Linear optics ID compensation on the 3 GeV storage ring at MAX IV Laboratory

A report presented for the degree of
Master of Science
by

Johann Schmand

Project duration: 8 months

Supervisor: Francesca Curbis
Co-supervisor: Magnus Sjöström
Department of Physics,
Division of Synchrotron Radiation Research

Abstract

MAX IV Laboratory is a modern synchrotron light source, providing synchrotron radiation used in a wide range of experiments. The production of synchrotron radiation used in experiments typically takes place in insertion devices (IDs), in which an oscillatory motion is induced in traversing electrons, stimulating the emission of synchrotron radiation. IDs consist of a periodic arrangement of magnets, and inevitably have an undesired impact on the electron beam. This report is an account of an attempt to implement and evaluate a first proof-of-concept compensation scheme on the 3 GeV ring for one of the lowest-order perturbations caused by the effects of IDs, namely the transverse beam size, largely governed by the betatron (or beta, β) functions.

The approach to eliminate perturbations in β functions relies on adjusting strengths of beam-focusing quadrupole magnets adjacent to the ID. In this way, the β functions in most of the storage ring can be restored to nominal values. Two different approaches were used to determine the quadrupole strength adjustments required for eliminating β function deviations. The first employs the use of so-called *kickmaps* to model the effects of IDs. It is an entirely theoretical approach, in which quadrupole adjustments are determined through numerical optimization on a model of the magnetic lattice comprising the storage ring. ID compensation using this approach has been studied at MAX IV Laboratory in the past, and has been used to develop and implement the ID compensation scheme on the 1.5 GeV ring.

The second approach to determining quadrupole adjustments was developed as part of this project, and consists of an optimization scheme on the online accelerator that iteratively measures β functions and calculates adjustments, eventually converging to the noise floor of the perturbation measurement. This approach is unconventional, and should be viewed as a proof of concept.

Evaluations of β function compensation were performed on quadrupole adjustments obtained from both methods. In this study, the kickmap-based compensation proved largely unsuccessful in eliminating perturbations in β functions, achieving reduction in root mean square RMS($\Delta\beta/\beta$) of no more than 47% for any one ID. In one of the initial tests, the online algorithm was used to reduce RMS($\Delta\beta/\beta$) by 90% in optimal conditions. The current implementation of the online compensation is known to be limited by BPM timing configuration and noise.

Table of contents

Acknowledgements

First and foremost, I would like to thank my supervisors Francesca and Magnus for their endless patience and support throughout this project. Thank you to Gustavo and Alfred for discussing ID compensation with me and teaching me a great deal on the matter. I want to thank Hamed, Mostafa and the whole ID team at MAX IV Laboratory for providing kickmaps and expert knowledge. A further thank you to the operators at MAX IV: Eric, Lena, Alexandra, Theo, Ali, Per, and Rutger for helping my clueless self perform the measurements I needed during long night shifts.

On a personal note, I want to thank my dear friends for their support and understanding in my time of absence during the writing of this report. Thank you, mom and dad, for being pillars to lean on whenever I need. Lastly, I want to thank my dear badminton friends for their willingness to complement any mental fatigue I may have felt during this project with physical fatigue.

List of acronyms and abbreviations

AT	Accelerator Toolbox (open-source software package developed by accelerator scientists)
BPM	Beam Position Monitor
EPU	Elliptically Polarizing Undulator
FB	Feedback
FF	Feedforward
ID	Insertion Device
IVU	In-vacuum undulator
IVW	In-vacuum wiggler
LOCO	Linear Optics from Closed Orbits (diagnostic tool using measured orbit response matrix)
PFS	Pole-face strip
RMS	Root Mean Square
SVD	Singular Value Decomposition
TBT	Turn-by-turn (accelerator jargon meaning "across" many revolutions in the storage ring)
	Relevant ring magnet denominations
DIP	Ring dipole bending magnets
DIPc	Pole-face strip magnets used to independently adjust quadrupole strengths of bending magnets
QDE	3 GeV-ring-specific horizontally defocusing pure quadrupoles installed at the each end of achromats (Quadrupole Defocusing End)
QF	3 GeV-ring-specific horizontally focusing pure quadrupoles installed in achromats (Quadrupole Focusing)
QFE	3 GeV-ring-specific horizontally focusing pure quadrupoles installed at the each end of achromats (Quadrupole Focusing End)
SQFO	1.5 GeV-ring-specific combined function quadrupole-sextupole installed at the ends of certain achromats. Horizontally focusing quadrupole component (Sextupole Quadrupole FOcusing)

1 Introduction

Synchrotron radiation, due to its ability to probe structures at small scales, has proven to be an asset to material sciences and other fields. At modern synchrotron radiation facilities, particularly electron storage rings, several experiments are provided with light simultaneously through beamlines. The production of synchrotron radiation for most beamlines takes place in insertion devices (IDs), installed in the long straight sections of the accelerator ring structure. IDs consist of periodic arrays of dipole magnets, which induce an oscillatory motion in the beam electrons passing through them, stimulating the emission of synchrotron radiation. The magnetic fields in physical IDs contain imperfections, which induce perturbations in the passing electron beam beyond the desired oscillation. These issues are amplified by the fact that during normal operation many IDs are in use simultaneously. The collective effects of many IDs can be cumulative, leading to even greater perturbations in beam properties. For these and more reasons, a scheme to compensate for the undesired effects of IDs on the electron beam is required for the operational integrity of the accelerator. This report constitutes an account of an attempt to compensate for one of the lowest-order perturbations caused by the effects of IDs, namely the transverse beam size.

The transverse size of an electron beam in a storage ring can be viewed as an envelope, which varies over the magnetic lattice comprising the accelerator. It can be thought of in analogy to rays of light passing through lenses, undergoing repeated focusing and defocusing. In the storage ring, beam focusing is typically achieved by the use of *quadrupole* magnets. The dependence of the beam size on position along the accelerator is characterized through the betatron (or beta, β) functions. The effects of IDs the on beam size can be seen as a focusing error, whereby the ID introduces beam focusing at a point of the magnetic lattice where it is undesired. To compensate for the focusing error caused by the presence of IDs, in this project, adjustments to the strengths of focusing and defocusing quadrupoles adjacent to IDs were found and tested. Specifically, the goal of this compensation was to restore nominal β functions in most of the accelerator in the presence of IDs.

At MAX IV Laboratory, the need for a compensation scheme for perturbations caused by IDs first became clear on the 1.5 GeV storage ring. Following the installation of a new ID, beam instabilities and low injection efficiencies became significantly more pronounced [1]. In response, among several other measures, a compensation for the β functions was implemented. The implementation of the β function compensation on the 1.5 GeV ring acts as the template for the compensation developed and evaluated on the 3 GeV ring in this project.

The 3 GeV ring, owing to its higher operating energy, stores a more stiff beam, which is affected less drastically by the presence of IDs. Despite the beam stiffness, degradation of key beam parameters such as lifetime and injection efficiency depending on ID configuration have been observed [1]. While ID compensation schemes were studied extensively in the design stage of the 3 GeV ring [2, 3], they were only partially implemented. Furthermore, there are 10 IDs currently installed on the 3 GeV ring, with one more having received funding, and studies ongoing for yet another. Moreover, the planned MAX4U lattice upgrade is expected to be in some senses more sensitive to the ID errors. For these reasons, ID compensation is necessary development on the 3 GeV ring. While severely limited in scope to compensation of β functions, the work presented in this report should be viewed as a step along the way.

Two different methods were employed in this project to determine the adjustments to focusing quadrupole strengths required for β function compensation on the 3 GeV ring. The first method is the same that was used to determine quadrupole adjustments for the β function compensation on the 1.5 GeV ring. It uses *kickmaps* to model the effects of IDs and calculates correction terms for quadrupole magnets numerically. These are calculated for several configurations of the ID. To account for all possible ID configurations, interpolation is used between configurations for which adjustments have been calculated. The second approach to determining quadrupole adjustments is more unconventional. A proof-of-concept implementation was developed and evaluated as part of this project. It is an "online" approach, which uses online measurements of β functions and a numerically determined response matrix to iteratively calculate and apply adjustments to quadrupole magnets on the accelerator. This method is still in development, and while initial steps were taken as part of this project to optimize the performance by reducing noise in β function measurements, many aspects pertaining to the online compensation have not yet been investigated. As part of the development process, the online compensation algorithm was tested on both storage rings at MAX IV Laboratory.

This report is divided into chapters as follows: Chapter 2 gives a brief theoretical background to transverse beam dynamics, accelerator magnets, and insertion devices. Chapter 3 describes the methodology behind the compensation schemes used, as well as the measurement techniques used to determine β functions. Chapter 4 begins with an evaluation of the kickmap-based corrections. It then covers some initial results of the online compensation algorithm, including a comparison to the kickmap-based compensation. Chapter 5 concludes this report and aims to give an outlook on how linear optics ID compensation on the 3 GeV ring could be approached moving forward based on the results of this work.

2 Background

Section 2.1 describes linear motion of particles in an electron storage ring, and how magnets commonly employed at storage rings affect both single-particle and multi-particle (or beam) dynamics. Furthermore, a brief account of the lowest-order magnetic perturbations and their effects is provided. Section 2.2 outlines the principles behind Singular Value Decomposition (SVD) and its use cases in this project. Section 2.3 introduces insertion devices and their working principles, followed by a description of two different types of insertion devices, as well as a qualitative argument as to how focusing errors can arise. Finally, a brief account of the MAX IV insertion devices is provided.

2.1 Transverse beam dynamics

This section is largely based on the treatment of transverse beam dynamics in an electron storage ring presented in [4], which describes some core concepts of linear transverse dynamics using simplifying assumptions. For a more comprehensive treatment of the subject, consider, for instance [5].

2.1.1 Coordinate system

A common choice electron storage rings is to define coordinates with respect to a reference orbit. The transverse plane can be described using cartesian coordinates (x, y), which represent transverse displacement orthogonal to the local reference orbit. In addition to the transverse plane, a longitudinal coordinate s can be defined along the reference orbit. The reference orbit is assumed to be in the (x, s)-plane. i.e. $y(s) = 0$ along the reference orbit. In general, the transverse coordinates depend on the longitudinal $(x, y) = (x(s), y(s))$. To describe the direction of particle motion, the derivatives

$$x'(s) = \frac{dx(s)}{ds} = \tan(\alpha_x), \quad y'(s) = \frac{dy(s)}{ds} = \tan(\alpha_y), \tag{2.1}$$

can be used. Here, α_x and α_y denote the angles of the projection of the particle's trajectory on the (x, s)-, (y, s)-planes, respectively. Using these coordinates, the transverse motion of a particle can be described using the position vectors

$$\vec{x}(s) = \begin{pmatrix} x \\ x' \end{pmatrix}, \quad \vec{y}(s) = \begin{pmatrix} y \\ y' \end{pmatrix}. \tag{2.2}$$

The planes spanned by (x, x') and (y, y') are called the horizontal and vertical *trace spaces*, respectively, and will be revisited later in this chapter. Furthermore, while for an individual particle, the coordinates x', y' are known as angular displacements, when discussing beam properties, x', y' are commonly referred to as *divergences*.

2.1.2 Lorentz force

The motion of a charged particle in an electromagnetic field is described by the Lorentz force $\vec{F}$

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}), \tag{2.3}$$

where q denotes the charge of the particle, $\vec{E}$ the electric field, $\vec{v}$ the particle velocity, and $\vec{B}$ the magnetic field. In any high-energy accelerator, to control particle motion in the transverse plane, only magnetic fields are employed. For simplicity's sake, only transverse magnetic fields will be considered, such that $\vec{v} \perp \vec{B}$. Thus, the Lorentz force equation is reduced to

$$\vec{F} = q\vec{v} \times \vec{B}. \tag{2.4}$$

2.1.3 Magnetic multipole expansion

To set up proper equations of motion using eq. 2.4, and to illustrate how different components of magnetic fields have different effects on the beam in a storage ring, it is worth mentioning a few key results of magnetic multipole expansion. The magnetic field can be expanded in the transverse coordinates (x, y)

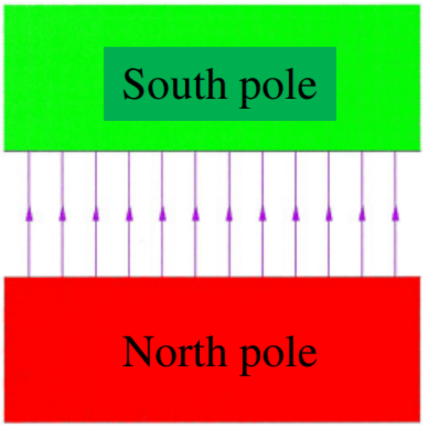

Figure 2.1: Schematic of the field lines produced by an upright dipole magnet. Reproduced from [6].

defined in Section 2.1.1. The result of this expansion, including up to linear terms in x, y, for each of the components B_x, B_y of the magnetic field is [4]:

$$\begin{aligned} B_x(x,y) &= B_0\left(-a_1 + \frac{b_2 y - a_2 x}{r_0}\right), \\ B_y(x,y) &= B_0\left(b_1 + \frac{a_2 y + b_2 x}{r_0}\right), \end{aligned} \tag{2.5}$$

where B_0 and r_0 are free parameters, and the expansion coefficients a_i, b_i can be used to divide the fields further into useful components. The coefficients b_i denote the *upright* expansion coefficients, whose associated magnetic fields have the useful property that the horizontal component vanishes in the (x, s)-plane, i.e. $B_x(x, y = 0) = 0$ in upright fields. In such a field, there is no vertical deflection of particles whose trajectories are in the (x, s)-plane, which is generally a desired property of magnets in a storage ring. The coefficients a_i are called *skew* coefficients. Throughout this work, upright magnetic fields are the main interest, and magnetic fields are assumed to be upright unless otherwise stated. By convention, the expansion coefficients are typically replaced by the *multipole strengths* s_n. The upright multipole strengths are given by:

$$s_n = \frac{q}{p}\left.\frac{\partial^{n-1} B_y(x,y)}{\partial x^{n-1}}\right|_{(x,y)=(0,0)} = \frac{q}{p}\frac{B_0}{r_0^{n-1}} b_n, \tag{2.6}$$

where p denotes the momentum of the reference particle, and the last equality shows the relation to the upright expansion coefficients b_i, and n denotes the order. The factor $1/p$ in equation 2.6 is the reason why the electron beam is said to be more "rigid" at higher energies. Of interest in this work are the two lowest-order terms s_1 and s_2. By convention, these are typically labeled $s_1 = \kappa$, $s_2 = k$, and are called *dipole* and *quadrupole* strengths, respectively. Dipole and quadrupole magnets are paramount to linear transverse dynamics in an electron storage ring, and warrant a brief description in this chapter. It should be noted that higher-order magnets, such as sextupoles and even octupoles are used in many storage rings. These are, however, mainly involved in nonlinear optics and thus outside the scope of this work.

Dipole magnet

The field generated by a dipole magnet is given by

$$\begin{aligned} B_x(x,y) &= 0, \\ B_y(x,y) &= B_0. \end{aligned} \tag{2.7}$$

The constant and homogeneous vertical field in eq. 2.4 leads to a deflection of particles in the horizontal plane at a constant bending radius ρ. Figure 2.1 shows the field lines of typical upright dipole. In a physical dipole magnet, there will be edge effects and other inhomogeneities, which are neglected in this simple treatment. The bending radius ρ of the deflection is related to the field to the field strength B_0 and the dipole strength κ via

$$\frac{1}{\rho} = \frac{q}{p} B_0 = \kappa. \tag{2.8}$$

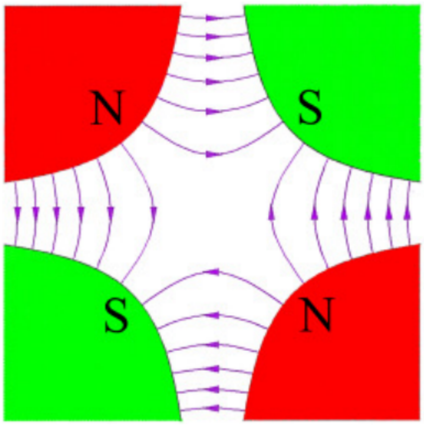

Figure 2.2: Schematic of the magnetic field lines produced by an upright quadrupole magnet. Reproduced from [6].

Quadrupole magnet

The field of an upright quadrupole magnet in terms of the quadruple strength k is given by

$$\begin{aligned} B_x(x,y) &= \frac{p}{q}ky, \\ B_y(x,y) &= \frac{p}{q}kx, \end{aligned} \tag{2.9}$$

where k denotes the quadrupole strength as in eq. 2.6. The field lines are illustrated in Figure 2.2. Equation 2.9 shows that the field is linear in the position displacements (x, y). This turns out to be a very useful property in accelerators. Inserting this field into the Lorentz force equation 2.4:

$$\vec{F} = q\vec{v} \times \vec{B} = pvk(-x\hat{e}_x + y\hat{e}_y), \tag{2.10}$$

where $\hat{e}_x$ and $\hat{e}_y$ denote the unit vectors along the coordinates x, y respectively. Equation 2.10 is an important result, as it shows that the field of a quadrupole causes a restoring force in one plane, depending on the sign of k. For this reason, quadrupoles are employed at accelerator facilities to focus the particle beam. Throughout this work, the sign convention for the quadrupole strength k is such that $k > 0$ implies a *horizontally* focusing quadrupole, i.e.

$$k > 0 \implies F_x \propto -x, \tag{2.11}$$

It should be noted that eq. 2.10 shows that a quadrupole will simultaneously act to focus the beam in one plane, and to defocus it in the other. For this reason, in beam focusing, quadrupole magnets are often installed in series with alternating signs of k along the accelerator. To obtain overall focusing in both planes, the quadrupoles need to have some separation. Such a periodic arrangement of quadrupoles is called a FODO lattice [5].

2.1.4 Equations of motion and betatron oscillations

With the definition of dipole and quadrupole strengths, the single-particle linear equation of motion can be constructed. Owing to the periodic arrangement of magnets in an electron storage ring, and assuming a *monochromatic* beam, meaning all particles have zero momentum deviation, the linear equations of motion can be expressed using Hill's equation [5]:

$$u''(s) + K(s)u(s) = 0, \tag{2.12}$$

where $u \in \{x, y\}$ denotes either of the transverse displacement coordinates. $K(s)$ is called the magnet parameter and is periodic along the magnetic lattice that makes up the accelerator structure:

$$K(s + L) = K(s), \tag{2.13}$$

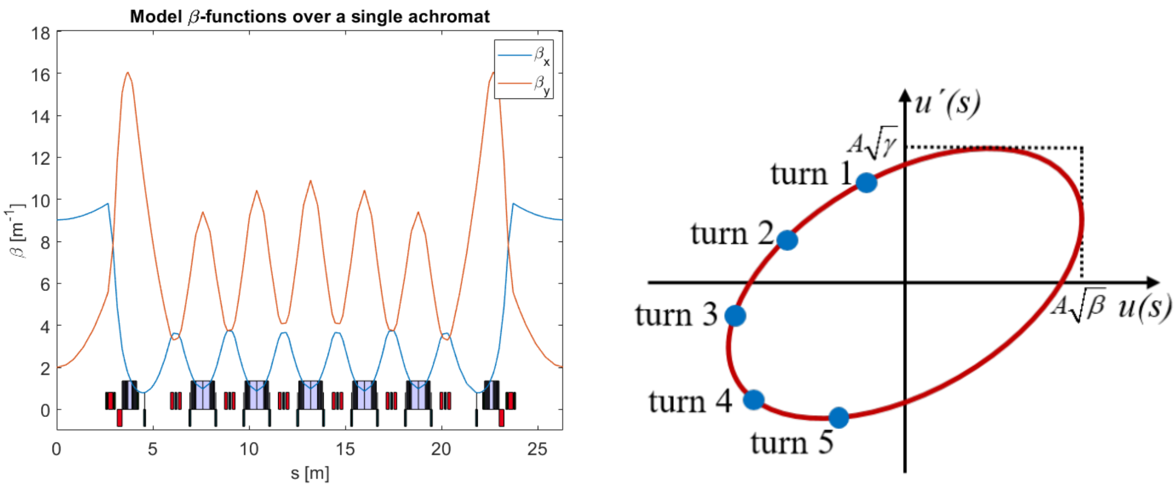

Figure 2.3: Left: β functions plotted over a model achromat. The magnets making up the achromat are indicated along the horizontal axis and are color coded: Red: pure quadrupoles. Black-violet-black: bending dipoles. Thin black lines: sextupoles. Right: Motion of a single particle in trace space at a fixed longitudinal coordinate s over a number of turns. The position of the particle shifts between turns, by a value determined by the tune Q_u. The shape and orientation of the ellipse are determined by the β-functions and derivatives thereof, while the area is determined by the particle's Courant-Snyder invariant A. Reproduced from [4], Figure 34.

where L denotes the length of the periodic structure. In principle, L could be the entire circumference, but owing to storage rings often being divided into a sequence of regular segments called *achromats*, L can commonly be an integer fraction of the circumference. The magnet parameter $K(s)$ takes a slightly different form in the two planes:

$$K(s) = \begin{cases} \dfrac{1}{\rho^2} - k(s), & \text{horizontal plane,} \\ k(s), & \text{vertical plane,} \end{cases} \tag{2.14}$$

with ρ the reference orbit bending radius from upright dipoles and k the quadrupole strength as before. The solution of eq. 2.12 can be written as

$$u(s) = A\sqrt{\beta_u(s)}\cos(\mu_u(s) + \varphi_0), \tag{2.15}$$

where the position-independent amplitude factor A and the initial phase φ_0 are determined by the individual particle. The parameter A^2 is known as the *Courant-Snyder invariant*, and remains constant for a given particle under the assumptions of zero transverse damping, and no excitations. The solution states that transverse particle motion is oscillatory around the reference orbit, and includes two parameters characterizing the oscillation:

- The beta (or betatron) functions $\beta_u(s)$, which are position-dependent functions governing the amplitudes of transverse oscillations,
- The phase advances $\mu_u(s)$, describing how the phase of the transverse oscillation changes as a function of the longitudinal coordinate s. The phase advances relate to the β functions through

$$\mu_u(s) = \int_0^s \frac{1}{\beta_u(\xi)} d\xi. \tag{2.16}$$

The transverse oscillations around the reference orbit are known as *betatron* oscillations, as they were first studied on betatron accelerators [4]. Figure 2.3 shows the β functions over a model achromat of the 3 GeV storage ring at MAX IV Laboratory. As can be seen, each of the β functions appears to oscillate over the length of the achromat. This behavior is a result of the periodic focusing and defocusing caused by the magnets in the achromat. The amplitudes of β_x and β_y appearing out-of-phase can be interpreted in terms of quadrupole focusing, as per eq. 2.10, stating that any one quadrupole is focusing in one plane and defocusing in the other. Consequently, β increases in one plane, and decreases in the other. To achieve beam stability overall, both focusing and defocusing magnets are necessary [4].

The behavior of β functions over one achromat can be generalized. Assuming an idealized storage ring, consisting of M identical achromats forming a closed loop, the period L of the magnet parameter K in eq. 2.14 is determined by the length of the achromat. Consequently, the β functions have an M-fold periodicity over the circumference of the storage ring.

2.1.5 Particle motion in trace space

Particles undergoing stable oscillatory transverse motion travel along ellipses in trace space (u, u'). In this linear treatment of particle dynamics, these planes are decoupled. The shape and orientation of these ellipses change with s, though the area for a given particle is constant, and is determined by the Courant-Snyder invariant, via Area $= \pi A^2$. An important parameter to characterize the motion of a particle along its trace space ellipses is the (betatron) *tune* Q_u, defined as

$$Q_u = \frac{1}{2\pi} \oint \frac{1}{\beta_u(s)} ds. \tag{2.17}$$

The tune Q_u describes the the number of oscillations along the phase space ellipse per turn (revolution) of the particle around the accelerator. The right-hand image in Figure[1] 2.3 illustrates this behavior, showing the trace space ellipse of a single particle at a fixed s over a number of turns. The tune Q_u determines how far along the ellipse the particle travels each turn (revolution). Moreover, the tune is a parameter critical to beam stability, and will be revisited later in this chapter.

2.1.6 Generalization of single-particle findings to beam properties

Thus far, only the single-particle case has been considered. To characterize the properties of a particle beam consisting of many particles, the parameter A in the single-particle solutions in eq. 2.15 can be put through a statistical treatment. Denote individual particles by a subscript i and assume that the average position deviation of particles is zero, i.e. $\langle u_i \rangle = 0$, where the brackets denote the statistical moment. The result of this treatment is that [4]

$$\epsilon = \frac{1}{2} \langle A_i^2 \rangle . \tag{2.18}$$

In equation 2.18, ϵ denotes the *emittance* of the beam in either the horizontal or vertical plane. The emittance is a paramount parameter in a synchrotron light facility, and represents, per eq. 2.18, an average of the amplitude factors A_i of individual particles. This, in turn, translates to a particle-averaged area spanned by the beam as a whole in trace space. Much like the Courant-Snyder invariants A^2, the emittances ϵ are invariant under the assumption of zero transverse damping and no excitations. Furthermore, in analogy with the single-particle transverse oscillations described in eq. 2.15, the amplitude of transverse oscillations $A_i\sqrt{\beta_u}$ can be generalized to

$$\sqrt{\epsilon \beta_u}, \tag{2.19}$$

the physical interpretation of which is the *beam size*, an envelope describing particle deviations from the reference as a function of the longitudinal coordinate. Figure[2] 2.4 illustrates how individual particle trajectories and transverse deviations relate to the beam envelope. It should be noted that in this simple illustration, all particles fall within the envelope defined by $\sqrt{\epsilon\beta}$. In a physical accelerator, this will not be the case, as the emittance is typically defined through a fraction of the trace space region occupied by particles.

2.1.7 Magnet perturbations

This section is dedicated to describing the effects of a small multipole strength perturbation in the two classes of magnets considered in this chapter: dipoles and quadrupoles. While such perturbations are

[1]License agreement
[2]License agreement

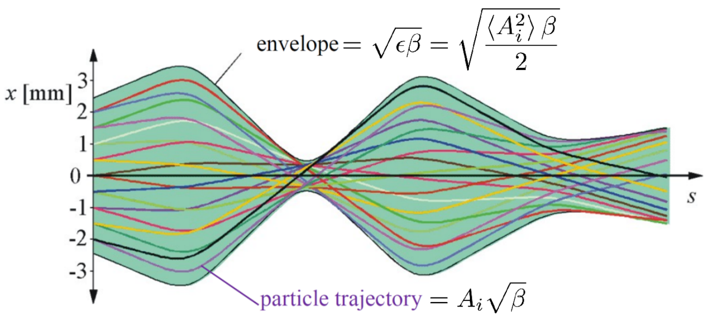

Figure 2.4: An illustration of the horizontal beam size envelope as a function of s. Individual particle trajectories within the envelope, and their oscillation amplitudes are indicated. Reproduced with modifications from [4], Figure 23.

ever-present in physical accelerators, they can also be used for measurements and corrections, as is described in chapter 3. For this reason, a brief overview of the relevant results is required.

Dipole perturbations

In a circular accelerator, a dipole field perturbation, expressed as an angular kick $\delta u'$, induces an oscillatory motion on particles. Let the perturbation be at longitudinal position $s = \xi$. Then, for a single particle, the oscillation can be expressed [5]:

$$u(s) = \sqrt{\beta(s)\beta(\xi)}\delta u' \sin\left(\mu(s) - \mu(\xi)\right), \tag{2.20}$$

where u denotes the transverse position deviation of the particle, β the betatron function, and μ the phase advance. In equation 2.20, it can be seen that the amplitude of the oscillation $u(s)$ depends on both $\beta(s)$ and $\beta(\xi)$, the values of the betatron function at the location of the observer, and the location of the error itself. As is explained in Section 3.4, this fact is exploited in pinger-based measurements of β functions. Commonly, on circular accelerators, dipole field perturbations are evaluated in terms of *closed orbit* distortions. The closed orbit is the closed trajectory on which the *center of charge* of the beam moves. The existence of a closed orbit stems from the periodicity of the magnetic lattice [4]. Introducing small perturbations in dipoles in the form of angular kicks $\delta u'$ at several points along the lattice, an equation describing the deviation in position of the closed orbit $u_c(s)$ can be derived [5]:

$$u_c(s) = \frac{\sqrt{\beta_u(s)}}{2\sin(\pi Q_u)} \oint \delta u'(\xi)\sqrt{\beta(\xi)}\cos(\mu(s) - \mu(\xi) + \pi Q_u)d\xi, \tag{2.21}$$

where the integral shows that all dipole errors $\delta u'$ over one revolution have been accounted for. This equation describes the deviation of closed orbit position after one revolution. An important consequence of eq. 2.21 is that from the fraction outside the integral, a resonance condition arises, namely that $Q_u \neq n,\ n \in \mathbb{Z}$. For integer tunes, the deviations $u_c(s)$ add coherently, leading to diverging betatron oscillation amplitudes. For this reason, circular accelerators cannot be operated at integer tunes.

Dipole perturbations in the form of kicks to the beam can be used for beam diagnostics. By kicking the beam using dipole corrector magnets and measuring the response in terms of closed orbit distortions in beam position monitors (BPMs) as per eq. 2.21, an *orbit response matrix* R_θ can be constructed. It should be noted that measurements of closed orbits require acquisition times significantly longer than the revolution period, to average out higher-order transverse beam motion [7]. Assume a total of N BPMs, indexed by $i = 1, ..., N$, and M dipole correctors, indexed by $j = 1, ..., M$. Denote the measured closed orbit distortion at BPM i by u_c^i, and the angular kick provided by corrector j by θ_j. Then, the orbit response matrix R_θ is given by

$$R_{\theta,ij} = \frac{\partial u_c^i}{\partial \theta_j}. \tag{2.22}$$

The accelerator optics and diagnostics tool *Linear Optics from Closed Orbits* (LOCO) [8] employs measurements of the response matrix. LOCO and its use case in this project is described in more detail in Section 3.3.

Quadrupole perturbations
Reasoning similar to that applied to dipole perturbations can be applied to quadrupole perturbations. Introducing quadrupole perturbations δk at various points of the magnetic lattice, there are two important results. The first is a constant tune shift ΔQ_u [5]:

$$\Delta Q_u = \frac{1}{4\pi} \oint \beta_u(s) \delta k(s) ds. \tag{2.23}$$

An important consideration regarding eq. 2.23 is that the tune shift ΔQ is dependent not only on the strength of the perturbation, but also the value of the β function at the location of the perturbation. The other result of interest is the deviation in β functions $\Delta\beta$ that arise as a consequence of quadrupole perturbations:

$$\Delta\beta_u(s) = \frac{\beta_u(s)}{2\sin(2\pi Q_{0,u})} \oint \delta k(\xi)\beta(\xi)\cos(2(\mu(s) - \mu(\xi) - \pi Q_u), \tag{2.24}$$

where $Q_{0,u}$ denotes the unperturbed tune, and Q_u denotes the perturbed tune, such that $Q_u = Q_{0,u} + \Delta Q_u$ with ΔQ_u as in eq. 2.23. The deviation in β functions $\Delta\beta$ is commonly referred to as *beta beat(ing)*. Similarly to eq. 2.21, beta beating depends on both the value of the β functions at the location considered, and the value of the β functions at the location of the perturbation.

Similarly to the orbit response matrix R_θ, a "β function response matrix" R_β can be constructed. Needless to say, this construction assumes a linear relationship between $\Delta\beta$ and quadrupole strengths. To construct R_β, quadrupole perturbations δk_l (with l indexing quadrupoles) can be introduced by varying the strength of quadrupoles in the magnetic lattice. The resulting perturbed beta functions β_i' can be compared to the unperturbed β_i to find $\Delta\beta_i = \beta_i' - \beta_i$. With δk_l and corresponding $\Delta\beta_i$, the beta function response matrix R_β can be constructed as

$$R_{\beta,il} = \frac{\Delta\beta_i}{\delta k_l}. \tag{2.25}$$

The physical interpretation of a matrix element $R_{\beta,il}$ is "the response (change) in β function measured in BPM i as a result of a quadrupole strength change δk_l in quadrupole magnet l". Within this project, R_β is used to correct beta beating. The mathematical formulation of, and solution to this problem is described in Section 2.2.1, and the process of determining R_β numerically is described in Section 3.5.2.

2.2 Singular Value Decomposition

This section is intended to present Singular Value Decomposition (SVD), and to briefly mention its applications within this project. References to more thorough treatments are provided throughout the text. SVD represents a factorization of any complex-valued $N \times M$ matrix A to:

$$A = USV^\dagger, \tag{2.26}$$

where U is a complex unitary $N \times N$ matrix, V is a complex unitary $M \times M$ matrix, S is an $N \times M$ diagonal rectangular matrix, and the dagger † denotes the conjugate transpose [9]. The diagonal elements of S, $\sigma_i = S_{ii}$ are real-valued and non-negative. They are called singular values, and can be arranged in decreasing order, such that

$$\sigma_1 > \sigma_2 > ... > \sigma_p, \quad p \leq M. \tag{2.27}$$

There are two applications of SVD within this project. All matrices used in this project are real-valued, so the conjugate transpose † can be replaced by the transpose ^{T}. The first application, which is described more practically in Section 2.2.1, pertains to solving a linear equation of the form:

$$\vec{y} = A\vec{x}, \tag{2.28}$$

where $\vec{x}$ is an *unknown* vector of length N, A is an $N \times M$ matrix, and $\vec{y}$ is a *known* vector of length M. In particular, A is not a square matrix. Therefore, the equation cannot be solved using the "normal" inverse A^{-1}. Instead, eq. 2.28 is reformulated to a least-squares problem, where the quantity

$$||A\vec{x} - \vec{y}|| \tag{2.29}$$

is to be minimized. For this type of problem, the Moore-Penrose pseudoinverse (henceforth pseudoinverse) is a useful tool [10]. The pseudoinverse of A, denoted here A^+, can be calculated from component matrices U, S, V in the decomposition in eq. 2.26 via

$$A^+ = VS^+U^T, \tag{2.30}$$

where S^+ denotes the pseudoinverse of S [9]. Owing to the rectangular diagonal structure of S, S^+ can be calculated by replacing the nonzero diagonal elements σ_i by their reciprocal $1/\sigma_i$ and transposing the resulting matrix. A^+ has the property that it minimizes the quantity

$$||AA^+\vec{y} - \vec{y}||. \tag{2.31}$$

Thus, by identifying $A^+\vec{y}$ with the unknown $\vec{x}$ gives the best solution to the initial linear equation 2.28.

The second application of SVD in this project pertains to direct, model-independent measurements of β functions by decomposing BPM data into spatial and temporal components, and is outlined in Section 2.2.2.

2.2.1 Using the Moore-Penrose pseudoinverse to determine quadrupole correction terms

As indicated in Section 2.1, the "β function response matrix" R_β as defined in eq. 2.25 can be used to determine quadrupole adjustments to eliminate β beating induced by the presence of an ID. Mathematically, by assumption, the problem can be formulated as a linear equation:

$$\vec{\beta} = R_\beta \vec{k}, \tag{2.32}$$

where $\vec{\beta}$ denotes a vector consisting of measured β function deviations (a vector containing all $\Delta\beta_i$ over BPMs i), and $\vec{k}$ a vector containing the unknown quadrupole adjustments. This problem is entirely analogous to the one presented in eq. 2.28, so skipping to the solution: the desired $\vec{k}$ can be found via

$$\vec{k} \approx R_\beta^+ \vec{\beta}, \tag{2.33}$$

where R_β^+ denotes the pseudoinverse of R_β. The implementation of the online compensation based on this approach to determining quadrupole adjustments is described in Section 3.5. The process of determining R_β is described in Section 3.5.2.

2.2.2 Measuring β functions using SVD on the beam history matrix

This section is dedicated to outlining the main result of a paper [11], which describes a way to perform model-independent β function measurements using SVD analysis on BPM data. This approach relies on observing the beam over a number of turns P. The beam history matrix B can be constructed via:

$$B_{ip} = b_{ip}, \tag{2.34}$$

where b_{ip} denotes the readout of BPM i during turn p. These measurements (called turn-by-turn, or TBT measurements) are different from measurements of the closed orbit distortions described in Section 2.1.7, as the acquisition times are significantly shorter, and the BPM trigger timings need to be precisely calibrated. The SVD of the beam history matrix can be interpreted as a decomposition into transverse eigenmodes of motion:

$$B = USV^T = \sum_{j\ modes} u_j \sigma_j v_j^T, \tag{2.35}$$

where the matrices $U = [u_1, ..., u_P]$ and $V = [v_1, ..., v_N]$ are comprised of the temporal and spatial eigenvectors characterizing the $j = 1, ..., N$ eigenmodes of transverse motion. The σ_j denote the amplitudes of said eigenmodes. The key assumption here is that transverse motion is dominated by the betatron oscillations described in Section 2.1.4. In the SVD of B, this is reflected by two singular values σ_i being significantly larger than others (the reason for there being two is that betatron oscillations are two-dimensional in phase space). Denoting these singular values $\sigma_\pm$ and their corresponding spatial eigenvectors $v_\pm$, the end result is that the β functions in the BPMs can be calculated:

$$\beta = C\left(\sigma_+^2 v_+^2 + \sigma_-^2 v_-^2\right), \tag{2.36}$$

where C denotes a constant. Equation 2.36 illustrates that the β functions in the BPMs can be determined by means of SVD, up to a constant. The practicalities of this kind of measurement is described in Section 3.4.1.

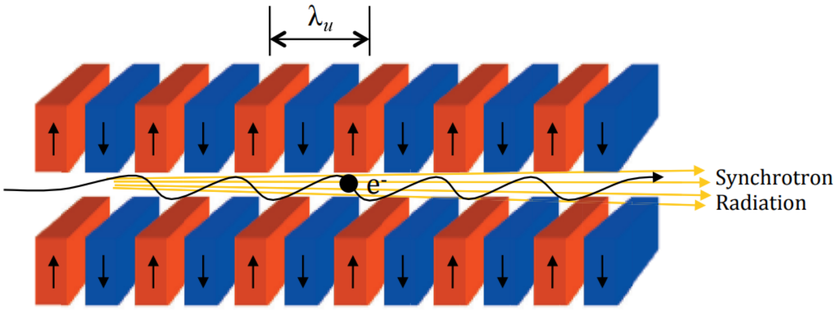

Figure 2.5: A schematic of the simplest kind of insertion device: a planar device, consisting of two parallel girders of dipoles with alternating polarities. The spatial period of magnets λ_u has been indicated. The oscillatory motion of an electron in the device stimulates the emission of synchrotron radiation. Reproduced from [3], Figure 5.1.

2.3 Insertion devices

Synchrotron radiation is the radiation emitted by highly relativistic charged particles undergoing acceleration perpendicular to the direction of motion [5]. Synchrotron radiation produced at modern sources can have useful properties, such as high collimation, high intensity, broad energy range, and tunable polarization [12]. Many modern electron storage rings employ the use of insertion devices (IDs) to produce synchrotron radiation. IDs usually consist of a periodic array of dipole magnets with a period λ_u. When the beam passes through the magnetic array, electrons are forced into oscillatory motion. The strong acceleration experienced by electrons in this motion leads to the emission of synchrotron radiation.

In electron storage rings, insertion devices are typically installed in long straight sections between achromats. This is to allow the device to span several tens of periods λ_u in space. In this work, the main interest in IDs lies in the magnetic geometry, and how the electron beam is affected by the forces applied in an ID. Despite its limitation in scope, this section aims to show that IDs do have a significant impact on the linear optics of the beam, and that these should be accounted for to restore the desired beam properties.

2.3.1 Types of insertion devices

Planar insertion devices

The simplest kind of insertion device is a *planar* ID, which consists of two girders of dipoles, arranged such that vertically opposite pairs of dipoles have alternating polarity along the axis of the ID. A schematic of such a device is provided in Figure 2.5. The girders are installed vertically above and below the (x, s)-plane, allowing the beam to pass through the gap. In the simplest approximation, the dipole fields are homogeneous and constant in the vertical, and particle motion is restricted to the (x, s)-plane. As per eq. 2.8, a particle traversing these alternating dipole fields will be deflected with a constant bending radius, whose curvature depends on the polarity of the dipole field. The net result of a particle traversing a series of dipole fields of opposite polarity is an oscillatory (or *undulatory*) motion. This motion resulted in the term *undulator* or *wiggler* to describe insertion devices of this kind. The strength of an undulator/wiggler is typically characterized through the undulator parameter K, given by [13]

$$K = \frac{eB_0\lambda_u}{2\pi\gamma m_0 c}, \tag{2.37}$$

where e denotes the elementary charge, B_0 the strength of the magnetic field, λ_u the spatial period of magnets, $\gamma = \sqrt{1 - v^2/c^2}$ the Lorentz factor, m_0 the electron rest mass, and c the speed of light.

Using K, a further distinction between IDs can be made. IDs with $K < 1$ are typically called *undulators*. In undulators, synchrotron radiation is emitted in a narrow cone, and radiation emitted in different undulator periods interferes, leading to a spectrum with coherent spectral lines [13]. IDs with $K > 1$ are typically called *wigglers*, and are conversely characterized by a larger cone of emitted synchrotron radiation and broader spectrum. It should be noted that in physical IDs, the distinction is oftentimes

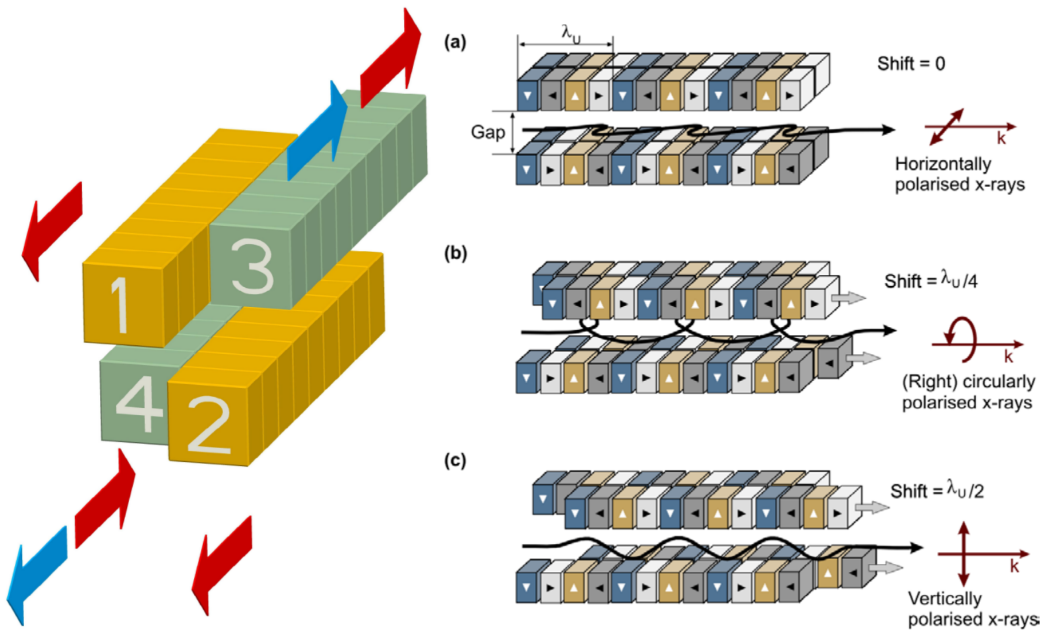

Figure 2.6: Left: Operation modes of an APPLE-II undulator: In the *elliptical* or *helical* mode, girders (1,2) are moved in one direction and girders (3,4) in the opposite, represented by the red arrows. In the *inclined* mode, girder 3 is moved in one direction, and girder 4 in the opposite, represented by the blue arrows. In the *universal* mode, the movements of the other two modes are combined. Reproduced from [15], Figure 1. Right: Different phases of the APPLE-II undulator in elliptical mode. a) Zero phase shift, giving horizontally polarized light, b) Phase shift of $\lambda_u/4$, giving rise to circularly polarized light, c) Phase shift of $\lambda_u/2$, giving rise to vertically polarized radiation. Reproduced from [16], Figure 8.

blurry, and IDs can display both undulator and wiggler properties.

Generally, wigglers have a stronger magnetic field B_0 than undulators, which bears significance within this project, as the strength of the magnetic field in an ID has a direct impact on beam optics. A typical planar ID has one adjustable parameter, namely the *gap*, defined as the vertical distance between the magnetic pole faces of the girders. The gap is an important parameter in the characterization of an ID, as the magnetic field strength B_0 in eq. 2.37 generally has a non-linear dependence on the gap. Thus, the beam may experience drastically different forces depending on the set gap. The field B_0 becomes stronger with decreasing gap. For this reason, the focus of this work is largely on IDs at their respective minimum-gap configurations.

Moreover, many planar IDs can be configured such that the gap at one end of the girders can be different to the other, effectively tilting each of the girders with respect to the horizontal plane. This parameter characterizing the gap difference at the ends of the ID is called *tapering*. Thus, a given planar ID (with given magnetic structure) is characterized by the adjustable gap and, if applicable, tapering. At the 3 GeV ring at MAX IV Laboratory, the magnet girders of all planar IDs are mounted inside the vacuum chamber containing the beam, giving rise to the terms in-vacuum undulator/wiggler (IVU/IVW).

Elliptically Polarizing Undulators

Elliptically Polarizing Undulators (EPUs) are insertion devices designed to be able to produce synchrotron radiation with controllable polarization. This can be achieved using a more complex magnet structure with more controllable parameters than that of planar IDs. Figure 2.6 shows a schematic of the APPLE-II type EPU [14], which represents the basic geometry of all EPUs at MAX IV Laboratory. The APPLE-II IDs at MAX IV Laboratory are installed outside the vacuum chamber. It consists of four girders of dipoles. Each of the girders is comprised of dipoles with different polarity, with an undulator period spanning four individual dipoles. The APPLE-II has three fundamental operating modes, defined by how girders are offset along the beam axis relative to each other. These modes (helical, inclined, universal), and how they correspond to girder movements, is shown in the left plot of Figure 2.6. The magnetic configuration of an APPLE-II EPU generally needs four parameters to be specified: the gap, the helical phase, the incline phase, and the mode. In helical mode, the incline phase is not used, and vice versa.

2.3.2 Beam focusing in insertion devices

This section is dedicated to establishing a qualitative understanding of the focusing effects of an insertion device on the beam in an electron storage ring within the framework of linear optics. To do so, consider the magnetic fields in a planar ID. Owing to the periodicity of the magnetic lattice, the magnetic field can be expanded in a Fourier series in terms of "wavenumbers" $\tilde{k} = 2\pi n/\lambda_u$, where λ_u denotes the magnetic period as before and n denotes the order of the expansion [12]. To first order, the field components along the x, y, s coordinate axes are:

$$\begin{aligned}
B_x(x,y,s) &= \frac{\tilde{k}_x}{\tilde{k}} B_0 \sinh(\tilde{k}_x x) \sinh(\tilde{k}_y y) \cos(\tilde{k}s), \\
B_y(x,y,s) &= B_0 \cosh(\tilde{k}_x x) \cosh(\tilde{k}_y y) \cos(\tilde{k}s), \\
B_s(x,y,s) &= -\frac{\tilde{k}}{\tilde{k}_y} \cosh(\tilde{k}_x x) \sinh(\tilde{k}_y y) \sin(\tilde{k}s),
\end{aligned} \tag{2.38}$$

where $\tilde{k}_{x,y}$ describe the field variation in the horizontal and vertical directions, respectively, and are subject to the condition $\tilde{k}_x^2 + \tilde{k}_y^2 = \tilde{k}^2$. From the fields in eq. 2.38, linear equations of motion may be derived [17]:

$$\begin{aligned}
x'' &= \frac{1}{2\rho^2} \frac{\tilde{k}_x^2}{\tilde{k}^2} x, \\
y'' &= -\frac{1}{2\rho^2} \frac{\tilde{k}_y^2}{\tilde{k}^2} y,
\end{aligned} \tag{2.39}$$

with $\rho = qB_0/p$ the dipole bending radius as before. In a typical planar insertion device, $\tilde{k}_x > 0$ and $\tilde{k}_y < 0$. In analogy to the Lorentz force equation for the upright quadrupole (eq 2.10), these values of $\tilde{k}_{x,y}$ imply that a typical planar ID acts *focusing* in the vertical plane, and *defocusing* in the horizontal. The vertical focusing characteristic of planar IDs is caused by electrons passing the longitudinal field B_s at an angle [17]. These field lines are particularly strong between the alternating poles of adjacent dipole magnets. Furthermore, as can be seen in eq. 2.38, the field B_s increases in magnitude for larger vertical deviations, representing that field lines (and hence focusing) become stronger closer to the poles. The horizontal defocusing is caused by a transverse field roll-off that varies with horizontal position along the ID axis [3]. The horizontal defocusing is typically much less pronounced than the vertical focusing.

The focusing error present in IDs can (in linear terms) be considered much like the quadrupole errors described in Section 2.1.7. A focusing error caused by an insertion device will induce β beating throughout the magnetic lattice, and cause a tune shift ΔQ. To account for the focusing errors caused by insertion devices, other quadrupole magnets, preferably close to the ID along the beam axis, can be used to restore beam optics [2]. Moreover, IDs cause closed orbit distortions (eq. 2.21), which can be corrected using corrector dipole magnets mounted close to the exits of the ID along the beam axis. Another effect of IDs is breaking the periodicity of the magnetic lattice. As mentioned in Section 2.1.4, in an ideal storage ring, β functions are periodic in the achromats. The focusing errors of IDs break that symmetry, which can have adverse effects on machine performance [2].

An important note is that the fields described in eq. 2.38 are an approximation of the magnetic field in planar IDs, which have the simplest possible geometry. The discussion above does not carry much weight for more complex magnet arrangements such as those that can be achieved in EPUs (Figure 2.6 for reference). As is discussed in later chapters, these IDs can have a more pronounced impact in the horizontal than the vertical. Furthermore, there are nonlinear effects pertaining to ID focusing that are neglected in this treatment, such as amplitude-dependent tune shift. ID magnetic fields can have skew components, which couples motion in the transverse plane [17].

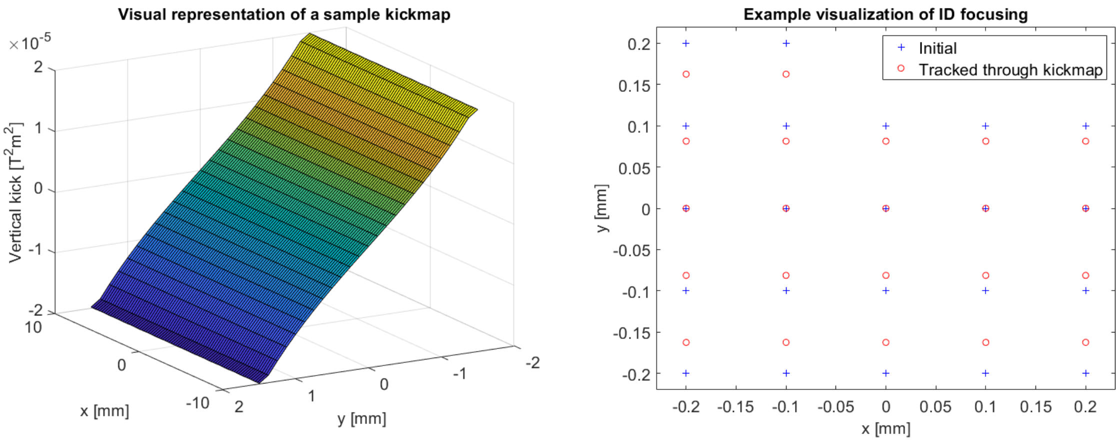

Figure 2.7: Left: a visualization of a sample kickmap representing the effects of planar ID in the vertical plane. Right: illustration of ID focusing using the kickmap in the left-hand plot in a simple tracking simulation. Note: for clarity, particles were tracked through a long drift section after the kickmap. The focusing is thus exaggerated compared to a physical system.

2.3.3 Modeling the dynamics of insertion devices using kickmaps

Due to the complexity of the magnetic fields in IDs, models are commonly used to describe the effects IDs have on optics. Kickmaps are a method of modeling the effects of IDs using dipole kicks [18]. In this work, kickmaps are used in simulation to represent the focusing error caused by IDs A kickmap is comprised of a transverse grid of points, each of which has an associated dipole kick. Each kick, at a point in (x, y) in the transverse plane, models the transverse kick an electron would experience when traversing the ID at (x, y). The dipole kicks are defined in one plane, thus a complete ID model representing transverse focusing requires two kickmaps, representing kicks in either plane. It should be noted that one pair of kickmaps (horizontal and vertical) represent a particular ID at a particular configuration in terms of the parameters gap, mode, helical and inclined phase described in Section 2.3.1. The left-hand plot in Figure 2.7 shows a visualization of a sample kickmap representing the effects of a planar insertion device in the vertical plane. The vertical dipole kicks in Figure 2.7 are largely independent of the horizontal displacement x, and are close to linear in y, which is in line with the expectations as discussed in Section 2.3.2.

Kickmaps are particularly useful in *tracking* studies, whereby a simulated beam is tracked through a model of the magnetic lattice of an accelerator. A simple tracking simulation is presented in the right-hand plot of Figure 2.7, where a "grid" of particles, regularly spaced in the transverse plane, has been tracked through the kickmap model in the left-hand plot followed by a drift section. The results of this experiment illustrate the vertical focusing expected in planar insertion devices. Vertical particle positions have approached the horizontal axis, with focusing being stronger for particles with larger initial displacements. The effects of tracking in the horizontal are barely appreciable. It should be noted that the horizontal kickmap was included in this simulation. However, horizontal kicks are around three orders of magnitude smaller than vertical kicks.

It should be emphasized that the applicability of kickmap modeling is not limited to planar IDs. EPUs can be modeled in much the same way. The main difference is in the "profile" of the kickmaps. While the vertical kicks representing a planar ID might form a plane with a constant inclination and a one-coordinate dependence, the kickmap representation of an EPU can form a highly irregular and more complex surface. Modeling EPUs using kickmaps has proven to be fruitful in previous studies regarding ID compensation at MAX IV Laboratory [18, 2]. Moreover, while not utilized in this project, kickmaps can be used to model nonlinear effects of IDs. The focusing error can be eliminated by applying corrections to adjacent quadrupoles, described in detail in Section 3.2. The kickmaps used in this work were calculated by the ID team at MAX IV Laboratory, using the RADIA software [19].

2.3.4 Overview of the insertion devices at MAX IV Laboratory

Figure 2.8 shows the beamlines with installed insertion devices on the two storage rings at MAX IV Laboratory. Table 2.1 indicates which type of insertion device (planar, EPU) is used at teach of the beamlines. The IDs on the 3 GeV ring are the main interest in this work. While all IDs have an effect on optics, the ones known to have the most significant are the IVW at the BALDER beamline, as well as the EPUs at the HIPPIE, SoftiMAX, and VERITAS beamlines.

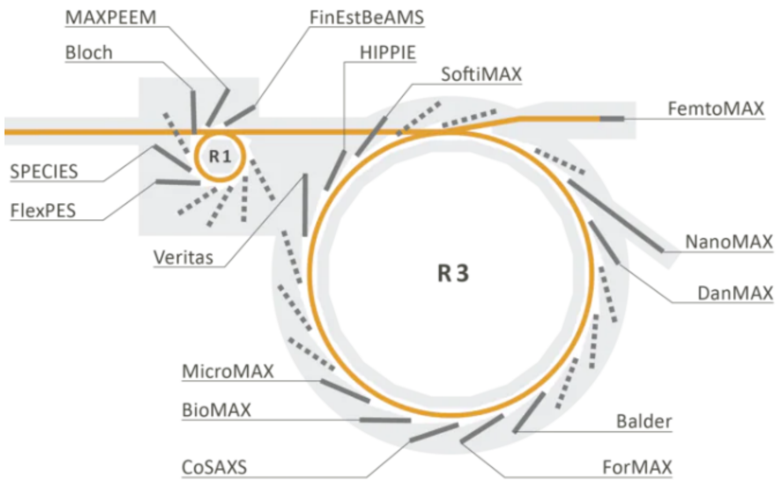

Figure 2.8: Schematic layout of the storage rings at MAX IV Laboratory. R1 and R3 denote the 1.5 and 3 GeV storage rings, respectively. The orange solid straight line near the top represents the linear accelerator. Beamlines with IDs mounted are represented by grey solid lines. Empty long straight sections are represented by dotted lines. The FemtoMAX beamline is connected to the linear accelerator and has been omitted from Table 2.1. Retrieved from [20].

Beamline name	ID type
3 GeV ring (R3)	
BALDER	IVW
BioMAX	IVU
CoSAXS	IVU
DanMAX	IVU
ForMAX	IVU
HIPPIE	EPU
MicroMAX	IVU
NanoMAX	IVU
SoftiMAX	EPU
VERITAS	EPU
1.5 GeV ring (R1)	
BLOCH	EPU
FinEstBeAMS	EPU
FlexPES	Planar
MAXPEEM	EPU
SPECIES	EPU

Table 2.1: A table of the MAX IV beamlines, indicating which type of ID is used. Legend: IVW: In-vacuum wiggler, IVU: In-vacuum undulator, EPU: Elliptically polarizing undulator. The planar ID at the FlexPES beamline is mounted outside the vacuum chamber.

3 Methods

This chapter is divided as follows: Section 3.1 introduces feedforward control schemes and beam position monitors, Section 3.2 explains the approach to linear optics ID compensation taken in this project. Section 3.3 is dedicated to the accelerator diagnostic tool *Linear Optics from Closed Orbits* and an evaluation of its performance pertaining to this project. Section 3.4 explains the working principles of the model-independent pinger-based β function measurements, and outlines an initial attempt at optimizing the measurement. The final Section 3.5 explains the working principles of the "online" correction algorithm to determine quadrupole adjustments.

3.1 Overview of central concepts

3.1.1 Feedforward control schemes

A feedforward (FF) control system is a protocol which seeks to eliminate deviations in a set of parameters A by applying a *known* adjustment to parameters B to eliminate the deviation [21]. To achieve this, a FF requires a precise model how a given deviation in A affects parameters B. This process is different to that of a feedback (FB), which relies on direct measurement of the deviations in parameters A to *calculate* a correction for parameters B. One of the fundamental advantages of a FF compared to a FB is speed. The FF process of receiving sensor values, associating them with an adjustment, and applying the adjustment is generally faster than a FB involving calculation. The speed provided by an FF scheme can constitute a significant advantage in systems with rapidly changing disturbances. On the other hand, the requirements on the internal FF model are stringent. Thus, the implementation of an FF control scheme requires a good understanding of the underlying processes.

In this project, the main use case for FF schemes is to eliminate β beating induced by the effects of IDs, by applying adjustments to flanking quadrupole strengths. The input sensor values to the FFs are the parameters used to describe ID configurations (gap, phase, mode) as described in Section 2.3.1. The output is the associated quadrupole adjustments Δk. The implementation of this compensation scheme, and the determination of adjustments Δk are described in more detail in Section 3.2.1.

3.1.2 Beam Position Monitors

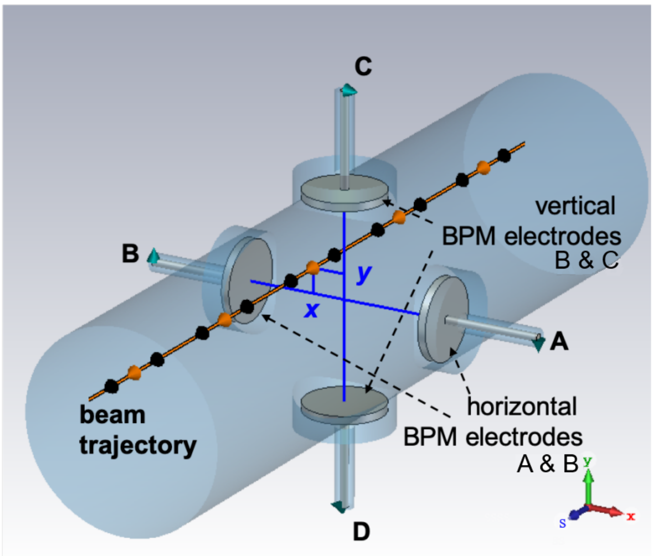

Figure 3.1: Schematic of a button BPM pickup. Pairs of electrodes are arranged opposite another in each of the transverse planes. The beam is indicated, showing its displacement (x, y) with respect to the reference orbit. BPMs are commonly installed such that the reference orbit lies at the intersection of normal lines drawn between opposite electrodes. Reproduced from [22], Figure 2.

This section is dedicated to a brief description of the working principles of Beam Position Monitors (BPMs). For a more comprehensive description, refer to, for instance [22, 7]. BPMs are beam diagnostic tools commonly employed at accelerator facilities. They provide non-destructive measurements of transverse beam position, defined as the displacement of the center of charge of the beam with respect to a reference (commonly the center of the vacuum chamber). Figure 3.1 illustrates the geometry of a button-type BPM pickup, consisting of four electrodes, arranged in opposing pairs facing the respective transverse axes. The BPMs installed on the 3 GeV ring at MAX IV Laboratory follow this four-electrode cross arrangement. However, with respect to Figure 3.1, the BPMs on the 3 GeV ring installed such that the axis between a pair of electrodes form an angle with the transverse axes, to avoid synchrotron radiation from interfering with horizontal electrodes [22].

To measure beam displacements, BPMs rely on image currents induced in the electrodes by the passing beam [22]. As the charged particle beam passes through the BPM, an image charge is induced on the electrodes. For a beam with position a displacement of $+x_0$, the amplitude of the signal inferred from the BPM at the *positive* end of the x-axis is greater than on the opposite end. Thus, in simple terms, using the relative amplitudes of signals inferred from opposing BPMs, the beam displacement x can be inferred. Measurements of vertical beam displacements are entirely analogous to horizontal measurements.

In a physical system, BPM measurements are always subject to noise. The amount of charge in the beam directly influences the amount of charge induced on electrodes, and thus the signal. For this reason, increasing the beam current is a surefire way to increase the signal-to-noise ratio in BPM measurements. Another important quantity pertaining to BPM measurements is the *position sensitivity* $S_{x,y}$. It defines the proportionality constant between the beam displacement, and the measured voltage difference between opposite electrodes [7]. In general, the position sensitivity $S > 0$, meaning that the signal difference increases in magnitude as a function of beam displacement. A consequence of this is that the signal-to-noise ratio is improved in case the beam has a large displacement. This property of BPMs is exploited in the pinger-based β function measurements, described in Section 3.4.

3.2 Working principles of ID compensation

This subsection is dedicated to a description of the general approach to ID compensation employed throughout this project. The main focus of this project is on the *local* compensation described in Subsection 3.2.1, whose objective is to eliminate β beating around the ring caused by an ID by finding correction terms applicable to flanking magnets. The second Subsection, 3.2.2 is dedicated to eliminating the tune shift caused by the combined effects of the ID and local compensation in an effort to restore the working point of the machine.

3.2.1 Local ID compensation to eliminate β beating

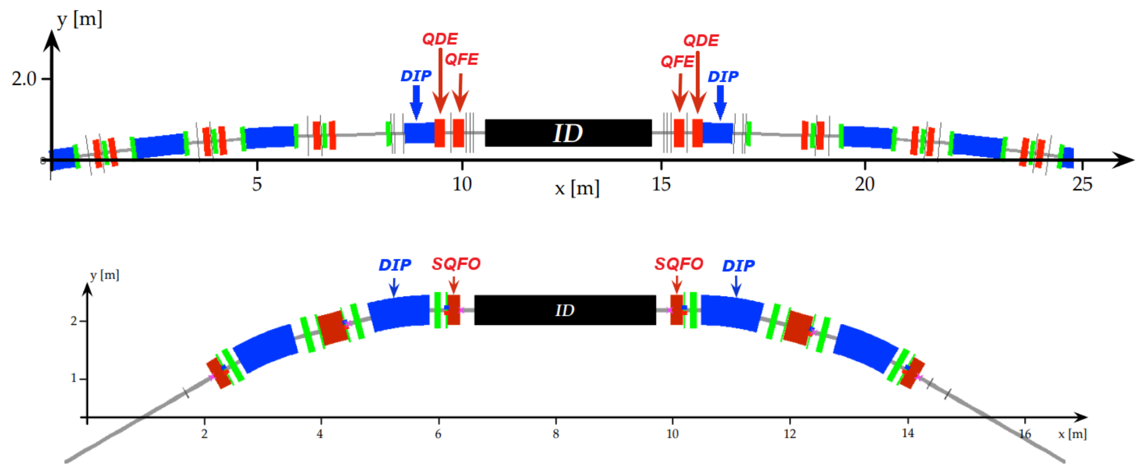

Figure 3.2: Schematics of IDs installed in long straight sections between achromats. Top: 3 GeV ring. Bottom: 1.5 GeV ring. The magnets are labeled by type: Blue: bending dipoles. Red: quadrupoles. Green: sextupoles. The magnets involved in local compensation of the central IDs have been labeled according to their family name. View of the horizontal plane. Reproduced with modifications from [2], Figures 1 & 2.

Overview of local ID compensation on the storage rings

The local compensation aims to make the ID transparent to the beam. Specifically, it aims to eliminate the focusing error caused by the ID by adjusting the quadrupole strengths of flanking magnets. The involvement of only adjacent magnets is why this compensation is referred to as *local*. In this way, the β functions can be restored throughout most of the ring, except in the region spanned by the outermost magnets involved in the compensation. Within the framework of this project, correction terms Δk, physically corresponding to changes in flanking quadrupole strengths, were found, and their effectiveness in restoring β functions throughout the rest of the ring evaluated.

To provide sufficient degrees of freedom to the compensation, two pairs of magnets with one magnet on either side of the ID are required. Furthermore, one pair should belong to a horizontally *focusing* family of magnets, whereas the other should belong to a horizontally *defocusing* family. The magnets involved in the compensation differ both between rings and IDs. Figure 3.2 shows schematics of IDs installed in long straight sections between achromats at either storage ring with the magnets involved in local compensation labeled. In the case of the 3 GeV ring, the magnet families involved in local compensation are QFE, which are horizontally focusing pure quadrupoles, and QDE, which are defocusing pure quadrupoles. For the 1.5 GeV ring, the magnet families involved are SQFO, which are horizontally focusing combined quadrupole-sextupoles, and DIP, which are combined-function dipole/quadrupole magnets, whose quadrupole field strength can be independently adjusted using pole-face strips (labeled DIPc).

A fundamental limitation of the local compensation is that the β functions inside and outside the region spanned by the outermost magnets (between QDE/DIP pairs on the 3 GeV and 1.5 GeV rings, respectively) cannot be restored simultaneously. Thus, the restoration of β functions in this region is forfeit when performing the compensation. In general, there are two solutions in terms of quadrupole strengths to the focusing error caused by the ID, which affect the beam inside the QDE/DIP region differently. One solution is an overall decrease in quadrupole strength. This solution is called under-focusing, as it results in an increased beam size in the ID. The other solution employs increased quadrupole strengths, called over-focusing, which results in a decreased beam size. The preferred solution throughout this project is the over-focusing, owing to a smaller beam size inside the ID itself being generally favorable to users at the beamlines.

Finding quadrupole strength adjustments

As mentioned in Section 2.3, planar IDs can be set to different *gaps*, and EPUs have even more configurable parameters, such as *phase* and *mode*. In terms of local ID compensation, each individual ID configuration corresponds to a different focusing error, meaning that the correction terms required to eliminate β beating are different for every ID configuration. This is an important consideration for the implementation of the local compensation. For the local compensation of a particular ID to be considered complete, correction terms for a number of different configurations should be found, particularly at small gaps, where the effects on optics are strongest.

Within the framework of this project, correction terms for the IDs on the 3 GeV ring were found in two different ways. The first was a numerical approach, relying on kickmaps to model the effects of IDs at different configurations, as described in Section 2.3.3. In this simulation, the IDs, at a given configuration, are represented by a single kickmap inserted into the model lattice of the storage ring at the appropriate longitudinal coordinate. The changes to β-functions caused by the presence of the ID kickmap are minimized using the Nelder-Mead simplex algorithm as described in [23]. As its optimization variables, the optimizer takes the gradient strengths of the flanking QDE/QFE quadrupoles. The goodness-of-fit is chosen to be the square sum of beta beats in both planes $\sum_n (\Delta\beta_{x,n}/\beta_{x,n})^2 + (\Delta\beta_{y,n}/\beta_{y,n})^2$ over every reference point n in the model lattice except the region between the outer flanking pair of quadrupoles (QDEs). This square sum is then normalized by the number of reference points in the model lattice.

This approach, using kickmaps to model the effects of ID on machine optics and subsequently finding correction terms through numerical optimization was taken in the implementation of local ID compensation for the IDs on the 1.5 GeV storage ring. This implementation has proven to be a success operationally, and so the same approach is applied in this project, which aims to perform the same kind of compensation for the IDs on the 3 GeV ring. It should be noted that this approach to determining quadrupole adjustments is entirely theoretical.

The second approach to determine correction terms is more unconventional, and was developed as a part of this project. It relies on model-independent β function measurements performed using pinger magnets. The working principles of these measurements are discussed in greater detail in section 3.4. These pinger-based β function measurements are still in development, and as such a characterization of the expected noise and a simple optimization of measurement parameters was carried out of this project. Using these measurements and a numerically determined "β function response matrix" R_β, correction terms are calculated and applied iteratively in an effort to gradually reduce β beating, eventually converging at a noise floor. The desired correction terms are taken to be the difference in quadrupole strength between the final and first iteration. As this approach represents, in essence, a feedback system on the machine itself, it was coined "online ID compensation algorithm" within the framework of this project. The algorithm is described in more detail in Section 3.5.

Implementation of local ID compensations on the accelerator

The implementation of the local compensation on the machine itself comes in the form of feedforward devices, configured for each ID individually. Each feedforward device has an associated FF table, which lists the determined correction terms given a particular ID configuration. Throughout this project, a typical FF table consists of 7-10 ID configurations and their associated correction terms for planar IDs, and around 60-80 configurations for EPUs. Table 3.1 shows the correction terms found using kickmaps for different configurations of the ForMAX ID. Planar devices generally have the simplest FF tables, since they only have one configurable parameter, namely the gap. EPUs have more configurable parameters. Within the framework of this project, however, only two EPU parameters were taken into consideration:

gap and helical phase.

Table 3.1: Feedforward table for the local ID compensation at the ForMAX beamline on the 3 GeV ring. The correction terms Δk were determined numerically using kickmaps to model the effects of the ID. The sign convention used in the numerical model of the ring is such that $k > 0$ corresponds to a horizontally focusing quadrupole. Hence, the correction terms $\Delta k_{QFE} > 0$ for the focusing QFE family, and $\Delta k_{QDE} < 0$ for the defocusing QDE family confirms that the corrections represent the desired *over-focusing* solution.

Gap [mm]	Δk_{QFE} [10^{-2}m^{-2}]	Δk_{QDE} [10^{-2}m^{-2}]
4.5	0.421	-1.175
5.5	0.280	-0.772
6.5	0.185	-0.516
8.0	0.102	-0.283
10	0.050	-0.133
12	0.025	-0.065
15	0.007	-0.023

EPU configurations being defined by more than one parameter is the reason why their associated FF tables are many times larger than those of planar IDs. Each combination of gap and phase has to be evaluated separately, and so for a desired granularity of 10 gaps and 8 phases, the FF table will consist of 80 rows. As can be seen from Table 3.1, the gaps listed do not cover every possible ID configuration. This is a property common to all determined FF tables, where a compromise has to be struck regarding the granularity of ID configurations covered. For ID configurations between the determined correction terms, linear interpolation is used.

Another feature of Table 3.1 is that only two correction terms are listed, one for each family. In general, the magnets in a QFE/QDE pair either side of a long straight are series-connected to the same power supply. Thus, any correction is applied to both magnets in a pair simultaneously. On the 3 GeV ring, there is one important exception to this scheme, which is the QDE/QFE pairs mounted either side of the IVW at the BALDER beamline. This ID is not centered around the center of the long straight it is mounted in, but has its center shifted downstream by a few meters. The QFE and QDE magnets flanking this ID were provided with individual power supplies with local compensation in mind. Due to the asymmetrical longitudinal position of the ID, the additional degrees of freedom provided by individual power supplies are necessary to be able to perform local compensation. Conceptually, in the process of determining correction terms, this only changes the number of variables in the numerical optimizer; the FF table for the IVW at the BALDER beamline has four columns of correction terms instead of two, one for each magnet involved.

It should be mentioned here that the situation is a little different on the 1.5 GeV ring. Consider the schematic provided in Figure 3.2. Each of the SQFO pairs either side of a long straight are series-connected, whereas each DIP magnet has an individual power supply. Thus, local ID compensation on the 1.5 GeV ring, for each ID, necessitates three correction terms. On the other hand, as each ID on this ring is centered on the center of their respective long straights, the correction terms found for the two DIP magnets are very similar.

As a final note on the implementation of local ID compensations on the storage rings, it would be prudent to describe why the desired implementation is in the form of a feedforward system, rather than a feedback system. The main reason is entirely practical; the available β function measurements are either entirely incompatible with normal operating conditions (as is the case for pinger-based measurements, see Section 3.4), or disturb the beam too much to be used during normal operation (as is the case for LOCO measurements, see Section 3.3). Even if these issues could be surmounted, speed is a point in favor of a feedforward system; in their current implementation, the feedforward devices are able to write correction terms to relevant actuators at a frequency of 10 Hz. While such a high frequency is redundant for IDs which are not changing gaps or phases, it allows for corrections to be applied during movement of IDs. Moving IDs, especially at small gaps, cause significant deviations in optics. A feedback system, on the other hand, is limited by the measurement frequency. In the context of local ID compensation, this translates to the time taken for a measurement of the β functions. As will be discussed in Section 3.4, a measurement of the β functions takes no less than 120 seconds using pingers, which is too slow to account for moving IDs.

Evaluating local ID compensations

In simulation, the objective of the local compensation is to restore model β functions. On the physical storage ring, the objective is reformulated slightly; instead of attempting to reproduce model optics, the objective of the compensation is to restore β functions to those of a prior *baseline* measurement. A *baseline* measurement in the context of this project refers to a measurement of β functions made at a machine configuration where every ID is set to its respective maximum gap. In this way, a baseline measurement is designed to represent the β functions in the ring in the absence of any insertion devices. Throughout this project, when evaluating the effectiveness of local ID compensations on the machine, the main parameter of interest is the *relative* β beating *with respect to a prior baseline measurement.* Mathematically,

$$\Delta\beta_u/\beta_u = (\beta_{u,measured} - \beta_{u,baseline})/\beta_{u,baseline},\ u \in \{x, y\},$$

where $u = x, y$ denotes either of the transverse planes. Thus, evaluating the effectiveness of a local compensation translates to determining the degree to which the compensation is able to restore the baseline β functions.

The evaluations are generally performed using three measurements:

i) Baseline measurement of β functions, acting as the "target" of the compensation as described above.

ii) β function measurement with the ID of interest at a particular configuration in the absence of any compensation.

iii) β function measurement, using the same ID configuration as ii), with the application of the compensation.

Using this three-step approach, it is possible to determine the extent to which the compensation is able to restore the baseline β functions. Furthermore, measurement ii) can be used to compare the effects on optics predicted from kickmaps to the effects of the ID on the machine itself. In this project, two methods were employed to determine the β functions in the ring. The more conventional method is LOCO, which is described in more detail in Section 3.3. The other method, which was developed during (though not as part of) this project was pinger measurements, described in Section 3.4.

3.2.2 Global compensation to eliminate tune shift

Another optics correction investigated in simulation as part of this project is the restoration of transverse tunes $Q_u,\ u \in \{x, y\}$. As described in eq. 2.17, the tunes are determined by the β functions around the ring. Even if the local compensation for a particular ID described in Section 3.2.1 restores the baseline β functions throughout most of the ring, the region between the outermost flanking defocusing magnets will be subject to altered β functions. Hence the tunes will shift with respect to the baseline, not only as a result of the direct effects of the ID, but also the local compensation. As the tunes are parameters critical to the stability of the beam, the global compensation aims to restore baseline tunes, and is meant to be used in conjunction with the local compensation. It should be noted that the prospect of a global ID compensation for the tunes in the form of a feedforward device was not investigated nearly as thoroughly as the local compensation. The main reason for this that there are already fully operational tune feedback systems on both rings. While a tune compensation in the form of a feedforward could, in principle, be advantageous during ID gap movement, during operation, the tune feedback system is sufficient. Moreover, as is discussed in Sections 4.1.3 and 4.1.4, the local compensation based on kickmaps proved insufficient. The development of model-independent pinger-based β function measurements (Section 3.4), which can be used to determine quadrupole adjustments via measurements (Section 3.5), opened a new avenue. For this reason, priorities within this project shifted to developing this online approach to determining quadrupole strengths. It should be noted that a global compensation scheme in the form of a feedforward system remains a possible development of linear optics ID compensation at MAX IV laboratory.

The label *global* compensation refers to the magnets involved in the compensation. Instead of adjusting two pairs of magnets by a relatively large amount (which would introduce significant β beating on its own), the global correction attempts to restore tunes by shifting the strengths of each magnet in two families by a relatively small amount. As in the case of the local compensation, this requires the involvement of one focusing and one defocusing family. On the 3 GeV ring, there are two potential candidate pairs of families. One is the aforementioned QFE/QDE pairs, the same involved in the local compensation. The other are the QF/DIP families, where QF is another family of horizontally focusing pure quadrupoles, and DIP, which, much like on the 1.5 GeV ring, are a family of horizontally defocusing multipoles. It should be noted here that while the usage of the QFE/QDE families in this way has been studied before [2], the implemented tune feedback system on the 3 GeV ring employs the QF/DIP families. This was done mostly for practical purposes, as the QF/DIP magnets were not involved in any control schemes. In the tune feedback system of the 1.5 GeV ring, the SQFO/DIP families are used, the same that are involved in the local compensation. An important note regarding the global compensation is that by adjusting the quadrupole strengths of two families to restore baseline tunes, β beating will be (re-)introduced. The β beating induced through the global correction is generally significantly smaller in magnitude than before the application of the local correction, so the combined local-global correction can be argued to be beneficial overall. Furthermore, the choice of magnet families to involve in the global correction has an impact on the envelope of β beating.

Finding correction terms for the global compensation was done numerically. In this simulation, IDs were modeled using kickmaps, and the local compensation was performed as described as in the previous Subsection 3.2.1. With the model in a post-local-compensation configuration, correction terms for the global compensation can be found. The optimization algorithm for this purpose employs a gradient-descent approach, constructing a Jacobian where each column represents the change in tunes caused by a small deviation in magnet strength applied to each magnet in the relevant family. For simplicity, the algorithm scales the quadrupole strength of each magnet in the respective families by a constant factor, including the magnets involved in the local compensation. Thus, the optimization algorithm finds the appropriate quadrupole strength scaling factors for each family.

The implementation of the global compensation on the machine would be much like that of the local compensation. The main difference is that the feedforward devices for the global compensations write correction terms to each actuator in two magnet families.

3.3 Linear Optics from Closed Orbits (LOCO)

3.3.1 Working principles of LOCO

Linear Optics from Closed Orbits (LOCO), originally described in [8], is a diagnostic tool for the linear optics in storage rings. It combines an orbit response matrix measurement, as outlined in Section 2.1.7, eq. 2.22 with an internal model of the magnetic lattice of the storage ring. From the internal lattice model, a model orbit response matrix R_θ^{mod} is calculated. The LOCO algorithm attempts to recreate the measured R_θ^{meas} by adjusting user-selected parameters in the model. For this purpose, a fit is performed to minimize the quantity χ^2:

$$\chi^2 = \sum_{i,j} \frac{\left(R_{\theta,ij}^{meas} - R_{\theta,ij}^{mod}\right)^2}{\sigma_i^2}, \tag{3.1}$$

where σ_i denotes the RMS noise of the i:th BPM measurement of closed orbit deviation, and j indexes the parameters involved in the fit. By performing this fit, an adjusted model lattice is created, whose associated response matrix minimizes χ^2. For linear optics considerations, dipole and quadrupole strengths are of primary interest, though in principle higher-order magnets (such as sextupoles, octupoles) can be included. For the LOCO fit to be able to account for IDs, these were represented as a series of three thin quadrupoles. The choice to not use kickmaps in the LOCO model was conscious, to reduce the overall reliance on them.

In this project, LOCO was mainly used to determine β beats on the 3 GeV ring. As explained in Section 3.2.1, this corresponds to making (at least) two measurements: one representing the *baseline*, and at least one other representing the *deviation*. LOCO determines β functions through numerical calculation from the adjusted model lattice obtained from the fit. For this reason, determining β functions using LOCO is a highly model-dependent approach. For convenience, throughout this report, LOCO-determined β functions are referred to as LOCO *measurements* of β functions, through strictly speaking the term measurement is misleading. Moreover, LOCO measurements are time-consuming. A typical orbit response matrix measurement takes around 20 minutes on the 3 GeV ring, and the subsequent fit another 5-10 (on powerful computers; significantly more if performed on a simple laptop).

3.3.2 Evaluation of LOCO as a tool to determine β functions

This section is dedicated to a discussion surrounding the effectiveness of LOCO as a tool to determine β functions on the ring, and to address some of the factors that could influence LOCO measurements of β functions.

Dependence on beam current

As mentioned in Section 3.1.2, BPM signal-to-noise can be increased by increasing the beam current. For this reason, an investigation was performed to see whether the beam current has a strong influence on LOCO measurements of β functions. Figure 3.3 shows different LOCO measurements of the β beats caused by BALDER at the minimum gap of 4.5 mm without any corrections. The left plot shows the low-current case at 3 mA, the right plot the higher-current case at 30 mA beam current. The behavior of β-beats in the vertical plane between the two measurements is similar, with peak amplitudes at just over 4% and a rather constant envelope. The vertical RMS β beating between the two measurements differs by about 0.08 percentage points, equivalent to about 3%. This difference is rather small, and suggests that increasing signal-to-noise in the BPMs by increasing the beam current does not markedly change the result. Moreover, consider the χ^2-values of said fits: the χ^2-values of the baseline measurements are 5.94 and 7.06 in the low- and high-current cases, respectively. The closed-gap measurements have corresponding values of 5.47 and 7.41, respectively. In general, a slightly higher χ^2 is to be expected in the higher-current case, as increasing signal-to-noise in the BPMs leads to smaller σ_i in eq. 3.1.

The case in the horizontal plane is different, where the low-current measurement gives a significantly larger overall beating at $> 0.6\%$ RMS compared to the high-current at $< 0.3\%$ RMS. From simulation (Table 4.1), only a very small beating is expected in the horizontal ($< 0.01\%$ RMS). If the kickmap is to be trusted, it could be that the higher signal-no-noise in the BPMs does significantly improve the horizontal measurement in particular, and that the measured, highly irregular β beating seen in the left plot of Figure 3.3 is an artifact produced in the LOCO fit as a result of large error margins in the horizontal σ_i. Despite the discrepancy in the horizontal measurement, the vertical measurement is of greater interest, and based on the similarity between the measured vertical β beating, the qualitative conclusion is that an increase in current from 3 mA to 30 mA does not significantly change (in particular, does not necessarily improve) measurement quality.

Dependence on baseline

Another factor pertaining to LOCO fits investigated is the dependence of data such as that presented in Figure 3.3 on baseline measurements. There were concerns that a machine at an uncorrected (to model optics) configuration, yielding a baseline measurement with large deviations in β functions with respect to the model, could yield systematically different results compared to a baseline close to model optics. In other words, this investigation was performed to see whether different baselines (with respect to each other) would significantly influence the resulting $\Delta\beta/\beta$-plots presented throughout this subchapter. Consider Figure 3.4, which in its top row shows two different baseline LOCO measurements. The top-left two plots display uncorrected horizontal (left) and vertical (right) β beats with respect to the *model*. The top-right plot shows horizontal and vertical β-beats in the case where optics have been corrected to resemble the model. The bottom row shows measurements of β-beats from BALDER at a 4.5 mm gap with respect to the baseline measurements above. The baseline measurements are quite different. The uncorrected case has peak deviations of $> 5\%$ and $> 4.5\%$ with respect to model optics in the horizontal and vertical, respectively. Notably, there are deviations in excess of 3.5% along most of the ring in the horizontal. The corrected case (top-right) displays much smaller deviations from model optics overall.

Despite this significant difference in baseline measurements, the ID measurements, using their respective

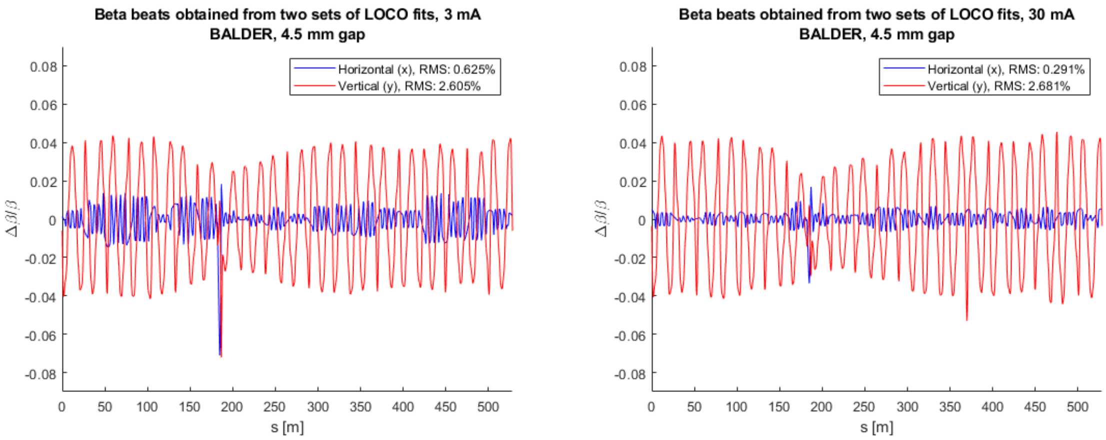

Figure 3.3: LOCO measurements of β-beats for BALDER at a gap of 4.5 mm. Left: 3 mA beam current; right: 30 mA. Note that data near the insertion device (around $s = 185$ m) has been interpolated and may not be accurate.

baseline measurements to calculate β beats, produce similar results. The RMS vertical β beats are 2.599% and 2.605% using the uncorrected and corrected baselines, respectively. In the horizontal, the corresponding values are 0.596% and 0.625%. Peak deviations are similar, at just above 4%, and the envelope does not vary significantly except in the immediate vicinity of the ID. Thus, there does not appear to be a significant difference between ID measurements depending on the baseline. This results speaks in favor of the approach taken to LOCO measurements throughout this project. Furthermore, it indicates a degree of robustness in LOCO measurements. In reality, the machine will never have the exact same optics (or otherwise) properties at two different points in time, even when cycling to the same configuration. Given this fact, the reproducibility of results (Figure 3.4, bottom row) across different machine configurations (Figure 3.4, top row) is a point of strength for this kind of measurement.

Reproducibility of kickmap effects using LOCO fitting

The last point that will be discussed in this section regarding LOCO as a tool to determine the effects of IDs on machine optics pertains to its internal model of the ID itself. Throughout this project, IDs, in the internal model used by LOCO for fitting, were represented by a longitudinal series of three thin pure quadrupoles. The main reason for this was to provide LOCO with sufficient degrees of freedom to attribute the linear optics effects of the ID to something in its model at the appropriate longitudinal coordinates. From this choice, an interesting question arises regarding whether LOCO's fitting procedure accurately reproduces the effects that the *kickmap* would have on the beam in its triple quadrupole representation.

This issue was investigated in a simple theoretical experiment, whereby a transverse grid of particles were tracked through (1) the LOCO-determined quadrupole triplet of a measurement of BALDER at a gap of 4.5 mm, and (2) the corresponding kickmap, in either case followed by a drift to emphasize the resulting particle positions. Figure 3.5 shows the results of this experiment. The red rings represent the final particle positions after being tracked through the kickmap. In this case, there is a clear focusing effect in the vertical, where particles with an initial vertical position of ± 0.2 mm have approached the horizontal axis to within ± 0.1 mm. Particles starting closer to the horizontal axis have undergone a similar focusing motion. The effects in the horizontal are small, with particles deviating only very slightly from their initial horizontal positions. The results of the kickmap tracking are in line with expectations from theory (Section 2.3.2).

The blue diagonal crosses represent final particle positions after being tracked through the LOCO model quadrupole triplet. The result is rather different compared to the kickmap. For one, the focusing in the vertical is not nearly as pronounced, and each of the five particle "columns" has seemingly tilted clockwise to some extent. The latter is likely a result of LOCO, in its fit, being explicitly allowed to fit not only *normal* quadrupole field components, but also *skew* components. It is clear from the figure that the LOCO model does not reproduce the strong vertical focusing observed in the case of the kickmap, and greatly

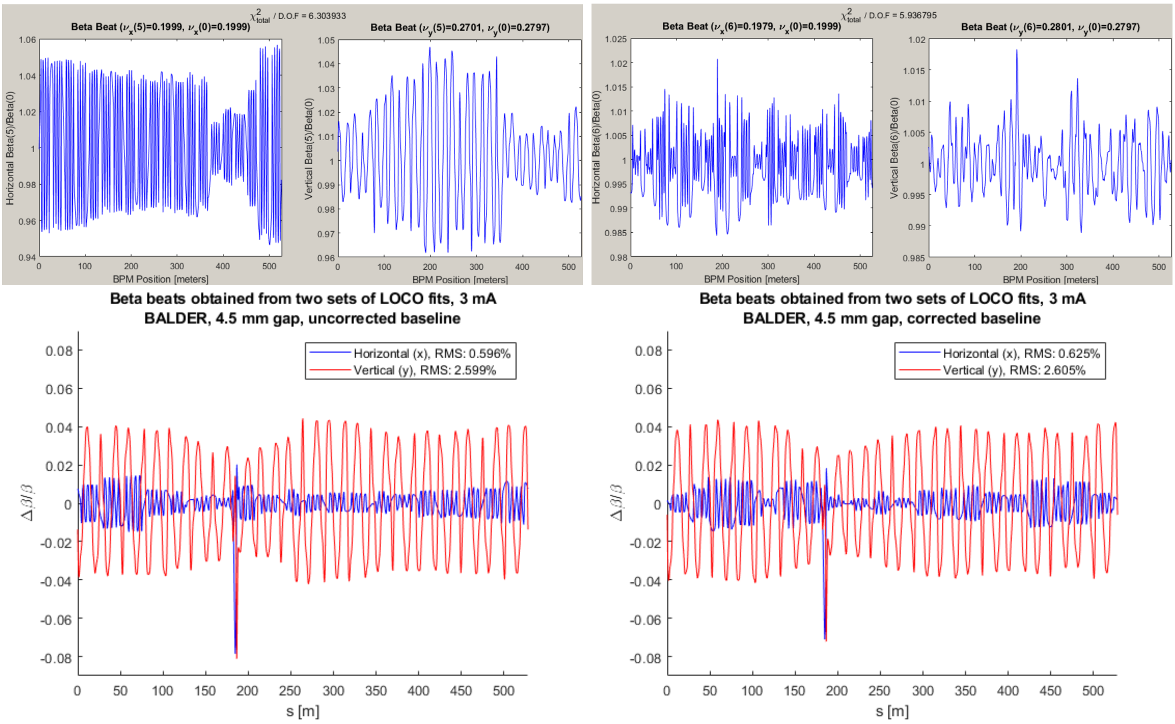

Figure 3.4: Top row: Different LOCO baseline measurements. Note that throughout this row, $\Delta\beta/\beta$ are calculated with respect to model optics. Bottom row: measurements of β-beating *with respect to the above baseline measurements* caused by BALDER at a gap of 4.5 mm.

exaggerates any potential skew quadrupole effects. It should be noted that the discrepancy between the results in this experiment has a multitude of possible explanations. One possibility is that modeling the kickmap as a series of pure quadrupoles is too simplistic. ID magnetic fields are known not only to have quadrupole components, but also octupole components. Furthermore, the LOCO fitting procedure is complex, with many degrees of freedom. Hence it is possible that effects that physically arise from an ID are attributed to a quadrupole (or other magnet) elsewhere in the ring. Regardless of the precise cause for this discrepancy, it is clear that LOCO's model dependence is a weakness within the framework of determining β-functions, especially in the presence of IDs. For this reason, and the development of the model-independent pinger β function measurements, the rest of this chapter is dedicated to pinger measurements, and their use in local ID compensation.

3.4 Pinger-based measurements of β functions

3.4.1 Working principles

The pinger-based measurement of β functions constitutes a (largely) model-independent and fast method of determining the β functions. As outlined in Section 2.2.2, it is based on the determination of the beam history matrix B followed by SVD analysis. The general measurement procedure is as follows:

i) The beam is kicked in one plane using the namesake pinger magnet to generate a transverse oscillation (as per eq. 2.20). To generate this kick, pinger magnets have an internal capacitor bank, whose voltage $V_{H,V}$ (H, V denoting horizontal, vertical) can be controlled. Changing the voltages $V_{H,V}$ corresponds to changing the strength of the kick.

ii) The disturbed beam traverses the ring a number of turns while the BPMs record the transverse position of the beam each turn.

iii) Steps i)-ii) are repeated a set number of times to give good statistics. The median of each BPM:s readout is taken as the final measurement value, to exclude outliers.

iv) Steps i)-iii) are repeated, in the other plane.

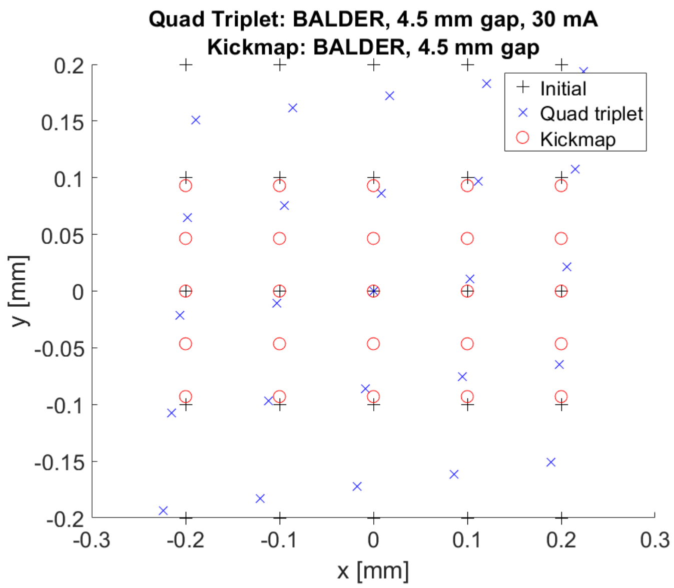

Figure 3.5: Simulation of a particle grid being tracked through a quadrupole triplet (referred to as "quad triplet") obtained from a LOCO fit from a measurement of BALDER at 4.5 mm gap, and the corresponding kickmap, respectively.

Once data has been gathered, SVD analysis gives the β functions times a scaling factor C. The scaling factor is the one model-dependent aspect of the pinger-based β function measurements. To obtain usable measurements, the constant C is determined through normalization with model β functions. The factor C is expected to be constant across measurements, so for measurements of $\Delta\beta/\beta$, it is not an issue.

Measuring the beam response matrix is somewhat involved, and requires both a precise calibration of the BPMs, and a specific beam circulating in the accelerator. In particular, the beam needs to consist of a short train of bunches (called a "few-bucket fill pattern"). The reason for this is that the pulse length of the pinger kick spans 1-2 turns, meaning that if a long train of bunches were kicked, they would experience different magnitude kicks. The need for precise calibration of the BPMs stems from a similar issue, namely that each BPM reading needs to be assigned to the correct turn. In its current implementation, the BPM timings are not perfectly accounted for, as is discussed in Section 3.4.2, though work is underway to resolve this issue.

3.4.2 Measurement reliability evaluation and resolution optimization

This section is dedicated to the study performed as part of this project on the pinger-based β function measurements pertaining to characterizing and optimizing the random error in β functions between measurements. The need for this type of study became apparent when testing the online local compensation algorithm on the 1.5 GeV ring, which is discussed in Section 4.2.2. The main result of this investigation is that practical configurations were found which resulted in BPM-averaged standard deviations of $\sigma(\Delta\beta/\beta) < 10^{-3}$ in both planes on the 3 GeV ring, and $< 2 \times 10^{-3}$ in both planes on the 1.5 GeV ring. Specifically, pinger voltages of $V_H = 500\,\text{V}$ and $V_V = 1500\,\text{V}$ with 25 pings for measurements on the 3 GeV ring, and $V_H = 600\,\text{V}$, $V_V = 1600\,\text{V}$ with no less than 20 pings on the 1.5 GeV ring. At these configurations, a single measurement of β functions takes about 130 s on the 3 GeV ring, and about 120 s on the 1.5 GeV ring. The complete set of configurations tested is presented in Figure 3.6.

According to this investigation, the standard deviation in measured β beating is expected to be $< 10^{-3}$ between measurements for the average BPM at the 3 GeV ring. To find this value, the median over a number of measurements was used, to reject outliers. This is an important result within the framework of the online optics optimization described in Section 3.5, as the method is reliant on pinger measurements. Specifically, the goodness parameters described here provide an estimated noise floor for measurements used in the calculation of correction terms, as well as measurements to evaluate their impact on ring optics. As part of the method development of pinger measurements of β functions, this optimization

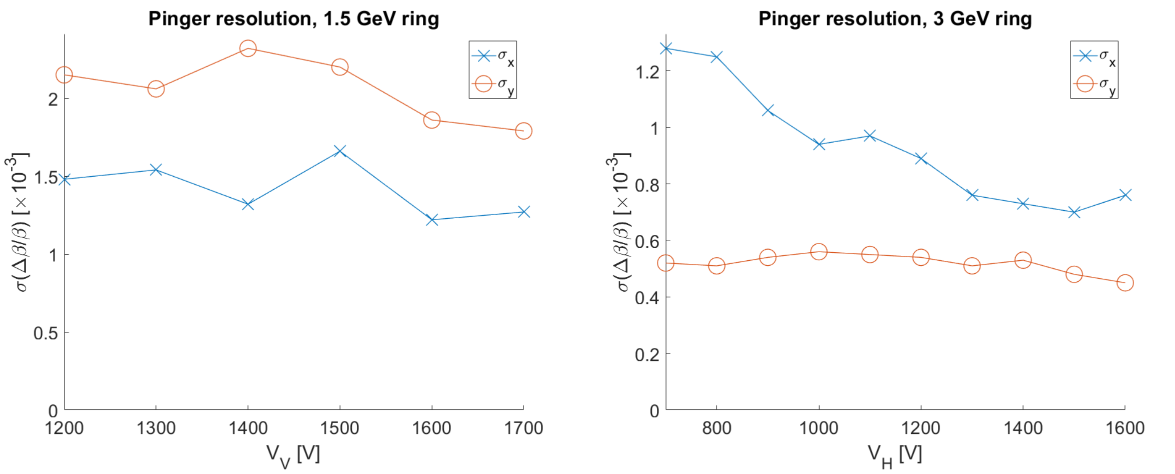

Figure 3.6: BPM-averaged $\sigma(\Delta\beta/\beta)$ in both planes as a function of: Left: vertical pinger voltage on the 1.5 GeV ring, 25 pings used. Right: horizontal pinger voltage on the 3 GeV ring, 20 pings used. Note that the vertical axes are not shared. Data from the region of BPMs subject to timing issues has been excluded.

was carried out on both storage rings. The parameters investigated were the kicker voltages V_H, V_V, as well as the number of kicks (pings) in each measurement N_P. The kicker voltages were investigated to determine a configuration to maximize signal in the BPMs, without damaging the beam, whereas the number of kicks was optimized to find a compromise between measurement time and statistics.

The general procedure of this investigation was to first perform a baseline measurement of β functions, followed by a further series of N_M β function measurements, each at the same kicker voltages V_H, V_V and number of pings N_P. For each of these N_M measurements, the β beats with respect to the baseline was calculated. This way, each BPM has an associated N_M values of β beats in each plane. To characterize the random error between measurements, the standard deviation $\sigma_{n,q}$, $q = x, y$ and n indexing the BPMs over these measurements is calculated for each BPM and plane. The final goodness parameter is chosen to be the mean of these standard deviations $\sigma_{n,q}$ over all BPMs. Thus, each tested configuration of kicker voltages has two associated goodness parameters, one for each plane.

There are a few important notes to make with regard to Figure 3.6. The first observation is that each of the rings has a generally markedly worse resolution in one particular plane. For the 1.5 GeV ring, it is the vertical, while for the 3 GeV ring, it is the horizontal. While this phenomenon is not fully understood, there is a physical reason why this behavior is expected. As per eq. 2.20, the strength of the kick provided by the pinger magnets is not only dependent on the voltage provided, but also on the *value* of the β-function in the relevant plane *at the location of the kickers*. Thus, for a given kicker voltage, one would expect a less powerful kick if the kicker is located in a region with a low β function value. The converse is also true. The strength of the kick determines the magnitude of the resulting oscillation of the beam, and thus influences the amount of signal that the BPMs receive, as mentioned in Section 3.1.2. In the model, the values of the β functions at the location of the kickers in the 1.5 GeV ring are $\beta_x = 5.94$ m at the horizontal kicker, and $\beta_y = 3.08$ m at the vertical. Correspondingly, on the 3 GeV ring, the values are $\beta_x = 1.17$ m at the horizontal kicker, and $\beta_y = 6.86$ m at the vertical. In this investigation, to compensate for this effect, the simplest approach was taken, namely to increase the voltage provided to the "weaker" pinger magnets. On the 3 GeV ring, increasing the kicker voltage markedly improved the resolution, as can be seen in Figure 3.6. The resolution improved from 1.77×10^{-3} at a horizontal kicker voltage of 500 V to 0.70×10^{-3} at 1500 V. The effect of increasing the "weak" vertical kicker voltage on the 1.5 GeV ring is not as clear, though there is some improvement, going from a vertical resolution of 2.15×10^{-3} to 1.79×10^{-3} at $V_V = 1200$ V and 1700 V, respectively.

In the effort of increasing resolution, increasing the kicker voltage has the added appeal over additional statistics in the form of more pings in that it does not increase the total time taken for a β function measurement. One of the main appeals of the pinger-based β function measurements is speed, and the

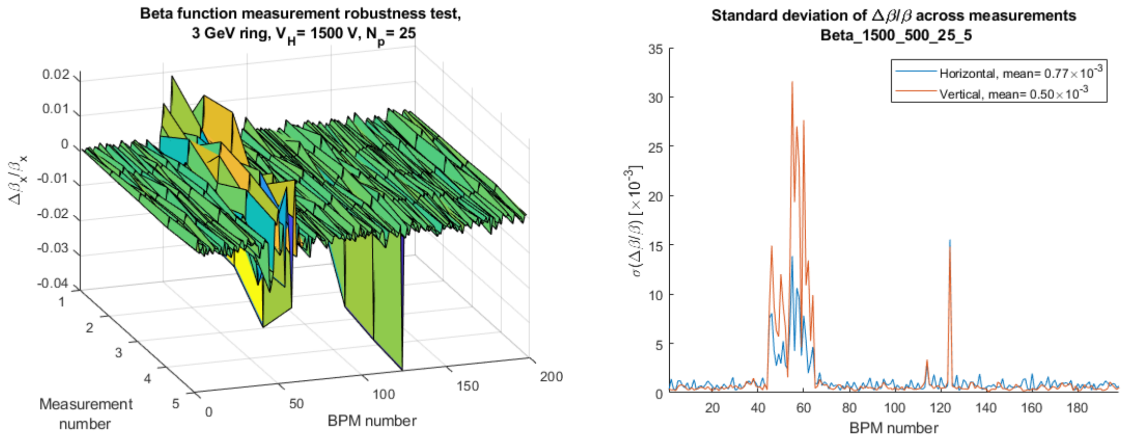

Figure 3.7: Resolution test of the pinger-based β function measurements. Left: β beats with respect to baseline over five measurements. Right: the corresponding per-BPM standard deviations over said five measurements. Note that for the calculation of the resolutions presented in Table 3.6, the timing affected region at BPMs indexed 45-66 and outliers 114 and 124 were excluded.

preferred configurations were chosen with this in mind. It should be noted that the main increase in measurement runtime as a result of a greater number of pings does not come from the time taken to physically ping the beam and wait for BPM readouts. This process happens at a rate of nearly 2 Hz. Instead, it is the subsequent data analysis and conversion of raw BPM data to usable β function measurements that adds significantly to the time taken overall. Nevertheless, with measurement times of no more than 150 s with good statistics, the pinger-based β function measurements outclass LOCO measurements in terms of speed, at around 30 and 8 minutes at the 3 GeV and 1.5 GeV rings, respectively. The measurement runtime could probably be reduced further, by rewriting parts of the analysis code or parallelizing the computation further, which was not purpose-made for pinger measurements.

Another point about the measurements themselves is that some BPMs had to be excluded from the analysis on which Figure 3.6 is based. This is because of a known timing issue in the BPM system, which is reflected in a region of subsequent BPMs providing systematically erroneous data. The BPM timing issue is present in β function measurements on both rings. The left surface plot in Figure 3.7 shows the measured horizontal β beats on the 3 GeV ring at the preferred horizontal pinger voltage of 1500 V. The region spanned by BPMs number 45-66 is characterized by large deviations in measured β functions compared to the baseline. The affected BPMs give erratic data, which is illustrated in the right plot of Figure 3.7, showing standard deviations per BPM over the five performed measurements. BPMs number 45-66 have associated standard deviations almost an order of magnitude larger than the median. Furthermore, BPMs number 114 and 124 have been identified as outliers and were also removed in the calculation of the resolution.

The presence of a region of BPMs providing erroneous data has implications for the online corrections described in the following section, 3.5. In particular, it means that on the 3 GeV ring, 22 out of the 198 BPMs, meaning just above 11%, provide incorrect data, which is used to calculate correction terms. This issue is more pronounced on the 1.5 GeV ring, where 8 out of the 37 BPMs fall into this region, corresponding to almost 22% of all BPMs. Importantly, the affected BPMs can vary between measurements. For this reason, in its first implementation, the online compensation algorithm does not account for these chunks of poor BPM data, which is a clear limitation. Ideally, the problem of BPM timing should be resolved in direct conjunction with the hardware configuration itself, though a more short-term solution could involve data manipulation to try to render the BPM data usable, or even simpler to appropriately exclude the data when calculating corrections. It should be noted here, however, that it was possible to perform at least partial online ID compensations in spite of these issues, as will be discussed in Sections 4.2.2 and 4.2.3.

3.5 Online correction algorithm to determine quadrupole adjustments

This section is dedicated to the description of the "online" β function correction algorithm developed as part of this project. The algorithm is iterative, and uses a numerically determined β function response matrix in conjunction the with pinger-based β function measurements described in Section 3.4 to calculate and apply correction terms to the relevant magnets. This section is divided into two subsections. The first, Section 3.5.1, is a description of the online correction algorithm, and the latter, Section 3.5.2 is a description of how the "β function response matrix" R_β is determined.

3.5.1 The online correction algorithm

The online ID correction algorithm is designed to perform *local* compensation only, that is, eliminating the β beating caused by the presence of an insertion device by adjusting the strengths of flanking quadrupoles. It is intended as a (more) model-independent approach to ID linear optics correction than the kickmap-based correction, which is achieved by relying on the model-independent pinger measurements of β functions. The implementation is that of an iterative optimization, and follows the procedure:

i) Perform a baseline β function measurement with fully open ID gaps using pingers and save the initial magnet strengths,

ii) Set the ID of interest to the desired gap and/or phase,

iii) Measure β functions and calculate the difference $\Delta\beta = \beta_{measured} - \beta_{baseline}$. In this step, measurements of β functions by BPMs inside the region spanned by the outermost flanking magnets involved in the compensation are discarded. In this first implementation of the online correction algorithm, data from the faulty BPM region is included in further calculations.

iv) Using the assumption that $\Delta\beta = R_\beta \Delta k$, calculate the correction terms Δk using the Moore-Penrose pseudoinverse of the β function response matrix, R_β^+, s.t. $R_\beta^+ \Delta\beta = \Delta k$, as described in Section 2.2.1. In other words, in this step, the focusing error physically caused by the ID is ascribed to quadrupole errors in its flanking magnets. Using the pseudoinverse, correction terms Δk to said flanking magnets are thus found to minimize $\Delta\beta$.

v) Apply the calculated correction $m\Delta k$ to the actuators controlling the flanking magnets, where $0 < m < 1$ denotes a scaling factor.

vi) Repeat steps iii)-v) a set number of iterations.

Ideally, the correction Δk calculated in step iv) completely eliminates β beating when applied. However, for stability in the implementation on the physical machine, it was decided that a scaling factor $0 < m < 1$ would be applied to the calculated correction. In particular, noise (as well as the faulty BPM region described in Section 3.4.2) in the measurement of β functions could translate to large errors in calculated Δk-values. With the inclusion of this scaling factor, the algorithm is expected to converge as the number of iterations increases, in the sense that as measured deviations $\Delta\beta$ decrease, as do the calculated correction terms Δk, both eventually settling at some noise floor. Furthermore, if the number of iterations chosen turns out to be insufficient, the algorithm can be run again, using the same initial baseline measurement as the "target" β functions. While a more dynamic stopping condition could be implemented, such as a sufficiently low RMS($\Delta\beta/\beta$), as part of the development process, it was decided that the algorithm be run a set number of iterations.

The online correction algorithm is decidedly different to the kickmap-based corrections owing to its model independence. It directly measures the affected β functions. This is principally a large advantage, both because modeling kickmaps requires significant effort, and because it eliminates any modeling issues as error sources. It should be noted, however, that there is one part of the online correction algorithm which in its current implementation remains model-dependent: the calculation of the "β function response matrix" R_β, which is discussed in more detail in the next Section, 3.5.2.

The desired implementation of local ID compensations is that of a feedforward system. Translating the results of the online optimization algorithm to a feedforward-table is simple; having performed the online

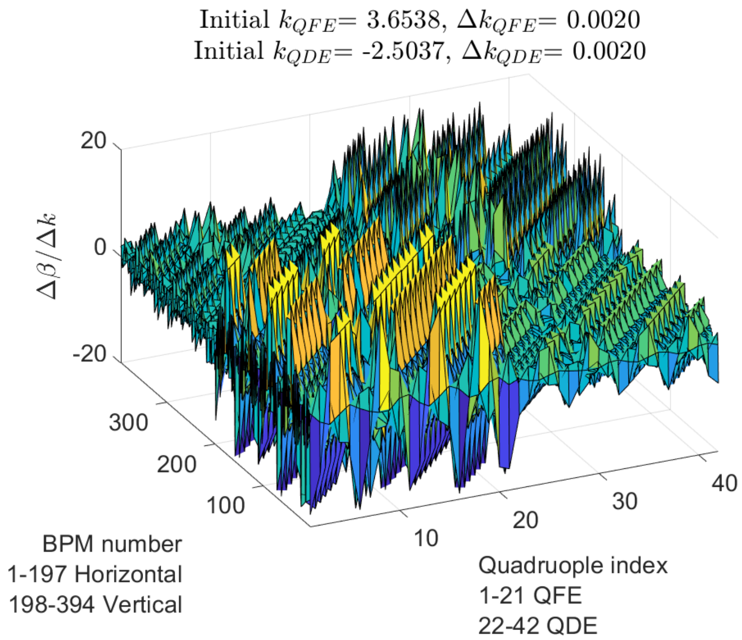

Figure 3.8: The β function response matrix R_β determined for the 3 GeV ring. Each column corresponds to one magnet actuator, and each row corresponds to one BPM, either in the horizontal or vertical. Note that the lattice model used for the calculation of R_β in this Figure, only 197 BPMs are present, as opposed to the 198 on the machine itself.

correction for an ID at a given configuration, the correction term to be put its feedforward table is simply taken to be the difference between the final and initial readout values of the flanking magnet strengths. Another consideration is time. The time required can likely be reduced further, by increasing the scaling factor m and reducing the number of iterations. As is discussed in Section 4.2.2, the convergence of the algorithm seems to be limited by noise in the β-function measurements, and thus further improvements to pinger measurements can likely be exploited to reduce the time taken for online corrections.

3.5.2 Determining the "β function response matrix" R_β

The β function response matrix R_β calculated for the 3 GeV ring is presented in Figure 3.8. It is calculated iteratively, by varying the quadrupole strength k of one of the magnet actuators by a value Δk at a time, and calculating the resulting difference $\Delta\beta$ between the new β functions and the model β functions in each of the BPMs. As both the horizontal and vertical planes are of interest, $2 \times N_{BPM} = 2 \times 198 = 396$ values of $\Delta\beta$ are determined in this way. The calculated differences $\Delta\beta$ are normalized by the applied variation in actuator strength Δk. The resulting array of 396 values of $\Delta\beta/\Delta k$ forms one column of R_β. This process is repeated over all actuators in the QDE/QFE families to form the complete matrix. As indicated at the top of Figure 3.8, the chosen variations Δk are less than 10^{-3} of the initial quadrupole strengths $k_{QFE/QDE}$. This is to ensure linearity of the response in β functions. Another consideration to be made is that throughout this report, the main figure of interest is *relative* β beating, that is $\Delta\beta/\beta$, while the matrix elements of R_β use *absolute* differences in β. The reason for this is that the compensation algorithm proved ineffective when matrix elements consisted of relative β beats. Likely this has to do with the fact that each matrix element, when expressed as $\Delta\beta/\beta$, naturally becomes normalized by a different number. This is a numerical issue, as the values of β functions can vary by almost an order of magnitude at different points in the lattice. The calculation of R_β for the 1.5 GeV ring is conceptually done the same way. The main differences between R_β between the two rings are that the 1.5 GeV ring uses the QF/DIPc families to perform the compensation, and that there are only 35 BPMs. R_β determined for the 1.5 GeV ring is presented in Figure 3.9.

On the 3 GeV ring, there are a total of 40 physical magnets in each of the QDE/QFE families. In general, these two magnets are connected to the same power supply, thus a variation applied to the actuator changes the two physical magnets in tandem. The exception to this are the two QFE and two QDE

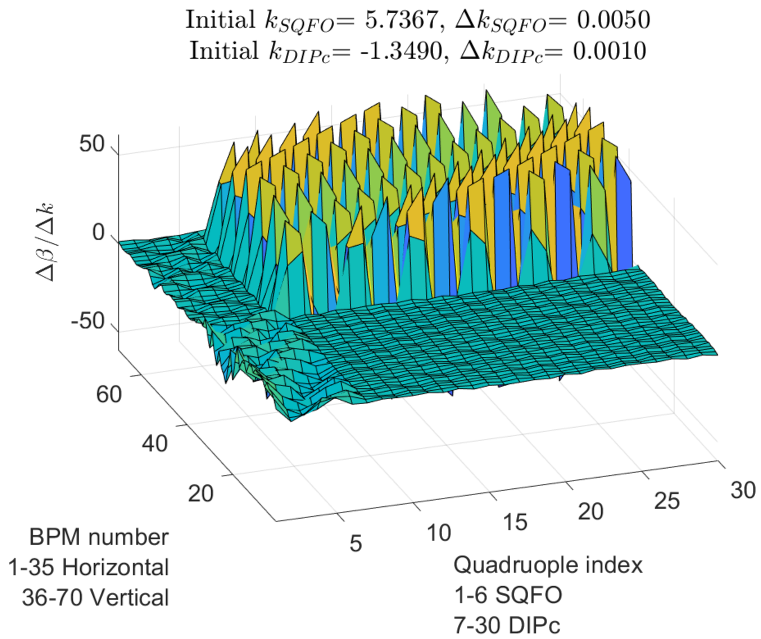

Figure 3.9: The β function response matrix R_β determined for the 1.5 GeV ring.

magnets flanking the IVW at the BALDER beamline. These four magnets each have an individual power supply, allowing their strengths to be varied independently. This was a conscious design choice, as due to the asymmetrical positioning of the BALDER IVW in its long straight, the degrees of freedom provided by four independent flanking quadrupoles are required to completely compensate for β beating. Thus, each magnet family has $19 + 2 \times 1 = 21$ actuators. The situation on the 1.5 GeV ring is somewhat different. There are only six actuators in the focusing SQFO family. Of these, five are of interest, as these correspond to power supplies to quadrupole pairs mounted to flank each of the five IDs on the ring. These five were installed to be able to perform ID compensation for each of the IDs. In the DIPc family (acting as defocusing quadrupoles), there are 24 actuators, corresponding to two per achromat. The more granular control was necessary over the PFS adjusting the quadrupole strength in bending dipoles was needed to enable correction to bare lattice optics.

One feature of R_β observed in both rings is that it is divided into quadrants based on the horizontal/vertical response of the focusing/defocusing quadrupole families. In Figures 3.8 and 3.9, it can be seen that the response in horizontal β_x is much greater when applying a variation to quadrupole strength in the *focusing* QFE/SQFO families of quadrupoles. Conversely, the vertical response is much greater when varying the strengths of the *defocusing* QDE/DIPc families. This behavior is expected from theory [4]. To illustrate this point further, the DIPc are installed in locations with large vertical β, and QFE/SQFO in locations with large horizontal β. As per equation 2.24, the value of the β function directly influences the resulting β beating. For clarity, some of the information in the calculated R_β is redundant. Only the columns representing magnet actuators of magnets flanking IDs are ever used in local β function compensation.

In principle, it would be possible to measure R_β rather than calculating it numerically using pinger β function measurements. The appeal of this approach is to obtain R_β determined entirely on the machine itself, which eliminates any sort of model dependence. However, there are a number of challenges with this approach. The first pertains to the issue of the faulty BPM regions mentioned in Section 3.4.2. Effectively, the presence of these faulty BPMs would render $2 \times 8 = 16$ and $2 \times 22 = 44$ rows in R_β on the 1.5 GeV and 3 GeV rings unusable. The second issue pertains to time. The quickest way of performing the required β response measurements would be to apply the variation Δk to one actuator, measure the response, and "reset" the variation, by applying $-\Delta k$ to the actuator, and repeating this process for the other actuators. The variation Δk applied on the physical machine would likely have to be significantly larger than the $< 10^{-3}|k|$ used in simulation, to be able to resolve the β response using pinger measurements. Due to magnet hysteresis, it is highly unlikely that the optics of the ring return to their original state after "resetting" the actuator strength. For this reason, any subsequent β response measurements should be assumed to be performed on slightly different machine optics than the first. While hysteresis might not be a significant issue for the first few measurements, the effects are cumulative, and it is likely that later measurements represent standard operating conditions increasingly worse. The simplest and arguably most reliable way to account for hysteresis is to cycle the machine as a whole, returning magnets (and thus optics) to a well-known configuration. This requires the beam to be dumped, and takes around 30 minutes on the 3 GeV ring. If the machine were to be cycled between every measurement of β response to quadrupole strength variation, it would have to be cycled 42 times, and the beam re-injected after each cycling. Thus a complete measurement of R_β would require in excess of 20 hours, which is unfeasible, especially given the length of the standard maintenance window of 24 hours weekly on each ring.

Following any hardware changes expected to change linear optics, R_β would have to be re-measured. Moreover, a deviation between the model-determined and real R_β is expected to result in slower (or no) convergence of the online algorithm. In the case of slower convergence, the issue can be circumvented by increasing the number of iterations. For these reasons, in this first test of online ID compensation, a numerically determined R_β was deemed sufficient.

4 Results

This chapter is divided as follows: Section 4.1 presents the results from simulated local compensation based on modeling the effects of IDs using kickmaps, as outlined in Section 3.2. Section 4.2 shows the results obtained from the online compensation, both in simulation and select machine tests. Section 4.3 is dedicated to a comparison between the approaches to finding quadrupoles: using kickmaps in simulation, and the online correction algorithm.

4.1 Kickmap-based approach

4.1.1 Simulated local compensation to eliminate β-beats

The numerical optimization described in Section 3.2.1 successfully reduces the goodness-of-fit to a value below a user-defined tolerance. Using this method, the deviations in beta functions can be reduced to a goodness-of-fit of $< 10^{-5}$. This result is consistent across all kickmaps tested. This a goodness-of-fit of $< 10^{-5}$ is not limited by numerical precision. In measurement, however, as will become clear later in this chapter, it is a smaller deviation than can be reliably measured on the machine using either method (LOCO or pingers). For this reason, pursuing even better corrections in simulation was not deemed worthwhile. An important note is that this near-perfect restoration of model beta functions is only achievable when neglecting all other parameters. For instance, the tunes $Q_{x,y}$ are altered in the process of changing quadrupole strengths, which is discussed further in Section 4.1.2.

The main results of the local compensation using kickmaps in simulation are presented in Table 4.1. In this table, the RMS β-beats caused by kickmaps representing each ID at the 3 GeV ring without any corrections are presented, as well as the determined gradient changes required to eliminate them. The kickmap configurations presented in this table correspond to those that cause the greatest deviations in beta functions. From Table 4.1, it becomes clear which IDs cause the greatest deviation from model optics. In the horizontal, the planar devices have very little effect on the beta functions, whereas the EPUs cause deviations in excess of 3% RMS. In the vertical, the situation is different. Here, each ID has noticeable effect. BALDER causes the largest deviations, with an RMS > 4%, followed by the EPUs, MicroMAX and ForMAX. The similarity between the two pairs SoftiMAX-VERITAS and BioMAX-NanoMAX are expected, since the physical insertion devices are identical. In simulation, this means that the IDs in these respective pairs are represented by the same kickmap. The small differences within these pairs are likely a consequence of the placement of the kickmaps in different straight sections in the model.

Using the results from these simulations, FF-tables were constructed, as described in Section 3.2.1. Evaluations performed on the machine of these FF-tables are discussed in Sections 4.1.3 and 4.1.4, using LOCO and pinger measurements, respectively.

4.1.2 Simulated global compensation to eliminate the tune shift

In simulation, the model tunes on the 3 GeV ring can be restored to within 4-decimal numerical precision using either pair of families (QDE/QFE or QF/DIPc). The results of this investigation are presented in Table 4.2, which, for the same kickmaps as in Table 4.1, lists the tunes with just the kickmap in place in the model (NC), as well as after the corresponding local correction has been performed (LC). The last two columns show the required gradient changes to eliminate the tune shift after the local compensation has been performed. The model tunes are $Q_x = 0.2000$, $Q_y = 0.2800$. The results are in line with those presented in Table 4.1, where IDs responsible for a large deviation in vertical beta functions cause a larger tune shift in the vertical. The same holds true for the horizontal. The reason for the omission of the QF/DIPc pair in Table 4.2 is that the DIPc magnets are multipoles, and the simple optimization algorithm used to find the correction terms only changes the quadrupole component of magnets used in the compensation. Physically, changing the strength of a magnet corresponds to changing the amount of current running through it. Changing the current through a multipole magnet, even the pole-face strips DIPc, cannot be assumed to affect only the quadrupole field component.

The global correction changes the overall optics of the ring, which effects new β-beats. These, however,

Table 4.1: RMS β-beats in simulation caused by the presence of a kickmap representing a particular ID at a minimum-gap configuration, as well as the gradient change required to correct said β-beats. The region between flanking QDE magnets has been omitted in the calculation of β-beats. Note that BALDER, due to its asymmetrical mounting, has two Δk-values listed for each family. These correspond to the upstream/downstream values, respectively. Moreover, "Phase" only applies to EPUs.

ID	Gap/Phase [mm]	RMS$\left(\frac{\Delta\beta_x}{\beta_x}\right)$ [%]	RMS$\left(\frac{\Delta\beta_y}{\beta_y}\right)$ [%]	Δk_{QDE} [%]	Δk_{QFE} [%]
BALDER	4.5	0.0025	4.1166	1.774/0.361	0.503/0.021
BioMAX	4.5	0.0008	0.8399	0.353	0.087
CoSAXS	4.5	0.0011	0.2860	0.123	0.030
DanMAX	4.0	0.0032	0.4115	0.176	0.043
ForMAX	4.5	0.0010	1.1378	0.469	0.115
HIPPIE	13.5/20	3.4649	1.3638	0.966	0.895
MicroMAX	4.5	0.0007	1.2521	0.514	0.127
NanoMAX	4.5	0.0008	0.8398	0.353	0.087
SoftiMAX	12.5/20	3.8025	1.4294	1.028	0.972
VERITAS	12.5/20	3.8020	1.4289	1.028	0.972

Table 4.2: Tune shifts in simulation caused by the presence of kickmaps. NC refers to the case where no corrections have been applied, and LC refers to the state of the simulation after the local correction. The two right-most columns indicate the required gradient changes to recover model tunes after the local correction has been performed.

ID	Gap/Phase [mm]	Q_x (NC)	Q_x (LC)	Q_y (NC)	Q_y (LC)	Δk_{QDE} [%]	Δk_{QFE} [%]
BALDER	4.5	0.2000	0.1995	0.2889	0.3009	-0.093	-0.024
BioMAX	4.5	0.2000	0.1998	0.2817	0.2858	-0.026	-0.007
CoSAXS	4.5	0.2000	0.1999	0.2805	0.2819	-0.009	-0.002
DanMAX	4.0	0.2000	0.1999	0.2808	0.2828	-0.013	-0.003
ForMAX	4.5	0.2000	0.1997	0.2824	0.2877	-0.035	-0.009
HIPPIE	13.5/20	0.1926	0.2011	0.2829	0.2878	-0.039	-0.015
MicroMAX	4.5	0.2000	0.1997	0.2827	0.2885	-0.038	-0.010
NanoMAX	4.5	0.2000	0.1998	0.2817	0.2858	-0.026	-0.007
SoftiMAX	12.5/20	0.1919	0.2012	0.2830	0.2881	-0.041	-0.016
VERITAS	12.5/20	0.1919	0.2012	0.2830	0.2881	-0.041	-0.016

are smaller than the deviations in beta functions without the local correction. Thus, by performing local and global compensations in combination, the machine's operating point can be restored, while the deviations in β-functions are drastically reduced. To illustrate this, consider Figure 4.1, which shows β-beats obtained from a kickmap representing BALDER at a minimum gap of 4.5 mm in four stages of compensation: no compensation, local compensation (LC), global compensation (GC) using either QFE/QDE or QF/DIPc. In this plot an interesting difference in the β-beats resulting from the global compensation becomes apparent: using QDE/QFE, there are positive deviations in the long straight sections, and a small negative dip on either end of an achromat, whereas using QF/DIPc there is little to no deviation in the long straights, followed by a "hill" throughout the dipoles.

The impact was not studied further, and it should be mentioned that while feedforward tables for a global compensation were generated for each ID, they were not tested extensively. There are a few reasons for this: first, the established tune feedback system works well and is employed in normal operation, meaning that the priority of establishing a tune-FF system is lower than establishing a β-function FF. The second reason is a change in priorities to the development of the pinger measurements, which allowed for a model-independent non-destructive approach to determine β-functions in the BPMs of each ring. Third, there were concerns pertaining to server overloads in the current implementation of FF-tables in the machine control system. In its current implementation, the IDFF-device would need to write to $2 \times 21 = 42$ actuators at a frequency of 10 Hz. With all tune-FF devices active simultaneously, this would likely lead to an overflow.

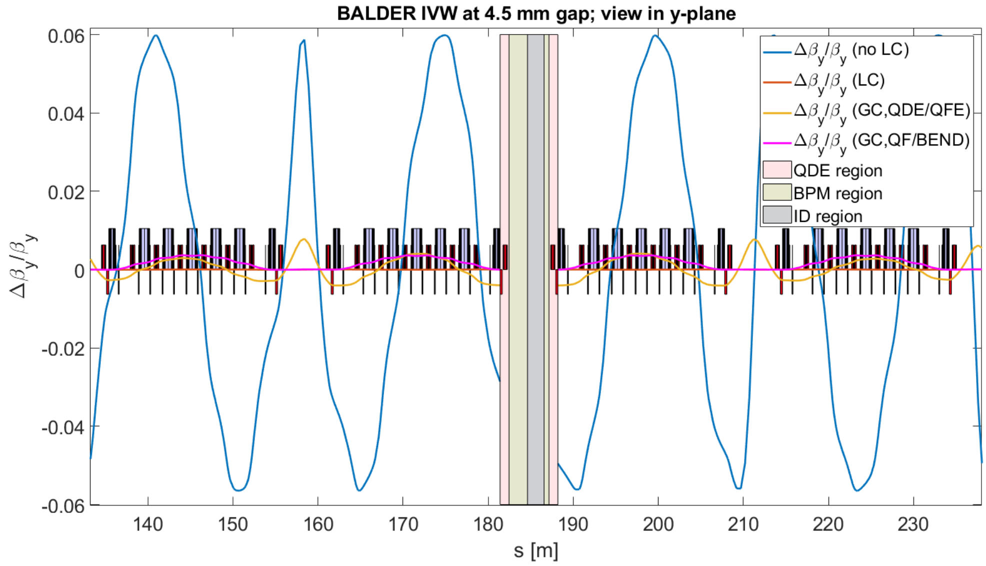

Figure 4.1: Plots of simulated β-beats caused by the presence of a kickmap representing BALDER at 4.5 mm gap, before and after different compensations. The different curves correspond to β-beats after different compensations. For clarity, the horizontal axis spans two achromats in either direction around the BALDER ID. Furthermore, relevant regions (between QDEs, between flanking BPMs and the extent of the ID itself) are been indicated. Note that the β-beats between the flanking QDE pairs have been omitted.

4.1.3 Evaluation of the kickmap-based local compensation on the 3 GeV ring using LOCO

This section is dedicated to the evaluation of the FF tables determined using kickmaps using the conventional optics diagnostics and correction tool, LOCO, as well as a discussion about the effectiveness of LOCO itself as a tool to determine β-functions, especially in the presence of IDs. The main results of the evaluation of the kickmap-based correction terms are presented in Table 4.3, presenting the RMS β beats for each ID at the minimum-gap configuration with and without the compensation active. Ideally, the compensation would entirely eliminate all deviations in beta functions, as it does in simulation. Averaging the change in RMS($\Delta\beta/\beta$) after and before applying the local compensation over all IDs gives -8.7% in the horizontal, and -23.8% in the vertical. While this corresponds to an improvement overall, the LOCO evaluation of the kickmap-based local compensation on the machine itself suggests that the correction terms found in simulation are insufficient to eliminate deviations in β functions in the ring caused by the presence of IDs.

There were issues with some of the measurements behind the data in Table 4.3. In particular, those marked with an asterisk * showed unexpectedly large deviations in beta functions. Likely, the systematic large deviations in β functions stem from an error in configuring the measurement and/or fit, as discussed in Section 3.3.2. One aspect is that all measurements presented in Table 4.3 were performed at a higher beam current of 30 mA compared to the typical low beam current of 3 mA used in LOCO measurements, in an attempt to increase signal-to-noise in the BPMs. While the increased current as such does not pose an immediate issue, it necessitates a number of adjustments to measurement parameters with respect to the typical 3 mA case, such as a reduced RF variation for the dispersion measurement. At higher currents, a large variation in RF frequency can lead to large beam losses, making the measurement unusable. While the exact causes of the poor measurements remain unknown, it is clear that the results should be viewed with skepticism. Unfortunately, due to machine time constraints and shifting priorities within this project, there was no opportunity to repeat these measurements.

Table 4.3: Evaluation of kickmap-determined corrections for the insertion devices using LOCO at the 3 GeV ring in the both planes. RMS values are given in per cent. NC refers to the case of no compensation, and LC refers to the case of local compensation. The columns "change" are calculated as the percentage difference between the cases of local and no compensation. IDs marked with an asterisk * represent measurements which are suspected to carry large systematic errors.

ID	Gap/Phase [mm]	$RMS\left(\frac{\Delta\beta_x}{\beta_x}\right)$ (NC)	$RMS\left(\frac{\Delta\beta_x}{\beta_x}\right)$ (LC)	Change [%]	$RMS\left(\frac{\Delta\beta_y}{\beta_y}\right)$ (NC)	$RMS\left(\frac{\Delta\beta_y}{\beta_y}\right)$ (LC)	Change [%]
BALDER	4.5	0.291	0.503	+73	2.681	1.266	-53
BioMAX*	4.5	1.490	1.155	-22	1.334	1.340	±0
CoSAXS*	4.7	1.595	1.476	-7	0.671	0.634	-6
DanMAX	4.4	0.975	0.963	-1	1.528	1.130	-26
ForMAX*	4.5	2.011	1.964	-2	1.296	1.031	-20
HIPPIE	14.0/20	3.022	1.904	-37	0.560	0.484	-14
MicroMAX*	5.0	1.995	2.405	+21	1.830	1.025	-44
NanoMAX*	4.5	4.327	2.327	-46	1.474	1.167	-21
SoftiMAX	12.5/20	4.288	2.998	-30	1.282	0.823	-36
VERITAS	12.5/20	3.242	2.066	-36	1.847	1.506	-18
Average change:				-8.7			-23.8

Table 4.4: Evaluation of kickmap-determined corrections for the insertion devices at the 3 GeV ring in the horizontal plane using *pinger* measurements of β functions. This table was placed on the same page as Table 4.3 for convenience.

ID	Gap/Phase [mm]	$RMS\left(\frac{\Delta\beta_x}{\beta_x}\right)$ (NC)	$RMS\left(\frac{\Delta\beta_x}{\beta_x}\right)$ (LC)	Change [%]	$RMS\left(\frac{\Delta\beta_y}{\beta_y}\right)$ (NC)	$RMS\left(\frac{\Delta\beta_y}{\beta_y}\right)$ (LC)	Change [%]
BALDER	4.5	0.528	0.522	-1	2.737	1.528	-44
BioMAX	4.5	0.241	0.278	+15	1.020	0.802	-21
CoSAXS	4.7	0.154	0.254	+65	0.977	0.924	-5
DanMAX	4.4	0.516	0.494	-4	1.235	1.129	-9
ForMAX	4.5	0.934	0.931	±0	1.032	0.699	-32
HIPPIE	14.0/20	2.629	1.393	-47	0.772	0.519	-33
MicroMAX	5.0	0.554	0.572	+3	0.880	0.656	-25
NanoMAX	4.5	1.871	1.873	±0	0.728	0.592	-19
SoftiMAX	12.5/20	3.570	1.871	-48	1.058	0.706	-33
VERITAS	12.5/20	2.860	1.562	-45	0.923	0.621	-33
Average change:				-6.2			-25.4

To illustrate how a local compensation might affect the β beats measured on the ring, consider Figure 4.2, which shows the effect of the kickmap-based correction terms on β beats caused by BALDER at a gap of 4.5 mm. These are the measurements that provided data for the first rows of Table 4.3. It is clear that the correction does reduce β-beating, particularly in the vertical. The horizontal case is different, however, showing an increase in RMS β beating of $> 70\%$. A possible partial explanation for this behavior is that measurement errors (BPM σ_i as in eq.3.1) in the horizontal are relatively more pronounced than in the vertical, reflected in the fit as a larger deviation. Nevertheless, the possibility of horizontal β beating increasing as a consequence of the application of the correction cannot be dismissed. The measurements were taken in succession, with no changes to the machine other than the application of the correction.

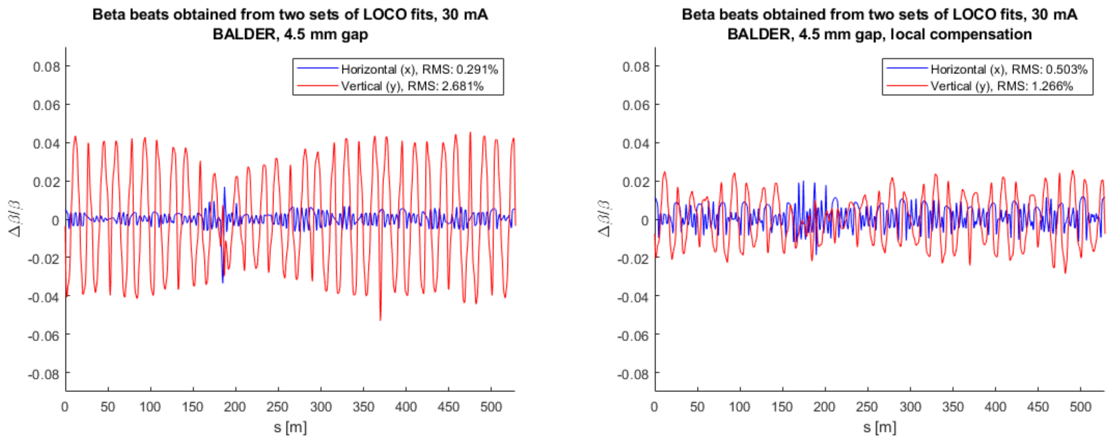

Figure 4.2: Combined LOCO measurements showing β beats in both planes before (left) and after (right) the application of kickmap-based correction terms to BALDER at 4.5 mm gap. Note that data in the region around $s = 185$ m has been interpolated, and may not be reliable.

The region around $s = 185$ m corresponds to the immediate vicinity of the relevant ID itself. Data in this region has been interpolated and may not be entirely reliable.

The left plot in Figure 4.2, the uncorrected case, provides an interesting point of comparison to Figure 4.1, in particular the latter's graph displaying $\Delta\beta_y/\beta_y$ before the application of any correction. This comparison is a direct comparison between expected behavior from theory and measurement. Of note are the peak amplitudes. In simulation, using kickmaps, peak amplitudes of almost exactly 6% are obtained, with peaks (and troughs) following constant envelopes across the entire lattice. The LOCO measurement shows peak amplitudes just in excess of 4% rather than 6%, and the envelope varies across the lattice, notably forming a "funnel" going either direction along s, centered around the ID itself. The observation regarding peak amplitudes of β beats will be revisited within the framework of pinger measurements, though the "funnel" visible in Figure 4.2 has a likely explanation. In the LOCO fits performed in Figure 4.2, as mentioned in Section 3.3, throughout this chapter, insertion devices are modeled as an array three thin pure quadrupoles in series along the longitudinal direction. The "funnel" in the amplitude envelope is likely an artifact produced in the LOCO fit owing to the fact that the representation of the ID is fundamentally different in the LOCO model.

Another observation that can be made in relation to Figures 4.2 and 4.1 is that the corrections found using kickmaps in simulation were such that they would correct β beats with peak amplitudes of 6%, whereas LOCO only measured 4%. Despite this, the correction, when applied to the physical machine, only reduced the observed $\Delta\beta_y/\beta_y$ by about a factor of 2. This indicates that the correction applied to quadrupole gradients on the machine was *insufficient* in terms of magnitude, rather than excessive, which is surprising, given that in simulation, the corrections should be sufficient for peak deviations of 6% rather than the measured 4%.

No matter the reason for this seeming insufficient magnitude of correction terms, an attempt was made to improve the correction. Under the assumption that the corrections still work in the linear regime, given that the first corrections yielded a reduction in vertical β-beating of 50% RMS, an attempt was made to simply double the correction terms, to see whether this would eliminate vertical beating entirely. Figure 4.3 shows LOCO measurements of corrected β beats. The left-hand plot is the same as the right-hand plot in Figure 4.2. The right-hand plot is the measurement resulting from doubling the correction terms obtained from simulation. While there is arguably some improvement after when doubling the correction, with vertical RMS β beating being reduced from 1.266% to 1.227%, it is clear that this change was unable to eliminate β beating altogether. The inability of both sets of correction terms to eliminate beating could be indicative of two effects: first, the optics have left the linear regime, meaning that there are higher-order effects at play, which would complicate finding correction terms considerably. Second, the kickmap does not represent the effects of the ID on the optics well. It should be noted that the two sets of measurements presented in Figure 4.3 were performed at different beam currents, 30 mA

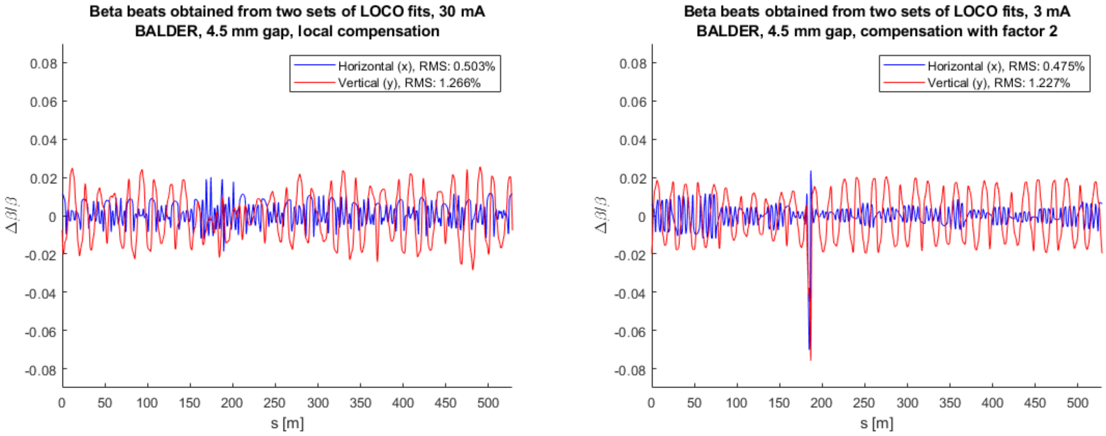

Figure 4.3: Corrected β-beats induced by BALDER at 4.5 mm gap measured using LOCO. Left: using correction terms found from simulation. Right: using the same correction terms scaled by a factor of 2.

in the case of the kickmap-true correction and 3 mA in the case of the doubled correction.

4.1.4 Evaluation of kickmap-based local compensation on the 3 GeV ring using pingers

This subchapter is dedicated to the evaluation of the correction terms determined from kickmaps in simulation (see Table 4.1 for reference) by the use of pinger measurements to determine β functions. It is intended as a complement to the previous section, 4.1.3, in an attempt to evaluate the same correction using two fundamentally different methods.

The main result of the evaluation of correction terms found using kickmaps is presented in Table 4.4, showing changes to RMS β beating in the horizontal and vertical following the application of the correction, respectively. The overarching conclusion of this pinger-based evaluation is similar to that of the LOCO-based evaluation, in that the average horizontal β beating decreases slightly at 6.2%, and the vertical β beating decreases by more, at 25.4%. The corresponding values using LOCO are 8.7% and 23.8%, respectively. Notably, the pinger-based evaluation suggests that there is little to no improvement in the horizontal plane except for the EPUs, which are known to cause the largest deviations in horizontal β functions at the used phase. In the case of CoSAXS, the measurement even suggests a large increase in β beat of 65% as a result of the compensation, though using relative errors might be unrepresentative given the small horizontal deviations caused by this ID. In the vertical plane, the pinger measurements indicate an improvement for each ID. This is similar to the LOCO evaluation, which also indicated an improvement in the vertical for every ID except for BioMAX, which showed no appreciable change. The overarching conclusion of the pinger-based evaluation of kickmap-based correction terms is that there is an appreciable decrease in β-beating for most IDs in the vertical plane. However, the pinger measurements indicate that in the horizontal, there is no appreciable change in β beating except for the EPUs, which each shows considerable improvement of $\geq 45\%$.

4.2 Online ID compensation using pinger measurements

4.2.1 Simulation

To test the functionality of the online local compensation algorithm as described in Section 3.5, simulations were performed before machine tests. Simulations of the online compensation algorithm were performed on models of both storage rings. Initial tests were performed in the absence of IDs; instead, quadrupole errors were introduced in the magnets involved in local ID compensation. The algorithm successfully corrected these errors and was able to restore model β functions. As a further test, the algorithm was tested on the model lattice of the 3 GeV ring with the inclusion of IDs modeled using the same kickmaps used to determine the correction terms presented in Table 4.1, in the hopes that the correction terms determined using the online algorithm would be identical. The main result of this investigation is that the online compensation algorithm is able to reproduce the correction presented in Table 4.1 to

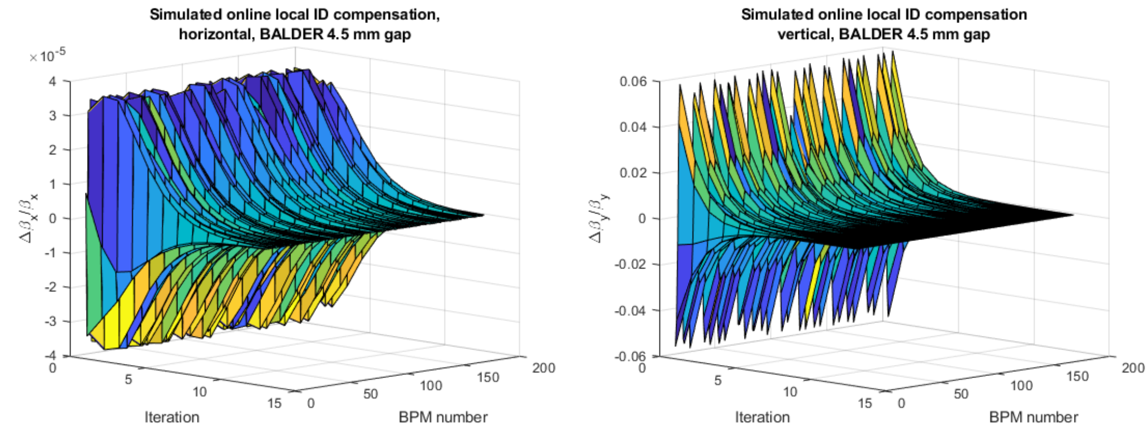

Figure 4.4: Simulated online compensation on the lattice model of the 3 GeV ring with a kickmap representing the IVW at the BALDER beamline. Note that the vertical are scaled differently between the plots. Left: Horizontal plane. Right: Vertical plane. The β beats in Iteration 1 are those determined before the application of any corrections. Data from the two BPMs within the region spanned by the outermost QDE pair has been omitted.

within the last decimal, the difference likely owing to numerical approximations.

To illustrate how the online compensation algorithm affects β beating over a number of corrective iterations, consider Figure 4.4, which shows a simulation of an online local compensation performed on a kickmap representing the IVW at the BALDER beamline on the 3 GeV ring. As can be seen in this idealized simulation of the algorithm, β beating, in both planes, are reduced each iteration, until eventually converging to RMS($\Delta\beta/\beta$)$< 10^{-4}$ in each plane. As mentioned in Section 3.5, the algorithm is designed to be iterative to provide a degree of resistance to noise in the β function measurements on the machine. This was achieved by scaling the each of the calculated corrections determined by a factor $0 < m < 1$. Scaling the corrections proved significant in simulation as well, as the algorithm is unable to converge if unaltered correction terms are applied directly. Remarkably, despite deviations in β_x being about three orders of magnitude smaller than deviations in β_y, the algorithm successfully reduces the horizontal β beating as well. Another feature visible in Figure 4.4 is that the vertical β beating appears to converge in fewer iterations than the horizontal, with respect to their respective maxima. The reason for this is not entirely known, though the relatively small magnitude of β function deviation compared to the vertical is likely part of the explanation.

Another difference between the online algorithm and the optimization based on the Nelder-Mead algorithm used to obtain the correction terms in Table 4.1 is the amount of "input data". While the Nelder-Mead simplex uses a goodness-of-fit based on the deviations in β functions in every reference point of the model lattice excluding the region between the flanking QDE magnets, the online algorithm, even in simulation, uses only data obtained from the reference points representing BPMs, also excluding the region between flanking QDE magnets. On the 3 GeV ring, the goodness-of-fit used in the Nelder-Mead simplex is based on > 7300 reference points, whereas the online algorithm uses a mere 195 in simulation (196 on the machine; the model is somewhat outdated). Despite this difference, the online algorithm is able to reproduce the results of Table 4.1.

Simulating the online compensation algorithm on the 1.5 GeV ring

The online algorithm was tested in simulation on the 1.5 GeV ring as well. During this project, machine time was generally more readily available on the 1.5 GeV ring, and as part of the development of the online compensation algorithm, machine tests were performed on this ring as well, which is discussed in Section 4.2.2. Prior to machine testing, it was deemed necessary to simulate the algorithm on this ring as well. Figure 4.5 shows a simulated online correction performed on the EPU at the BLOCH beamline, at 15 mm gap and 23.7 mm phase, a configuration known to have a large impact on optics.

The algorithm reduces RMS($\Delta\beta/\beta$) from 11% to 0.12% in the horizontal plane, and from 7.9% to 0.09%

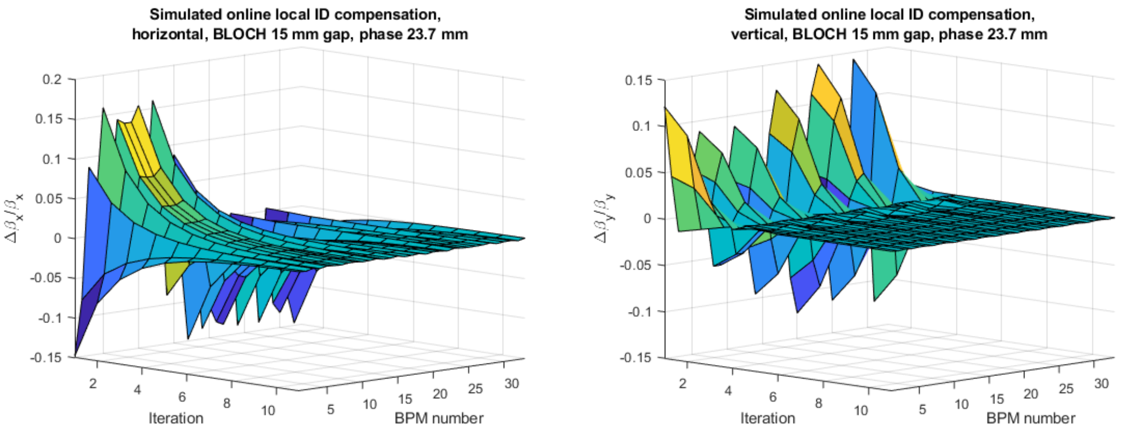

Figure 4.5: Simulated online compensation on the lattice model of the 1.5 GeV ring with a kickmap representing the EPU at the BLOCH beamline. Note that the vertical axes are **not** shared between the plots. Left: Horizontal plane. Right: Vertical plane. The β beats in Iteration 1 are those determined before the application of any corrections. Data from the two BPMs within the region spanned by the outermost DIPc pair has been omitted.

in the vertical. An improvement of around two orders of magnitude is a similar result to that obtained for the 3 GeV ring. A similar behavior to that seen in Figure 4.4 is the convergence of β beating in the vertical in fewer iterations than in the horizontal. This difference cannot be explained by an order-of-magnitudes difference in pre-correction β beating. A possible explanation comes from the response matrix presented in Figure 3.9, where it can be seen that the vertical-DIPc quadrant has significantly higher peaks than the horizontal-SQFO quadrant. The calculated correction terms would in this explanation be "weighted" more strongly towards correcting the vertical β deviations.

Overall, a reduction of RMS β beating of around two orders of magnitude can be considered a favorable result for the algorithm. The final RMS β beats of 0.12% and 0.09% in the horizontal and vertical are worse than what can be achieved using the Nelder-Mead simplex described in Section 4.1.1. However, these values are of the same magnitude as the resolutions determined for pinger measurements in Figure 3.6 (around 0.1%). Given the measurement resolution, the simulated compensation was deemed sufficient. Moreover, the ability of the algorithm to reproduce the corrections presented in Table 4.1 to within the last decimal speaks for its viability.

4.2.2 Testing on the 1.5 GeV ring

This section is dedicated to machine tests of the online local ID compensation algorithm performed on the EPUs at the BLOCH and FINEST beamlines at the 1.5 GeV ring. These IDs are known to have a large impact on optics. An important limitation of the measurements presented in this section is that, as mentioned in Section 3.4.2, they were performed before the optimization of pinger-based β function measurements. In fact, the online correction measurements presented in this section prompted the investigation into the stability and resolution of pinger-based β function measurements. In particular, vertical β function measurements appear noisy compared to the horizontal, and the compensation consequently ineffective. Due to time constraints and machine availability, repeat measurements using the optimized pinger configurations determined in Section 3.4.2 were not possible.

Testing on the EPU at the BLOCH beamline

In this test, the EPU at the BLOCH beamline was configured to a gap of 14 mm and a phase of 0 mm. Figure 4.6 shows the measured β beats as functions of the number of corrective iterations. It can be seen that the deviations in β functions caused by the ID at this configuration are more significant in the vertical than in the horizontal. Furthermore, the data is rather noisy, especially in the vertical. Nevertheless, particularly in the vertical, it appears as though the peak deviations in β functions are reduced somewhat as the number of iterations are increased. It should be noted that only two BPMs have been excluded from these plots, rather than the 8 that fell in the region subject to timing issues mentioned

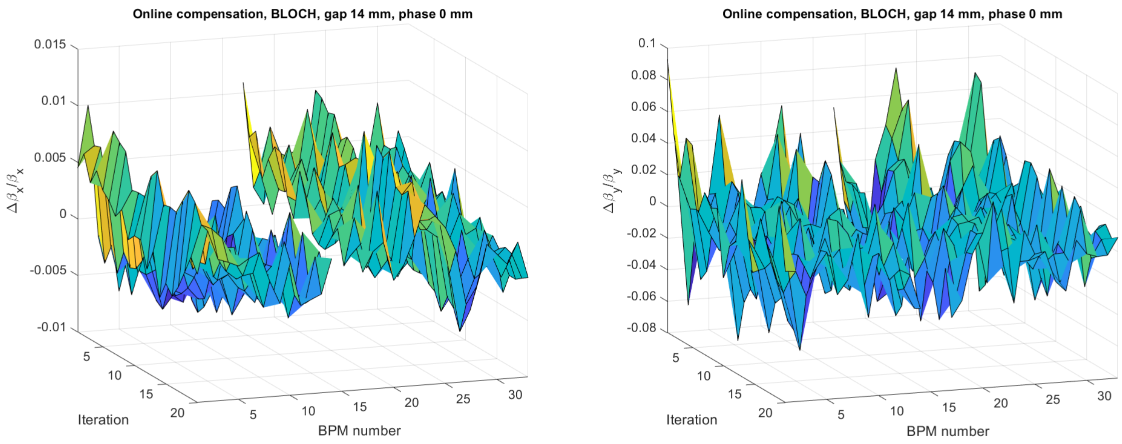

Figure 4.6: Online local ID compensation performed on the EPU at the BLOCH beamline at 14 mm gap and a helical phase of 0. Left: horizontal plane. Right: vertical plane. Note that the vertical axes are not shared between the plots, and that BPMs number 15, 16 have been identified as being subject to timing issues and omitted.

in Section 3.4.2. The reason for the stems from a preliminary BPM calibration which was unfortunately lost during an update.

To quantify the effects of the online compensation algorithm, consider Figure 4.7, showing RMS($\Delta\beta/\beta$) in each plane as a function of the number of iterations. The errors indicated in the horizontal are determined from Figure 3.6. The vertical error is an estimate, as the vertical pingers were configured to $V_V = 600\,\mathrm{V}$ during the BLOCH measurement, lower than any of the investigated voltages in Figure 3.6. In the right-hand plot, the vertical errors indicated are 0.35%, an overestimate compared to all values presented in Figure 3.6. In Figure 4.7, it can be seen that RMS β beating does decrease overall over the first seven iterations. In the subsequent iterations, no marked improvement can be discerned. An exponential fit $\Delta\beta/\beta = ab^n + c$, where n denotes the number of iterations, a the initial deviation over noise, b the relative decrease in RMS β beating each iteration, and c denoting a constant representing the noise floor, was performed on each data set. The fits overlap with data remarkably well in the vertical. The seeming exponential behavior of the RMS β beating indicates that the β function response is indeed linear in the correction terms Δk. In this test, the correction terms were scaled by a factor of 0.2 to increase overall stability, as mentioned in Section 4.2.1. Thus, a fit parameter $b = 1 - 0.2 = 0.8$ is expected. In the vertical case, $b = 0.63$ is found. The reason for this is not entirely known, but could relate to the relatively high noise floor.

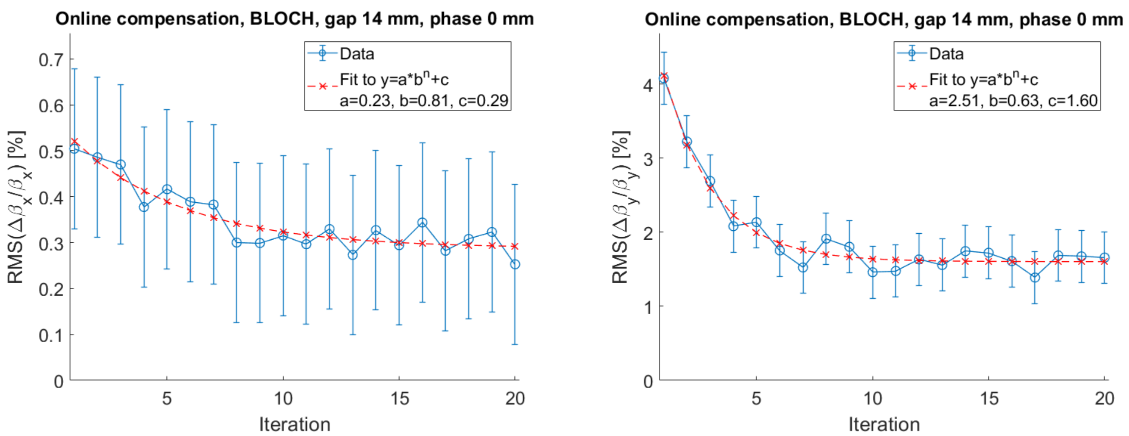

Figure 4.7: RMS β beating as a function of the number of corrective iterations during online local compensation of the EPU at the BLOCH beamline. Left: horizontal plane. Right: vertical plane. Note that the vertical axes are not shared between the plots.

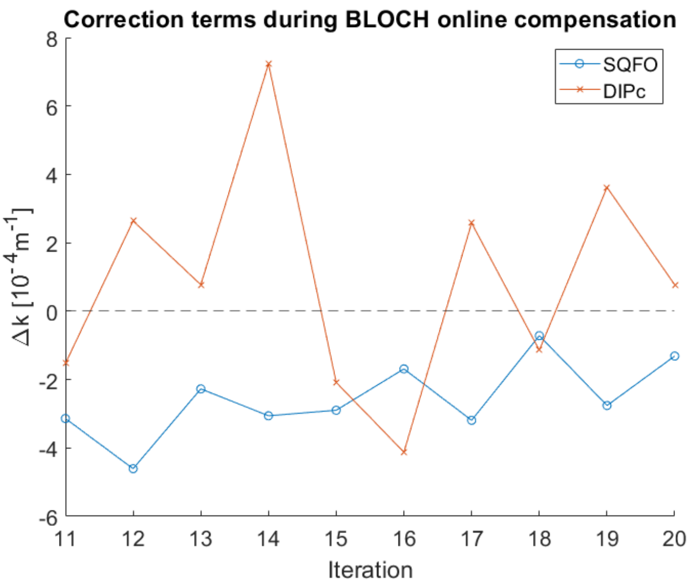

Figure 4.8: Determined correction terms over the last ten iterations of the online compensation presented in Figure 4.6. The dashed line represents the horizontal axis.

Notably, in the latter half of iterations, the RMS β beating does not monotonously decrease, or even stay constant. Instead, it appears to increase and decrease repeatedly. This alternating behavior is particularly apparent in the horizontal, and is interpreted as the convergence of the algorithm to the noise floor. Nevertheless, the online local compensation algorithm was able to reduce RMS β beating from initial values of 0.50% and 4.08% in the horizontal and vertical, respectively, to values of 0.29% and 1.60% obtained from the fit, corresponding to decreases of 42% and 60%. While the online compensation algorithm does improve β beating significantly in this case, its inability to reduce vertical RMS β beating below 1.6% was deemed unsatisfactory, and was investigated further. In particular, an understanding of whether noise in the vertical β function measurements was the limiting factor was deemed necessary. To address this matter, the correction terms found each iteration were investigated.

Figure 4.8 shows the correction terms Δk determined and applied by the algorithm over the last ten iterations, i.e. where no improvement is discernible. What stands out is that the correction terms for the magnets mainly involved in the horizontal compensation (see response matrix in Figure 3.9 for reference), SQFO, do not change sign. The corrections for the DIPc magnets, whose β function response is mainly in the vertical, do change sign several times. The interpretation of the alternating behavior of Δk_{DIPc} presented in Figure 4.8 is that this noise floor has been reached, supporting the notion that noise in the vertical β function measurement is the limiting factor of the compensation algorithm in this case. It should be emphasized that while noise cannot be avoided, reducing noise appears worthwhile in the effort of improving the effectiveness of the online compensation. Within the framework of this project, this discovery led to the optimization of β function measurements described in Section 3.4.2.

Testing on the EPU at the FINEST beamline

The other machine test of the online ID compensation algorithm performed on the 1.5 GeV ring was performed on the EPU at the FINEST beamline, at a configuration of 14 mm gap and 0 phase. There is a difference compared to the compensation test on the BLOCH beamline, in that the number of pings used in the β function measurement was increased from 10 to 15 in both planes. This increase in statistics is expected to result in less noise. Figure 4.9 shows the measured β beating as functions of the number of iterations. Comparing to Figure 4.6, the measurement appears much less noisy in the horizontal plane, likely owing to the improved statistics. Furthermore, it can be seen that this ID configuration has significant impact on horizontal β functions, with initial peak deviations in excess of 7%. Owing to the improved statistics (and likely the large deviations), it is visible that the online algorithm is able to reduce β beating. In the zero-phase configuration, the FINEST EPU is expected to behave much like a

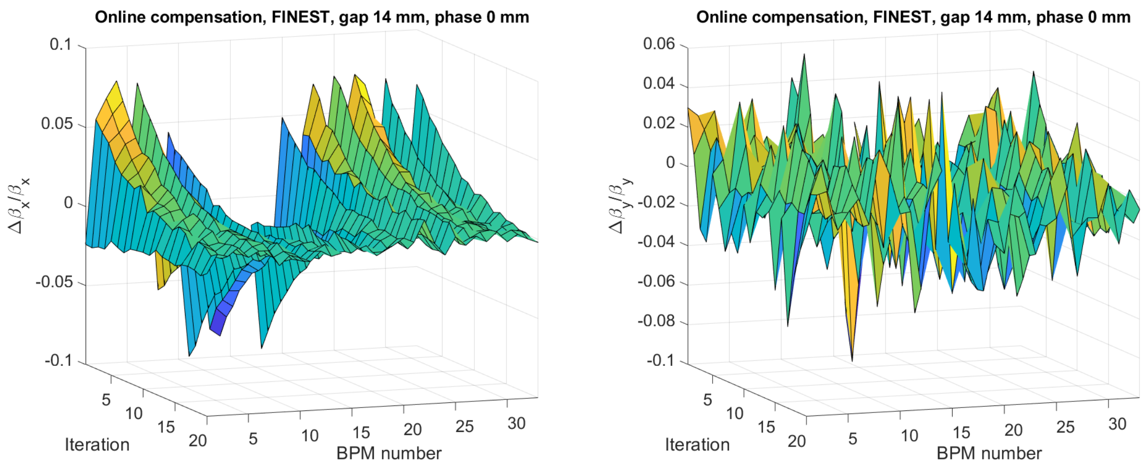

Figure 4.9: Online local ID compensation performed on the EPU at the FINEST beamline. Left: horizontal plane. Right: vertical plane. BPMs number 15, 16 have been omitted.

planar ID (Section 2.3.1), leading to small horizontal β beating. The cause of observed large deviations in horizontal β functions is not known.

To quantify the improvement, consider Figure 4.10, showing RMS($\Delta\beta/\beta$) as a function of the number of iterations. The horizontal case is particularly promising, where RMS β is reduced by 91%, from 5.73% to 0.48% in the last iteration. In the vertical, on the other hand, no real improvement is visible. It should be noted that this measurement appears limited by the number of iterations; it is not as clear that the algorithm has reached convergence, particularly in the horizontal. The constant fit parameter $c = 0.19$, while the last measurement of RMS($\Delta\beta_x/\beta_x$) is 0.48%. This discrepancy indicates that convergence has not been reached.

In the vertical, the compensation appears to be almost entirely ineffective. The RMS β beating seems to fluctuate around 2%, reaching a value of 1.76% in the final iteration. This value is slightly higher, though comparable to what was found in Figure 4.7, which reached an RMS($\Delta\beta/\beta = 1.63\%$), suggesting that the vertical compensation is at its noise limit. While the increased number of pings from 10 to 15 should improve statistics by a factor of $\sqrt{15/10} \approx 1.22$, the beam current had decreased to 2.4 mA during the FINEST measurement from 2.9 mA during the BLOCH measurement, a decrease of 17%. The decrease in current directly affects the signal-to-noise in the BPMs. Owing to the improvement in statistics of roughly 20% and a worsening of signal-to-noise of about 17%, it is assumed that the overall measurement

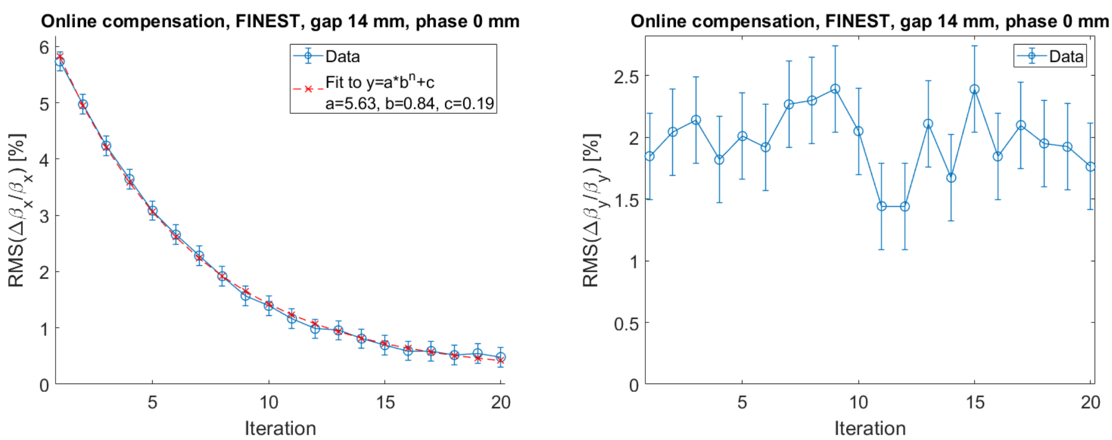

Figure 4.10: RMS β beating as a function of the number of corrective iterations n during online local compensation. Left: horizontal plane. Right: vertical plane. Note that the vertical axes are not shared between the plots.

quality is comparable to that of the BLOCH measurement. One factor that could influence the quality of β function measurements that should be mentioned are the capacitor banks in the pinger magnets, particularly whether these collect the same amount of charge between pings. An investigation into this detail of measurement hardware is outside the scope of this project. Given these two measurements on the 1.5 GeV ring, it appears that the overall noise floor for the RMS vertical β beating is no less than 1.6-1.8% at this initial pinger configuration.

As a final note in this section, it should be emphasized that reaching RMS($\Delta\beta_x/\beta_x$)$<$ 0.5% given an initial deviation of almost 6% can be considered a success for the online compensation algorithm. The results in the vertical plane are less impressive. The two most likely explanations are that either the noise floor was reached, or something went wrong with the vertical β function measurement. It would seem advisable to repeat the online compensation on the FINEST EPU. Nevertheless, the BLOCH measurement shows that the compensation is sufficiently robust to achieve a degree of compensation in the vertical, provided that deviations are sufficiently large compared to noise. Furthermore, identifying noise in measurements of β_y as the limiting factor is a useful result in understanding the capabilities of the online compensation algorithm, and led to a first attempt at optimizing the algorithm, as described in Section 3.4.2. Lastly, the FINEST measurement provided the insight that a correction scaling factor of 0.2 might be too low to reach convergence of the algorithm in 20 iterations. In the interest of time, it is worth investigating whether a similar degree of compensation can be achieved with a larger scaling factor.

4.2.3 Testing on the 3 GeV ring

In this section, the only (limited by machine time availability) machine test of the online compensation algorithm performed on the 3 GeV ring is discussed. It was carried out on the IVW at the BALDER beamline, at the minimum 4.5 mm gap. In contrast to the measurements presented in Section 4.2.2, these tests were carried out after the noise optimization described in Section 3.4.2. Moreover, the scaling factor was increased from 0.2 to 0.5. For this ID, the vertical plane is of particular interest (see Tables 4.3/ 4.4 for reference); Figure 4.11 shows measured vertical RMS β beats as functions of the number of iterations. A comparison can be made between Figure 4.11, and the right-hand plot in Figure 4.4, the simulated online compensation modeling the ID as a kickmap. The most striking difference is the peak deviation. Using the kickmap representation, peak deviations of just under 6% are expected, whereas the measured peak deviations do not exceed 5%. This discrepancy between model and measurement raises raises some questions regarding the accuracy of the kickmap to model the effects of the ID on optics.

To quantify the improvement, consider Figure 4.12, showing the RMS β beats during the compensation, as

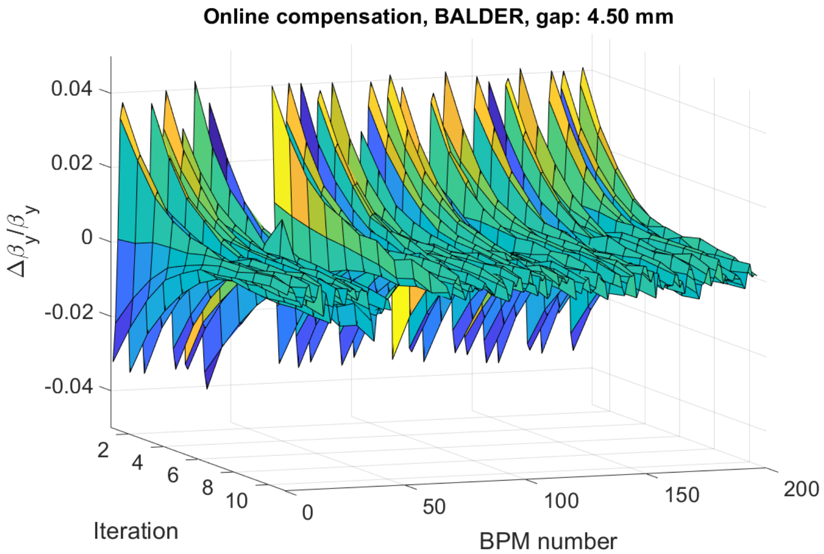

Figure 4.11: Vertical β beating during online compensation performed on the IVW at the BALDER beamline. The faulty BPM region spans BPMs numbered 45-66. Two more outliers have been omitted.

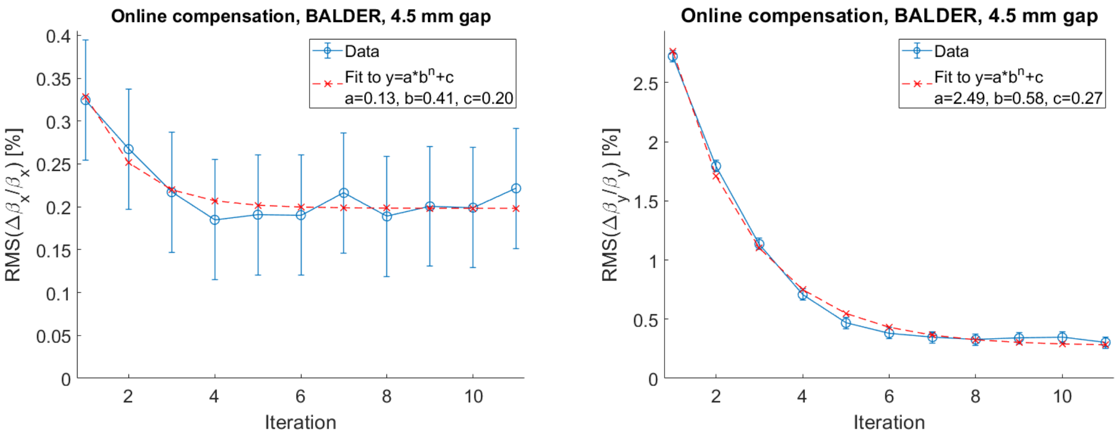

Figure 4.12: RMS β beating as a function of the number of corrective iterations during online local compensation. Left: horizontal plane. Right: vertical plane. Note that the vertical axes are not shared.

functions of the number of iterations. The errors indicated in this Figure were obtained from Figure 3.6, using a pinger configuration of $V_H = 1500\,\mathrm{V}$ and 25 pings. Exponential fits were carried out as in Section 4.2.2. In the vertical, the algorithm appears to converge well within the performed 11 iterations, to a value of RMS($\Delta\beta_y/\beta_y = 0.30\%$), down from an initial 2.72%, corresponding to an improvement of 89%. Measurements indicate an improvement in the horizontal as well, from 0.32% to 0.22% in the last iteration, a reduction of 31%. Regarding the fits, as the scaling was set to 0.5, a base of $b = 0.5$ is expected. In the vertical, the fit suggests a base of $b = 0.58$, whereas the horizontal fit gives $b = 0.41$. It should be noted that the horizontal fit likely suffers from the fact that the initial deviations of 0.33% are less than a factor 2 greater than the noise floor of 0.2%, as determined from the fit. The results obtained in this study are promising for the online compensation algorithm. Furthermore, the discrepancy in $\Delta\beta_y/\beta_y$ between the kickmap-based simulation and model-independent measurement constitutes a point that should be investigated further.

4.3 Comparison of kickmap-based and online-based corrections in experiment

This section is dedicated a comparison of the effectiveness of the kickmap-based and online local ID compensations employed in this project. Due to limitations in machine time availability, the only ID on which both methods were tested was the IVW at the BALDER beamline. Figure 4.13 shows three pinger measurements of vertical β beating performed with the BALDER IVW at a minimum gap of 4.5 mm. The three cases covered are: no compensation, kickmap-based compensation, and online compensation. Referring to the legend of Figure 4.13, the "Kickmaps" measurement is that presented in Table 4.4. The "No correction" and "Online" measurements correspond to iterations 1 (meaning before the application of any quadrupole strength adjustments) and 11 (the final iteration), respectively, of Figure 4.11.

From Figure 4.13, it is clear that the online compensation algorithm outperforms the kickmap-based compensation. As mentioned in Section 4.2.3, RMS($\Delta\beta_y/\beta_y$) is reduced from an initial 2.72% to 0.30% using the online algorithm, compared to 1.53% using the kickmap-based compensation. While, as discussed in Section 3.4.2, it is known that BPM noise is a limiting factor of the pinger-based β function measurements, a remaining post-correction RMS($\Delta\beta/\beta$) of 1.53% is significantly more than can be explained by noise alone, indicating that there is a discrepancy between the effects of the ID predicted by the kickmap, and the effects observed on the machine. This discrepancy should be investigated further.

Another interesting observation lies in considering the quadrupole adjustments $\Delta k_{QFE/QDE}$ determined by the two different methods to eliminate the β beating. Table 4.5 shows the quadrupole adjustments as determined through the two different methods. As can be seen in Table 4.5, the determined quadrupole adjustments are rather similar between the two methods. Specifically, there is good agreement in the determined Δk_{QFE}, and both methods find that the upstream Δk_{QDE} needs to be subject to the largest change. The similarity between the determined quadrupole adjustments indicates that the discrepancy

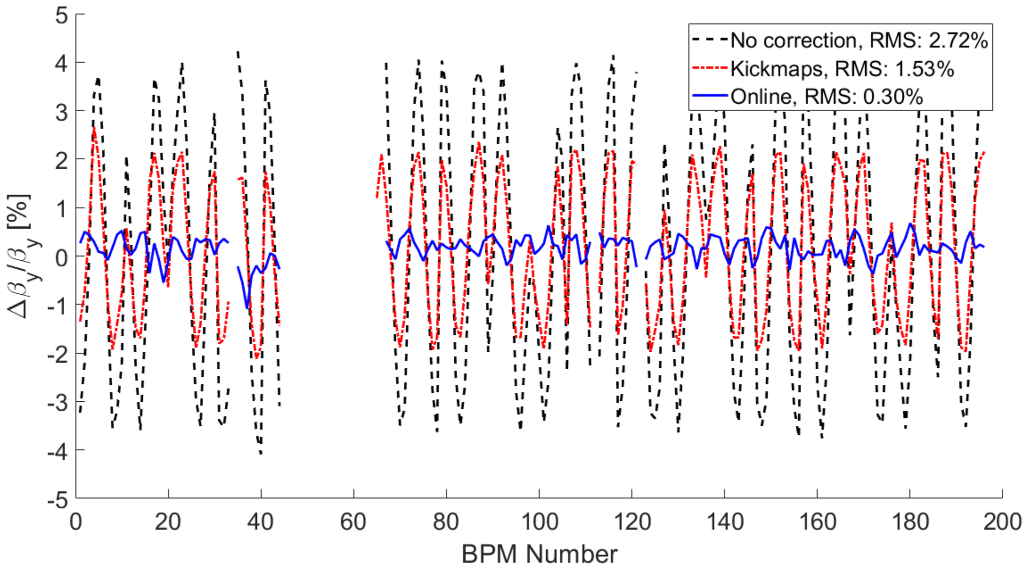

Figure 4.13: Comparison of the kickmap-based and online compensations on the IVW at the BALDER beamline at 4.5 mm gap. The faulty BPM region (approximately numbers 45-66, slight variation between measurements) and outliers have been omitted.

between kickmap-predicted ID focusing and that observed on the machine is rather small.

Table 4.5: Quadrupole adjustments determined for the four quadrupoles involved in the local compensation for the IVW at the BALDER beamline using the two different methods. Upstream/downstream magnets are denoted us/ds, respectively.

Method	Δk_{QFE} (us/ds) [%]	Δk_{QDE} (us/ds) [%]
Kickmap	0.50/0.02	1.77/0.36
Online	0.50/-0.02	1.85/0.68

5 Conclusions and Outlook

The goal of this project was to implement and test a linear optics ID compensation scheme on the 3 GeV ring at MAX IV Laboratory. Specifically, the project was initially limited in scope to implementing and testing feedforward solutions to eliminate deviations in β functions and betatron tunes caused by the presence of IDs. The problem was approached using two methods. In the first approach, the effects of IDs were modeled through the use of kickmaps. Correction terms were calculated through numerical optimization performed on a model lattice with the inclusion of kickmaps, which were used in a physics model of the accelerator to compute quadrupole corrections to nullify the ID perturbations. This was done for both a local compensation using flanking quadrupoles to eliminate β beating, as well as for a global compensation to restore model tunes, to be used in conjunction with the local compensation. In this way, for each ID on the 3 GeV ring, two feedforward devices were devised, one for each compensation. Due to impracticalities in the implementation, and the well-established tune feedback system on the 3 GeV ring, the global compensation was not pursued beyond simulation.

The effectiveness of the kickmap-based local ID compensation was evaluated using both the model-dependent, conventional accelerator optics diagnostics and calibration tool LOCO (in Section 4.1.3), and the pinger-based model-independent β function measurements (in Section 4.1.4). The general conclusion reached is the same for both methods: the kickmap-based corrections do reduce β beating, but not to a satisfactory degree. It should be noted that the feedforward-devices configured in this stage were used during routine operation of the 3 GeV ring during the first quarter of 2026, with no discernible improvement to overall performance [24]. They were later disabled during troubleshooting of unrelated issues to reduce complexity, with no appreciable deterioration of performance, and have remained disabled since. As discussed throughout Sections 4.1.3 and 4.1.4, during the evaluation of the performance of the kickmap-based local compensation, the ability of kickmaps to model the effects of IDs came into question.

For this project, the insufficient quality of kickmap-based quadrupole corrections was not a significant roadblock, as β function compensation could be pursued using another approach. However, it is an important result for current developments and research at MAX IV Laboratory more generally. Currently, the MAX4U lattice upgrade (Conceptual Design Report [25]) is in the technical design phase. As part of the development process, significant effort is being put into developing ID compensation on the new lattice. To model the effects of IDs in simulations using the new lattice, there is some reliance on kickmaps. The discrepancy in β beating between kickmap predictions and measurements on the current accelerator performed as part of this project should be investigated. A further point warranting this investigation is that, fundamentally, kickmaps can be used to model higher-order effects, meaning that their use case goes well beyond the β-function-limited scope of this project.

The other method employed to perform β function compensation in this project was the iterative online algorithm described in Section 3.5. Developed as a part of this project, initial tests were performed on both storage rings. The results obtained from testing on the 1.5 GeV ring, as described in Section 4.2.2, pointed to BPM noise as a limiting factor of the pinger-based β function measurement, and consequently, the online compensation algorithm. Thus, the study presented in Section 3.4.2 was performed to find reasonably fast yet better-optimized measurement configurations in terms of pinger kick strength and number of pings, in the effort of reducing noise. Unfortunately, due to time constraints, repeats of the measurements presented in Section 4.2.2, with improved measurement configurations, could not be performed. While further testing is needed, the initial results of the online compensation algorithm can be considered a success, as in all cases tested, the measured ID-induced β beating could be reduced to the noise floor. Particularly the test on the BALDER IVW presented in Section 4.2.3 indicates the potential of the approach, with β beating being reduced to $< 0.4\%$ RMS in each plane.

As a first proof-of-principle, it should be emphasized that a number of different steps can be taken to improve the online compensation algorithm. The most obvious improvement would be a further reduction of noise in the pinger-based β function measurements. While the investigation presented in Section 3.4.2 resulted in *better* resolution, it is possible that simply increasing pinger voltages further could improve measurement quality. There is an (yet unknown) upper limit, however, to the kick strength that can be applied without losing the stored electron beam. Moreover, it could be worth investigating whether more statistics in the form of more kicks is worthwhile, despite increasing the time required for a measurement. As mentioned in Section 3.4.2, the dominant contribution in measurement time comes from data processing. Reducing the time taken for data processing could be achieved through code optimization and/or further parallelization of computations.

Another striking issue with the current pinger measurements of β functions (and consequently, the online compensation algorithm) is the faulty BPM region. In its current implementation, the online compensation does not address this issue, simply using the faulty BPM data to calculate quadrupole adjustments. Ideally, this issue should be dealt with through further work on BPM timing configurations, which is ongoing. If, however, addressing the issue at the source proves unfeasible, an attempt could be made to salvage usable data in post-processing. As a crude workaround, a simple noise analysis could be implemented as part of the online compensation algorithm to discard faulty BPM data. As a final note on potential developments of the online compensation algorithm, the model-independence of the approach could be extended to the "β function response matrix" R_β by measuring it on the machine rather than relying on a model calculation. This, however, as discussed in Section 3.5.2, would incur a large cost in time. In combination with noise inevitably having an effect on the measured response matrix elements, it is not immediately obvious that measuring R_β is a worthwhile pursuit.

Nevertheless, in the immediate future, based on the results of this project, the online approach seems to be the most promising in the pursuit of minimizing β beating caused by IDs. Most of the steps necessary to generate FF tables based on online compensation have been taken in this project; the implementation is limited mainly by the availability of the accelerators for the required measurements.

References

[1] M. Sjöström (2026), private communication

[2] S. Leemann, H. Tarawneh, in *6th International Particle Accelerator Conference* (2015), p. TUPJE038, **DOI** 10.18429/JACoW-IPAC2015-TUPJE038

[3] Various, *Detailed Design Report on the MAX IV Facility* (MAX IV Facility, 2011)

[4] W. Hillert, *Transverse linear beam dynamics* (2021), 2107.02614, https://arxiv.org/abs/2107.02614

[5] H. Wiedemann, *Particle Accelerator Physics*, Graduate Texts in Physics (Springer, Berlin, Germany, 2015), ISBN 978-3-540-49045-6, 978-3-319-18316-9, 978-3-319-30759-6, 978-3-319-18317-6

[6] C. Mühle, in *Beam Injection, Extraction and Transfer* (CERN Accelerator School, 2017)

[7] P. Forck, P. Kowina, D. Liakin, Synchrotron Radiation News **1** (2008)

[8] J. Safranek, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **388**(1), 27 (1997), ISSN 0168-9002, **DOI** https://doi.org/10.1016/S0168-9002(97)00309-4

[9] A. Ben-Israel, T. Greville, *Generalized Inverses: Theory and Applications*, CMS Books in Mathematics (Springer New York, 2006), ISBN 9780387216348

[10] J.C.A. Barata, M.S. Hussein, Brazilian Journal of Physics **42**(1-2), 146–165 (2011), ISSN 1678-4448, **DOI** 10.1007/s13538-011-0052-z

[11] C.X. Wang, V. Sajaev, C.Y. Yao, Phys. Rev. ST Accel. Beams **6**, 104001 (2003), **DOI** 10.1103/PhysRevSTAB.6.104001

[12] J. Bahrdt (2006), **DOI** 10.5170/CERN-2006-002.441

[13] K. Wille (1998), **DOI** 10.5170/CERN-1998-004.61

[14] F. Hannon et al., in *European Particle Accelerator Conference* (2004)

[15] J. Bahrdt, G. Wüstefeld, Physical Review Special Topics - Accelerators and Beams **14** (2011), **DOI** 10.1103/PhysRevSTAB.14.040703

[16] C. Schmitz-Antoniak, Reports on Progress in Physics **78**, 062501 (2015), **DOI** 10.1088/0034-4885/78/6/062501

[17] J. Safranek, in *Lecture notes* (USPAS, 2015)

[18] E. Wallén, S. Leemann, in *Proceedings of 2011 Particle Accelerator Conference* (2011), p. TUPJE038

[19] J. P. Elleaume, O. Chubar, in *Proceedings of the PAC97 Conference* (1997), pp. 3509–3511

[20] G. Rift (2024), web, accessed 2026-05-02, https://www.maxiv.lu.se/beamlines-accelerators/beamlines/

[21] in *Practical Process Control for Engineers and Technicians*, edited by W. Altmann, D. Macdonald, S. Mackay (Newnes, Oxford, 2005), pp. 142–146, ISBN 978-0-7506-6400-4

[22] M. Wendt, *Bpm systems: A brief introduction to beam position monitoring* (2020), 2005.14081, https://arxiv.org/abs/2005.14081

[23] J. Lagarias et al., SIAM Journal on Optimization **9**, 112 (2006), **DOI** 10.1137/S1052623496303470

[24] M. Sjöström (2026), private communication

[25] Various, *MAX4U - Conceptual Design Report* (MAX IV Facility, 2025)