and no additional contributions from dispersion were introduced. Following the introduction of this systematic quadrupole error, the tunes were re-matched to their nominal values using the MQT circuits employed operationally for tune control. This approach has previously been demonstrated to reproduce accurately the β-beating observed in 2022 as a consequence of an energy error [168]. The study was performed for the HL-LHC round-optics configuration at end-of-squeeze ($\beta^* = 15$ cm).

An energy offset of 2×10^{-4} in $\Delta p/p$ was considered, corresponding to the upper range of energy deviations regularly encountered during LHC commissioning between 2022 and 2025. When applied to the HL-LHC model, this perturbation generated a peak β-beat of approximately 50%, as shown in Fig. 8.2 (only the vertical plane is presented, as the horizontal plane exhibited a comparable magnitude). By comparison, standard LHC optics corrections routinely reduce the residual β-beat to the level of approximately 7% [125]. It should be noted that the peak β-beating before optics corrections was approximately 256%.

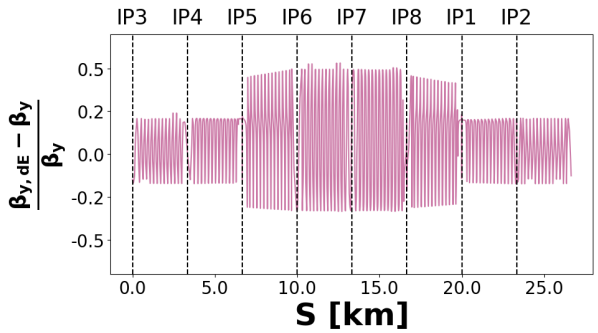


Figure 8.2: Simulated vertical β-beating (where the β-function with the quadrupole trim is defined to be $\beta_{y,\mathrm{dE}}$, and the β-function without it is β_y) for HL-LHCB1 vs s, corresponding to a $\Delta p/p$ offset of $\approx 2\times10^{-4}$ for round-optics.

The principal quadrupolar sources of the β-beating were investigated by applying strength shifts individually to different quadrupole classes. Figure 8.3 presents the peak β-beating in the vertical plane of HL-LHCB1 caused by the IR and arc magnet powerings.

As expected, the dominant contribution originates from the IR triplet magnets (Q1 to Q3, shown in green in Fig. 8.3), with the largest effect arising from the Q2 magnets (seen in Fig. 3.9). The main arc quadrupoles (QF and QD, shown in blue) also contribute at a non-negligible level. In contrast, the quadrupoles located in the matching sections were found to have a consistently minimal impact for both beams and in both transverse planes.

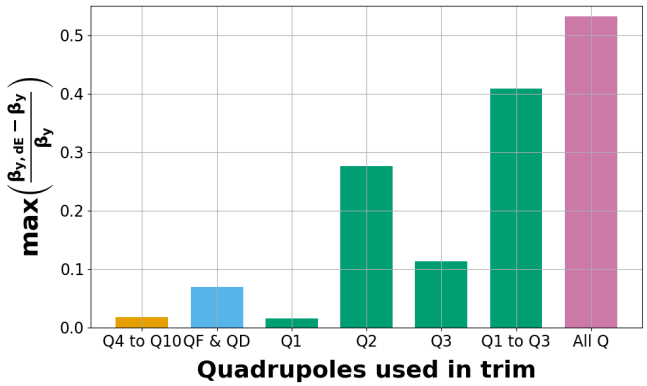


Figure 8.3: Simulated maximum vertical β-beating (same definition as in Fig. 8.2) for HL-LHCB1 for a $\Delta p/p$ offset of $\approx 2\times10^{-4}$, for various quadrupole types in the HL-LHC, where orange signifies quadrupoles in the matching sections, blue the main arc quadrupoles, green the triplet magnets, and pink is all the quadrupoles.

Other operational scenarios for the HL-LHC include the use of flat-optics [170], with a β^* configuration of 30/7.5 cm as mentioned earlier. In this configuration, significantly larger β-

beating is observed compared to the round-optics case (see Fig. 8.4), reaching values exceeding 60% for $\Delta p/p = 2 \times 10^{-4}$.

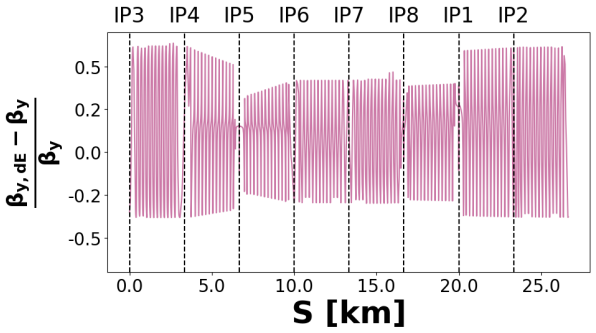


Figure 8.4: Simulated vertical β-beating (where the β-function with the quadrupole trim is defined to be $\beta_{y,\mathrm{dE}}$, and the β-function without it is β_y) vs s, for HL-LHCB1 corresponding to a $\Delta p/p$ offset of $\approx 2 \times 10^{-4}$ for flat-optics.

Given the typical magnitude of energy errors observed during LHC Run III commissioning, primarily arising from orbit set-up, the persistence of such effects in the HL-LHC era would be expected to induce significant variations in the accelerator optics throughout the commissioning period. This would represent a considerable challenge for optics commissioning.

These observations strongly motivate the development of mitigation strategies. One such approach under investigation is the automatic identification of energy shifts directly from optics measurements [172–175], with the aim of improving control over this variability during operation. In addition, for the 2025 commissioning campaign, the OP team introduced improved regulation of the average orbit corrector strengths during the orbit set-up process, with the objective of reducing the magnitude of such energy offsets. This included disabling the feedback systems only during the beam-based measurements, rather than for the entire measurement shift.

8.2 Impact of terrestrial tides on HL-LHC optics

Given the pronounced sensitivity of HL-LHC optics configurations to energy offsets, the impact of additional sources of energy drift was also considered. Among these, terrestrial tides constitute one of the most significant contributors in the LHC. The gravitational deformation of the Earth's crust induces variations in the accelerator circumference, therefore modifying the path length

of the circulating particles. The resulting mismatch between the RF frequency and the altered circumference leads to shifts in the beam energy [176].

Extensive studies of tidal effects on beam energy were previously carried out for the Large Electron-Positron Collider (LEP) [177], where such effects were of particular importance, notably in the context of precision measurements of the Z boson. In contrast, for proton-proton operation in the LHC, tidal effects on beam energy are generally less critical. During standard operation, energy drifts induced by terrestrial tides are also effectively mitigated by the LHC orbit, energy, and radial feedback systems (the radial loop) [178], which act to regulate the RF frequency and adjust the effective path length to compensate for variations in the accelerator circumference.

During optics commissioning, energy drifts induced by terrestrial tides have not historically constituted a significant concern. Optics commissioning is typically performed in shifts of 8 to 16 hours, and individual measurements can, in some cases, extend over several hours. During such periods, the energy and orbit feedback systems are often disabled for extended intervals in order to avoid interference with beam excitations required for optics measurements [179].

In this context, it is of interest to assess whether the combined effect of tidal energy variations and the increased optics sensitivity of the HL-LHC at very low β^* could lead to appreciable β-beating variations on the time scale of typical optics measurements. Figure 8.5 shows the expected variation of $\Delta p/p$, as reconstructed from the measured radial loop feedback at flat-top.

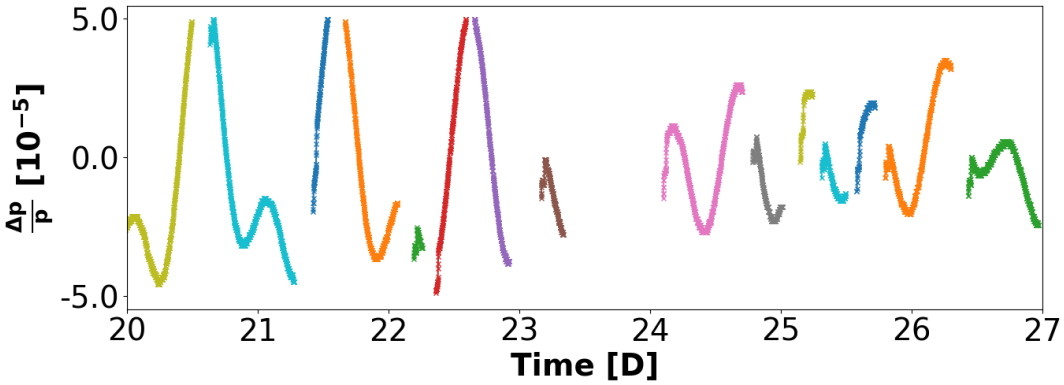


Figure 8.5: $\Delta p/p$ vs s, calculated from the LHCB1 RF frequency data over a one-week period in August in 2024, filtered on flat-top periods; the different colours signify separation of shifts.

From Fig. 8.5, the peak-to-peak variation in $\Delta p/p$ is observed to be of the order of 10^{-4}. The rate of change of $\Delta p/p$ over time scales relevant to optics commissioning was evaluated using

both radial loop measurements and simulations of tidal effects performed with the LHC steering program: Yet Another Steering Programme (YASP) [180]. Fluctuations of the $\Delta p/p$ were of similar magnitude in both the simulated and measured data.

Figure 8.6 presents the predicted rate of change of $\Delta p/p$ per hour, obtained from simulations over several months. A clear cyclical behaviour is observed, with the highest rates corresponding to periods of spring tides. Figure 8.7 shows the distribution of the expected variation in $\Delta p/p$ due to simulated tidal effects over a 200-day period, for several time scales relevant to optics commissioning.

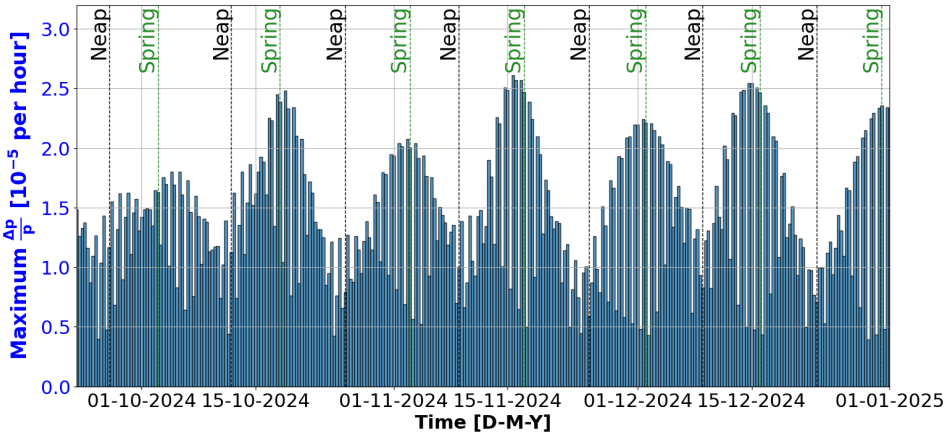


Figure 8.6: Maximum change in YASP-predicted $\Delta p/p$ over an hour vs time. Spring (green) and neap (black) tide times taken from [8].

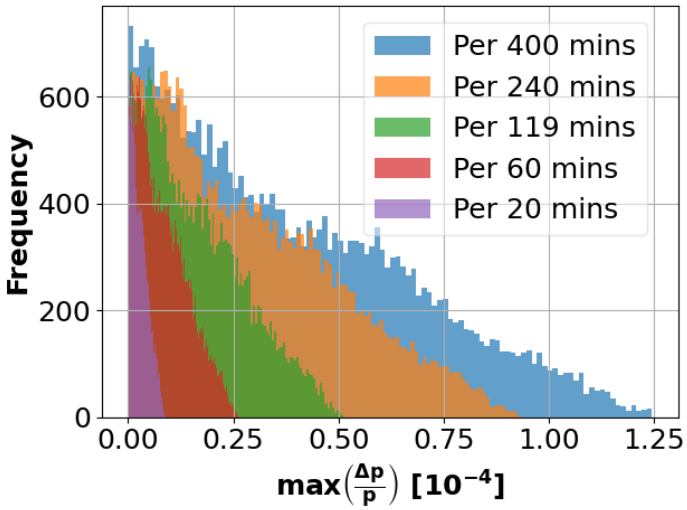


Figure 8.7: Distribution of the predicted tidal variation in $\Delta p/p$, predicted on several time intervals relevant to optics commissioning.

With regard to the impact of terrestrial tides, the scenario previously presented in Fig. 8.4 may represent an overly pessimistic case. Modelling an energy shift on the reference orbit, as is appropriate for orbit set-up induced offsets, does not introduce any sextupolar contribution. A more representative description of tidal effects corresponds to an off-momentum orbit in a lattice of modified effective length, leading predominantly to chromatic β-beating, as illustrated in Fig. 8.8.

In this scenario, sextupoles act to mitigate the maximum off-momentum β-beating, which also tends to be more localised around the experimental IRs. Nevertheless, during optics commissioning, local corrections in the IRs remain of primary importance, and residual errors in these areas continue to be of significant concern.

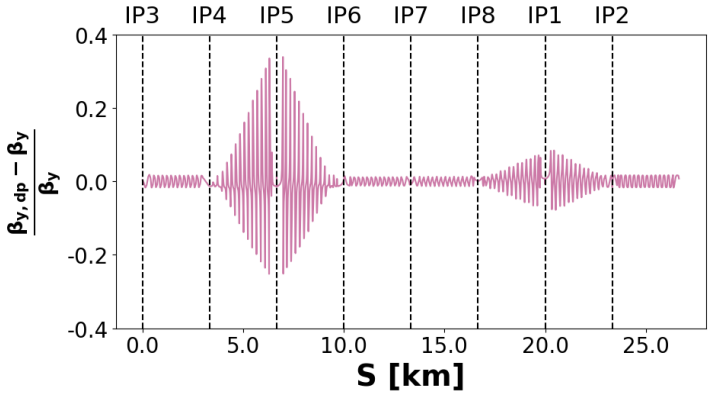


Figure 8.8: Simulated vertical off-momentum β-beating (where the β-function with the RF frequency trim is defined to be $\beta_{y,\mathrm{dp}}$, and the β-function without it is β_y) vs s, for HL-LHCB1 corresponding to a $\Delta p/p$ offset of $\approx 2\times10^{-4}$ for flat-optics.

Figure 8.9 presents the variation of the peak off-momentum β-beating as a function of $\Delta p/p$ for both flat- and round- HL-LHC optics configurations. By combining this chromatic β-beating with the expected distribution of $\Delta p/p$ variations shown in Fig. 8.7, it is possible to estimate the magnitude of tide-induced peak optics variations. The resulting distribution, evaluated over time scales relevant to optics commissioning, is shown in Fig. 8.10, and indicates a non-negligible impact.

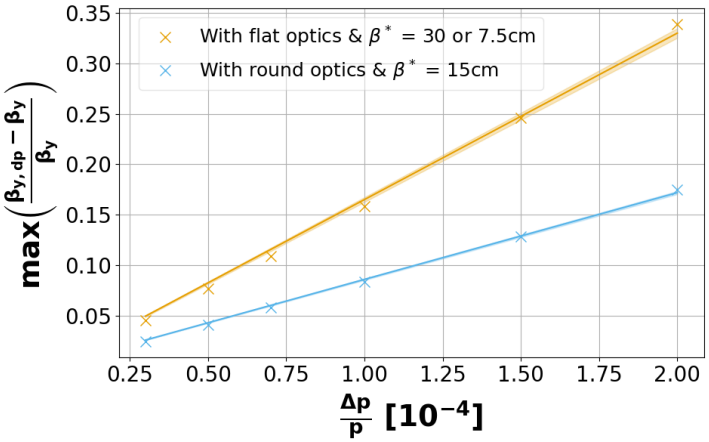


Figure 8.9: Predicted maximum vertical off-momentum β-beating (same definition as for Fig. 8.8) HL-LHCB1 as a function of $\Delta p/p$ where the lines are a linear fit of the points.

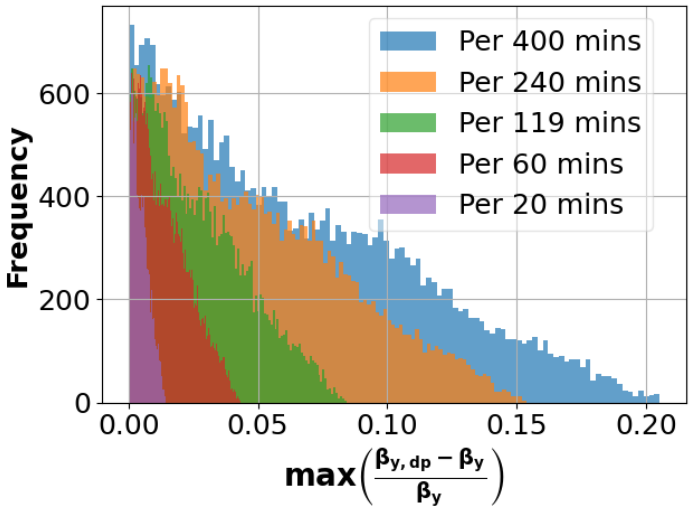


Figure 8.10: Frequency distribution of the predicted maximum change in Δchrom β_y/β_y for HL-LHCB1 flat-optics over several different time intervals.

SbS (see Section 5.4 for further details) constitutes the primary analysis tool used to determine local optics corrections in the LHC IRs. It is therefore of interest to assess whether this revised modelling of energy errors can be observed in the reconstructed phase advance and, consequently, whether it is of the level where it could spoil the measurements used to determine the local IR correction. This can be investigated by systematically comparing the phase advance obtained under off-momentum conditions with that corresponding to the nominal (on-momentum) case.

Figure 8.11 presents the SbS analysis of the simulated phase advance, evaluated from a BPM

located to the left of IP5 in a flat-optics configuration. The results indicate that a momentum variation corresponding to that expected over a period of 2 hours during peak spring tide conditions induces a deviation in the reconstructed phase advance of a magnitude comparable to that typically addressed through optics corrections. A duration of approximately 2 hours is representative of a standard SbS measurement and analysis cycle.

The resulting peak phase deviation, of the order of 0.01 2π, lies at the threshold at which corrections would generally be considered in local optics analysis. This is comparable to phase errors observed and corrected during commissioning, for example a value of 0.03 2π reported in recent SbS measurements [181].

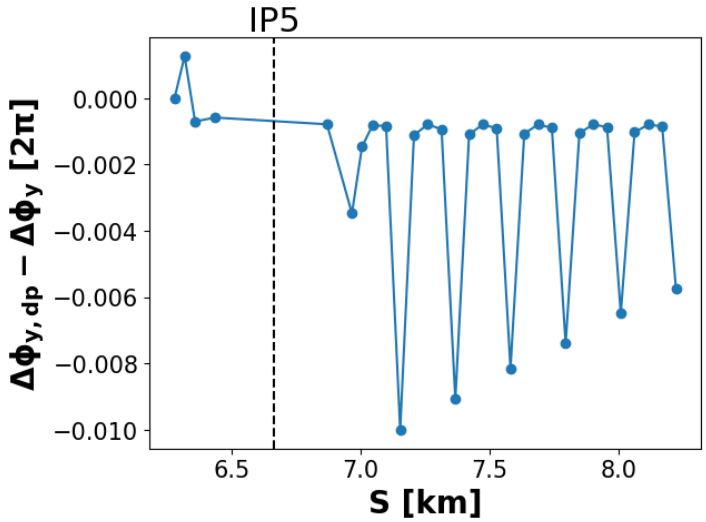


Figure 8.11: Simulation of the SbS phase advance vs s, starting from BPM.10L5, expected to be generated over a 2 hour period at the worst time during the spring tide with flat-optics. The line is included only to guide the eye.

In light of these findings, operational procedures during commissioning have been revised to ensure that the radial loop is not left disabled for extended periods during optics measurements. In addition, recent updates to the operational tools for the ACD now allow the automatic toggle of feedback systems during beam excitation, reducing the risk of unintended omission of the procedure.

Looking ahead to HL-LHC commissioning, it will remain essential to maintain this increased level of attention during optics measurements, in particular with regard to controlling radial drifts induced by terrestrial tides.

8.3 Conclusion

Repeated observations of optics variations during LHC commissioning, arising from energy errors induced by closed-orbit set-up, have raised concerns for operation in the HL-LHC era, where the reduced β^* is expected to amplify these effects. Simulations of energy shifts representative of those typically encountered in the LHC, when applied to HL-LHC end-of-squeeze optics, indicate that the resulting β-beating can be substantial. This highlights the need for improved control of orbit-induced energy errors in HL-LHC operation, with dedicated studies currently under way.

In view of the increased sensitivity of HL-LHC optics to energy offsets, the potential impact of energy variations induced by terrestrial tides was also investigated. Historically, such effects have not been actively controlled during optics commissioning. The analysis indicates that the resulting β-beating variations can become significant on time scales relevant to optics measurements. Consequently, optics measurement procedures at very low β^* should be adapted to ensure more careful compensation of tidal drifts during optics commissioning.

CHAPTER 9

CONCLUSION

By measuring and correcting forced-RDTs using ACD-driven oscillations, this work establishes a comprehensive and experimentally validated methodology for the identification, characterisation, and correction of non-linear optics errors in the LHC. The results have direct implications for future HL-LHC operation, as well as for other accelerators operating in strongly non-linear regimes.

The thesis first exploited forced-RDTs to quantify and benchmark normal sextupolar resonances in the LHC at injection. Good agreement was observed between measurements and the model; however, an unexpectedly large normal sextupolar resonance was identified, whose impact had previously been neglected. It was subsequently demonstrated that simulation-driven correction strategies based on free-RDTs can be reliably translated into measurable improvements in the forced resonance behaviour, and overall beam performance.

The successful operational implementation of sextupolar resonance corrections, resulting in enhanced beam lifetime and dynamic aperture, confirms the practical relevance of these methods. Furthermore, the extension of measurement capabilities to dynamically evolving accelerator conditions constitutes a significant methodological advance, enabling non-linear optics studies to be performed under non-static operational scenarios for the first time.

A detailed investigation of Landau Octupole-driven octupolar resonances at injection energy

revealed previously unobserved discrepancies with the magnetic model, ultimately traced to hardware-related effects consistent with circuit hysteresis, significant only at injection energy. The introduction of a novel non-linear Segment-by-Segment (SbS) diagnostic approach inspired by an existing linear SbS method enabled the localisation of problematic circuits with minimal beam time, illustrating the potential of advanced beam-based techniques for efficient accelerator diagnostics. These findings underline the importance of incorporating hardware effects into optics models, particularly in regimes where accelerator reproducibility is critical.

A long-standing challenge in collider operation is the control of non-linear errors in the experimental insertions, where large β-functions enhance the sensitivity to higher-order magnetic field errors. This challenge was addressed through the development of a novel correction methodology targeting higher-order field errors in the interaction regions via RDT feed-down. This represents a substantial achievement and was successfully validated through beam-based measurements.

By exploiting feed-down mechanisms, the method enables the effective compensation of non-linear errors without requiring the direct measurement of higher-order RDTs or detuning coefficients. Consequently, the time required to characterise higher-order errors, such as dodecapolar components, is substantially reduced, rendering their inclusion in routine yearly commissioning both feasible and practical.

The practical impact of this methodology was established through the successful correction of skew octupolar errors in the ATLAS insertion, leading to a significant reduction in the strength of the $3Q_y$ resonance in the operational configuration, for which skew octupole feed-down constitutes a principal contribution.

Improved control of such resonances is of considerable importance for both LHC and HL-LHC operation, as they have been associated with collimator hierarchy breaking, which had a notable operational impact, for example during the 2024 run. Furthermore, the observation of dodecapolar feed-down effects, together with the benchmarking of the corresponding magnetic model, confirms both the sensitivity and robustness of the measurement framework, extending its applicability to higher-order field components.

Beyond local field errors, this thesis presents the first experimental validation of direct correction strategies for long-range beam-beam driven non-linear RDTs using common-aperture

magnets. This approach was also applied to the linear regime for the correction of long-range beam-beam induced β-beating.

The demonstrated capability to extract, model, and correct these effects constitutes a significant step towards integrating the correction of long-range beam-beam interactions into routine optics commissioning, therefore extending beyond the traditional single-particle paradigm.

Collectively, the results presented here have established forced non-linear RDT analysis as a standard and reliable tool for LHC commissioning and operation. This capability has already enabled higher-order studies, improved dynamic aperture, and facilitated new correction schemes. These developments open the way towards a unified framework in which both non-linear and beam-beam effects can be simultaneously measured and corrected using beam-based techniques.

Finally, this work highlights emerging challenges for HL-LHC operation, in particular the increased sensitivity of the optics to energy offsets and external perturbations, such as terrestrial tidal effects, during commissioning at extremely low β^*. The illustrated impact of these phenomena on β-beating underscores the need for improved control strategies and adapted measurement procedures in future commissioning campaigns.

In summary, this thesis establishes a comprehensive experimental and methodological framework for the enhancement of non-linear optics control in the LHC and HL-LHC through the application of forced-RDTs. Many synchrotrons, such as SuperKEKB and ESRF, have expressed interest in the direct measurement of resonance strengths using RDTs; however, these measurements have traditionally been highly challenging. This work demonstrates that forced-RDTs constitute a powerful and practical tool for the benchmarking and optimisation of non-linear optics, extending to very high orders and remaining effective even for strongly non-linear accelerator configurations.

The techniques developed in this thesis not only enhance present accelerator performance but also provide essential tools to ensure the robustness and operational flexibility required for future high-luminosity hadron collider operation.

Several avenues for future work naturally arise from these results. A normal sextupole correction scheme for implementation during the next commissioning period should be devised. In addition, the projected increase in DA resulting from the normal sextupole correction should

be investigated.

The non-linear SbS methodology should further developed and applied in other studies, and the approach should be directly implemented within the MAD-NG framework.

The higher-order feed-down correction could similarly benefit from calculating the corrector response directly within MAD-NG. Such an implementation would provide greater flexibility in generating and applying corrections during machine operation, therefore improving its practicality for use in the control room.

For the LRBB corrections, further investigation is required to develop a method capable of correcting the f_{2002} RDT. In addition, the origin of the observed discrepancy between simulated and measured responses in the linear optics regime should be understood.

Finally, the terrestrial tides study could be extended by investigating the effects of physically increasing the lattice length to emulate the deformation induced by terrestrial tides. Such a study would help establish whether the observed changes in the chromatic optics provide an accurate representation of the underlying physical effects.

BIBLIOGRAPHY

[1] E. H. Maclean, R. Tomás, F. Schmidt *et al.*, "Measurement of nonlinear observables in the Large Hadron Collider using kicked beams," *Phys. Rev. ST Accel. Beams*, vol. 17, p. 081002, Aug. 2014. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.17.081002

[2] E. Lopienska, "The CERN accelerator complex, layout in 2022." 2022. [Online]. Available: https://cds.cern.ch/record/2800984

[3] AC Team, "Diagram of an LHC dipole magnet," 1999. [Online]. Available: https://cds.cern.ch/record/40524

[4] J.-L. Caron, "LHC quadrupole cross section." 1998, AC Collection. Legacy of AC. Pictures from 1992 to 2002. [Online]. Available: https://cds.cern.ch/record/841485

[5] L. Fiscarelli, M. Pentella, G. Dima *et al.*, "Magnetic measurements at 1.9 K on the LMBRD01 (D2 cold mass with MBRD2)," Nov. 2025, WP3 Meeting. [Online]. Available: https://indico.cern.ch/event/1603268/contributions/6755778/attachments/3176608/5649811/LFiscarelli_MM_LMBRD1_19112025_a.pdf

[6] L. Fiscarelli, F. Beltron, S. Izquierdo Bermudez *et al.*, "Field quality of MQXFB magnets," Sep. 2025, 15th HL-LHC Collaboration Meeting. [Online]. Available: https://indico.cern.ch/event/1559978/contributions/6664968/

[7] E. Maclean, R. Tomás, T. Persson *et al.*, "Review of linear and nonlinear optics measurements in the CERN LHC," in *16th International Particle Accelerator Conference proceedings*, no. FRYD2, 2025, pp. 3250–3255. [Online]. Available: https://cds.cern.ch/record/2952834/files/document.pdf

[8] S. Wright. (2024) Spring and Neap tides 2024. Accessed 13 February 2025. [Online]. Available: https://tides.today/en/journal/spring-tides-2024

[9] L. Deniau, "MAD-NG Reference Manual," 2023, accessed 18 March 2026. [Online]. Available: https://madx.web.cern.ch/madx/releases/madng/html/index.html

[10] S. Myers and H. Schopper, Eds., *Particle Physics Reference Library . Volume 3: Accelerators and Colliders*. Springer, 2020.

[11] H. Wiedemann, *Particle Accelerator Physics: Third Edition*. Springer Berlin, Heidelberg, 2007, DOI: 10.1007/978-3-540-49045-6/COVER.

[12] E. H. Maclean, "Modelling and correction of the non-linear transverse dynamics of the LHC from beam-based measurements," Ph.D. dissertation, Hertford College, Oxford, 2014. [Online]. Available: https://cds.cern.ch/record/1951379

[13] R. Wolf, "Field Error Naming Conventions for LHC Magnets," CERN, Geneva, Tech. Rep. LHC-M-ES-0001, 2001. [Online]. Available: https://lhc-div-mms.web.cern.ch/tests/MAG/FiDeL/Documentation/lhc-m-es-0001-30-00.pdf

[14] H. Grote, G. Roy, and F. Schmidt, "The MAD-X Program User's Reference Manual," 2022. [Online]. Available: https://mad.web.cern.ch/mad/webguide/manual.html

[15] E. Courant and H. Snyder, "Theory of the alternating-gradient synchrotron," *Annals of Physics*, vol. 3, no. 1, pp. 1–48, 1958. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0003491658900125

[16] S. Guiducci, "Chromaticity," in *CAS - CERN Accelerator School : 5th General Accelerator Physics Course*, 1994. [Online]. Available: https://cds.cern.ch/record/398300

[17] CERN, SixTrack Homepage. Accessed 13 March 2026. [Online]. Available: https://lhcathome.web.cern.ch/projects/sixtrack

[18] G. Iadarola *et al.*, "Xsuite: An Integrated Beam Physics Simulation Framework," in *HB2023*, vol. HB2023, no. TUA2I1, 2024, pp. 73–80. [Online]. Available: https://cds.cern.ch/record/2882300/files/document.pdf

[19] A. J. Dragt, *Lie Methods for Nonlinear Dynamics with Applications to Accelerator Physics*. University of Maryland, 2005. [Online]. Available: https://inspirehep.net/files/904b6c446c2b8f728b8e5168e104a925

[20] L. Deniau, J. Gray, A. Teska *et al.*, "Methodicalacceleratordesign/mad-ng: v1.1.12," Apr. 2026. [Online]. Available: https://doi.org/10.5281/zenodo.19454632

[21] L. Deniau, "MAD-NG, a standalone multi-platform tool for linear and non-linear optics design and optimisation," CERN ATS Seminar 4 Dec 2025. [Online]. Available: https://indico.cern.ch/event/1615996/attachments/3186419/5670242/ld_2512-mad_ats_slides.pdf

[22] F. S. Carlier, R. Tomás, E. H. Maclean *et al.*, "First experimental demonstration of forced dynamic aperture measurements with LHC AC dipoles," *Phys. Rev. Accel. Beams*, vol. 22, p. 031002, Mar. 2019. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.031002

[23] R. Meller, A. Chao, J. Peterson *et al.*, "Decoherence of Kicked Beams," Fermilab, Tech. Rep. SSC-N-360, 1987. [Online]. Available: https://lss.fnal.gov/archive/other/ssc/ssc-n-360.pdf

[24] G. Russo, "Precise tune determination and split beam emittance reconstruction at the CERN PS synchrotron," Ph.D. dissertation, Universität Frankfurt, 2023. [Online]. Available: https://inspirehep.net/files/c1ffb1594f869b6b1cfa34e46ae054cf

[25] A. Wolski, *Beam Dynamics in High Energy Particle Accelerators*. Imperial College Press, 2014. [Online]. Available: https://www.worldscientific.com/doi/abs/10.1142/p899

[26] H. Goldstein, C. P. Poole, and J. L. Safko, *Classical Mechanics*, 3rd ed. Addison-Wesley, San Francisco, 2002.

[27] É. Forest, *Beam Dynamics A New Attitude and Framework*. CRC Press, 1998. [Online]. Available: https://doi.org/10.1201/9781315138176

[28] A. Franchi, "Studies and Measurements of Linear Coupling and Nonlinearities in Hadron Circular Accelerators," Ph.D. dissertation, Universität Frankfurt/M., 2006. [Online]. Available: https://inspirehep.net/files/ae76ed6644ae0aeb0e5a9926346e8f31

[29] R. Tomás, "Direct Measurement of Resonance Driving Terms in the Super Proton Synchrotron (SPS) of CERN using Beam Position Monitors," Ph.D. dissertation, Universitat de València, 2003. [Online]. Available: https://repository.cern/records/bjrvf-ns620

[30] A. Bazanni, E. Todesco, and G. Turchetti, "A normal form approach to the theory of nonlinear betatronic motion," CERN, Tech. Rep. 02, 1994. [Online]. Available: http://cds.cern.ch/record/262179?ln

[31] W. Herr and B. Muratori, "Concept of luminosity," in *CAS - CERN Accelerator School: Intermediate Accelerator Physics*, 2006, pp. 361–378. [Online]. Available: https://cds.cern.ch/record/941318/files/p361.pdf

[32] J. Coupard, H. Damerau, A. Funken *et al.*, "LHC Injectors Upgrade, Technical Design Report," CERN, Tech. Rep. CERN-ACC-2014-0337, 2014. [Online]. Available: https://cds.cern.ch/record/1976692

[33] W. Herr, "Features and implications of different LHC crossing schemes," CERN, Geneva, Tech. Rep. LHC-Project-Report-628, CERN-LHC-Project-Report-628, 2003. [Online]. Available: https://cds.cern.ch/record/604005

[34] O. S. Brüning, J. Poole, P. Collier *et al.*, *LHC Design Report*. CERN, 2004, vol. 1. [Online]. Available: https://cds.cern.ch/record/782076

[35] L. Rossi, "The LHC Main Dipoles and Quadrupoles toward Series Production," *IEEE Trans. Appl. Supercond.*, vol. 13, no. LHC-Project-Report-623, CERN-LHC-Project-Report-623, pp. 1221–1228, 2003. [Online]. Available: https://cds.cern.ch/record/599389

[36] V. I. Ptitsyn, S. Tepikian, and J. Wei, "BNL-built LHC magnet error impact analysis and compensation," in *18th Biennial Particle Accelerator Conference Proceedings*, no. TUP43, 1999, pp. 1575–1577. [Online]. Available: https://cds.cern.ch/record/419550

[37] B. Bellesia, F. Bertinelli, P. Fessia *et al.*, "Influence of Copper Wedge Dimensions on Field Harmonics in the Main LHC Dipoles," CERN, Geneva, Tech. Rep. LHC-Project-Report-630, CERN-LHC-Project-Report-630, 2003. [Online]. Available: https://cds.cern.ch/record/611583

[38] T. H. B. Persson, Y. Inntjore Levinsen, R. Tomás *et al.*, "Chromatic coupling correction in the Large Hadron Collider," *Phys. Rev. ST Accel. Beams*, vol. 16, p. 081003, Aug. 2013. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.16.081003

[39] T. H. B. Persson, "Beam-Based Error Identification and Correction Methods for Particle Accelerators," Ph.D. dissertation, Chalmers University of Technology, Jun. 2014. [Online]. Available: https://repository.cern/records/zdagb-cpx34

[40] N. Aquilina, M. Lamont, M. Strzelczyk *et al.*, "Chromaticity decay due to superconducting dipoles on the injection plateau of the Large Hadron Collider," *Phys. Rev. ST Accel. Beams*, vol. 15, p. 032401, Mar. 2012. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.15.032401

[41] M. Haverkamp, "Decay and Snapback in Superconducting Accelerator Magnets," PhD Thesis - Research External, Graduation UT, University of Twente, Netherlands, Oct. 2003. [Online]. Available: https://s3.cern.ch/inspire-prod-files-5/534c4c00cc03ff9a3878d8b2c2069263

[42] CERN, "WISE homepage," accessed 12 March 2026. [Online]. Available: http://wise.web.cern.ch/WISE/

[43] CERN, "FiDeL Group Documentation on the Magnetic Model of the LHC." [Online]. Available: https://web.archive.org/web/20250606214413/https://lhc-div-mms.web.cern.ch/tests/MAG/Fidel/

[44] CERN, "LHC The Guide FAQ," Apr. 2021. [Online]. Available: https://cds.cern.ch/record/2809109/files/CERN-Brochure-2021-004-Eng.pdf

[45] V. Balagura, "Van der Meer scan luminosity measurement and beam–beam correction," *Eur. Phys. J. C*, vol. 81, no. 1, p. 26, 2021.

[46] S. van der Meer, "Calibration of the effective beam height in the ISR," CERN, Geneva, Tech. Rep., 1968. [Online]. Available: https://cds.cern.ch/record/296752

[47] T. Mastorides, C. Rivetta, J. D. Fox *et al.*, "RF system models for the CERN Large Hadron Collider with application to longitudinal dynamics," *Phys. Rev. ST Accel. Beams*, vol. 13, p. 102801, Oct. 2010. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.13.102801

[48] L. Bottura, M. Lamont, E. Todesco *et al.*, "Pre-Cycles of the LHC Magnets during Operation," CERN, Geneva, Tech. Rep. CERN-ATS-2010-174, 2010. [Online]. Available: https://cds.cern.ch/record/1283477

[49] J. Wenninger, "Functional specification LHC modes," CERN, Geneva, Tech. Rep. LHC-OP-ES-0022 rev 2.1, 2019. [Online]. Available: https://edms.cern.ch/ui/file/1070479/2.1/LHC-OP-ES-0022-V2.1.pdf

[50] R. Bailey and P. Collier, "Standard Filling Schemes for Various LHC Operation Modes," CERN, Geneva, Tech. Rep. LHC-PROJECT-NOTE-323, 2003. [Online]. Available: https://cds.cern.ch/record/691782

[51] J. Wenninger, "LHC Report: operation with holes," Sep. 2017.

[52] A. Boccardi, M. Gasior, R. Jones *et al.*, "First Results from the LHC BBQ Tune and Chromaticity Systems," CERN, Geneva, Tech. Rep. LHC-Performance-Note-007, CERN-LHC-Performance-Note-007, 2009. [Online]. Available: https://cds.cern.ch/record/1156349

[53] M. Gasior and R. Jones, "High Sensitivity Tune Measurement by Direct Diode Detection," in *7th European Workshop on Beam Diagnostics and Instrumentation for Particle*

Accelerators proceedings, no. CERN-AB-2005-060, Jun. 2005. [Online]. Available: https://cds.cern.ch/record/895142/files/ab-2005-060.pdf

[54] T. Persson and R. Tomás, "Improved control of the betatron coupling in the Large Hadron Collider," *Phys. Rev. ST Accel. Beams*, vol. 17, p. 051004, May 2014. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.17.051004

[55] R. Calaga and R. Tomás, "Statistical analysis of RHIC beam position monitors performance," *Phys. Rev. ST Accel. Beams*, vol. 7, p. 042801, Apr. 2004. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.7.042801

[56] A. A. Nosych, U. Iriso, and A. Olmos, "Overview of the geometrical non-linear effects of button BPMs and methodology for their efficient suppression," in *Proceedings of IBIC14*, Monterey, 2014. [Online]. Available: https://proceedings.jacow.org/IBIC2014/papers/tupf03.pdf

[57] A. Nosych, "Geometrical non-linearity correction procedure of LHC Beam Position Monitors," CERN, Tech. Rep. 1342295, 2014. [Online]. Available: https://edms.cern.ch/ui/#!master/navigator/document?D:1894463902:1894463902:subDocs

[58] E. C. Giraldo, C. Boccard, L. Jensen *et al.*, "BPM system upgrades and temperature regulation," LHC Optics Measurement and Corrections review 17-18th Jun 2013. [Online]. Available: https://indico.cern.ch/event/246159/contributions/533854/

[59] E. B. Holzer, B. Dehning, E. Effinger *et al.*, "Beam Loss Monitoring for LHC Machine Protection," *Physics Procedia*, vol. 37, pp. 2055–2062, 2012, Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011). [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1875389212019281

[60] S. Morales Vigo, B. Salvachua, C. P. Welsch *et al.*, "Beam lifetime monitoring using beam loss monitors during LHC Run 3," in *14th International Particle Accelerator Conference (IPAC 2023) proceedings*, no. THPL086, May 2023, pp. 4645–4648. [Online]. Available: https://cds.cern.ch/record/2886324

[61] D. Belohrad, L. Jensen, O. Jones *et al.*, "The LHC Fast BCT system: A comparison of Design Parameters with Initial Performance," CERN, Geneva, Tech. Rep. CERN-BE-2010-010, 2010. [Online]. Available: https://cds.cern.ch/record/1267400

[62] M. Krupa and M. Gasior, "A New Wall Current Transformer for Accurate Beam Intensity Measurements in the Large Hadron Collider," *MDPI Energies*, vol. 16, no. 21, p. 7442, 2023. [Online]. Available: https://cds.cern.ch/record/2882888

[63] M. Sapinski and T. Kroyer, "Operational limits of wire scanners on LHC beam," in *13th Beam Instrumentation Workshop (BIW08)*, no. TUPTPF066, May 2008, pp. 274–278. [Online]. Available: https://proceedings.jacow.org/BIW2008/papers/tuptpf066.pdf?n=BIW2008/papers/TUPTPF066.pdf

[64] E. Bravin, G. Burtin, A. Fisher *et al.*, "First Beam Measurements with the LHC Synchrotron Light Monitors," in *1st International Particle Accelerator Conference proceedings*, no. MOPE057, 2010, pp. 1104–1106. [Online]. Available: https://cds.cern.ch/record/1271742

[65] D. Butti, E. Bravin, S. Burger *et al.*, "Status of the Synchrotron Radiation Telescope at the Large Hadron Collider," in *14th International Beam Instrumentation Conference (IBIC 2025)*, 2025, p. 342. [Online]. Available: https://cds.cern.ch/record/2954936

[66] E. Carlier, L. Ducimetiere, and E. Vossenberg, "A Kicker Pulse Generator for Measurement of the Tune and Dynamic Aperture in the LHC," in *Conference Record of the 2006 Twenty-Seventh International Power Modulator Symposium*, 2006, pp. 463–466. [Online]. Available: https://ieeexplore.ieee.org/document/4216236

[67] R. Miyamoto, M. Cattin, J. Serrano *et al.*, "Signal quality of the LHC AC Dipoles and its impact on beam dynamics," in *1st International Particle Accelerator Conference proceedings*, no. CERN-ATS-2010-063, 2010, pp. 1104–1106. [Online]. Available: http://cds.cern.ch/record/1269893

[68] J. Serrano and M. Cattin, "The LHC AC Dipole system: an introduction," CERN, Tech. Rep. CERN-BE-Note-2010-014, 2010. [Online]. Available: https://cds.cern.ch/record/1263248

[69] R. Tomás, "Adiabaticity of the ramping process of an AC dipole," *Phys. Rev. ST Accel. Beams*, vol. 8, p. 024401, Feb. 2005. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.8.024401

[70] D. Valuch, G. Kotzian, W. Höfle *et al.*, "ADT operation. A cook book of black magic," 2016, lBOC Meeting No. 61. [Online]. Available: https://indico.cern.ch/event/527889/contributions/2168749/attachments/1278254/1897561/ADT_operation.pdf

[71] D. Boussard, W. Höfle, and T. P. R. Linnecar, "The LHC Transverse Damper (ADT) Performance Specification," CERN, Geneva, Tech. Rep. CERN-SL-Note-99-055-HRF, 1997. [Online]. Available: http://cds.cern.ch/record/702559

[72] T. H. B. Persson, G. Baud, J. M. Coello De Portugal Martinez Vazquez *et al.*, "Linear coupling dependence on intensity and a next step towards a feedback (MD1850)," CERN, Geneva, Tech. Rep. CERN-ACC-2017-0008, 2017. [Online]. Available: https://cds.cern.ch/record/2255976

[73] B. Salvachua, R. W. Assmann, R. Bruce *et al.*, "LHC collimation cleaning and operation outlook," in *4th Evian Workshop on LHC beam operation proceedings, Evian-les-Bains, France, 17 - 20 Dec*, 2012, pp. 155–160. [Online]. Available: https://cds.cern.ch/record/2302435

[74] K. Paraschou and X. Buffat, "Impact of phase-space distortions on the collimator hierarchy," 2024, LBOC meeting 167. [Online]. Available: https://indico.cern.ch/event/1420698/contributions/5973654/attachments/2870529/5025491/20230504_LBOC_phase_space_distortions.pdf

[75] R. Tomás, X. Buffat, J. Coello *et al.*, "Imperfections and corrections," in *CERN Accelerator Course on Numerical Methods for Analysis, Design and Modelling of Particle Accelerators proceedings*, 2020, 31 pages, contribution to the CAS - CERN Accelerator School: Numerical Methods for Analysis, Design and Modelling of Particle Accelerators, 11-23 November 2018, Thessaloniki, Greece. [Online]. Available: https://cds.cern.ch/record/2923746

[76] L. Malina, "Harpy: A Fast, Simple and Accurate Harmonic Analysis with Error Propagation," in *13th International Particle Accelerator Conference proceedings*, no. WEPOMS035, 2022, pp. 2326–2329. [Online]. Available: https://proceedings.jacow.org/ipac2022/papers/wepoms035.pdf

[77] M. Aiba, S. Fartoukh, A. Franchi *et al.*, "First β-beating measurement and optics analysis for the CERN Large Hadron Collider," *Phys. Rev. ST Accel. Beams*, vol. 12, p. 081002, Aug. 2009. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.12.081002

[78] P. Castro, "Luminosity and beta function measurement at the electron-positron collider ring LEP," Ph.D. dissertation, Universitat de València, Mar. 1996. [Online]. Available: https://repository.cern/records/eny2v-4y338

[79] A. Wegscheider, A. Langner, R. Tomás *et al.*, "Analytical n beam position monitor method," *Phys. Rev. Accel. Beams*, vol. 20, p. 111002, Nov. 2017. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.111002

[80] E. J. Høydalsvik, T. H. B. Persson, and R. Tomas García, "Evaluation of the closest tune approach and its MAD-X implementation," CERN, Tech. Rep. CERN-ACC-NOTE-2021-0022, 2021. [Online]. Available: https://cds.cern.ch/record/2778887

[81] S. White, E. Maclean, and R. Tomás, "Direct amplitude detuning measurement with AC dipole," *Phys. Rev. ST Accel. Beams*, vol. 16, p. 071002, Jul. 2013. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.16.071002

[82] A. Franchi, L. Farvacque, F. Ewald *et al.*, "First simultaneous measurement of sextupolar and octupolar resonance driving terms in a circular accelerator from turn-by-turn beam position monitor data," *Phys. Rev. ST Accel. Beams*, vol. 17, p. 074001, Jul. 2014. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.17.074001

[83] O. S. Brüning, F. Schmidt, W. Fischer *et al.*, "Comparison of measured and computed dynamic aperture for the SPS and the HERA proton ring," *Part. Accel.*, vol. 54, pp. 223–235, 1996. [Online]. Available: https://cds.cern.ch/record/316947

[84] L. C. Teng, "Error Analysis for the Low β Quadrupoles of the Tevatron Collider / Location and Strengths of Correction Dipoles in the Low β Insertion," FERMILAB, Tech. Rep. FERMILAB-TM-1097, Mar. 1982. [Online]. Available: https://inspirehep.net/files/f8c2171cde88615cf40e155280e7e70f

[85] M. Benedikt, F. Schmidt, R. Tomás *et al.*, "Driving term experiments at CERN," *Phys. Rev. ST Accel. Beams*, vol. 10, p. 034002, Mar. 2007. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.10.034002

[86] F. Pilat, S. Binello, P. Cameron *et al.*, "Non-linear Effects in the RHIC Interaction Regions: Measurement and Correction," in *PAC 2003 proceedings*, May 2003, pp. 1703–1705. [Online]. Available: https://proceedings.jacow.org/p03/PAPERS/TPPB041.PDF

[87] W. Fischer, F. Schmidt, and R. Tomás, "Measurement of sextupolar resonance driving terms in RHIC," CERN, BNL, Tech. Rep. CERN-AB-2003-032-ABP, CERN-SL-2002-050-AP, BNL-C-AD-AP-81, Oct. 2003. [Online]. Available: https://cds.cern.ch/record/618884

[88] C. Caliari, A. Oeftiger, and O. Boine-Frankenheim, "Beam-based identification of magnetic field errors in a synchrotron using Deep Lie map networks," *Phys. Rev. Accel. Beams*, vol. 28, p. 024601, Feb. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.28.024601

[89] Y. Ohnishi, H. Koiso, A. Morita *et al.*, "Dynamic Aperture Optimization in SuperKEKB," in *Proceedings of HF2014, Beijing, China*, 2015, pp. 73–78. [Online]. Available: https://proceedings.jacow.org/HF2014/papers/frt1b2.pdf?n=HF2014/papers/FRT1B2.pdf

[90] A. Ropert and L. Farvacque, "Non-linear optics studies at the ESRF," in *PACS2001. Proceedings of the 2001 Particle Accelerator Conference (Cat. No.01CH37268)*, 2001, pp. 1702–1704 vol.3. [Online]. Available: https://proceedings.jacow.org/p01/PAPERS/TPPH019.PDF

[91] H. Grote, F. Schmidt, and L. H. A. Leunissen, "LHC Dynamic Aperture at Collision," CERN, Geneva, Tech. Rep. LHC-PROJECT-NOTE-197, 1999. [Online]. Available: https://cds.cern.ch/record/691988

[92] J. Poole, "Field quality issues for LHC magnets: Analysis and perspectives for quadrupoles and separation dipoles," CERN, Geneva, Tech. Rep. CERN-AB-2004-014-ADM, 2004. [Online]. Available: https://cds.cern.ch/record/712881

[93] O. Brüning and L. Rossi, *The High Luminosity Large Hadron Collider*, O. Brüning, Ed. World Scientific, 2024. [Online]. Available: https://cds.cern.ch/record/2892725

[94] X. Buffat, F. S. Carlier, J. M. Coello De Portugal Martinez Vazquez *et al.*, "Optics Measurement and Correction Strategies for HL-LHC," CERN, Geneva, Tech. Rep. CERN-ACC-2022-0004, 2022. [Online]. Available: https://cds.cern.ch/record/2808650

[95] M. Hofer, "Studies on the nonlinear beam dynamics for future particle colliders," Ph.D. dissertation, TU Wien, 2023. [Online]. Available: https://inspirehep.net/files/18125202432b1c44c9d5d6c93ee40a85

[96] P. Raimondi, "FCC lattice evolution," Jun. 2023, FCC Week 2023. [Online]. Available: https://indico.cern.ch/event/1242395/contributions/5417974/attachments/2673162/4634460/FCC%20lattice%20evolution%20tuning%20workshop.pdf

[97] E. H. Maclean, R. Tomás, M. Giovannozzi *et al.*, "First measurement and correction of nonlinear errors in the experimental insertions of the CERN Large Hadron Collider," *Phys. Rev. ST Accel. Beams*, vol. 18, p. 121002, Dec. 2015. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.18.121002

[98] E. H. Maclean, R. Tomás, F. S. Carlier *et al.*, "New approach to LHC optics commissioning for the nonlinear era," *Phys. Rev. Accel. Beams*, vol. 22, p. 061004, Jun. 2019. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.061004

[99] F. S. Carlier, "A Nonlinear Future - Measurements and corrections of nonlinear beam dynamics using forced transverse oscillations," Ph.D. dissertation, University of Amsterdam, 2020, presented 08 May 2020. [Online]. Available: https://cds.cern.ch/record/2715765

[100] G. Vanbavinckhove, M. Aiba, R. Bartolini *et al.*, "First Measurements of Higher Order Optics Parameters in the LHC," in *2nd International Particle Accelerator Conference proceedings*, 2011, pp. 2073–2075. [Online]. Available: https://cds.cern.ch/record/1382076

[101] K. Paraschou, H. Bartosik, G. Iadarola *et al.*, "Injection optics - ABP discussion 9 March," 2023. [Online]. Available: https://indico.cern.ch/event/1264211/contributions/5309523/attachments/2608404/4505931/20230314_LBOC_injection_optics.pdf

[102] R. Tomás, F. Carlier, F. Chudoba *et al.*, "Optics for Landau damping with minimized octupolar resonances in the LHC," *Journal of Instrumentation*, vol. 19, no. 05, p. T05010, May 2024. [Online]. Available: https://doi.org/10.1088/1748-0221/19/05/T05010

[103] E. Maclean, K. Paraschou, R. Tomás *et al.*, "3Qy resonance correction at LHC injection," in *16th International Particle Accelerator Conference proceedings*, no. WEPM008, 2025, pp. 1964–1967. [Online]. Available: https://cds.cern.ch/record/2952717

[104] R. Tomás, "Normal form of particle motion under the influence of an AC dipole," *Phys. Rev. ST Accel. Beams*, vol. 5, p. 054001, May 2002. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.5.054001

[105] J. W. Dilly, M. Le Garrec, E. H. Maclean *et al.*, "Report from LHC MD 6863: Amplitude Detuning Corrections from Feed-Down," CERN, Tech. Rep. CERN-ACC-NOTE-2023-0028, 2023. [Online]. Available: https://cds.cern.ch/record/2913600

[106] E. H. Maclean, R. Tomás, F. S. Carlier *et al.*, "Commissioning of the nonlinear chromaticity at injection for LHC Run II," CERN, Tech. Rep. CERN-ACC-NOTE-2016-0013, 2016. [Online]. Available: https://cds.cern.ch/record/2121333

[107] G. Valentino, R. Bruce, S. Redaelli *et al.*, "Anomaly Detection for Beam Loss Maps in the Large Hadron Collider," *J. Phys. Conf. Ser.*, vol. 874, no. 1, p. 012002, Jul. 2017. [Online]. Available: https://dx.doi.org/10.1088/1742-6596/874/1/012002

[108] M. Hofer and R. Tomás, "Impact of two-dimensional decoherence on the measurement of resonance driving terms," *Journal of Physics: Conference Series*, vol. 2687, no. 6, p. 062006, Jan. 2024. [Online]. Available: https://doi.org/10.1088/1742-6596/2687/6/062006

[109] M. Hofer, F. S. Carlier, J. W. Dilly *et al.*, "MD 3603: Dynamic Aperture with uncorrected dipole b3," CERN, Tech. Rep. CERN-ACC-NOTE-2020-0049, 2020. [Online]. Available: https://cds.cern.ch/record/2733022

[110] M. Le Garrec, "LHC effective model for optics corrections : measurements and corrections of high-order non-linear optics," Ph.D. dissertation, Goethe-Universität Frankfurt, 2025. [Online]. Available: https://repository.cern/records/wtztf-mv419

[111] N. Sammut, L. Bottura, and J. Micallef, "Mathematical formulation to predict the harmonics of the superconducting Large Hadron Collider magnets," *Phys. Rev. ST Accel. Beams*, vol. 9, p. 012402, Jan. 2006. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.9.012402

[112] N. Ryckx, S. Redaelli, M. Lamont *et al.*, "Combined Ramp and Squeeze at the Large Hadron Collider," in *3rd International Particle Accelerator Conference proceedings*, V. Suller, Ed., 2012, pp. 157–159. [Online]. Available: https://proceedings.jacow.org/IPAC2012/papers/moppc016.pdf

[113] M. McAteer, Y. Levinsen, E. Maclean *et al.*, "Magnet polarities in the LHC," CERN, Tech. Rep. CERN-ACC-NOTE-2014-0012, 2014. [Online]. Available: https://cds.cern.ch/record/1667590

[114] W. Herr, "Introduction to Landau Damping," in *CAS 2013 proceedings*, no. arXiv:1601.05227, 2014, Trondheim, Norway, 18-29 Aug 2013. [Online]. Available: https://cds.cern.ch/record/1982428

[115] H. Bartosik, M. Carlà, and K. Cornelis, "Reconfiguration of SPS Landau octupole circuits to minimise second order chromaticity," CERN, Tech. Rep. CERN-ACC-2019-065, 2019. [Online]. Available: https://cds.cern.ch/record/2696265

[116] R. Miyamoto, A. Jansson, M. J. Syphers *et al.*, "Nonlinear dynamics studies in the Fermilab Tevatron using an AC dipole," in *Particle Accelerator Conference (PAC09) proceedings*. Brookhaven National Laboratory, May 2009. [Online]. Available: https://www.osti.gov/biblio/965878

[117] J. Gareyte, J.-P. Koutchouk, and F. Ruggiero, "Landau damping dynamic aperture and octupole in LHC," CERN, Geneva, Tech. Rep. LHC-Project-Report-91, CERN-LHC-Project-Report-91, 1997, revised version number 1 submitted on 2003-08-21 14:12:02. [Online]. Available: https://cds.cern.ch/record/321824

[118] J. Barranco, X. Buffat, T. Pieloni *et al.*, "Beam-beam effects, octupoles and Landau damping," Oct. 2017, Combined WP2 and WP3 coordination Meeting. [Online]. Available: https://indico.cern.ch/event/669849/contributions/2740676/attachments/1537807/2410440/FFC_eurocircol_landau_beam_beam.pdf

[119] F. Van der Veken, A. Abramov, R. Bruce *et al.*, "Loss Map Validation Restart 2022," Oct. 2022, lHC Collimation Working Group #268. [Online]. Available: https://indico.cern.ch/event/1208387/contributions/5081745/attachments/2524327/4341263/LossMaps_restart.pdf

[120] R. Tomás, A. Calia, M. D'Andrea *et al.*, "Mitigation of losses at injection protection devices in the CERN LHC," in *14th International Particle Accelerator Conference proceedings*, vol. 2023, no. MOPL019, 2023, pp. 547–550. [Online]. Available: https://cds.cern.ch/record/2886034

[121] F. Van der Veken, Collimation Team, and OP team, "Feedback from Run 3 on Collimation Losses: Impact of Optics Errors and Beam–Beam Effects," Sep. 2025, 15th HL-LHC Collaboration Meeting. [Online]. Available: https://indico.cern.ch/event/1559978/contributions/6664918/attachments/3146820/5586885/HLLHC%20-%20Non-Linearities%20in%20Collimation.pdf

[122] K. Paraschou, H. Bartosik, L. Deniau *et al.*, "Emittance Growth From Electron Clouds Forming in the LHC Arc Quadrupoles," in *HB2023*, vol. 2023, 2024, pp. 487–490. [Online]. Available: https://cds.cern.ch/record/2901227

[123] H. García Morales, J. Cardona, and R. Tomás García, "Comparison of Segment-by-Segment and Action-Phase-Jump Techniques in the Calculation of IR Local Corrections in LHC," in *12th International Particle Accelerator Conference proceedings*, 2021, pp. 632–635. [Online]. Available: https://inspirehep.net/files/71cacc4bcead8771d3c48be5eeb30295

[124] T. Persson, J. Cardona, F. Carlier *et al.*, "Optics Correction Strategy for Run 3 of the LHC," in *13th International Particle Accelerator Conference proceedings*, 2022, pp. 1687–1690. [Online]. Available: https://cds.cern.ch/record/2845780

[125] R. Tomás, T. Bach, R. Calaga *et al.*, "Record low β beating in the LHC," *Phys. Rev. ST Accel. Beams*, vol. 15, p. 091001, Sep. 2012. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.15.091001

[126] T. Persson, F. Carlier, J. C. de Portugal *et al.*, "LHC optics commissioning: A journey towards 1% optics control," *Phys. Rev. Accel. Beams*, vol. 20, p. 061002, Jun. 2017. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.061002

[127] R. Tomás, M. Bai, R. Calaga *et al.*, "Measurement of global and local resonance terms," *Phys. Rev. ST Accel. Beams*, vol. 8, p. 024001, 2005. [Online]. Available: https://journals.aps.org/prab/abstract/10.1103/PhysRevSTAB.8.024001

[128] A. Franchi, T. Beier, M. Kirk *et al.*, "A method to measure the skew quadrupole strengths in the SIS-18 using two BPMs," in *9th European Particle Accelerator Conference proceedings*, 2004, pp. 1957–1959. [Online]. Available: https://proceedings.jacow.org/e04/PAPERS/WEPLT052.PDF

[129] O. Brüning and L. Rossi, Eds., *The High Luminosity Large Hadron Collider*. World Scientific, Mar. 2024. [Online]. Available: https://cds.cern.ch/record/2892725

[130] F. C. Pilat, Y. Luo, N. Malitsky *et al.*, "Beam-Based Non-linear Optics Corrections in Colliders," in *21st IEEE Particle Accelerator Conference proceedings*, 2005, pp. 601–605. [Online]. Available: https://cds.cern.ch/record/928146

[131] W. Fischer, J. Beebe-Wang, Y. Luo *et al.*, "RHIC Proton Beam Lifetime increase with 10- and 12-pole Correctors," in *1st International Particle Accelerator Conference proceedings*, no. THPE099, 2010, pp. 4752–4754. [Online]. Available: https://proceedings.jacow.org/IPAC10/papers/thpe099.pdf

[132] B. Parker *et al.*, "The SuperKEKB Interaction Region Corrector Magnets," in *7th International Particle Accelerator Conference*, no. TUPMB041, 2016, pp. 1193–1195. [Online]. Available: https://proceedings.jacow.org/ipac2016/papers/tupmb041.pdf

[133] N. Ohuchi, Y. Arimoto, K. Akai *et al.*, "SuperKEKB beam final focus superconducting magnet system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1021, p. 165930, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168900221008949

[134] N. Ohuchi *et al.*, "Final-focus Superconducting Magnets for SuperKEKB," in *9th International Particle Accelerator Conference proceedings*, no. FERMILAB-CONF-18-452-TD, Jun. 2018, pp. 1215–1219. [Online]. Available: https://inspirehep.net/files/d9187939936d407c5d7a6d852c063aba

[135] E. Cruz Alaniz, A. Seryi, E. H. Maclean *et al.*, "Non linear field correction effects on the dynamic aperture of the FCC-hh," in *8th International Particle Accelerator Conference proceedings*, no. TUPVA038, 2017, pp. 2143–2146, on behalf of EuroCirCol WP3. [Online]. Available: https://cds.cern.ch/record/2281162

[136] M. Giovannozzi, R. De Maria, F. Van Der Veken *et al.*, "Field quality and DA: latest results from simulation studies," Nov. 2016, 6th HL-LHC Collaboration Meeting. [Online]. Available: https://indico.cern.ch/event/549979/contributions/2263210/

[137] G. Apollinari, I. Béjar Alonso, O. Brüning *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC)*, ser. CERN Yellow Reports: Monographs. Geneva: CERN, 2017. [Online]. Available: https://cds.cern.ch/record/2284929

[138] A. Verdier and E. Wildner, "Feed-down effect in dipole alignment," in *PACS2001. Proceedings of the 2001 Particle Accelerator Conference*, vol. 3, 2001, pp. 1693–1695. [Online]. Available: https://cds.cern.ch/record/512466/files/tpph016.pdf

[139] J. Dilly, "Corrections of high-order nonlinearities in the LHC and High-Luminosity LHC beam optics," Ph.D. dissertation, Humboldt-Universität zu Berlin, 2024. [Online]. Available: https://inspirehep.net/files/82e196d221988f8a3c9f22e926a735a4

[140] J. Dilly, V. Ferrentino, M. Le Garrec *et al.*, "First operational dodecapole correction in the LHC," *Phys. Rev. Accel. Beams*, vol. 26, p. 121001, Dec. 2023. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.26.121001

[141] J. Dilly, F. Carlier, J. Gray *et al.*, "Status of the CERN Optics Measurement and Correction Analysis Tools," in *16th International Particle Accelerator Conference proceedings*, 2025, pp. 1972–1975. [Online]. Available: https://cds.cern.ch/record/2953142/files/document.pdf

[142] OMC team, OMC team GitHub. [Online]. Available: https://github.com/pylhc/omc3

[143] D. Jacquet, R. Gorbonosov, and G. Kruk, "LSA - the High Level Application Software of the LHC - and Its Performance During the First Three Years of Operation," in *14th International Conference on Accelerator & Large Experimental Physics Control Systems*, 2013. [Online]. Available: https://cds.cern.ch/record/1696961

[144] V. Shiltsev, Y. Alexahin, V. Lebedev *et al.*, "Beam-beam effects in the Tevatron," *Phys. Rev. ST Accel. Beams*, vol. 8, p. 101001, Oct. 2005. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.8.101001

[145] W. Fischer and Y. Luo, "Beam-beam observations in the RHIC," in *ICFA Mini-Workshop on Beam-Beam Effects in Hadron Colliders 2013*. CERN, Oct. 2014, pp. 19–25. [Online]. Available: http://cds.cern.ch/record/1957061

[146] Y. Funakoshi, Y. Ohnishi, K. Ohmi *et al.*, "Beam-beam related issues at SuperKEKB," in *ICFA mini workshop: Beam-Beam Effects in Circular Colliders BB24*, 2024.

[Online]. Available: https://indico.global/event/9305/contributions/90640/attachments/41513/77761/BeamBeamIssuesFunakoshi.pdf

[147] M. Vogt, "Beam-beam effects in HERA," in *ICFA Mini-workshop on Beam-Beam in Hadron Colliders 2013*, 2013, presented at the ICFA Mini-Workshop on Beam-Beam in Hadron Colliders, CERN, Geneva, Switzerland, 18-22 March 2013. [Online]. Available: https://indico.cern.ch/event/189544/contributions/1464633/attachments/267387/374410/vogtm_y2013m03d18.pdf

[148] D. Kaltchev and W. Herr, "Analysis of long-range studies in the LHC - Comparison with the model," in *ICFA Mini-workshop on Beam-Beam in Hadron Colliders 2013*, 2014, pp. 93–100, comments: 8 pages, contribution to the ICFA Mini-Workshop on Beam-Beam Effects in Hadron Colliders, CERN, Geneva, Switzerland, 18-22 Mar 2013. [Online]. Available: https://cds.cern.ch/record/1957062

[149] W. Herr, "Beam-beam effects and Dynamic β^*," in *LHC Lumi Days 2012*, 2012. [Online]. Available: https://indico.cern.ch/event/162948/contributions/1417430/attachments/191879/269237/S3_WH.pdf

[150] K. Ohmi, Y. Funakoshi, K. Hirosawa *et al.*, "Optics Aberration at IP and Beam-beam Effects," in *62nd ICFA Advanced Beam Dynamics Workshop on High Luminosity Circular e^+e^- Colliders*, no. TUOBB01, 2019, pp. 66–70. [Online]. Available: https://proceedings.jacow.org/eefact2018/papers/tuobb01.pdf

[151] M. P. Crouch, "Luminosity Performance Limitations Due to the Beam-Beam Interaction in the Large Hadron Collider," Ph.D. dissertation, University of Manchester, 2018. [Online]. Available: https://research.manchester.ac.uk/en/studentTheses/luminosity-performance-limitations-due-to-the-beam-beam-interacti/

[152] T. Ieiri, Y. Funakoshi, T. Kawamoto *et al.*, "Measurement of dynamic beam-beam effects," *Phys. Rev. ST Accel. Beams*, vol. 8, p. 124401, Dec. 2005. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.8.124401

[153] F. Carlier, R. Tomás, and X. Buffat, "First measurement of resonance driving terms from head-on beam-beam interaction in the LHC," *Journal of Instrumentation*, vol. 20, no. 10, p. P10005, Oct. 2025. [Online]. Available: https://doi.org/10.1088/1748-0221/20/10/P10005

[154] P. Gonçalves Jorge, J. Barranco García, X. Buffat *et al.*, "Measurement of Beta-Beating Due to Strong Head-on Beam-Beam Interactions in the LHC," in *8th International Particle Accelerator Conference*, no. TUPVA030, 2017, pp. 2121–2124. [Online]. Available: https://cds.cern.ch/record/2289678

[155] W. Herr and T. Pieloni, "Beam-Beam Effects," CERN, Tech. Rep. arXiv:1601.05235, 2014, 29 pages, contribution to the CAS - CERN Accelerator School: Advanced Accelerator Physics Course, Trondheim, Norway, 18-29 Aug 2013. [Online]. Available: https://cds.cern.ch/record/1982430

[156] W. Herr, X. Buffat, R. Calaga *et al.*, "Long Range Beam-beam Effects in the LHC," in *ICFA Mini-workshop on Beam-Beam in Hadron Colliders 2013*, 2014, pp. 87–92, presented at the ICFA Mini-Workshop on Beam-Beam in Hadron Colliders, CERN, Geneva, Switzerland, 18-22 March 2013. [Online]. Available: https://cds.cern.ch/record/1953724

[157] S. Fartoukh, S. Kostoglou, H. Bartosik *et al.*, "2023 Machine Configuration for Optimal Performance," in *Joint Accelerator Performance Workshop 5-8 Dec 2022*, 2022. [Online]. Available: https://indico.cern.ch/event/1194548/contributions/5093929/attachments/2560367/4412969/JAPW22_2023_configuration.pdf

[158] D. Mirarchi, "Observations on collimator hierarchy breakage below 36 cm," Apr. 2024, LBOC meeting 165. [Online]. Available: https://indico.cern.ch/event/1405768/contributions/5925146/attachments/2843311/4970776/BrokenHierarchy.pdf

[159] T. Persson, "Long Range Beam-Beam investigation using Weak-Strong beams in the LHC," in *ICFA mini workshop: Beam-Beam Effects in Circular Colliders BB24*. CERN, 2024. [Online]. Available: https://indico.global/event/9305/contributions/90651/attachments/41527/77800/beam-beam%20workshop.pdf

[160] T. H. B. Persson, X. Buffat, F. S. Carlier *et al.*, "LHC MD #12263 - Weak-Strong investigation of Long-Range Beam-beam," CERN, Tech. Rep. CERN-ACC-NOTE-2025-0011, 2025. [Online]. Available: https://cds.cern.ch/record/2930597

[161] A. Poyet, A. Bertarelli, F. Carra *et al.*, "First experimental evidence of a beam-beam long-range compensation using wires in the Large Hadron Collider," *Phys. Rev. Accel. Beams*, vol. 27, p. 071003, Jul. 2024. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.27.071003

[162] CERN, "Beam-Beam Wire Compensation Review Meeting 14-15 October," 2024. [Online]. Available: https://indico.cern.ch/event/1437020/

[163] G. Sterbini *et al.*, "The BBWC Experimental Program in LHC and Potential for the HL-LHC," 2024, Beam-Beam Wire Compensation Review Meeting 14-15 October. [Online]. Available: https://indico.cern.ch/event/1437020/contributions/6047557/attachments/2946519/5178321/BBWC_Review-10.pdf

[164] G. Sterbini *et al.*, "Status of the Beam-Beam Long-Range Wire Compensation Studies," in *ICFA mini workshop: Beam-Beam Effects in Circular Colliders BB24*, 2024. [Online]. Available: https://indico.global/event/9305/contributions/90650/attachments/41517/77771/ICFA_EPFL_2024-4.pdf

[165] T. Gläßle, H. B. GmbH, Y. Levinsen *et al.*, "CPyMAD Manual," 2026. [Online]. Available: https://hibtc.github.io/cpymad/

[166] G. Iadarola, R. De Maria, S. Łopaciuk *et al.*, "Integrated beam-beam modeling with Xsuite," in *ICFA mini workshop: Beam-Beam Effects in Circular Colliders BB24*, 2024. [Online]. Available: https://indico.global/event/9305/contributions/90673/attachments/41538/77824/012_Xsuite_bb_workshop.pdf

[167] D. Mirarchi, M. Hostettler, G. Ripellino *et al.*, "LHC Luminosity Calibration and Monitoring Working Group 4 August 2025," 2025. [Online]. Available: https://indico.cern.ch/event/1575154/contributions/6636183/attachments/3114061/5521094/202508_LLCMWG_LumiImbalance-1.pdf

[168] F. Carlier, A. C. Ojeda, R. De Maria *et al.*, "LHC Run 3 optics corrections," in *14th International Particle Accelerator Conference proceedings*, vol. 2023, no. MOPL015, 2023, pp. 535–538. [Online]. Available: https://cds.cern.ch/record/2884562

[169] J. Wenninger, "Orbit corrector magnets and beam energy," CERN, Geneva, Tech. Rep. CERN-SL-Note-97-06-OP, 1997. [Online]. Available: https://cds.cern.ch/record/703742

[170] R. Tomás, N. Mounet, Y. Angelis *et al.*, "Towards a High Luminosity LHC with even higher performance," in *16th International Particle Accelerator Conference proceedings*, no. MOPM008, 2025, pp. 278–281. [Online]. Available: https://meow.elettra.eu/81/pdf/MOPM008.pdf

[171] R. De Maria, R. Bruce, D. Gamba *et al.*, "High Luminosity LHC optics and layout HLLHCV1.4," in *10th International Particle Accelerator Conference proceedings*, no. MOPMP019, 2019, pp. 468–471. [Online]. Available: https://cds.cern.ch/record/2696117

[172] J. Gray and F. Carlier. (2024, Oct.) Energy Correction via Global Corrections. OMC meeting on MD planning, LRBB b3 correction. [Online]. Available: https://indico.cern.ch/event/1467197/contributions/6180274/attachments/2947653/5180423/OMC_15_05_24_Energy_correction_via_global_corrections.pdf

[173] J. Gray, "Beam energy discrepancy: What can the optics mismatch reveal?" Sep. 2025, 15th HL-LHC Collaboration Meeting. [Online]. Available: https://indico.cern.ch/event/1559978/contributions/6664897/

[174] J. Gray. (2026) Energy Error GUI. [Online]. Available: https://indico.cern.ch/event/1663136/#9-energy-error-gui-and-2026-co

[175] J. Gray, K. Skoufaris, A. Oeftiger *et al.*, "Understanding Energy-Induced Optics Distortions in the LHC," in *[submitted to 17th International Particle Accelerator Conference]*, no. WEP5041, 2026.

[176] J. Wenninger, "Large storage ring sensitivity to tides & large-scale perturbations," 2021. [Online]. Available: https://indico.cern.ch/event/982987/contributions/4191304/attachments/2192132/3705229/LHC.TidesAndQuakes.Feb21.pdf

[177] L. Arnaudon, B. Dehning, A. Hofmann *et al.*, "Effects of terrestrial tides on the LEP beam energy," *Nuclear Instruments and Methods in Physics Research Section A*, vol. 357, pp. 249–252, 1995. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/0168900294015260

[178] J. Wenninger and E. Todesco, "Large Hadron Collider momentum calibration and accuracy," *Phys. Rev. Accel. Beams*, vol. 20, no. CERN-ACC-2017-0007, p. 081003, 2017. [Online]. Available: https://cds.cern.ch/record/2254678

[179] R. Tomás, O. Brüning, M. Giovannozzi *et al.*, "CERN Large Hadron Collider optics model, measurements, and corrections," *Phys. Rev. ST Accel. Beams*, vol. 13, p. 121004, Dec. 2010. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.13.121004

[180] J. Wenninger, "YASP Steering Program User Guide," 2005. [Online]. Available: https://web.archive.org/web/20221121014248/https://jwenning.web.cern.ch/documents/YASP/YASP-user-guide.pdf

[181] Y. Angelis, "LHC - Proton Physics Optics commissioning outcome," Mar. 2026, OMC Meeting on LHC commissioning, numerical deconv, OMC3 error bars. [Online]. Available: https://indico.cern.ch/event/1663136/contributions/6992444/attachments/3242181/5784012/OMC_ppCommis26_Angelis.pdf