CERN-ACC-NOTE-2020-0045
joel.daniel.andersson@cern.ch

Orbit Correction Studies on the HL–LHC Layout and Optics V1.5

Joel D. Andersson, Riccardo De Maria, Davide Gamba

Abstract

In this report three separate studies on the topic of orbit correction are presented for the HL–LHC V1.5 layout and 15 cm β^* optics for protons. It is shown that the constraints put on the error correction for the assumed machine imperfections distribution, together with the implementation of the desired orbit knobs, are compatible with the available corrector strengths. Results based on a simplified model of the present LHC orbit feedback are presented, showcasing its efficiency in maintaining beam collisions and the inherent orbit stability expected in LHC and in HL–LHC. Finally, necessary short-term beam position monitor stability for adequate position-based correction of beam separation is investigated and estimated to be under one micrometre.

Geneva, Switzerland
July 13, 2020

Contents

1 Introduction

This document accounts for beam orbit studies performed on the HL–LHC version 1.5 optics [1, 2]. Three complementary studies are being reported in this document. Firstly, results from studying the orbit corrector budget will be discussed, that is, how the closed orbit in the machine can be corrected and all orbit knobs be implemented while the orbit corrector usage lies within given constraints. Secondly, a simple model of the orbit feedback system used in LHC will be used to estimate its impact on the beam orbit stability at Interaction Points (IPs). Lastly, estimates on the necessary beam position monitor (BPM) stability to correct adequately orbit at Interaction Point 1 and 5 (IP1 and IP5) will be provided.

Unless otherwise stated, all studies have been performed for proton beams at 7TeV and for the fully squeezed optics with $\beta^* = 15cm$ in IP1 and IP5.

This document will be oriented towards results and less so towards the underlying assumptions and technical details. For a more detailed account of the treatment, see the corresponding Master's thesis [3]. The analysis was conducted using a Python package POCKPy [3], designed for closed orbit analysis of HL–LHC.

2 Orbit Corrector Budget

The objective is to implement a set of orbit knobs and to correct for closed orbit perturbation caused by machine errors, both of which require the use of orbit correctors. Each orbit corrector has, by design, an upper limit on its strength, hence it makes sense to talk of an *orbit corrector budget*, which enforces a constraint to be respected across both problems. This orbit corrector budget problem will be stated, analyzed, and solved for HL–LHC optics, version 1.5. Unless otherwise specified, computations will be performed assuming the nominal collision energy of 7 TeV where the correctors are the least effective in kicking the beam.

All analyses are based on *orbit response matrices* for the machine imperfections and orbit correctors, $\mathrm{RM_e}$ and $\mathrm{RM_c}$, respectively. In response matrix jargon, the set of machine imperfections are stored in a vector $\boldsymbol{e}$ and the set of orbit corrector strengths are stored in a vector $\boldsymbol{c}$. The problem of error correction can therefore be written as:

$$\boldsymbol{r} = \mathrm{RM_e}\boldsymbol{e} + \mathrm{RM_c}\boldsymbol{c} \tag{1}$$

where the closed orbit residual $\boldsymbol{r}$ is retrieved by finding a suitable correction $\boldsymbol{c}$ given the machine imperfections $\boldsymbol{e}$. Similarly, the problem of implementing an orbit knob can be stated as:

$$\boldsymbol{k} = \mathrm{RM_c}^{\dagger}\boldsymbol{c} \tag{2}$$

where † denotes taking a subset of rows, corresponding to the positions in the accelerator where the closed orbit is to be set to a given value, and subset of columns, corresponding

to the correctors that are available for implementation, and $\boldsymbol{k}$ encodes the requisite values at the subset of rows.

2.1 Boundary conditions and specifications

For the investigation of the orbit corrector budget the Octant 5 was investigated, from Q25 to Q25. Because the interaction regions in Octant 1 and 5 are the main area affected by the HL–LHC upgrade, and since they are symmetric once the horizontal and vertical planes are exchanged, it suffices to consider one of them. In treating the error correction, element errors are assigned to applicable elements up to Q20 on either side of IP5, assuming uniform distributions tabulated in Table 1.

Table 1: Assumed distributions of element errors. All distributions are taken to be uniform distributions on intervals specified by the values tabulated. The error of any given element is taken to be independent of all other elements and errors, unless otherwise stated.

Element type	Error type	Uniform distribution boundaries
Dipole	Relative field error	± 0.002
	Roll along beam axis	± 0.5 mrad
	Horizontal misalignment	± 0.5 mm
	Vertical misalignment	± 0.5 mm
Quadrupole	Relative field error	± 0.002
	Roll along beam axis	± 1 mrad
	Horizontal misalignment	± 0.5 mm
	Vertical misalignment	± 0.5 mm
BPM	Horizontal misalignment	± 0.5 mm
	Vertical misalignment	± 0.5 mm

The values assumed in Table 1 are compatible with studies on previous versions of the HL–LHC layout and optics, e.g. [4, 5]. Longitudinal misalignment of elements has been withheld from the analysis due to its small impact on the result and the complexity of including it in the computations with the current analytical framework. A complete and more accurate evaluation of the expected misalignment is being carried out by the HL–LHC Alignment Working Group. The acceptance criteria specified in [6] for the MBRD magnets (D2) for the roll are ± 2 mrad, instead of ± 0.5 mrad specified in Table 1. This should not have a major impact on the results discussed in this document, as such an error can be easily compensate by order of 10 mTm strength on the nearby orbit correctors.

The BPM misalignment will be strongly correlated with the misalignment of neighbouring quadrupoles and dipoles to which most BPMs are mechanically attached. Based on [7], it is possible to associate almost all BPMs with a dipole or quadrupole, and then

proceed to posit some relations between their transverse misalignments. Table 2 shows the mapping between BPMs and the corresponding magnetic element, based on [7]. Subsequent assumptions on how misalignments are interrelated between linked elements is discussed in Section 2.2.

Table 2: Mapping between BPM and corresponding quadrupole or dipole, based on inferences from [7]. BPM elements mapping to 'NONE' implies that the element is assumed to be not connected to any quadrupole or dipole.

BPM	Quadrupole/Dipole
BPMQSTZA.1	MQXFA.1
BPMQSTZB.A2	MQXFB.A2
BPMQSTZB.B2	MQXFB.B2
BPMQSTZB.A3	MQXFA.3
BPMQSTZB.B3	NONE
BPMQSTZB.4	MBXF.4
BPMQSTZW.4	NONE
BPMQBCZ[AB].4	MBRD.4
BPMYA.4	MQY.4
BPM(R).5	MQML.5
BPM(R).6	MQML.6
BPM(R).7	MQM.7
BPM.8	MQML.8
BPM.9	MQMC.9
BPM.10	MQML.10
BPM.($\geq$11)	MQ.$\geq$11

For the error correction, the following constraints are put on the residual orbit:

- Zero residual orbit at IP5.
- Zero residual orbit at the Crab Cavities (CCs), ACFGA.[AB][LR]5.B[12].
- No orbit leakage at the end of the accelerator segment (i.e. zero residual orbit and angle at Q25), in practice enforced by setting zero orbit at Q20.
- 2*RMS residual orbit along the accelerator segment within the limits on the available aperture (see Sec. 2.1.1).

The zero orbit at IP5 is enforced because of the collisions, whereas the zero orbit at the CCs follows from beam loading considerations of those devices [8]. To enforce the zero orbit and angle, BPMs close to the specified element positions were corrected to zero. If two BPMs are corrected to zero, and no source of closed orbit perturbation is situated between them, it then follows that the perturbation between these BPMs is zero. Enforcing

no orbit leakage at the extremities of the segment under consideration is in line with the assumption that each octant can be treated independently.

To enforce a certain length of the orbit knobs, it was chosen to limit the allowed orbit corrector to use by the optimisation algorithm. The knobs to be implemented, together with the respectively allowed correctors, were:

- `IP_CROSSING`: Implement a crossing angle (up to ± 295 μrad) in the vertical plane for both beams in IP5. Allowed correctors: every corrector from IP5 up to and including those in Q4 on each side of IP5.
- `IP_SEPARATION`: Implement a separation of the beams (up to ± 0.75 mm) at IP5 in the horizontal plane. Allowed correctors: every corrector up to Q4 on each side of IP5, excluding all correctors of type `MCBY`.
- `CC_MOVE_B1`/`CC_MOVE_B2`: Implement an orbit at the CCs (up to ± 500 μm) in the horizontal and vertical plane, independently for each beam. Allowed correctors: all correctors up to Q6 on each side of IP5.
- `IP_OFFSET_REMOTE`: Implement an offset of both beams (up to ± 2 mm) in the horizontal and vertical plane at IP5 when offsetting every quadrupole from Q4 on the left to Q4 on the right of IP5 by ± 2 mm, and Q5 by ± 1 mm. Allowed correctors: every corrector from Q4 to Q8 on each side of IP5, excluding all correctors of type `MCBRD`.
- `IP_OFFSET_CORR`: Implement an offset of both beams (up to ± 0.5 mm) in the horizontal and vertical plane at IP5. Allowed correctors: all correctors up to Q8, excluding all correctors of type `MCBRD`.
- `LUMISCAN_B1`/`LUMISCAN_B2`: Implement an orbit at IP5 (up to ± 100 μm) in the horizontal and vertical plane for each beam, independently. Allowed correctors: every corrector from Q4 to Q5, on each side of IP5.

The corrector budget, i.e. how much corrector strength in Tm is allotted to each corrector, is given in Table 3. The corrector strength allocated for the knob implementation is the total strength reduced by the amount needed for orbit correction (i.e. 2 times the RMS corrector strength needed to correct a uniform distribution of errors, see LHC [11]).

2.1.1 APERTURE

The available aperture is computed using MAD-X [12] and following the approach from [13] where the phase advance between the TCTs and the dump kickers determines the needed aperture around IPs. The rationale is that if a horizontal dump kicker 'misfires', the impact will be an orbit perturbation with local impact dependent on phase-advance. Using the aperture values tabulated for different phase advances directly provides estimates on the necessary aperture for a given dump kicker failure around the IPs.

Table 3: Orbit corrector budget in HL–LHC. Corrector names are provided as regular expressions [9] and corrector strength is given in units of Tm [10].

Corrector name	**Corrector strength limit** [Tm]
`MCB[HV]`	1.895
`MCBC[HV].9`	1.895
`MCBC[HV].[78]`	2.8
`MCBC[HV].[56]`	2.1
`MCBY[HV].[AB]?4`	2.25
`MCBRD[HV]`	5.0
`MCBXFA[HV].3`	4.5
`MCBXFB[HV].[AB]2`	1.87

In the approach taken here, we use the tabulated values from [13] for the region delimited by
`TCTPH.6[LR][15].B[12]` for IP1 and IP5, and by `TCTPH.4[LR][28].B[12]` at IP2 and IP8, whereas for the remaining machine circumference, a constant aperture limit of $19.4\,\sigma$ is posited. As there are a total of fifteen dump kickers per beam, we take the approach of letting the aperture limit at a given position correspond to the maximum value across all kickers. The resulting aperture limit can be seen in Fig. 1.

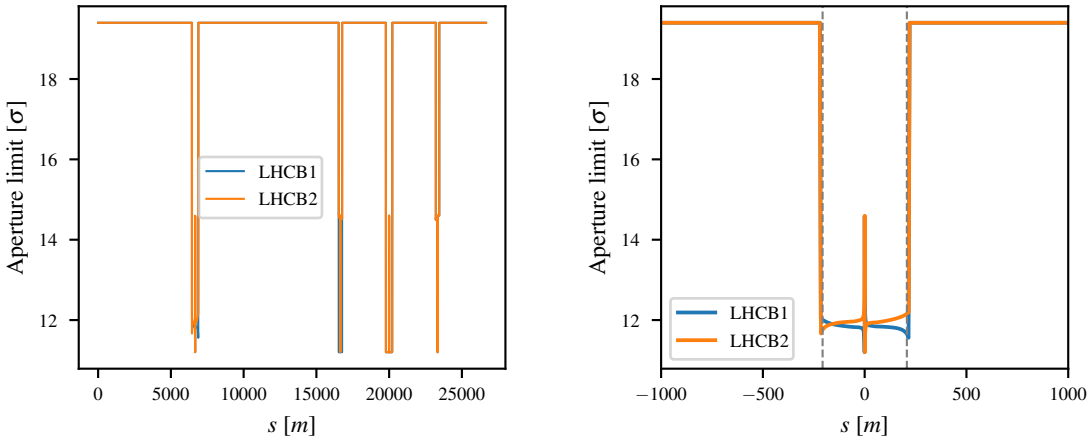

(a) Full machine.

(b) Around IP5.

Figure 1: Aperture limit shown globally (a) and around IP5 (b). $s = 0$ corresponds to IP3, and the Gray dashed vertical lines in (b) correspond to Q5.

With a lower limit on aperture, it is possible to compute the maximum allowed closed orbit in HL–LHC using the `APERTURE` command in MAD-X. In practice, this is done by scanning the added closed orbit (`COR` parameter in MAD-X), where the available orbit

budget at a location is equal to the maximum closed orbit value such that the minimum aperture is still attained. Doing this for the 15cm β^* round collision optics, where the crossing knobs at IP1 and IP5 were set to $\pm 295\,\mu$rad, results in the closed orbit limits shown in Fig. 2. This limit on the orbit becomes the effective bound on the error cor-

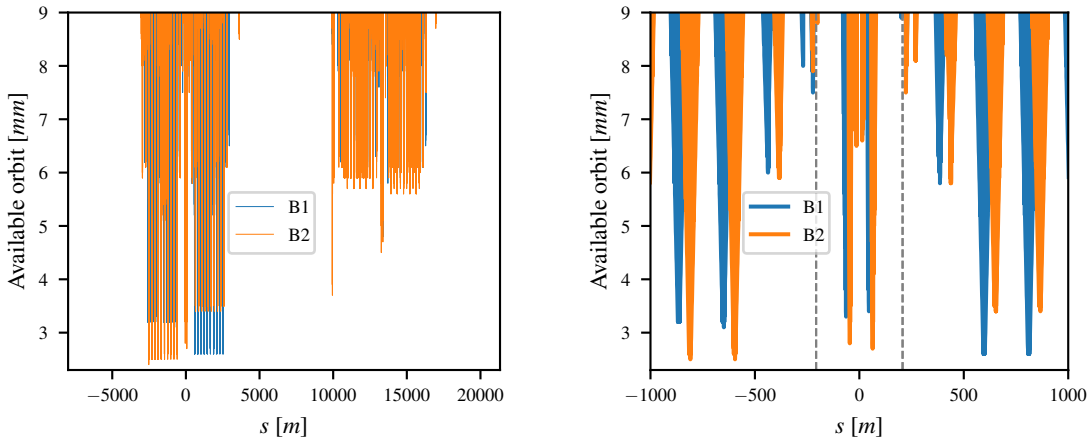

(a) Full machine.

(b) Around IP5.

Figure 2: Maximum allowed radial closed orbit shown globally (a) and around IP5 (b) for 15cm β^* round collision optics. $s = 0$ corresponds to IP3. Gray dashed vertical lines correspond to Q5.

rection, where we require that the 2*RMS residual after error correction is less than the limit provided in Fig. 2b. As regards the asymmetry between IP1 and IP5 in Fig. 2a, it is caused by the choice of crossing plane. Since IP5 has a vertical crossing plane and dipole apertures have a greater margin in the horizontal plane, the net result is a worsening of maximum tolerable radial orbit distortion budget around IP5 compared to that around IP1. This is also the motivation why the corrector budget is studied for Octant 5 and not Octant 1: the error correction in Octant 5 has tighter orbit constraints.

2.2 Error correction

For the error correction, two methods have been implemented in POCKPy:

1. Linear correction based on a weighted pseudoinverse of the corrector response matrix.
2. Constrained convex optimization.

The first method constitutes a direct approach and has been used before [5]. However, it has the drawbacks of not being adaptable (the same mapping will be applied to the most pathological as well as the most easily corrected machine realisations) and of lacking functionality for mapping weights and singular value cutoffs to given constraints. Granted

these limitations, the linear approach is still preferable as it is used in practice, and allows for direct computation of the residual orbit and corrector strength covariance matrices. For these reasons, the linear correction is the method that will be used in this study. More information and use of the convex optimization are reported in [3].

2.2.1 CHOOSING A LINEAR CORRECTION

For the linear correction, the choice of parameter values is important. It is possible to pseudoinvert the corrector response matrix, evaluated at each BPM, directly with all singular values and define it as the correction strategy. This is unfeasible in practice for the reasons that:

1. Some BPMs have to be corrected down to zero (i.e. at the IPs and CCs).
2. Some regions of the machine may have tighter orbit constraints, and thus need to be corrected more aggressively (e.g. the triplet areas).
3. Regions with fewer BPMs will be less favoured in the least-squares solution if they are not weighted.
4. The smallest singular values form a basis in orbit space that is 'expensive' to correct, and may be dominated by numerical noise.

To address these issues we weight the BPMs in the machine and constrain the number of singular values used. By providing additional weight to BPMs the residual orbit at these positions is prioritized in the least-squares minimization, and if a weight is much larger than most other weights, the result is that such a BPM will be corrected to zero where possible. To make the parameter search a tractable problem, we limit the weights to three cases: zero orbit, triplet and arc BPMs.

To enforce zero orbit constraints, we weight the corresponding BPMs as 1.0, and the remaining BPMs in the arc as 10^{-6}. The last parameters, the weight of the triplet BPMs and the number of singular values, were determined by exploring the corresponding two-dimensional parameter space. The process for choosing these parameters is covered in Appendix A, with the final result being a triplet weight of 1.25×10^{-6} and 47 out of 57 singular values (per plane).

An alternative approach to weighting differently the BPMs of the triplet and those of the arc, is to perform two independent orbit corrections: one global correction over the whole machine, one more local dedicated to each IR, and by scanning the number of singular values for both corrections similarly to what explained in Appendix A. This latter approach is that normally used in operation in the LHC [14], and it is expected to be equivalent to that used in this study.

2.2.2 ERROR CORRECTION RESULTS

For the error correction results shown here, we first consider the following three cases, where the distributions involved are those of Table 1:

1. Only BPM errors are present.
2. Only dipole and quadrupole errors are present.
3. BPM, dipole, and quadrupole errors are present.

In the third case, we make the assumption that BPMs attached to any other element (as per Table 2) share misalignment with that element. BPMs that are not attached to any other element have a misalignment that is independent of all other misalignments. The first case reveals the impact of BPMs error, independently from any actual source of magnetic error, i.e. it shows the residual orbit and corrector strength usage just due to attempting to correct to zero a false non-zero orbit. The second case considers only misalignments of the magnetic elements, assuming to have perfect BPMs perfectly aligned on the ideal orbit. The fact that most BPMs are mechanically linked to a nearby magnet (see Table 2) introduces a fundamental correlation between part of the BPM error and magnetic element misalignments, which is modelled by the third case. The most complete and probably more realistic scenario is a combination of those three cases. Assuming those cases are independent from each other, the RMS residual orbit and corrector usage can then be computed by opportune scaling and adding in quadrature of the single effects.

Correcting for these three cases leads to the arc residual orbits shown in Fig. 3.

In the arc, the error correction evidently allows for a great margin with respect to the constraints on orbit. The same result is shown for the triplet in Fig. 4. Here the orbit is, while permissible, greater than expected, and strongly affected by the BPM errors. The underlying reason for the spike in the residual orbit entering the triplet can be traced to the zero-orbit constraint at the IP. In the current formulation of the error correction, we correct to the center of the BPMs next to the IP, but granted that these positions are perturbed away from the ideal orbit, correcting them down to zero will not necessarily be attractive.

The correction strategy must thus be modified to account for this added complexity. Here we consider the two following alternatives:

1. We do not focus on correcting at the IP, i.e. the closest BPMs have equal weights as every other BPM in the triplet.
2. We assume that the BPM error for the closest BPMs to the IP is zero, and we still correct there with high weight.

The first of these two options is intuitive: correcting for the IP caused the orbit to explode, therefore not correcting at the IP could be an interesting alternative. Indeed, this is also the strategy that is typically adopted in the LHC for the first steering in the Interaction Region (IR), while the problem of finding collision is addressed at a second stage, independently

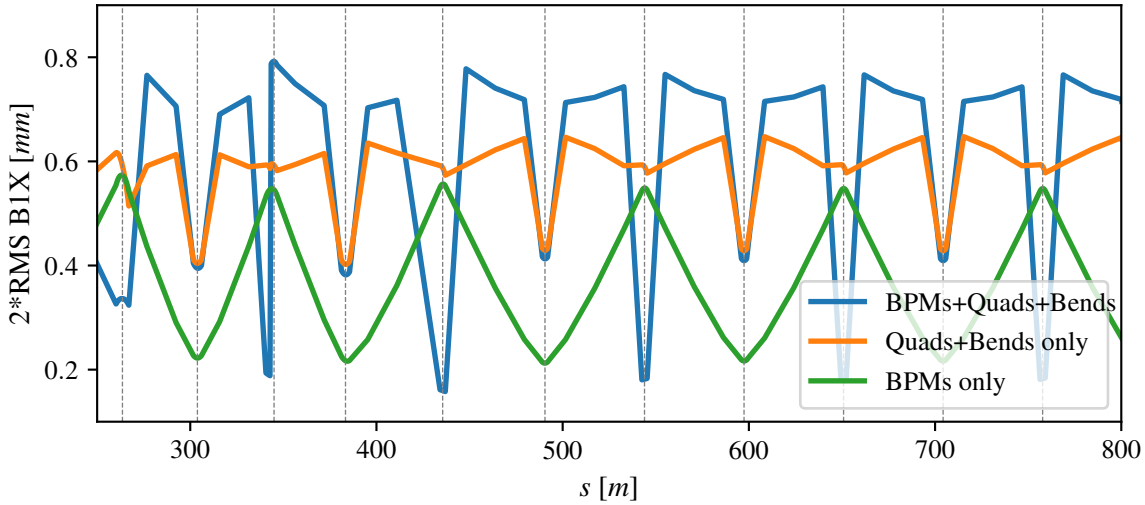

(a) Orbit given with respect to actual transverse center of elements.

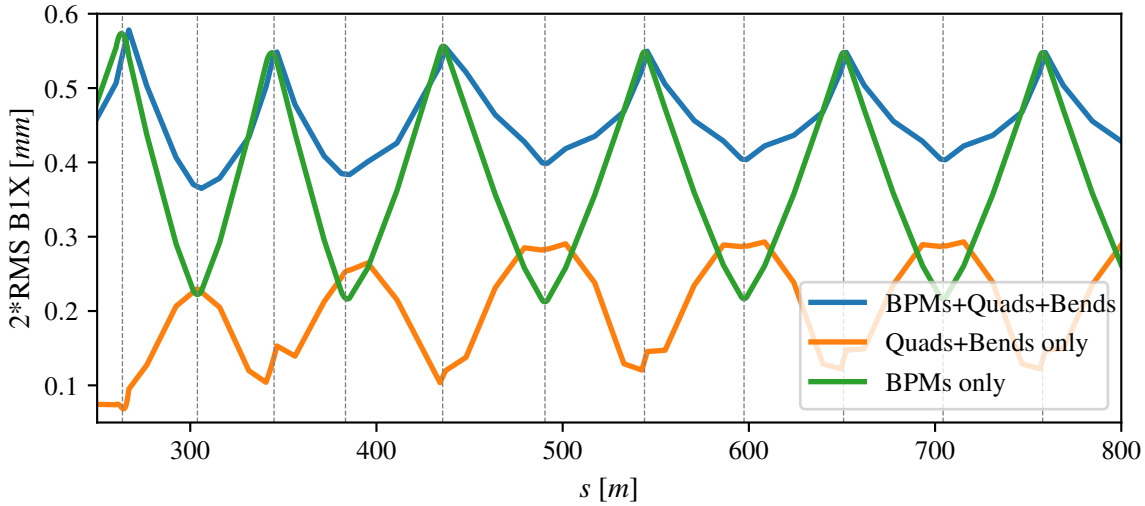

(b) Orbit given with respect to ideal transverse center of elements.

Figure 3: Residual orbit in the right-hand side arc shown for Beam 1, evaluated at all dipoles and quadrupoles after applying the linear correction. Each line represents a different scenario, i.e. assuming that only BPM errors are present (green), only quadrupole and dipole errors (orange) and dipole and quadrupole errors fully correlated with BPM (blue). The gray dashed vertical lines correspond to the position of Q7 to Q17.

from the bare orbit correction [14]. The second variant follows from the assumption that, with experience, one will be able to find the misalignment of the BPMs next to the IP with respect to the ideal orbit passing by the ideal IP. This comes out the same as saying, in this framework of analysis, that these BPMs have no alignment error.

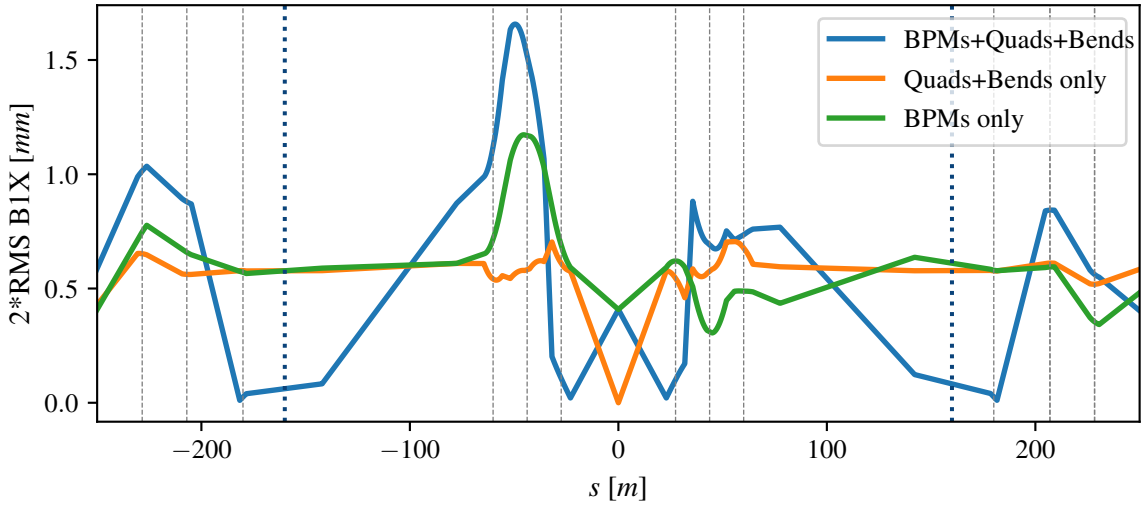

(a) Orbit given with respect to actual transverse center of elements.

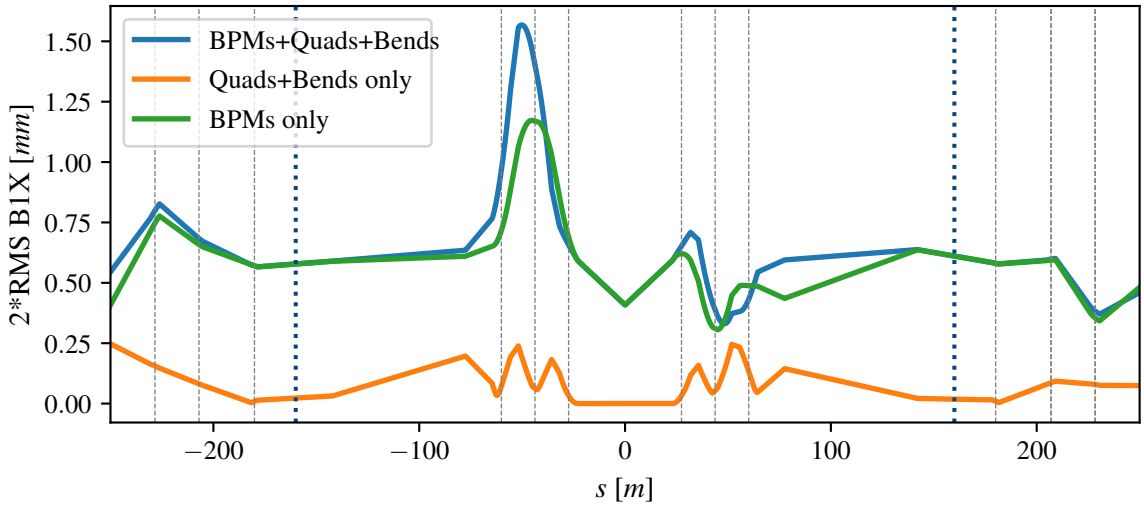

(b) Orbit given with respect to ideal transverse center of elements.

Figure 4: Horizontal residual orbit in the triplet shown for Beam 1, evaluated at all dipoles and quadrupoles after applying the linear correction. Each line represents a different scenario, i.e. assuming that only BPM errors are present (green), only quadrupole and dipole errors (orange) and dipole and quadrupole errors fully correlated with BPM (blue). The gray dashed vertical lines correspond to Q1 to Q6 on each side of the IP, and the thicker blue dashed lines correspond to the CCs.

Considering the most realistic case of having all quadrupoles and dipoles error, and BPMs attached to the nearby magnet (as per Table 2), the result of these two variants compared to the initial scenario are shown in Fig. 5. As expected, assuming to know the

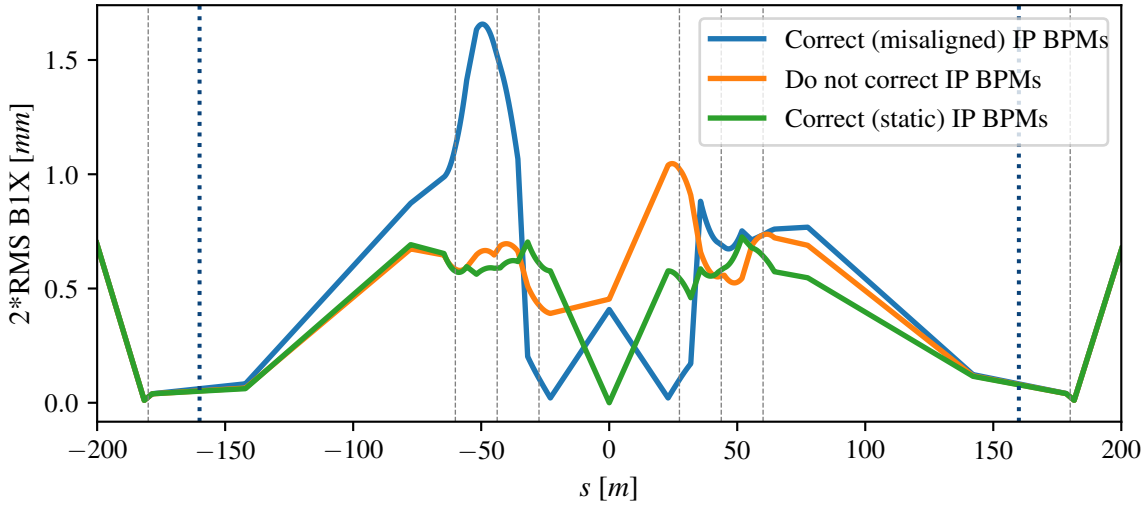

(a) Orbit given with respect to the actual transverse center of elements.

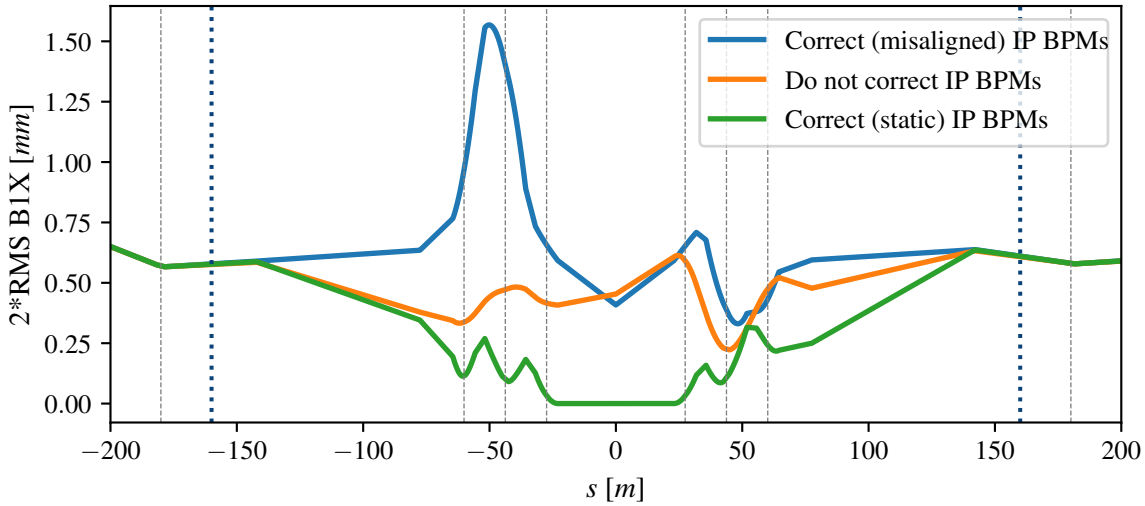

(b) Orbit given with respect to the ideal transverse center of elements.

Figure 5: Horizontal residual orbit in the triplet shown for Beam 1, evaluated at all dipoles and quadrupoles after applying three different variations of the linear correction. Each line represents a different strategy for correcting at the BPMs next to the IP, i.e. forcing the beam going through the center of those BPMs, even if they are misaligned with respect to the ideal orbit (blue), not forcing the beam to go through the center of those BPMs (orange) and assuming to know and forcing the ideal target at those BPMS (green). The gray dashed vertical lines correspond to Q1 to Q4 on each side of the IP, and the thicker blue dashed lines correspond to the CCs.

alignment information of the IP BPMs leads to the best orbit.

The corrector strength used to achieve the correction is showcased in Fig. 6. The

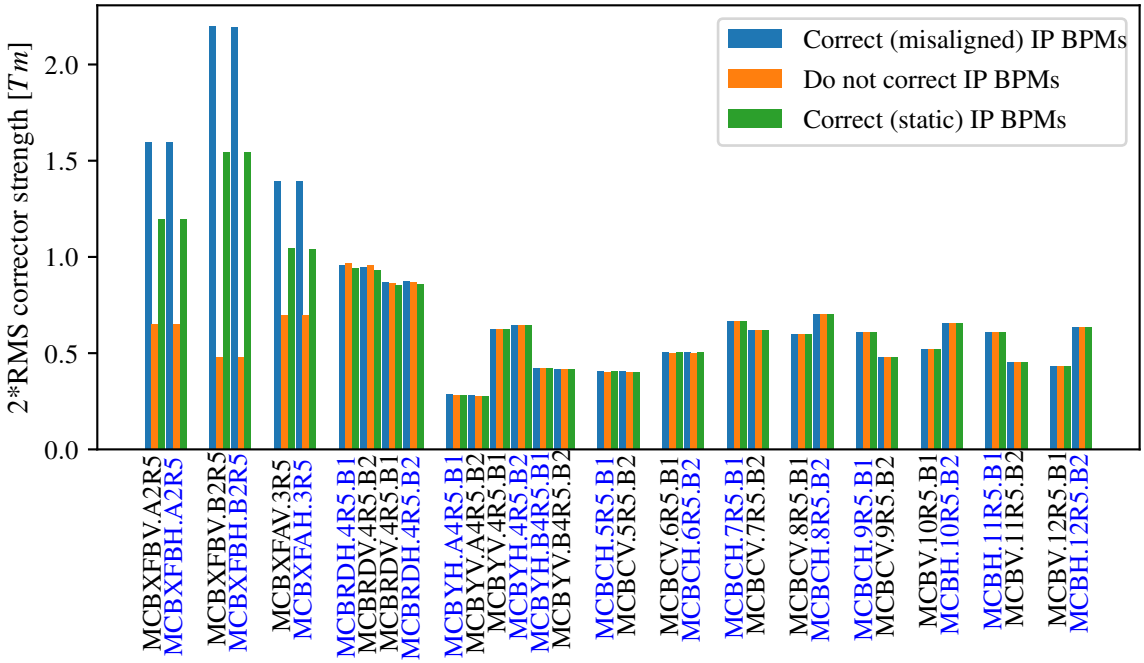

Figure 6: Corrector strength usage shown for correctors on the right-hand side of the IP when the three different variations of the linear correction strategy are employed.

result of the corrector usage matches our expectations. If we insist on correcting to the center of the IP BPMs, ignoring their alignment, then we are almost surely attempting to implement a nonphysical orbit and we will use additional corrector strength to this end. Not enforcing zero orbit at the IP is, as expected, a cheaper option, and enforcing zero orbit at the IP BPMs where they are static lies in between the other two cases.

It should be noted that the orbit in the CCs shown in Fig. 5 is not necessarily desirable. So far, to minimize the orbit at the CCs, the surrounding BPMs are corrected to zero. However, similarly to the case of the IP, these BPMs are assumed to move and therefore correcting them to zero will necessarily imply a non-zero orbit at the cavities. Alternatively, we can assume to have these BPMs and CCs to be on the ideal reference orbit, i.e. they are treated as static. The residual orbit resulting from combining this setup with static IP BPMs is shown in Fig. 7. Knowing the alignment of the BPMs next to the CCs improves the residual orbit at the CCs considerably. The impact of this approach on the corrector budget usage is shown in Fig. 8. As for the case of the IP, the residual orbit between Q3 and Q6, as measured with respect to the ideal orbit, is driven to a large extent by an excessive corrector strength usage from forcing the orbit through the center of the BPMs next to the CCs in case they are misaligned. Unlike the correction of the IP, where we assume that the alignment of the adjacent BPMs will be known at some

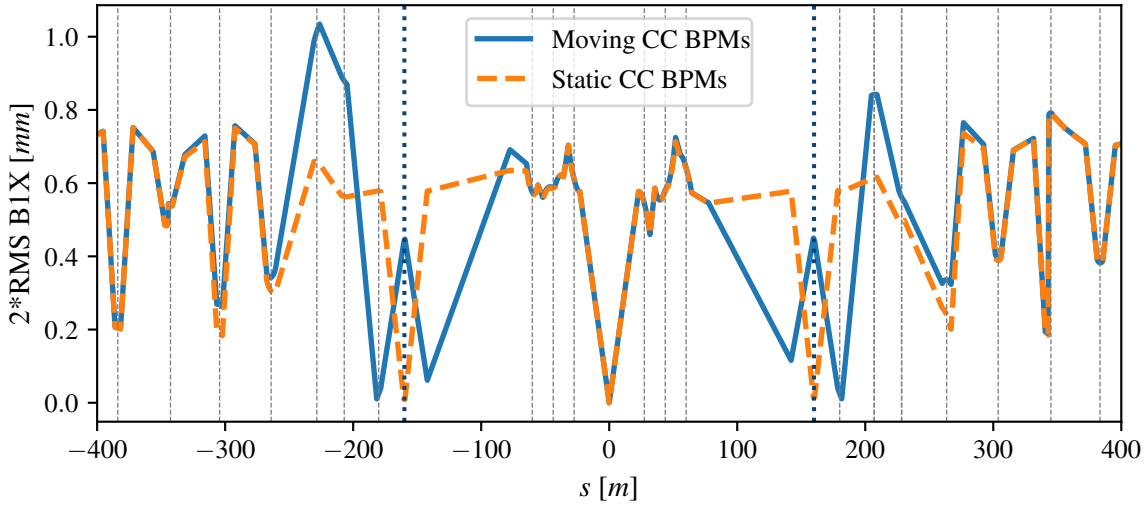

(a) Orbit given with respect to the actual transverse center of elements.

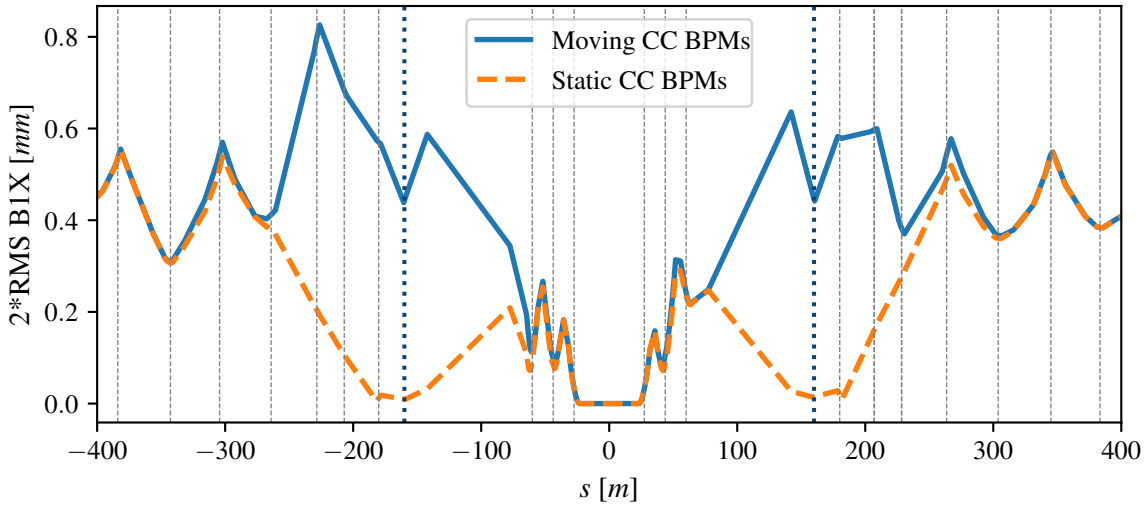

(b) Orbit given with respect to the ideal transverse center of elements.

Figure 7: Horizontal residual orbit for Beam 1, evaluated at all dipoles, quadrupoles and CCs, assuming either that the BPMs next to the CCs are static on the ideal reference or moving together with the associated magnet. The gray dashed vertical lines correspond to Q1 to Q10 on each side of the IP, and the thicker blue dashed lines correspond to the CCs.

point, correcting for the orbit at the CCs is not a settled question, as the CC themselves could be used as quasi-BPM (measuring the beam-induced Radio Frequency power) and they could be remotely aligned around the beam. Therefore, for the remaining studies we keep the pessimistic scenario of "Moving CC BPMs" in Figs 7 and 8. Naturally, should

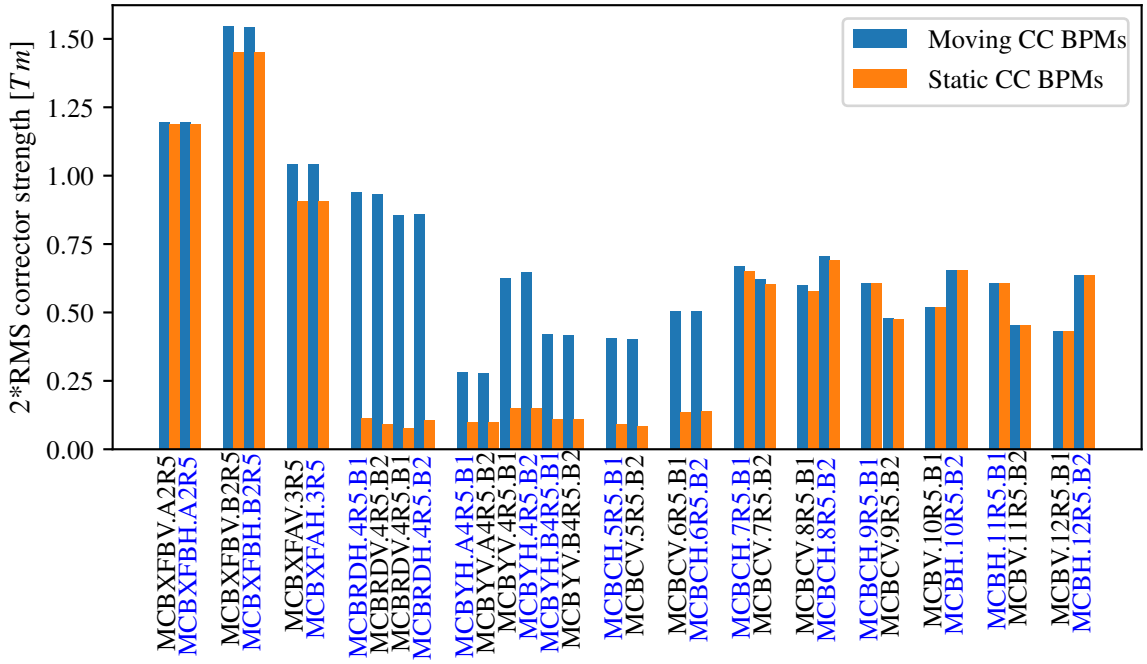

Figure 8: Corrector strength usage shown for correctors on the right-hand side of the IP when the BPMs next to the CCs are assumed either static or moving.

the "Static CC BPMs" scenario be the most realistic one, the corrector usage would drop considerably for the correctors in Q4 to Q6, which would allow for greater margins in the corrector budget and a crossing knob with less orbit in the CCs.

Summarizing, for subsequent studies we make the assumption that IP BPM alignment information is available, but not for the CC BPMs, resulting in the expected corrector strength usage shown in Fig. 8 for moving CCs.

2.3 Knob implementation

The implementation of `IP_CROSSING` was retrieved from the v1.5 round optics (choosing the crossing knob to be a 80%/20% combination of the "long" and "short" versions defined in the MAD-X model [2]), and the remaining knobs were implemented using the convex optimizer in POCKPy. The optimizer was run with the goal of minimizing the L^2-norm of the total corrector usage, with increased weights on the `MCBY` correctors, resulting in Fig. 9. As the figures demonstrate, all knobs can be implemented within the provided bounds on corrector strength. The corresponding orbits of all the implemented knobs can be seen in Appendix C, but the crossing knob in particular can be seen in Fig. 10. The crossing knob used, achieves a beam separation at the CCs which is less than 1 mm, and achieves a beam separation in the triplet which is above $12\,\sigma$.

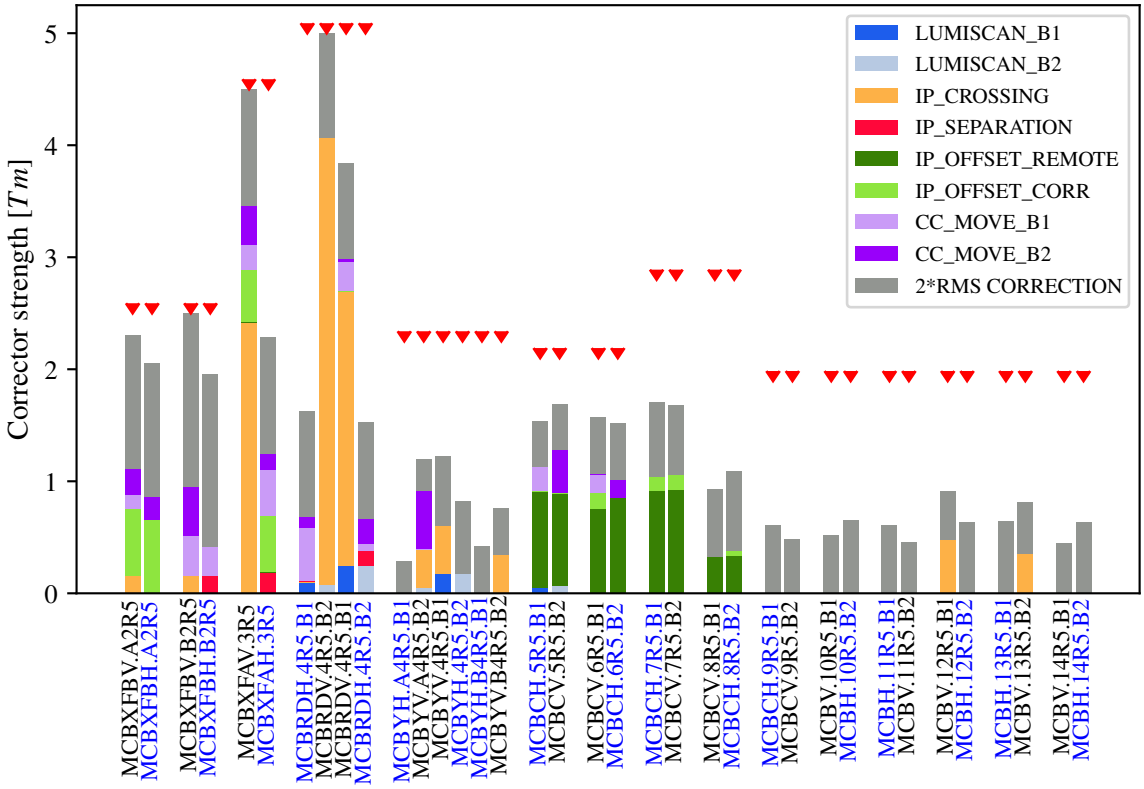

Figure 9: Bar plot of the corrector strength used for implementing all knobs and a 2*RMS error correction. Red triangles represent the strength limit for each corrector. Shown for all correctors up to Q14 on the right-hand side of IP5.

Detailed numerical tables summarising the corrector budget used by knobs and residual orbit are reported in Appendix B.

2.4 Corrector failure scenarios

In this section we will investigate whether individual orbit corrector failures still result in an operable machine. The process for assessing this will proceed as follows: remove one orbit corrector at the time and verify that the orbit corrector budget still holds. This provides an estimate of how dispensable each corrector is.

As a first step, each corrector was withheld for the error correction and a linear error correction was performed with the same scaling, except that the withheld corrector was weighted as zero (thus not being used in the correction). For this correction, it was assumed that the BPMs closest to the IP were static and corrected down to zero, while the BPMs closest to the CCs were supposed to move together with their associated magnet. Only the correctors up to Q18 were tested. The resulting 2*RMS residual for Beam 1 in

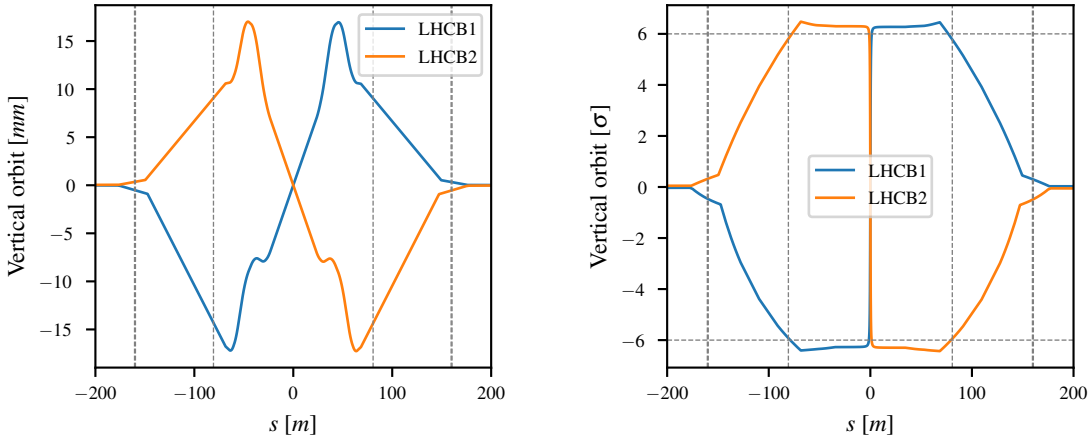

(a) Crossing knob in millimeters.

(b) Crossing knob in beam sigmas.

Figure 10: Orbit of the `IP_CROSSING` knob. Gray dashed vertical lines show the position of the CCs and D1, and gray horizontal lines correspond to 6σ.

the vertical plane is shown in Fig. 11.

For the correctors showcased in the given beam and plane, it is only the removal of `MCBXFAV.3L5` that causes the residual orbit to violate the aperture constraints. The same analysis conducted for the other beam and planes resulted in that the Q3 correctors could not be removed without violating the aperture constraints. In essence, from a strictly error correction point of view, the Q3 correctors are the only essential ones.

No detailed study was performed in case of loss of more correctors at the same time. In first approximation, one can assume that loosing two correctors simultaneously (with exception of Q3 correctors) is tolerated if the correctors are "far enough" from each other or they act on orthogonal planes.

To further investigate the robustness of the orbit correction, the previously defined knobs were subsequently implemented by the same routine as before: reserving the 2*RMS corrector strength employed in the error correction, then using the remaining corrector strength to implement the required knobs. Before discussing the results of this endeavour, it should be stated that this is an artificial approach, in the sense that the orbit knobs as defined presume all correctors being available. If one corrector is removed, one or more orbit knobs may become impossible to implement within the constraints on corrector usage. If such a corrector were to fail, in practice the orbit knobs would have to be redefined. A second comment to these results is that, while knobs were implemented as advertised, the orbits themselves were not checked in detail. In practice, even provided a guarantee that an orbit knob functions as defined, the overall orbit shape can still be undesirable to implement in the machine. Bearing these caveats in mind, the corrector failures that resulted in unfeasible knob implementations were: all `MCBRD` correctors, in Q4 `MCBY[HV].4[LR]5.B[12]`, and all Q3 correctors.

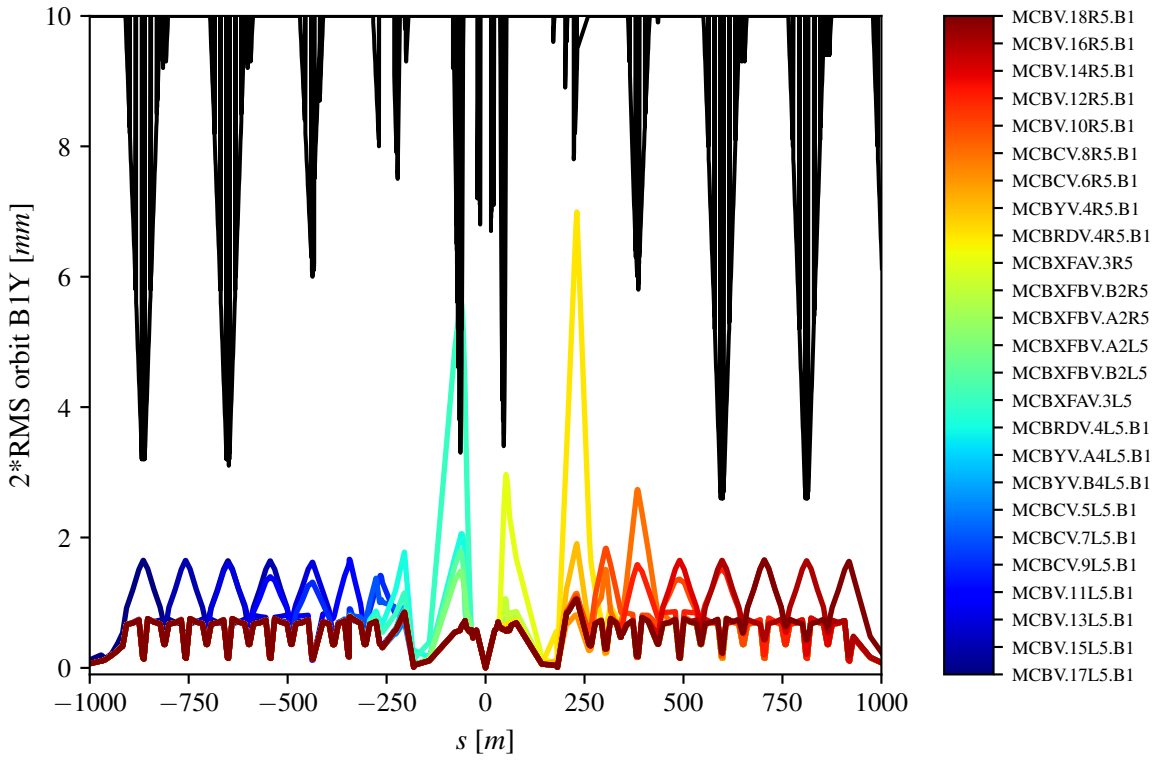

Figure 11: Plot of 2*RMS residual orbits, for Beam 1 in vertical plane, when removing a single corrector, where the upper black curve corresponds to the maximum tolerable radial orbit distortion in the machine as reported in Section 2.1.1. Each curve of a different color corresponds to the residual orbit achieved without the given corrector. The orbit is provided at the actual transverse center of all quadrupoles and dipoles.

2.4.1 Q9 CORRECTOR FAILURES

In practice, the correctors that deserve particular attention are those in Q9, as these correctors, in Octant 1 and 5, will have absorbed enough radiation dose to run a risk of failure in Run 4 [15]. For this reason, the failure of the Q9 corrector is here taken into account separately, beginning with the residual orbit shown in Fig. 12.

The increment in orbit from losing the Q9 correctors is still well-within the constraints on orbit in the arc. The updated corrector budget plot is shown in Fig. 13. As no orbit knob uses the Q9 correctors, the only effect of their failure is shifting around corrector strength to implement the new error correction. This results in additional corrector strength usage for the Q5 and Q7 correctors, but it constitutes such a small increment that the previous knob implementation does not need to be recomputed.

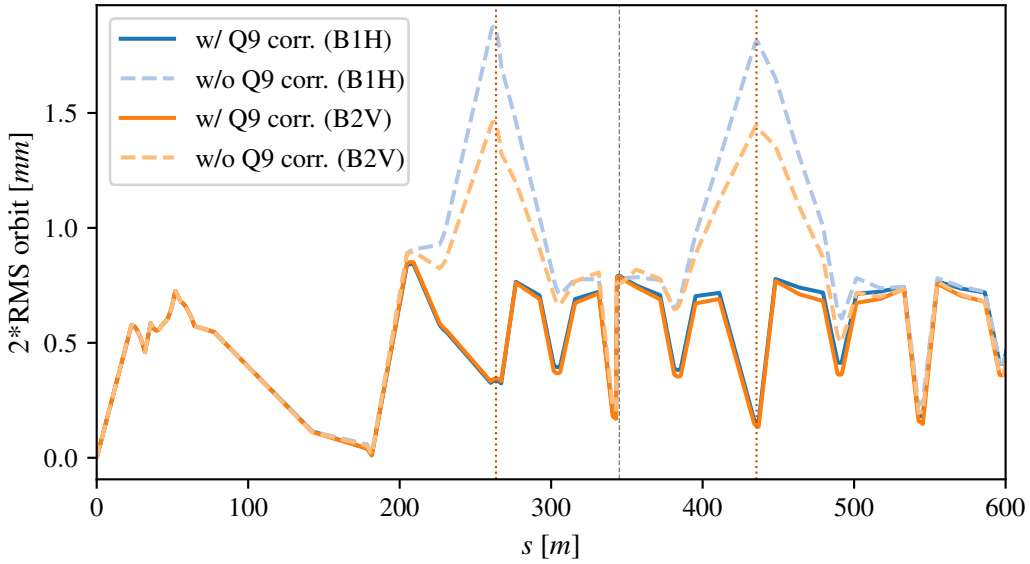

Figure 12: Plot of the 2*RMS residual orbit, with and without using the right-hand side Q9 correctors, shown for both combinations of affected beams and planes. Orbit is given with respect to the center of quadrupoles and dipoles, and the three vertical lines in the plot correspond to Q7, Q9 and Q11.

2.5 Flat optics

Using the same aperture limit as that of Fig. 1, the maximum tolerable radial orbit distortion for the alternative "flat optics" [16] is reported in Fig. 14. The maximum tolerable radial orbit distortion shown here is worse than that of the round optics, due to the larger beam size. It is worth pointing out that the region considered for the corrector budget (up to Q25 on each side of IP5) does not differ noticeably between the round and flat optics, with the only notable difference being an exchange of crossing planes between IP1 and IP5 (IP5 has a horizontal crossing plane in the flat optics). Because the quadrupole strengths are the same for the two optics, the error correction yields the same results for both optics, and by symmetry the corrector budget can be extended from the round optics to the flat optics. Importantly, even though the flat optics yield a degradation of maximum tolerable radial orbit distortion in the machine below the 2mm threshold set in [13], the 2*RMS error correction provides an orbit (see for example Fig. 13) which still allows for a considerable margin with respect to the closed orbit available in the machine (Fig. 14).

2.6 Conclusion

Under the given assumptions, the orbit corrector budget holds. There is a sufficient amount of orbit corrector strength allotted for error correction and implementing the

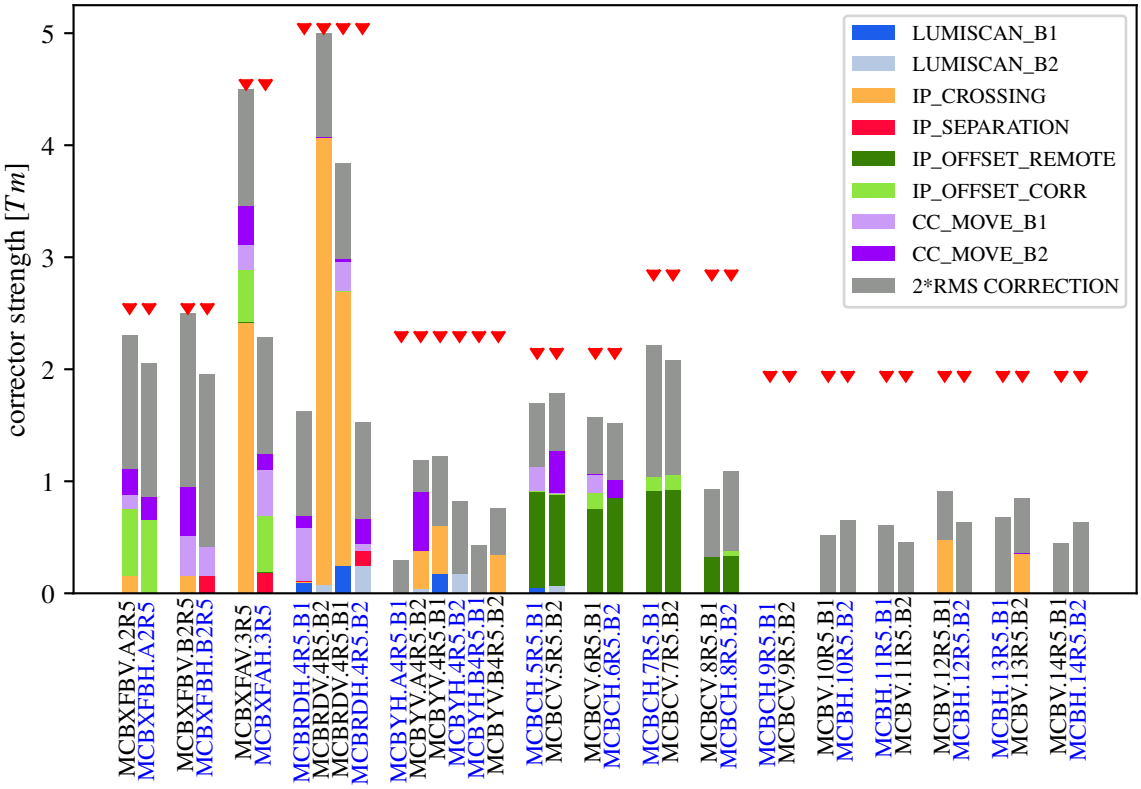

Figure 13: Bar plot of the corrector strength used for implementing all knobs and a 2*RMS error correction where no Q9 correctors were used. Red triangles represent the strength limit for each corrector. Shown for all correctors up to Q14 on the right-hand side of IP5.

listed knobs for the HL–LHC optics version 1.5. In all cases investigated for the error correction, even the most pessimistic one, the residual orbit with respect to the center of the elements is well-within the maximum tolerable radial orbit distortion as provided by aperture computations. The orbit corrector budget is also quite robust in the sense that individual corrector failures do not in general result in a machine violating the budget, with the exception of correctors in Q3 and Q4. In particular, the case of failure of the most irradiated Q9 correctors would result in an error correction with good margins and without affecting implementation of the various orbit knobs.

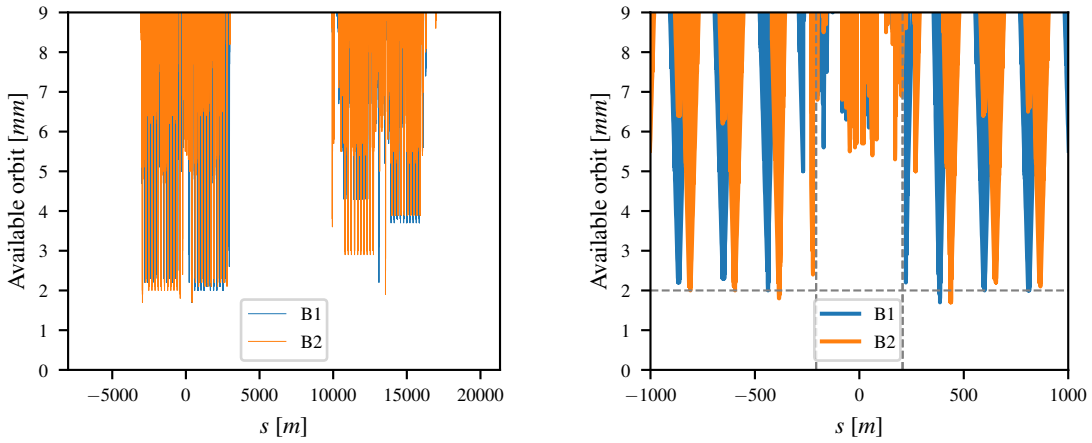

(a) Full machine.

(b) Around IP5.

Figure 14: Maximum allowed closed orbit for the flat optics shown globally (a) and around IP5 (b). $s = 0$ corresponds to IP3. Gray dashed vertical lines correspond to Q5.

3 Orbit Feedback System and Orbit Stability at IPs

Currently in LHC there is a feedback system employed for stabilizing the closed orbit during operation, referred to as the *Orbit Feedback System* (OFB) [17, 18]. It is run throughout the different phases of the machine cycle, from beam injection to dump at the end of collisions. Here the focus will be on the phase of Stable Beams, when beams are in collision.

During Stable Beams the global orbit is not a concern: the order of the usual closed orbit perturbation accrued over Stable Beams is too small to cause any beam loss. Exceptions to this are abrupt events (e.g. great ground motion, or equipment faults), but they are outliers and not the topic of this analysis. The issue is instead that the typical perturbation causes the beams to drift away from each other at the collision points, which could lead to luminosity loss. The object of analysis is therefore not the usual global orbit, but rather the beam separation at the collision points.

HL–LHC will also profit of the same orbit feedback system for closed orbit stabilization, which will be eventually consolidated [19]. As a first approximation of the feedback system in HL–LHC, we will transfer the OFB as defined in LHC to HL–LHC and compare the two implementations. This allows for some insight to be gained into how the OFB impacts luminosity and how this can be expected to scale for HL–LHC.

3.1 Model assumptions

The orbit feedback was modeled as a simple linear correction using the available orbit correctors [3]. The signal $\boldsymbol{y}$ measured at BPMs can be expressed as:

$$\boldsymbol{y} = \mathrm{B}\boldsymbol{e} + \boldsymbol{w}, \tag{3}$$

where $\boldsymbol{e}$ represents machine errors that, under the effect of a transfer matrix B, induce a closed orbit perturbation and $\boldsymbol{w}$ is noise added to the signal, i.e. BPM reading errors. Assuming a corrector response matrix, A, the corrector strength used for a given signal can be written as

$$\boldsymbol{c} = \mathbf{pinv}(\mathrm{A})\boldsymbol{y} = \mathbf{pinv}(\mathrm{A})\mathrm{B}\boldsymbol{e} + \mathbf{pinv}(\mathrm{A})\boldsymbol{w}, \tag{4}$$

where $\mathbf{pinv}(\mathrm{A})$ is the "pseudoinverse" of A which is computed by using the Singular Value Decomposition (SVD) of A. Depending on the number of singular values retained for the inversion, the needed corrector strength and the actual residual orbit after correction can vary drastically. The following data and assumptions were then used:

1. The LHC OFB does not use any corrector that is shared between the beams [20]. For HL–LHC, all shared correctors are assumed to be included in the OFB.
2. The OFB uses 40 out of 520 singular values (per transverse plane) for the pseudoinverse of the corrector response matrix [20].

3. All closed orbit perturbation taking place during collision is, or can be treated as being, caused by transverse movement of quadrupoles.
4. All transverse quadrupole movement is Independent and Identically Distributed (I.I.D.) in each dimension, which is what typically assumed for ground motion studies, e.g. [21].
5. All BPM reading errors are I.I.D.. in each dimension, assuming that those errors are dominated by the BPM signal acquisition chain.
6. As reference values, $\sigma_{\text{quad}} = 0.3$ μm is taken as the RMS quadrupole movement during a typical collision time of 12 hours and $\sigma_{\text{BPM}} = 20$ μm as the RMS BPM error for one reading [22].
7. LHC optics uses: $\beta^* = 30$ cm, $\varepsilon_{\text{N}} = 2.5$ μm and $E = 6.5$ TeV.
8. HL–LHC optics uses: $\beta^* = 15$ cm, $\varepsilon_{\text{N}} = 2.5$ μm and $E = 7$ TeV.

Under these assumptions, the RMS beam separation at an IP can be shown to satisfy the following equation [3]:

$$\sigma_{\text{d}} = \sqrt{(a \times \sigma_{\text{quad}})^2 + (b \times \sigma_{\text{BPM}})^2}\,, \tag{5}$$

where a and b are functions of the response matrices and the number of singular values employed in the pseudoinverse. For the residual orbit at an arbitrary point, there exists an analogous equation where a and b are functions of longitudinal position. Equivalently, if the orbit is not corrected there exists a linear relation as per:

$$\sigma_{\text{d}} = c \times \sigma_{\text{quad}}\,, \tag{6}$$

where c is a function of only the error response matrix.

3.2 Global residual orbit

Given the previous assumptions, an example result for the residual orbit in the full machine can be seen in Fig. 15 where LHC and HL–LHC are compared when using 40 and 400 singular values for the linear correction strategy in the OFB. The plots shown in Fig. 15 are normalized in units of beam size. For HL–LHC, the β-functions around IP1 and IP5 are greater than for LHC, which has the effect that quadrupole misalignments in these regions have a greater impact on the closed orbit at every other position, including the IPs. In addition, the same values of the normalized emittance for the two machines and the slightly different beam energy further contribute to the difference between the machines. Overall, the results shown in 15a are expected: there is a worsening of the orbit at IP1 and IP5 where the beam size has been shrunk in HL–LHC. To quantify why the beams are not corrected more aggressively, see Fig. 15b where the same scenario is plotted but for 400 singular values. In essence, the errors in the BPM readings propagate to the orbit, worsening it substantially, unless the number of singular values is adequately constrained for Stable Beams.

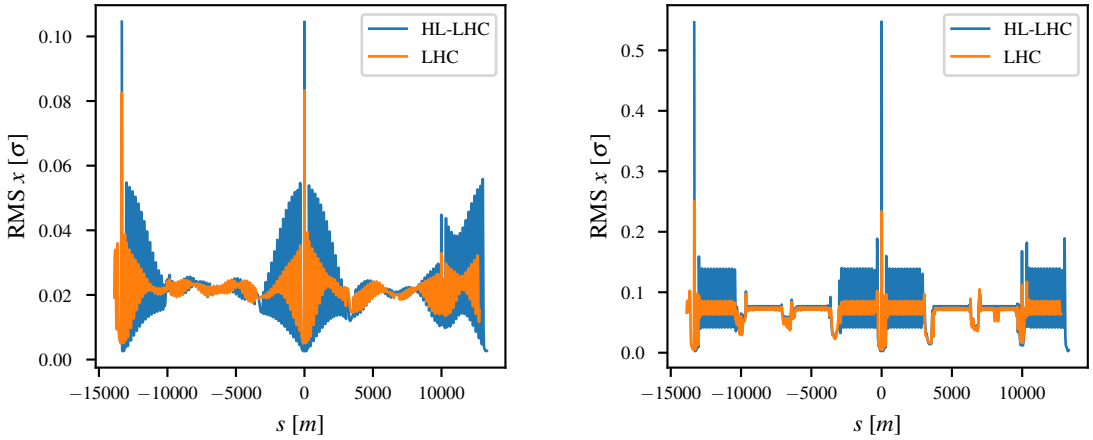

(a) 40 singular values.

(b) 400 singular values.

Figure 15: RMS horizontal residual orbit (with respect to the ideal reference) for Beam 1 at the end of collision using the OFB, measured in beam size and compared between LHC and HL–LHC for $\sigma_{\text{quad}} = 0.3$ μm, $\sigma_{\text{BPM}} = 20$ μm for 40 (a), and 400 (b) singular values used. $S = 0$ corresponds to IP5 and the leftmost peak corresponds to IP1 (note the difference in vertical scale).

3.3 Beam separation at IP1 and IP5

Figure 16 shows the beam separation at IP1 and IP5 in LHC for 40 singular values, and different values of the RMS quadrupole movement and BPM noise using Eq. (5). Note that even for perfect BPMs, the OFB with 40 singular values is only able to reduce the beam separation by roughly 20 % for a given magnitude of the quadrupole error. This is a direct result of the correction strategy being global and unweighted. If the OFB is to have the ability to considerably correct for beam separation, then it needs to weight the corrector response matrix and possibly change the number of singular values. In practice however, this added capability to correct for beam separation would necessarily propagate more BPM errors, and so the real bottleneck is the level of noise in the BPMs and not so much the choice of correction strategy.

It should be emphasized that the x-axes in Fig. 16 should not be directly associated with time. Even if the assumption is made that the RMS quadrupole error increases linearly with time (contrary to e.g. a Brownian motion where it is proportional to the square-root of time), it would be a stretch to claim any continuity. Rather, the graphs should be seen as scans over possible quadrupole errors by the end of collision and what the respective impact on beam separation would be.

For a given beam orbit separation d at the IP, the relative luminosity loss (L/L_0) is

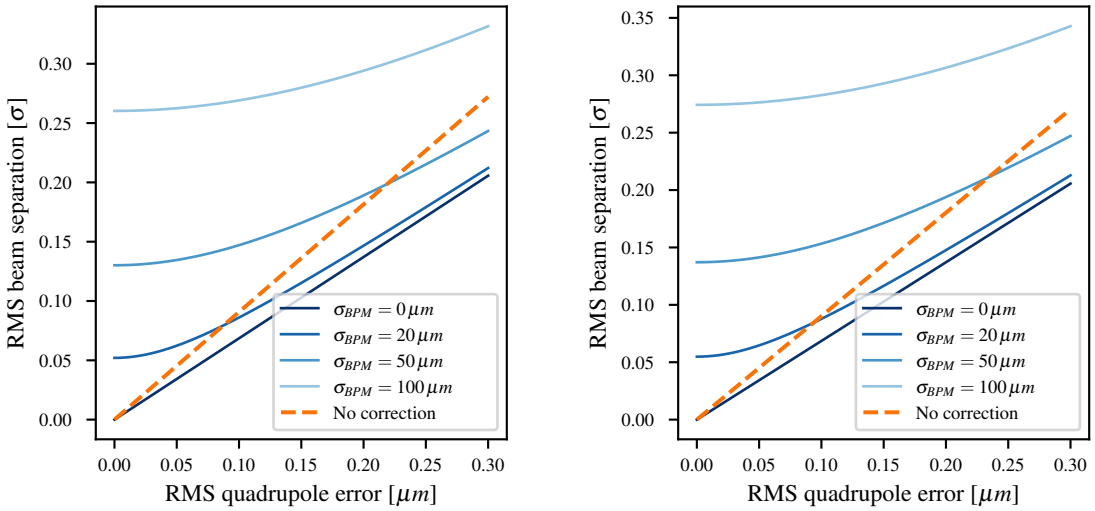

(a) IP1 (b) IP5

Figure 16: RMS beam separation measured in beam size at IP1 (a) and IP5 (b) using the OFB with 40 singular values, plotted versus σ_{quad} for several values of σ_{BPM}.

given by Eq. (7):

$$\frac{L}{L_0} = 1 - \exp\left(-\frac{d^2}{4\sigma^2}\right). \tag{7}$$

Therefore, a beam separation of 0.1 beam σ corresponds to about 0.25% luminosity loss and 0.2 beam σ corresponds to about 1% luminosity loss. To be noted in Fig. 16 that with the present OFB running and an assumed BPM noise of the order of $20-50\mu$m, then one would loose of the order of 0.25% luminosity even if the quadrupoles would not move. In reality, the OFB is typically run with a gain much lower than one[1]. This is equivalent to averaging over several orbit acquisitions, effectively reducing the BPM noise and therefore reducing the minimum achievable RMS beam separation with OFB on in Fig. 16. Additionally, other effects than orbit could play a role in the 0.1% luminosity-loss/-noise level, therefore the contribution of the OFB to luminosity loss is probably not observable. If higher OFB gain would be used in LHC, and according to the results presented here, a luminosity jitter would be visible.

Fig. 17 shows a comparison between LHC and HL–LHC for the same settings. From the perspective of beam separation at IP1 and IP5, HL–LHC performs worse, but comparably, and, as expected, there is no notable difference between the two IPs. In terms of just maintaining collision, beams are expected to remain in collision for the whole duration of a typical fill (i.e. even after the quadrupole-displacement-equivalent errors

[1]For the LHC OFB a gain of one is defined as the gain required for damping a static orbit error in 12-14 seconds from 100% to about 10% [14]. In this document, instead, a gain of one means to dump a static orbit error in a single correction iteration.

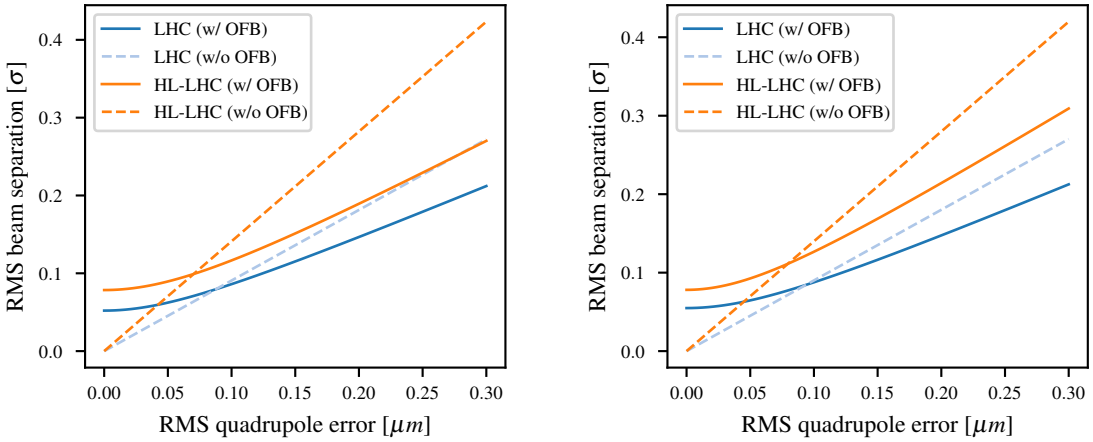

(a) IP1 (b) IP5

Figure 17: RMS beam separation at IP1 (a) and IP5 (b) using the OFB with 40 singular values as a function of the RMS quadrupole error, measured in beam σ's for $\sigma_{\mathrm{BPM}} = 20\,\mu$m. Here plotted against the benchmark of using no correction.

have grown up to $0.3\,\mu$m), even without OFB. In this case, at the end of an HL–LHC fill the beams would have separated by $0.4\,\sigma$ which, according to Eq. (7), would give a 4 % instantaneous luminosity loss by the end of collision for HL–LHC, which would not be acceptable.

Clearly, the actual dynamics during an HL–LHC fill will be much more complex, also due the expected operational scenarios including luminosity steps [23]. Already in LHC, luminosity scans are used during the fill to re-establish head-on collision, effectively cancelling any accumulated quadrupole-equivalent accumulated errors. A key property of the luminosity scans is that they can be performed on top of the OFB. In other words, if a luminosity scan is performed with the OFB active, the OFB does not cancel the orbit knob. This is a consequence of the OFB using a global correction, and the luminosity scan modifying the orbit localized around IPs. Due to the availability of the luminosity scan, operating virtually independently of the OFB, beam separation can and is made smaller than what is predicted here for LHC.

Lastly, it is worth mentioning that for Run 1 (2010-2013), LHC was run with the OFB inactive during collision and still proved stable [20]. This was however for other parameters; the collision energy was lower at 3.5 TeV, the beam sizes at collision were bigger, and no Achromatic Telescopic Squeezing (ATS) optics [24] was used, so that iteration of LHC was overall more resilient to closed orbit perturbation. Nevertheless, there are grounds for believing that LHC could remain operational even without the OFB during collision. This could be verified with a dedicated Machine Development slot during Run3.

3.4 Beam separation at IP2 and IP8

The procedure for determining the beam separation at IP2 and IP8 is analogous to that of IP1 and IP5, but it is not the quantity of interest from the perspective of luminosity. IP2 and IP8 undergo offset leveling, in the horizontal and vertical plane, respectively, during collision. Due to the luminosity's dependence on beam separation, the plane being leveled will be more sensitive to orbit perturbation, and so the quantity of interest is not the beam separation considered before, but rather the perturbation in the leveled plane. Here the following optical conditions were assumed (proton run):

1. LHC optics: $\beta^* = 10m$, and $3m$ for IP2 and IP8, respectively.
2. HL–LHC optics: $\beta^* = 10m$, and $1.5m$ for IP2 and IP8, respectively.

The nominal separation between the two beams was considered to be 4σ and 2σ for IP2 and IP8, respectively.

Once more using Eq. (5), it is possible to compute the RMS perturbation of the orbit in their respective leveling planes for IP2 and IP8 in LHC, which is done in Fig. 18. In

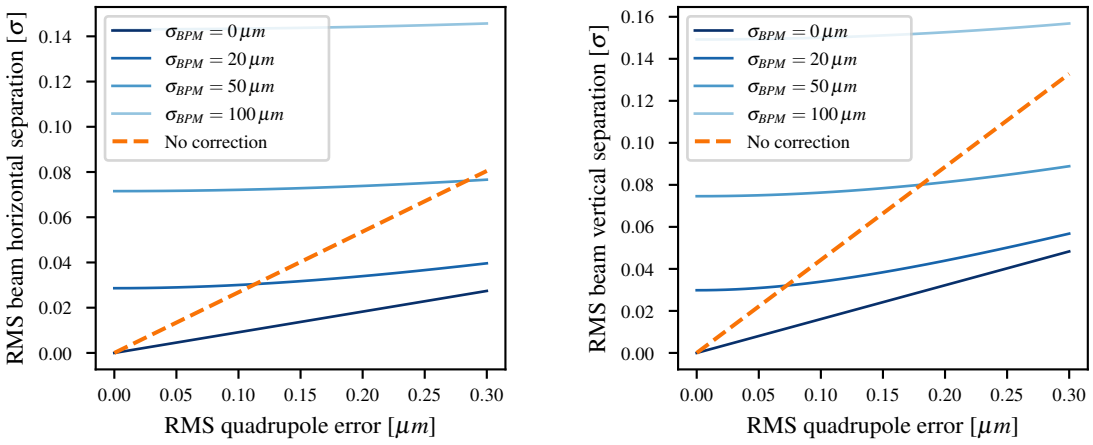

(a) IP2 (b) IP8

Figure 18: RMS beam separation measured in beam size at IP2 (a) and IP8 (b), in their respective leveling plane, versus σ_{quad} for several values of σ_{BPM}. The OFB is assumed to run with 40 singular values.

relation to IP1 and IP5, IP2 and IP8 appear to be more sensitive to BPM errors. The efficacy of the OFB appears greater for IP8, where under the assumption of an RMS BPM error of $50\,\mu$m, it is capable of improving the orbit already for an RMS quadrupole movement of $0.2\,\mu$m versus the corresponding $0.3\,\mu$m of IP2.

By the same procedure, a comparison between the OFB performance between LHC and HL–LHC is provided in Fig. 19. For head-on collisions, fluctuations in beam separation take the expression of direct loss in luminosity whereas for offset-leveled collisions,

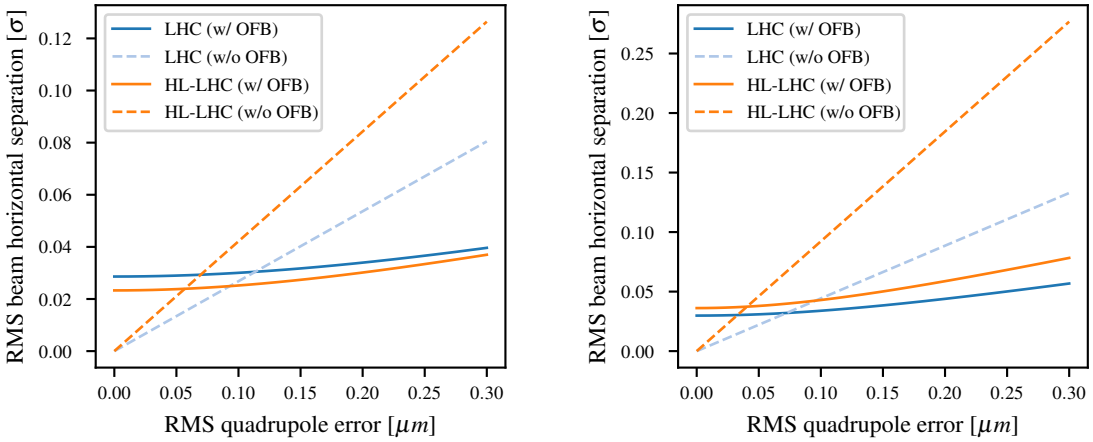

(a) IP2

(b) IP8

Figure 19: RMS beam separation at IP2 and IP8, in their respective leveling plane, measured in beam σ's as a function of the RMS quadrupole error, assuming $\sigma_{\text{BPM}} = 20$ μm and using the OFB with 40 singular values. Here plotted against the benchmark of using no correction.

fluctuations in separation can result in either increased or reduced luminosity. More importantly, for a given RMS beam separation in the leveling plane, the RMS luminosity loss will be proportional to the target beam separation, and this latter changes over a fill and may differ between fills. In summary, translating the results in Fig. 19 to luminosity is non-trivial and will not be performed here.

3.5 Conclusion

The results derived here indicate that the OFB used in LHC improves the beam separation stability at IP1 and IP5 during collision, but not significantly. This result largely translates to IP2 and IP8 where the OFB is somewhat more effective. Transferring the OFB as currently implemented for the LHC to HL–LHC leads to similar results, with a general worsening of beam separation driven in part by a reduced beam size at collision points Eq. (7). If one were to improve OFB performance in stabilising the orbit separation at the IPs, one would need to rely on better BPM stability. How much better is the argument of the next study.

4 BPM Stability Requirements for Improved Orbit Stability at IPs

The integrated luminosity over a fill can be viewed as a benchmark for the performance of a collider. In the previous chapter it was shown that the present BPM and OFB performances together with the remarkable stability of the LHC allow for maintaining collision with no significant losses of luminosity. Among others, potential new sources of closed orbit perturbation that could worsen the situation in HL-LHC are the new IR1 and IR5 triplets mechanical design, which could be more sensitive to mechanical vibrations, and the extensive use of *luminosity levelling* for those IPs. The latter will be achieved by changing in steps the β-function (β^*-*levelling*) and/or the beam separation (*orbit-levelling*) at the IPs. At each step, non-negligible closed orbit perturbation might be induced. Because of this, using a luminosity scan might be necessary to bring the beams back into head-on collision. Performing a luminosity scan takes in the order of a minute [14], and in this process some luminosity is inevitably lost as the luminosity itself is used as the signal to optimize over. During LHC Run 2, a few β^*-levelling steps were performed, and at times, no luminosity scans were necessary to recover head-on collision condition, therefore the associated integrated luminosity loss was probably negligible. In HL–LHC, the levelled luminosity will be increased nominally by a factor five while the number of bunches is kept constant [25]. This will increase the event pile-up and so the β^*-levelling used in HL–LHC will be more involved, with many more levelling steps, thus potentially increasing the number of luminosity scans performed during collision and thus the probability of losing a significant amount of integrated luminosity.

One alternative to luminosity scans, and in general to stabilise the orbit at the IPs independently of external sources of orbit perturbation, could be to deploy a more aggressive OFB that make use of the BPMs closest to the IP and correcting them to the design orbit using orbit correctors. Such an approach would require highly accurate and stable BPMs close to the IPs. The subject of this study is therefore to estimate how accurate these BPMs would have to be in order to find and maintain optimal collision, without, for example, continuously make use of luminosity scans.

4.1 Model assumptions

As luminosity scans are performed to bring beams into head-on collision, any local IP correction employed as a substitute has to reduce the beam separation sufficiently. As a benchmark of this, we use $0.1\,\sigma$ beam separation, corresponding to a 0.25 % instantaneous luminosity loss. In other words, if a local IP correction can achieve a $0.1\,\sigma$ RMS beam separation, it is considered a valid substitute for luminosity scans.

To study these corrections, the following setup was used and model assumptions were made:

1. Only Octant 5 is considered from Q25 to Q25, with quadrupole errors from Q20 to Q20.

2. All closed orbit perturbations are, or can be treated as being, caused by transverse movement of quadrupoles.
3. All quadrupoles' transverse movements are I.I.D. in each plane.
4. All BPM reading errors are I.I.D.. in each plane.
5. As a reference value, $\sigma_{\text{quad}} = 0.3\,\mu$m is taken as the RMS quadrupole movement to correct for during collisions [22].
6. HL–LHC optics for protons: $\beta^* = 15$ cm, $\varepsilon_{\text{N}} = 2.5\ \mu$m and $E = 7$ TeV.

The choice of considering only one octant is mainly due to computational reasons. The underlying assumption is that local errors are assumed to be the dominating source of errors at the corresponding IP. In practice, the whole machine errors will contribute to the worsening of the orbit at one IP. However, the orbit perturbation will be dominated by the movement of the triplet quadrupoles of IP1 and IP5.

For the extent of this study, pairs of BPMs will be used where one pair is defined as per Fig. 20. In other words, the ith BPM pair is the pair of BPMs formed by taking the

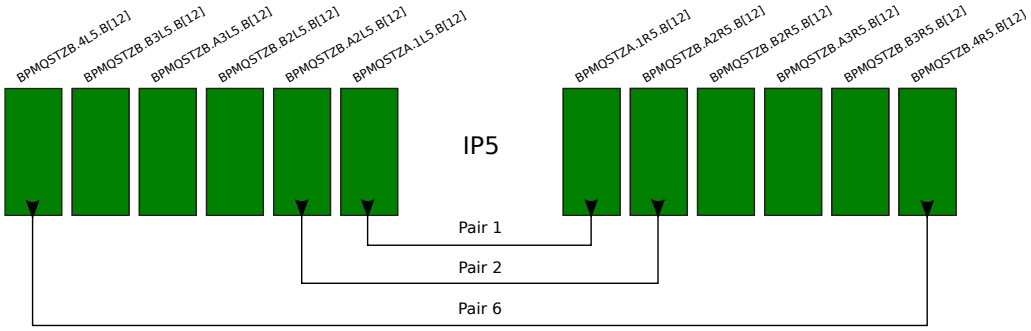

Figure 20: Definition of the BPM pairs in the interaction region around IP5.

ith BPM away from the IP, on each side of the IP and for both beams. Each single BPM is assumed to be able to measure the beam position for both counter rotating beams with the same accuracy.

4.2 Single BPM pairs

Under these assumptions, Fig. 21 shows the residual RMS beam separation at IP5 when using a single pair of BPMs as a function of RMS BPM error. Each point in the graph is computed using the optimal number of singular values for the orbit correction for that particular configuration, which also explains why all curves saturate to the level of not correcting at all. In those cases, the scanning procedure finds more beneficial to use zero singular values, i.e. not correcting at all, than attempting any correction that would just worsen the beam separation. No strategy on how to find the best number of singular

values in practice is provided. In fact, this aims to represent the best-case scenario one could achieve scanning over the number of singular values. In practice this is what is typically done also empirically in LHC [14]. If one uses only one BPM pair at a time, the trend is clear: the closer the BPM pair is to IP5, the more effective it is at correcting the beam separation. The underlying reason has to do with the number of error sources between each BPM pair. The closest BPM pair performs the best as there is no error source between the BPMs, and the further away a BPM pair is the more error sources. Because of this relation, the correction of the closest pair is in fact independent of the level of closed orbit perturbation; if the closed orbit is corrected to zero at these BPMs, it directly follows that the closed orbit at IP5 is zero too. In the same manner, correcting to zero at a BPM pair when there is an error source in between will invariably lead to a residual at the IP. This is most clearly seen in Fig. 21 at the leftmost points where the BPMs are perfect: pair 1 corrects to zero beam separation, every other pair does not with a performance worsening with distance from the IP.

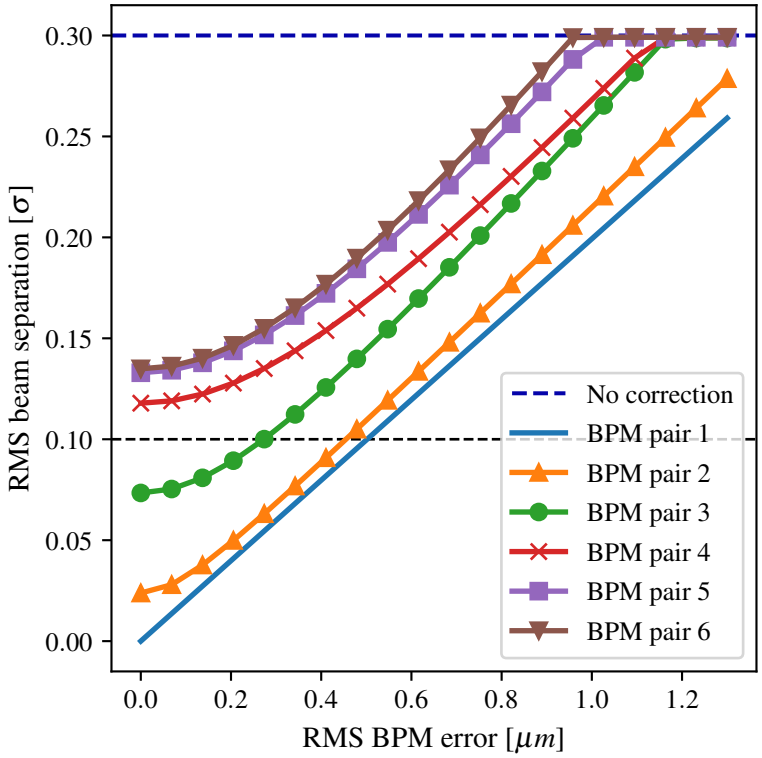

Figure 21: RMS beam separation at IP5 after local correction as a function of RMS BPM error for $\sigma_{\text{quad}} = 0.3$ μm, plotted for different single pairs of BPMs used for the correction. The dashed blue line is the expected beam separation assuming no correction. The dashed black line corresponds to 0.25% luminosity loss.

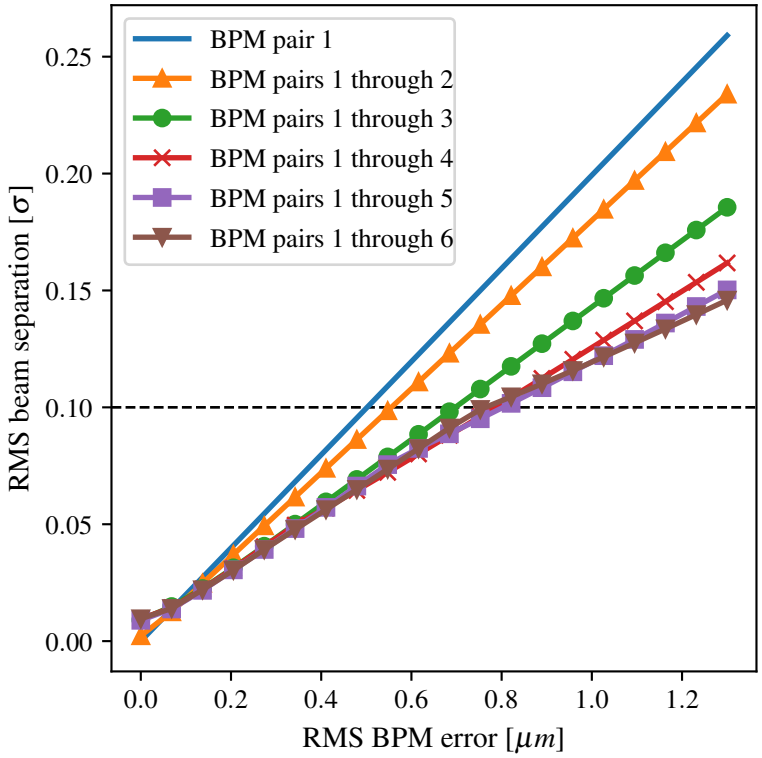

Figure 22: RMS beam separation at IP5 after local correction as a function of RMS BPM error for $\sigma_{\text{quad}} = 0.3$ μm, plotted for different sets of BPM pairs used in the correction. The dashed black line corresponds to 0.25% luminosity loss.

4.3 Multiple BPM pairs

If one has access to several imperfect BPMs close to each other with no sources of orbit perturbation in between, then it is possible to reduce the level of noise in the readings based on the collective reading being nonphysical. A naive schematic of this is shown in Fig. 23. The figure assumes that any physical orbit must be a straight line for the region. In such a system, the orbit can be estimated based on the line with the least-squares distance to every measurement, i.e. the error would scale with a $1/\sqrt{n}$ law, where n is the number of BPMs.

In practice, in an accelerator having several BPMs with no other elements in between is normally not affordable. Instead, a single BPM is placed in between two magnets that could be source of orbit perturbation. Still, one could imagine that not all orbits are physically possible, and so using several BPMs instead of a single one per side of the IP should provide some benefit. This principle is used in Fig. 22 where groups of BPM pairs are used in the correction. Here it becomes evident that one can improve the result the single BPM next to the IP up to about 40% by using all the six BPMs available in

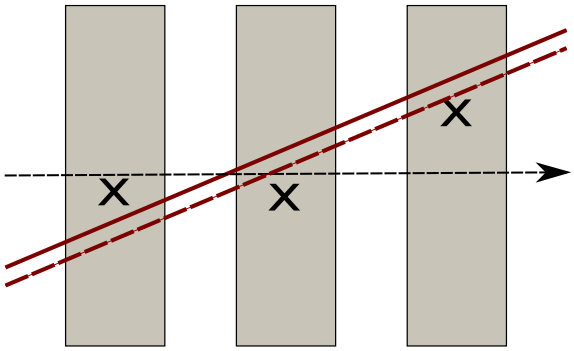

Figure 23: Schematic of beam measurement by imperfect BPMs. The gray blocks correspond to BPMs, the full red line corresponds to the actual closed orbit, the black crosses correspond to the BPM reading and the red dashed line corresponds to the reconstruction of the closed orbit based on the readings.

the IR for a given quadrupole and BPM errors, which is compatible with the $1/\sqrt{n}$ law mentioned before. Note that the gain saturates already after using the first four BPM pairs. This can be understood by the fact that the BPMs in the IR are interleaved by quadrupoles, which are sources of orbit distortion, therefore one cannot discern between all combination of BPM noise and quadrupole movements. To be stressed, once again, that each point in Fig. 22 is obtained by scanning over the number of singular values used for the correction and by choosing the number that minimize the beam separation at the IP. In a real machine this could be found by keep adding a singular value to the error correction until perturbations on the luminosity signal are seen.

The results from Fig. 21 and 22 are summarized in Table 4. The table provides the tolerated RMS BPM error that still allow for maintaining collision with a luminosity loss below 0.25% for the assumed quadrupole error of $\sigma_{\text{quad}} = 0.30\,\mu$m, and assuming a dedicated local OFB. This stability would need to be kept over the full collisions period, which is in the order of ten hours, or over a shorter period after which a luminosity scan could be performed to re-set the reference orbit.

It is worth noting that the correction performs worse when including the sixth pair of BPMs. The underlying reason is likely to be the fact that the trade-off between utilizing more information and prioritizing the correction of the new BPM: when the sixth BPM pair is added, we manage to filter away more BPM errors, but we also correct more at the BPM which is detrimental to the IP orbit. Looking at Fig. 22, it is evident that adding the sixth BPM pair pays off for greater BPM errors, but this payoff becomes tangible at beam separations greater than what we are here interested in.

Table 4: Table over tolerated RMS BPM error to achieve a beam separation at IP5 corresponding to a 0.25% instantaneous luminosity loss for $\sigma_{\text{quad}} = 0.3\ \mu$m for different BPMs employed in the correction.

BPM pairs used	**Required RMS BPM stability [μm]**
1	0.50
2	0.45
3	0.27
4	N/A
5	N/A
6	N/A
1 through 2	0.56
1 through 3	0.70
1 through 4	0.79
1 through 5	0.80
1 through 6	0.76

4.4 Finding Collision

A separate use of local IP correction that has not been mentioned is as a means of finding a luminosity signal. If HL–LHC is not operated for an extended period of time, or some equipment fault takes place between two fills, then the closed orbit in the machine might be large enough that collision cannot be found after the ramp. Luminosity scans are also performed in these scenarios, but finding collision can take considerably longer than a minute since there is no luminosity signal. Once a luminosity signal is found, head-on collision can be achieved in around a minute [14]. If the stability of the closest BPM pair around IPs is good enough, it would be possible to correct down to the design orbit at these locations and a luminosity signal would be attained instantly.

A beam separation of $4\,\sigma$ still gives a sufficient luminosity signal to be able to quickly find head-on collision [14]. As correction using the closest BPM pair is independent of the closed orbit in the machine, the result from Table 4 can be scaled directly to compute the necessary BPM stability to find always collision:

$$\text{Required BPM stability } \sigma_{\text{BPM}} \text{ to find collision} = 40 \times 0.50\,\mu\text{m} = 20\,\mu\text{m} \tag{8}$$

where the factor 40 is the quotient between $4\,\sigma$ and $0.1\,\sigma$. The timescale of this stability dictates how useful the local IP correction can be to find collision. If it is less than a day, then it is not very useful, whereas if it is longer than a day then one can possibly save time using this approach. Note that the value computed here is for the $\beta^* = 15cm$ optics, which is normally expected to be in use at the end of collision. For a higher β^* optics, one should expect to relax the obtained value by the square root of the β^* ratio.

4.5 Conclusions

Based on the results presented, the BPMs closest to the points of collisions would need a stability in the order of 0.5 micrometer on a timescale of around ten hours to be used as a means of steering and maintaining beams into head-on collision (luminosity loss below 0.25%; 1 micrometer BPM stability for ensuring luminosity loss below 1%) using a dedicated local OFB. The BPM pairs most effective at correcting orbit at the points of collision are the closest ones, with their effectiveness degrading with distance. Note that the stability requirement is given assuming an instantaneous orbit correction, without giving specific boundary for the OFB bandwidth. Depending on the source of BPM acquisition instability, one could achieve the desired stability by averaging the BPM acquisition over time at the expense of the maximum bandwidth of the OFB. Additionally, using up to the first four BPMs next the IP on each side could be beneficial.

If one is interested in using local IP correction as a means of finding collision, then the minimum BPM stability is around twenty micrometers for whatever timescale is considered.

The results presented here are compatible with those of a previous study on older HL–LHC optics and under different assumptions [26]. For that study, the necessary BPM stability to replace luminosity scans during collision was estimated to be $\pm 2\mu m$ (criterium: luminosity loss smaller than 1%), and $\pm 44\mu m$ to find back collision after a fill (criterium: find 1% of luminosity at the beginning of the fill). For finding collision after a long stop, the operational experience is that $20 - 50\mu m$ BPM stability is enough to find back collision [14], which is also compatible with the value of $20\mu m$ suggested by this study.

As to the larger question of whether it is worthwhile to implement such a local OFB in HL-LHC, the following items have to be performed:

1. Assess if a degradation of machine stability is expected in HL–LHC compared to LHC.
2. Investigate the feasibility of micrometer stability for the BPMs near the collision points.
3. If it proves feasible to upgrade the BPMs to this precision, quantify the cost of doing so, which can be favorably compared to integrated luminosity.

Acknowledgments

Research supported by the HL–LHC project.

Appendices

A Search for error correction parameters

In this appendix, the process and heuristics for choosing the error correction parameters are covered. As stated in the error correction Sec. 2.2, the weights for equality constraints was set to 1.0 and those for arc BPMs to 10^{-6}, and subsequently the number of singular values were 47 out of 57 singular values (per plane). To evaluate a given linear correction defined by these parameters, the following was performed:

1. Compute the correction strategy based on a weighted pseudoinverse of the corrector response matrix.
2. Use the correction strategy to compute the RMS residual orbit and corrector strength usage for the errors considered (BPMs assumed to have no misalignment errors).
3. Analyze the L^2-norm of the residual orbit in the triplet, arc and the L^2-norm of the corrector strength usage.

Using this approach, it is possible to plot the performance of a given correction strategy, for a given number of singular values, as a function of the triplet BPM weights.

As a second heuristic, we want as few singular values as possible in the correction strategy. The idea behind such a rule is that if n singular values are sufficient to adequately correct for errors, then the $n+1$ singular value might improve the orbit slightly, but will typically do so at an additional expense of corrector strength. Moreover, too many singular values allow for the correction to reconstruct nonphysical orbits, e.g. noise returned by BPMs. For these reasons, the minimal number of singular values possible will be used.

As mentioned, for the correction strategy the number of singular values was to be kept as few as possible and the only other parameter entering into consideration was that of the weight on triplet BPMs. Scanning these two parameters, the final correction strategy came out as using 47 out of 57 singular values and a weight on the triplet BPMs equal to 1.25×10^{-6}. A plot demonstrating the aptitude of these values is shown in Fig. 24 What the plot shows is that both 46 and 47 singular values are sufficient for adequately correcting the triplet, but 46 singular values are insufficient for adequately correcting the arc. For 47 singular values and a triplet weight of 1.25×10^{-6}, the arc together with the triplet are both adequately corrected, and so 47 singular values are the minimum number of singular values necessary, and for the given choice of triplet weight it provides a valid correction strategy as per our qualifications.

As a last comment, note that these parameter values are expected to change if alterations of (or imperfections in) the current optics take place. Small changes in the optics may significantly perturb the singular-value decomposition of the response matrices, and

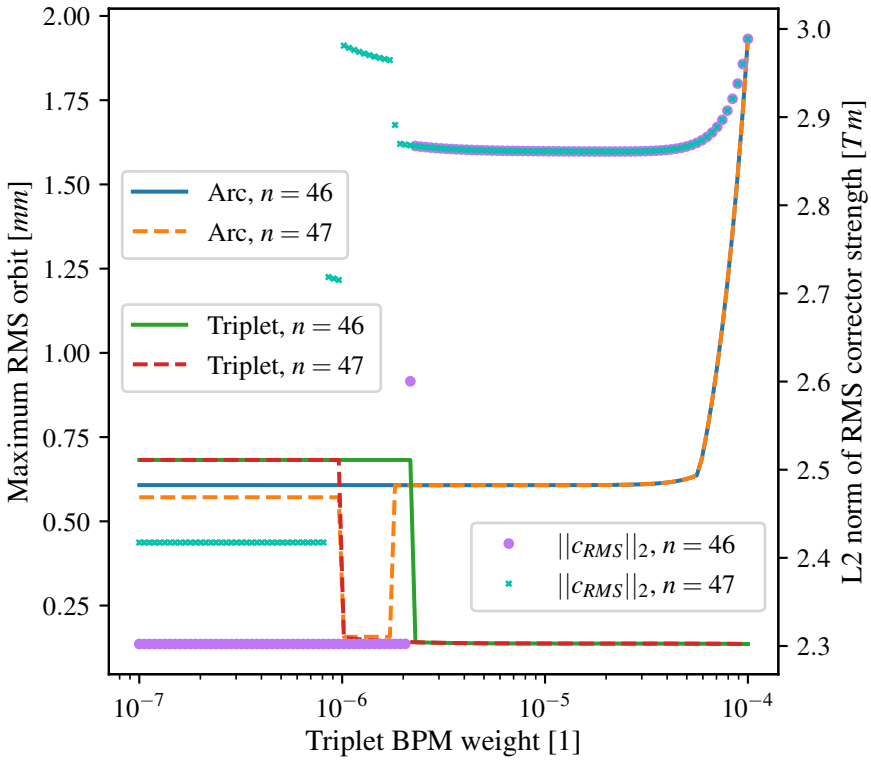

Figure 24: Triplet weight scan for 46 and 47 singular values plotting the maximum RMS residual orbit, across each beam and dimension, and the L^2-norm of the RMS corrector strength usage for a linear correction.

so these values should not be considered immutable but rather approximately optimal for the given model of HL–LHC.

B Numerical Results from Error Correction and Knob Implementation

The corrector budget is summarized in Table 5 and the orbit knobs, together with the error correction, are summarized in Table 6. From the perspective of critical apertures and margins of corrector strength, numerical values of the RMS residual orbit have been requested. These are provided in Table 7 and 8.

Table 5: Table of the corrector strength budget, shown for all correctors up to Q15. The values shown are taken as the maximum of the absolute corrector strength used across all matched correctors and are given in units of Tm. The 'SUM' column refers to the maximum over all absolute-value sums for the given pattern.

CORRECTORS	LUMI_B1 [±0.1 mm]	LUMI_B2 [±0.1 mm]	IP_CROSS [±295 μrad]	IP_SEP [±0.75 mm]	OFF_REMOTE [±2 mm]	OFF_CORR [±0.5 mm]	CC_B1 [±0.5 mm]	CC_B2 [±0.5 mm]	ORBIT [2σ]	SUM	LIMIT
^MCBXFB[HV].A2	0.00	0.00	0.16	0.00	0.00	0.67	0.25	0.11	1.20	2.07	2.50
^MCBXFB[HV].B2	0.00	0.00	0.16	0.16	0.00	0.00	0.35	0.45	1.55	2.50	2.50
^MCBXFA[HV].3	0.00	0.00	2.42	0.19	0.00	0.50	0.41	0.21	1.04	4.50	4.50
^MCBRD[HV].4	0.25	0.00	3.99	0.14	0.00	0.00	0.10	0.47	0.94	4.85	5.00
^MCBY[HV].[AB]?4	0.04	0.18	0.43	0.00	0.00	0.00	0.00	0.53	0.65	1.10	2.25
^MCBC[HV].5	0.00	0.07	0.00	0.00	0.88	0.00	0.44	0.00	0.43	1.14	2.10
^MCBC[HV].6	0.00	0.00	0.00	0.00	0.85	0.14	0.18	0.00	0.52	1.20	2.10
^MCBC[HV].7	0.00	0.00	0.00	0.00	0.91	0.14	0.00	0.00	0.67	1.45	2.80
^MCBC[HV].8	0.00	0.00	0.00	0.00	0.33	0.05	0.00	0.00	0.71	1.06	2.80
^MCBC[HV].9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.61	0.61	1.90
^MCB[HV].10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.66	0.66	1.90
^MCB[HV].11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.61	0.61	1.90
^MCB[HV].12	0.00	0.00	0.48	0.00	0.00	0.00	0.00	0.00	0.64	0.91	1.90
^MCB[HV].13	0.00	0.00	0.41	0.00	0.00	0.00	0.00	0.00	0.64	0.84	1.90
^MCB[HV].14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.64	0.64	1.90
^MCB[HV].15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.64	0.64	1.90

Table 6: Table of the absolute orbit, with respect to the ideal orbit, shown for a subset of elements. The values shown correspond to the maximum value taken over all the selected elements, both beams and planes and sides of the IP, and are given in units of mm. The 'SUM' column refers to the maximum over all absolute-value sums for the given pattern.

ELEMENTS	LUMI_B1 [±0.1 mm]	LUMI_B2 [±0.1 mm]	IP_CROSS [±295 μrad]	IP_SEP [±0.75 mm]	OFF_REMOTE [±2 mm]	OFF_CORR [±0.5 mm]	CC_B1 [±0.5 mm]	CC_B2 [±0.5 mm]	ORBIT [2σ]	SUM
^TAXS	0.10	0.10	5.89	0.75	2.00	0.50	0.00	0.00	0.00	8.49
^MQXFA.[AB]1	0.12	0.12	11.06	0.91	2.00	0.61	0.00	0.00	0.13	13.92
^MQXFB.[AB]2	0.16	0.16	16.88	1.17	2.00	1.07	0.23	0.22	0.32	20.36
^MQXFA.[AB]3	0.10	0.10	17.12	0.78	2.00	1.08	0.34	0.32	0.31	20.56
^MBXF	0.02	0.02	14.87	0.42	2.00	1.08	0.31	0.29	0.35	18.32
^TAXN	0.14	0.14	4.46	0.12	2.00	0.76	0.48	0.46	0.52	8.17
^MBRD	0.18	0.18	1.91	0.04	2.00	0.72	0.58	0.55	0.64	5.74
^MCBRD	0.20	0.20	0.86	0.01	2.00	0.70	0.63	0.59	0.69	4.74
^ACFGA	0.13	0.15	0.60	0.00	2.00	0.66	0.50	0.50	0.50	4.28
^MCBY[HV].[AB]?4	0.09	0.11	0.04	0.00	2.00	0.60	0.56	0.58	0.58	3.71
^MQY.4	0.08	0.09	0.00	0.00	2.00	0.59	0.48	0.51	0.58	3.60
^TCLMB.4	0.10	0.11	0.14	0.00	2.00	0.61	0.54	0.56	0.49	3.76
^TCLMC.5	0.02	0.03	0.00	0.00	2.00	0.45	0.21	0.20	0.66	3.29
^MCBC[HV].5	0.00	0.00	0.00	0.00	2.02	0.43	0.17	0.16	0.68	3.27
^MQML.5	0.01	0.02	0.00	0.00	2.01	0.43	0.19	0.18	0.71	3.26
^TCLMC.6	0.00	0.00	0.00	0.00	2.09	0.34	0.06	0.06	0.83	3.23
^MCBC[HV].6	0.00	0.00	0.00	0.00	2.08	0.33	0.00	0.00	0.86	3.23
^MQML.6	0.00	0.00	0.00	0.00	2.10	0.33	0.04	0.04	0.85	3.24
^MCBC[HV].7	0.00	0.00	0.00	0.00	0.66	0.11	0.00	0.00	0.59	1.19
^MQM.[AB]7	0.00	0.00	0.00	0.00	0.72	0.10	0.00	0.00	0.58	1.25
^MCBC[HV].8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.59	0.59
^MQML.8	0.00	0.00	0.00	0.00	0.07	0.01	0.00	0.00	0.58	0.61
^MCBC[HV].9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55
^MQMC.9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	0.52
^MQM.9	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55
^MCB[HV].10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.59	0.59
^MQML.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.57	0.57
^MCB[HV].11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.58	0.58
^MQ.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.56
^MQTLI.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.57	0.57
^MCB[HV].12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.56
^MQT.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	0.52
^MQ.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55
^MCB[HV].13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.56
^MQT.13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	0.52
^MQ.13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55
^MCB[HV].14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.56
^MQ.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55
^MCB[HV].15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.56
^MQ.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.55	0.55

Table 7: Maximum RMS orbit (taken with respect to the center of elements) in mm per applicable element granted the IP BPMs are treated as static, and the CC BPMs assumed to be moving. Each value tabulated is taken as the maximum across both beams, dimensions and all matched elements.

ELEMENTS	BPMs only	Quads+Bends only	Quads+Bends+BPMs
^MQXFA.[AB]1	0.00	0.35	0.35
^MQXFB.[AB]2	0.14	0.35	0.36
^MQXFA.[AB]3	0.16	0.35	0.34
^MBXF	0.14	0.30	0.35
^MBRD	0.32	0.29	0.06
^MQY.4	0.29	0.29	0.02
^MQML.5	0.34	0.31	0.45
^MQML.6	0.40	0.34	0.53
^MQM.[AB]7	0.29	0.31	0.19
^MQML.8	0.29	0.31	0.20
^MQMC.9	0.27	0.30	0.24
^MQM.9	0.27	0.30	0.40
^MQML.10	0.28	0.30	0.20
^MQ.11	0.28	0.30	0.20
^MQT.12	0.27	0.32	0.39
^MQ.12	0.27	0.30	0.21
^MQ.13	0.27	0.30	0.21
^MQ.14	0.27	0.30	0.21
^MQ.15	0.27	0.30	0.21

Table 8: Maximum RMS orbit (taken with respect to the ideal orbit) in mm per element granted the IP BPMs are treated as static, and the CC BPMs assumed to be moving. Each value tabulated is taken as the maximum across both beams, dimensions and all matched elements.

ELEMENTS	BPMs only	Quads+Bends only	Quads+Bends+BPMs
^TAXS	0.00	0.00	0.00
^MQXFA.[AB]1	0.00	0.06	0.06
^MQXFB.[AB]2	0.14	0.12	0.16
^MQXFA.[AB]3	0.16	0.12	0.16
^MBXF	0.14	0.10	0.17
^TAXN	0.26	0.03	0.26
^MBRD	0.32	0.02	0.32
^MCBRD	0.34	0.00	0.34
^ACFGA	0.25	0.00	0.25
^MCBY[HV].[AB]?4	0.29	0.01	0.29
^MQY.4	0.29	0.01	0.29
^TCLMB.4	0.25	0.01	0.25
^TCLMC.5	0.32	0.04	0.33
^MCBC[HV].5	0.33	0.05	0.34
^MQML.5	0.34	0.05	0.36
^TCLMC.6	0.39	0.08	0.42
^MCBC[HV].6	0.41	0.09	0.43
^MQML.6	0.40	0.09	0.42
^MCBC[HV].7	0.27	0.14	0.29
^MQM.[AB]7	0.29	0.15	0.29
^MCBC[HV].8	0.29	0.11	0.29
^MQML.8	0.29	0.12	0.29
^MCBC[HV].9	0.26	0.09	0.27
^MQMC.9	0.27	0.09	0.26
^MQM.9	0.27	0.10	0.27
^MCB[HV].10	0.27	0.13	0.29
^MQML.10	0.28	0.13	0.29
^MCB[HV].11	0.27	0.13	0.29
^MQ.11	0.28	0.13	0.28
^MQTLI.11	0.28	0.13	0.28
^MCB[HV].12	0.27	0.14	0.28
^MQT.12	0.27	0.14	0.26
^MQ.12	0.27	0.14	0.27
^MCB[HV].13/14/15	0.27	0.14	0.28
^MQT.13	0.27	0.14	0.26
^MQ.13/14/15	0.27	0.15	0.28

C Knob Orbits

In this section, all the knob orbits implemented in the orbit corrector budget study are showcased.

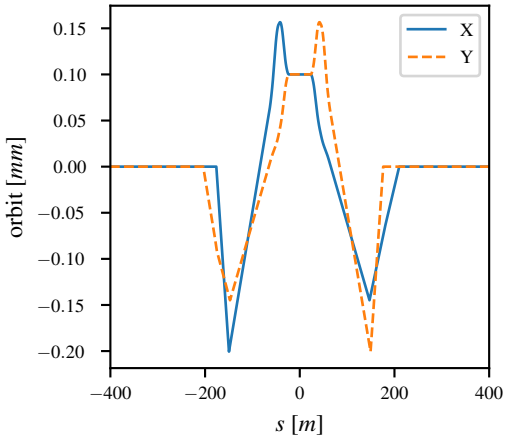

(a) Beam 1.

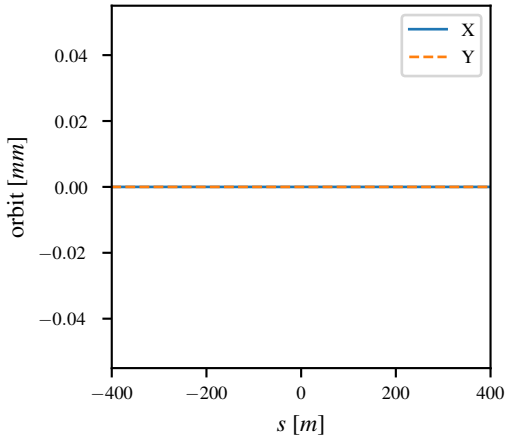

(b) Beam 2.

Figure 25: Plot of the `LUMISCAN_B1` knob, shown around `IP5`.

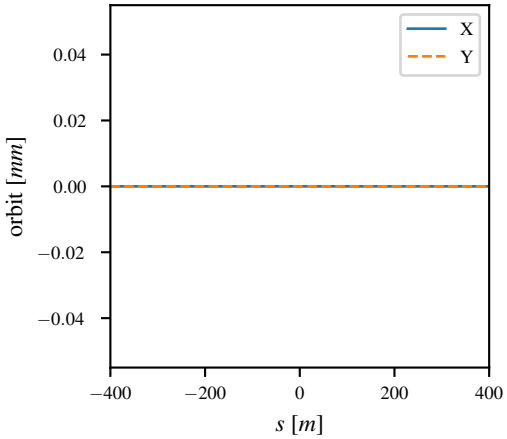

(a) Beam 1.

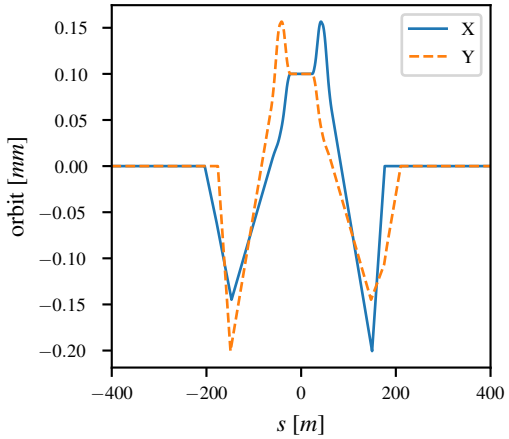

(b) Beam 2.

Figure 26: Plots of the `LUMISCAN_B2` knob, shown around `IP5`.

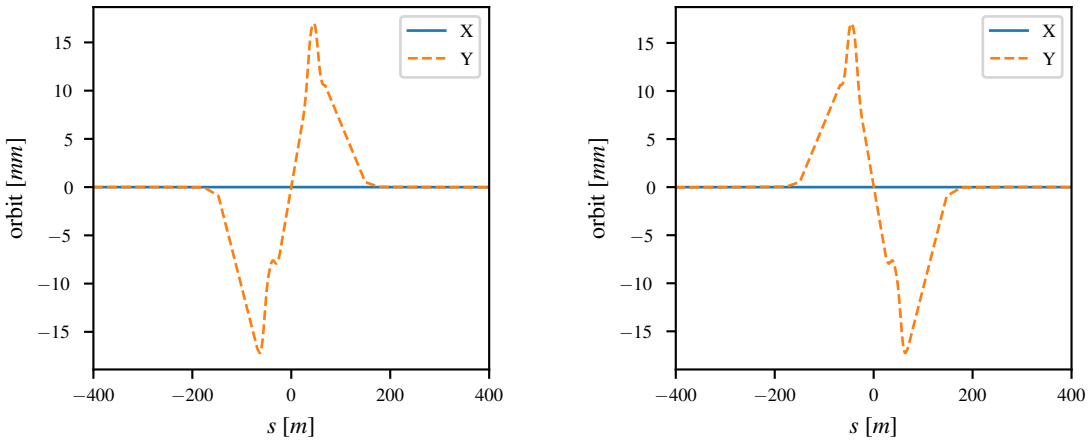

(a) Beam 1.

(b) Beam 2.

Figure 27: Plots of the `IP_CROSSING` knob, shown around `IP5`.

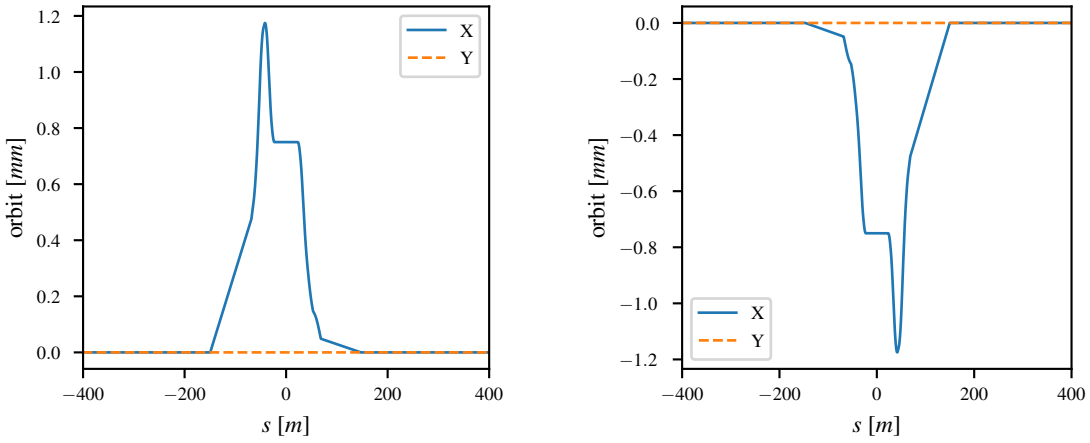

(a) Beam 1.

(b) Beam 2.

Figure 28: Plots of the `IP_SEPARATION` knob, shown around `IP5`.

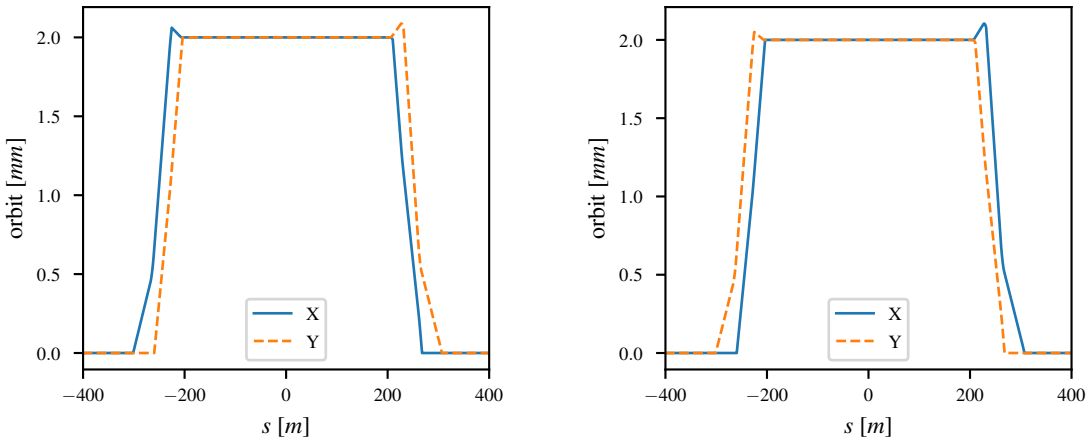

(a) Beam 1.

(b) Beam 2.

Figure 29: Plots of the `IP_OFFSET_REMOTE` knob, shown around `IP5`.

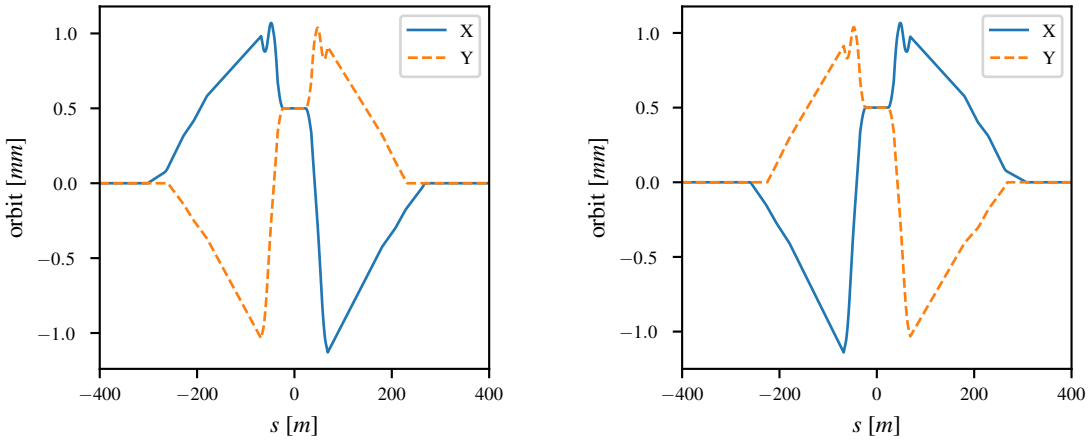

(a) Beam 1.

(b) Beam 2.

Figure 30: Plots of the `IP_OFFSET_CORR` knob, shown around `IP5`.

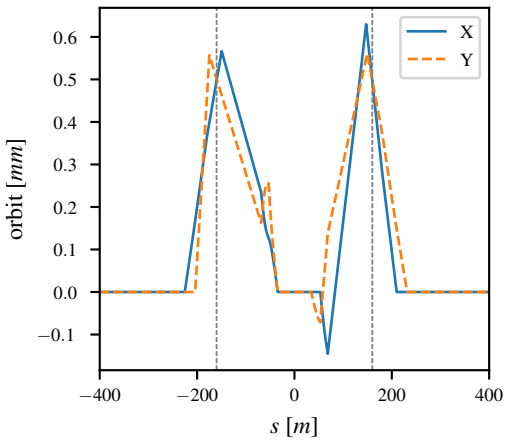

(a) Beam 1.

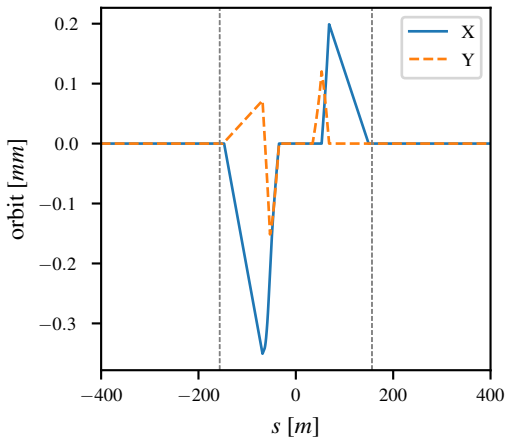

(b) Beam 2.

Figure 31: Plots of the `CC_MOVE_B1` knob, shown around `IP5`. Gray dashed vertical lines correspond to the location of the CCs.

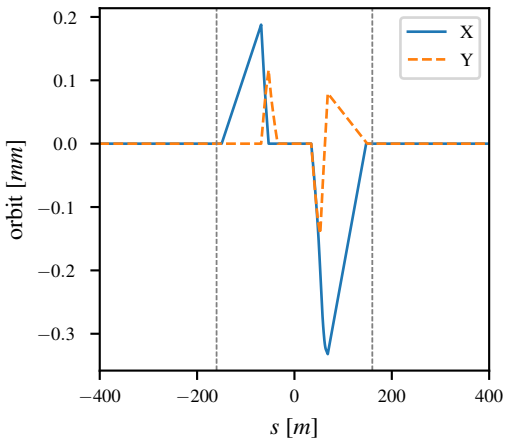

(a) Beam 1.

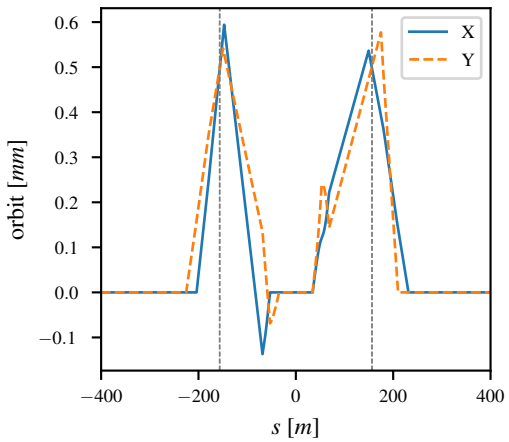

(b) Beam 2.

Figure 32: Plots of the `CC_MOVE_B2` knob, shown around `IP5`. Gray dashed vertical lines correspond to the location of the CCs.

D OFB Beam Separation Formulas

All the results on beam separation in the OFB study were based on Eqs. (5) and (6), repeated below:

$$\sigma_{d,corrected} = \sqrt{(a \times \sigma_{\mathrm{quad}})^2 + (b \times \sigma_{\mathrm{BPM}})^2}\,,$$
$$\sigma_{d,raw} = c \times \sigma_{\mathrm{quad}}\,.$$

where a and b determine the scaling from RMS quadrupole movement and RMS BPM error to orbit with the OFB active, and c the scaling from RMS quadrupole movement to orbit with the OFB inactive. The values of a, b and c utilized in the study are tabulated in Table 9.

Table 9: Table of the a, b and c values in (5) and (6) to reproduce the radial beam separation in IP1 and IP5, and the horizontal/vertical separation in IP2/IP8. All values assume the global correction strategy with 40 singular values, barring the use of shared correctors in LHC and allowing them in HL–LHC.

Machine	Scaling factors IP	a	b	c
LHC	IP1	6.85e+05	2.60e+03	9.07e+05
	IP2	9.14e+04	1.43e+03	2.68e+05
	IP5	6.84e+05	2.74e+03	9.01e+05
	IP8	1.61e+05	1.49e+03	4.43e+05
HL–LHC	IP1	8.62e+05	3.92e+03	1.41e+06
	IP2	9.57e+04	1.16e+03	4.21e+05
	IP5	9.97e+05	3.90e+03	1.40e+06
	IP8	2.31e+05	1.81e+03	9.22e+05

By plugging in the values from the table into Eqs. (5) and (6), it is possible to extend the results presented for the performance of the OFB to different assumptions in the RMS quadrupole and BPM errors. This could prove useful if rough estimates of the beam separation during Stable Beams are required.

E Element Naming Conventions

As element names occur in the final results, it is worth to spend a little time on the element naming convention used in LHC and HL-LHC. Figure 33 shows a section of LHC. The

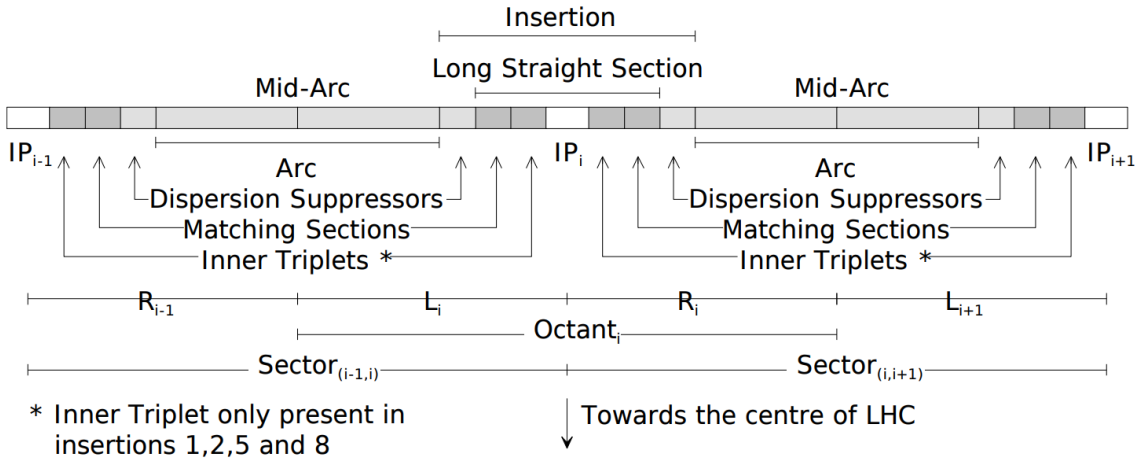

Figure 33: In-depth view of a segment in LHC [27].

key take-aways from Figure 33 for the present purpose are:

1. Each octant is divided into two *half-arcs* surrounding an *insertion*.
2. Each octant is also divided into a left side and a right side.
3. The center-point of some octants is the *interaction point*, `IP` for short, with their surrounding region referred to as *interaction region*, abbreviated `IR`.

From the perspective of lattice definitions, there are eight `IP`s, but this is only for notational ease. An interaction point in the strict sense is a point where the two beams collide, and is only a feature of octant 1, 2, 5 and 8 where experiments are run. When an `IP` or `IR` is referred to in this document, it is taken for granted that it applies to one of these octants. What all octants nevertheless have in common is that they all have a long straight section in the middle as part of the insertion. The arc can be perceived to be roughly uniform across LHC whereas the long sections differ from octant to octant.

As the base pattern is a FODO lattice (alternatingly focusing and defocusing quadrupoles), the machine can be broken up into half-cells containing one quadrupole each. In doing so, each half-cell is given a number, where the i^{th} quadrupole away from the center of its octant is associated with the i^{th} half-cell. With this in mind, the general naming convention can be summarized as follows:

```
<TYPE><SPECIAL>.<EXTRA><HALF_CELL><LR><OCTANT>.B<12>
```

- TYPE: Entry specifying the type of element. See Table 10 for example entries.
- SPECIAL: Optional entry which can be used to specify sub-type of element, e.g. H or V to signify if a corrector is acting on the horizontal or vertical plane.
- EXTRA: Optional entry used to separate between otherwise identically named elements. E.g. A, B, C to separate between three bending magnets in the same half-cell.
- LR: Entry specifying which side of the closest IP the element is on. Assumes either L or R.
- OCTANT: Entry specifying the octant the element is a part of. Valid entries are 1 to 8.
- 12: Entry specifying which beam the element is part of. Either 1 or 2, unless the element is shared between the two beams in which case the element name ends with the OCTANT entry.

Table 10: Table over prefixes for different element types.

Element type	**Prefix**
Bending magnet	MB
Quadrupole	MQ
Orbit corrector	MCB
BPM	BPM
Crab cavity	ACFCA
Drift	DRIFT

For example, the element MQ.25L5.B1 is a quadrupole on the left side of IP5, in the 25^{th} half-cell and for Beam 1. The special identifier can be used in multiple ways, for example MQML.10R1.B1 is a different type of quadrupole in half-cell 10, on the right side of IP1 for beam 1. Here the special identifier describes the type of quadrupole. For MCBH.21R5.B1, the special identifier H signifies that it is a horizontal orbit corrector. Two other horizontal orbit correctors are MCBYH.A4R5.B1 and MCBYH.B4R5.B1. In this case the two horizontal orbit correctors share type MCBY, octant, side of IP, half-cell and beam, which is why they make use of the extra specifiers A and B to tell them apart.

Note that there are elements which skip the appendage of .B<12>. These correspond to elements which are common to both beams, which can only happen in the IR. This is due to the fact that when two beams are brought to collision they pass through the same equipment close to the point of collision. Examples of this are MCBXFAV.3L5, a shared vertical corrector, and MBXF.4L5, a shared bending magnet. This has an important impact on the analysis, since powering a shared corrector impacts both beams, same thing goes for imperfections in shared bending magnet.

References

[1] R. De Maria, "HL–LHC Optics v1.5 Update," in *154th HiLumi WP2 Meeting*, July 2019. https://indico.cern.ch/event/836669/contributions/3507590/.

[2] R. De Maria, "HL–LHC MAD-X optics V1.5 Tag v0.1," January 2020. https://github.com/lhcopt/hllhc15/.

[3] J. Andersson, "A linear framework for orbit correction in the high-luminosity large hadron collider," Master's thesis, Lund University, 12 2019. ISSN: 1404-6342.

[4] M. Fitterer, S. Fartoukh, M. Giovannozzi, and R. De Maria, "Crossing scheme and orbit correction in IR1/5 for HL-LHC," in *Proceedings of the 6th International Particle Accelerator Conference*, (Richmond, VA, USA), p. TUPTY036, May 2015.

[5] D. Gamba and R. D. Maria, "IP Orbit Correction Update for HL-LHC," in *Proceedings of the 9th International Particle Accelerator Conference, IPAC18*, (Vancouver, BC, Canada), p. THPAF039, May. 2018.

[6] E. Todesco *et al.*, "Acceptance Criteria of INFN-Genova in-kind Contribution (Recombination Dipole D2)," Tech. Rep. EDMS n. 2051868 v.0.1, CERN, Geneva, May 2020.

[7] "HL-LHC Lay-out IP5 Drawings - Front View - Optics 1.5," Tech. Rep. EDMS n. 1687089 v.1.5, CERN, Geneva, Dec 2019. NOTE: Requires permission to view.

[8] R. Calaga, "Crab cavity operational aspects," in *25th HiLumi WP2 Task Leader Meeting*, (Geneva), March 2014. https://indico.cern.ch/event/307357/contributions/1675652/.

[9] Jan Goyvaerts, "Regular Expression Tutorial – Learn How to Use Regular Expressions." https://regular-expressions.info.

[10] G. Apollinari, I. Béjar Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, and L. Tavian, "HL–LHC Technical Design Report," Tech. Rep. EDMS n. 1723851 v.0.71, CERN, Geneva, 2016.

[11] R. D. Maria *et al.*, "Can we simplify HL-LHC circuits?," in *LHC Performance Workshop*, January 2017. https://indico.cern.ch/event/580313/contributions/2359640/.

[12] G. R. L. Deniau, H. Grote and F. Schmidt, "MAD-X: User's Reference Manual Version 5.05.02." http://madx.web.cern.ch/madx/, 2019.

[13] R. Bruce *et al.*, "Updated parameters for HL-LHC aperture calculations for proton beams," Tech. Rep. CERN-ACC-2017-0051, CERN, Geneva, Jul 2017.

[14] Jörg Wenninger. Private Communication, July 2019.

[15] D. Gamba *et al.*, "Impact of radiation on orbit correctors and their possible failure," in *162th HiLumi WP2 Meeting*, Nov 2019. https://indico.cern.ch/event/858514/contributions/3638798/.

[16] R. De Maria *et al.*, "High Luminosity LHC optics and layout HLLHCV1.4," in *Proceedings of the 10th International Particle Accelerator Conference, IPAC19*, (Melbourne, Australia), p. MOPMP019, May 2019.

[17] R. Steinhagen, J. Andersson, L. Jensen, R. Jones, and J. Wenninger, "LHC Orbit Stabilisation Tests at the SPS," in *Proceedings of the Particle Accelerator Conference 2005*, (Knoxville, TN, USA, USA), pp. 886 – 888, June 2005.

[18] A. Boccardi *et al.*, "Commissioning and Initial Performance of the LHC Beam-Based Feedback Systems," in *Proceedings of the 1st International Particle Accelerator Conference*, (Kyoto, Japan), p. WEPEB041, May 2010.

[19] L. Grech, D. Alves, S. Jackson, J. Wenninger, and G. Valentino, "Feasibility of hardware acceleration in the lhc orbit feedback controller," in *Proceedings of ICALEPCS2019*, (New York, NY, USA), p. MOPHA151, October 2019.

[20] J. Wenninger, "Operation and Configuration of the LHC in Run 2," Tech. Rep. CERN-ACC-NOTE-2019-0007, CERN, Mar 2019.

[21] D. Gamba, R. Corsini, M. Guinchard, M. Schaumann, and J. Wenninger, "Estimated Impact of Ground Motion on HL-LHC Beam Orbit," in *Proceedings of the 9th International Particle Accelerator Conference*, p. THPAF040, Jun. 2018.

[22] A. Gorzawski and J. Wenninger, "Orbit Drifts at the LHC Interaction Points During the Squeeze," Tech. Rep. CERN-ACC-NOTE-2015-0013, CERN, May 2015.

[23] E. Metral *et al.*, "Update of the HL-LHC operational scenarios for proton operation," Tech. Rep. CERN-ACC-NOTE-2018-0002, CERN, Jan 2018.

[24] S. Fartoukh, "Achromatic telescopic squeezing scheme and application to the lhc and its luminosity upgrade," *Phys. Rev. ST Accel. Beams*, vol. 16, p. 111002, Nov 2013.

[25] B. J. Holzer *et al.*, "Optics Design and Lattice Optimisation for the HL-LHC," Tech. Rep. CERN-ACC-2013-0134, CERN, May 2013.

[26] M. Fitterer and R. De Maria, "BPM Tolerances for HL-LHC Orbit Correction in the Inner Triplet Area," Tech. Rep. CERN-ACC-2015-0176, CERN, May 2015.

[27] R. Saban, "Equipment Naming Conventions in LHC," Tech. Rep. EDMS n. 103369 v.1.0, CERN, Geneva, Sep 1999.