


Development of Interaction Point Feedback Systems for CERN's Future Circular Lepton Collider



John Patrick Thomas Salvesen

Oriel College, University of Oxford

Thesis submitted in fulfilment of the requirements for the degree of Doctor of
Philosophy at the University of Oxford
Hilary Term, 2026



UNIVERSITY OF
OXFORD

Abstract

CERN’s Future Circular Electron–Positron Collider (FCC-ee) is a proposed high-luminosity circular lepton collider, aimed at delivering unprecedented precision in measurements of the Z, W, Higgs and top quark sectors. Achieving the required luminosity performance demands nanometre-scale control of the relative beam position at the interaction points (IPs). Dynamic perturbations arising from ground motion, magnet vibrations and beam–beam interactions introduce orbit fluctuations that reduce effective luminosity and require active stabilisation. Fast beam-based feedback systems operating at the IP are therefore essential to preserve luminosity and ensure robust operation.

This thesis presents the development of analytic and simulation frameworks for the design and optimisation of IP feedback systems in circular lepton colliders, with particular emphasis on FCC-ee. An analytic model of multi-IP feedback in circular machines is derived, explicitly incorporating the coupling between beam–beam deflections, lattice transport and feedback correction. The formulation enables identification of the physical mechanisms governing stability and gain limits, and provides a computationally efficient tool for parameter optimisation in machines with one, two or four IPs.

Large-scale beam–beam tracking simulations and dedicated machine development studies are used to investigate the practical implementation of IP feedback at SuperKEKB, an operating luminosity frontier lepton collider. At SuperKEKB, the iBump beam–beam deflection feedback system is analysed using operational data to characterise its stability, correction behaviour and operational tuning. Separately, the iBump system is modelled in Xsuite with a realistic lattice description, and its response to driven disturbances is studied as a function of error frequency. Together, these studies provide insight into the behaviour and limitations of IP feedback in a high-luminosity collider environment.

For FCC-ee, performance requirements and tolerances that must be met by the IP feedback are established. A feedback architecture based on integral control is developed and evaluated. Full lattice tracking simulations incorporating beam–beam interaction and feedback correction are performed to assess luminosity recovery under both single-frequency perturbations and realistic models of final-focus quadrupole vibrations. The results demonstrate that appropriately optimised IP feedback can significantly mitigate dynamic orbit perturbations and preserve luminosity under realistic disturbance conditions.

For the minus signs that caused disproportionate stress

Acknowledgements

I would like to express my deepest gratitude to my university supervisor, Professor Philip Burrows, for selecting me for this DPhil and for giving me the opportunity to undertake this research. His guidance, insight, and steady support throughout this journey have been invaluable.

I am equally grateful to my CERN supervisor, Dr Frank Zimmermann, for his mentorship during my time at CERN and for his continued encouragement and support. His perspective and enthusiasm for accelerator physics have greatly influenced this work.

My sincere thanks go to my colleagues in the FCC collaboration for their support, discussions, and collaborative spirit throughout my DPhil. The breadth of expertise and openness within the FCC community have made it a stimulating and rewarding environment in which to work.

I am also deeply thankful to my colleagues at KEK. In particular, I would like to thank Dr Ryuichi Ueki for his supervision and support during my secondments to KEK. I am grateful to EAJADE for supporting my repeated visits, which were instrumental to this work.

At CERN, I have been fortunate to work alongside many supportive colleagues. Special thanks go to my office mates — Tuuli, Elias, Peter, Satya, Erik, Antonia, and Frank — for the engaging discussions and for creating such a positive and uplifting working atmosphere. I would also like to thank Emily Howling and Sasha Horney, my JAI colleagues at CERN, for sharing this journey and for their friendship and support throughout.

My time at Oriel College has been one of the great privileges of this DPhil. I am grateful for the friendships formed there that have accompanied the academic work. Floreat Oriel!

Finally, and most importantly, I would like to express my deepest appreciation to my parents, family, partner, and friends for their constant love, encouragement, and belief in me. Your support has meant more than I can adequately express.

Contents

1	Introduction	1
1.1	The Standard Model	1
1.2	Particle Accelerators	3
1.3	CERN’s Future Circular Lepton Collider	6
1.4	SuperKEKB	9
1.5	Thesis Outline	12
2	Accelerator Physics and Feedback Systems	14
2.1	Beam Dynamics	14
2.2	Interaction Region Optics	18
2.3	Beam–Beam Effects	21
2.4	Oscillation Modes and Resonances	24
2.5	Accelerator Imperfections	28
2.6	Control Systems	31
2.7	Accelerator Feedback Systems	38
2.8	Modelling of Particle Accelerators	47
3	Analytic IP Feedback Modelling	50
3.1	Model Overview	50
3.2	Error Modelling	55
3.3	Benchmarking	61
3.4	Feedback Modelling	65
3.5	One IP Feedback Model	73
3.6	Two IP Feedback Model	79
3.7	Four IP Feedback Model	86
3.8	Summary	95
4	Operational Performance of IP Feedback at SuperKEKB	98
4.1	SuperKEKB IP Feedbacks	98
4.2	SuperKEKB Machine Studies Overview	104
4.3	SuperKEKB Collision Commissioning	107
4.4	Varied Current MD	113
4.5	Constant Current MD	121
4.6	Historical iBump Data	124
4.7	Summary	127

5	Modelling SuperKEKB and the iBump Feedback System	129
5.1	The SuperKEKB Optical Lattices	129
5.2	Modelling SuperKEKB Optics with Xsuite	133
5.3	Radiation and Beam–Beam Modelling for Multi-Turn Tracking	143
5.4	iBump Feedback Modelling	152
5.5	Multi-Turn Particle Tracking Simulations	157
5.6	Summary	165
6	Modelling IP Feedback for FCC-ee	167
6.1	The FCC-ee Lattice	167
6.2	FCC-ee IP Feedback Signals	169
6.3	FCC-ee IP Feedback Tolerances	181
6.4	Radiation and Beam–Beam Modelling for Multi-Turn Tracking	196
6.5	Feedback Modelling	205
6.6	Multi-Turn Particle Tracking Simulations	209
6.7	Summary	228
7	Conclusions	230
7.1	Summary	230
7.2	Further Work	233

1.1.1 Fermions and Bosons

Fermions

Fermions are a class of elementary particle characterised by half-integer spin. Within the SM, fermions are divided into leptons and quarks. There are six leptons and six quark flavours, arranged into three generations. For each fermion, a corresponding antiparticle exists with identical mass and opposite quantum numbers.

Leptons comprise charged leptons and neutrinos, with the former interacting via the electromagnetic and weak forces, while the latter interact only via the weak interaction. Quarks participate in electromagnetic, weak, and strong interactions.

Bosons

Bosons are a class of elementary particle characterised by integer spin. Within the SM, bosons include the spin 1 vector gauge bosons and the spin 0 scalar Higgs boson.

The gauge bosons arise from the local gauge invariance of the SM under the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$ where the subscripts denote colour (strong interaction), left-handed weak isospin, and weak hypercharge, respectively [3, 4]. Locality implies that the symmetry holds at all points in space–time. This structure gives rise to eight gluons associated with the strong interaction, the photon associated with electromagnetism, and the $W^\pm$ and Z bosons associated with the weak interaction.

In addition to the gauge bosons, the SM contains a scalar Higgs boson responsible for electroweak symmetry breaking and mass generation.

1.1.2 The Electroweak Force

The electromagnetic and weak interactions are unified within the SM as a single electroweak interaction. This unification is realised through a locally gauge-invariant theory based on the symmetry group $SU(2)_L \times U(1)_Y$ [3].

A defining feature of the weak interaction is its chiral nature: only left-handed fermions and right-handed anti-fermions participate in charged-current weak interactions. This structure leads to the violation of parity symmetry, a phenomenon observed experimentally in weak processes [5].

1.1.3 The Strong Force

The strong interaction is described by quantum chromodynamics (QCD), a non-Abelian gauge theory based on the local symmetry group $SU(3)_C$. Quarks carry colour charge and interact via the exchange of gluons, which themselves carry colour [4]. A defining feature of QCD is confinement, which prevents the observation of isolated quarks or gluons.

1.1.4 The Higgs Mechanism

In the SM, the electroweak sector is formulated as a locally gauge-invariant theory, which forbids explicit mass terms for both gauge bosons and fermions [6–8]. Mass generation is instead achieved through spontaneous breaking of the electroweak symmetry via the Higgs mechanism. In this framework, a scalar Higgs field acquires a non-zero vacuum expectation value, leading to the breaking of the $SU(2)_L \times U(1)_Y$ symmetry to the electromagnetic $U(1)_{EM}$ subgroup. This results in the weak gauge bosons acquiring mass. The photon remains massless and a single physical scalar particle arises: the Higgs boson. Fermion masses arise through Yukawa interactions with the Higgs field, with the coupling strengths determining the observed fermion mass hierarchy.

1.1.5 Physics Beyond the Standard Model

The SM has been remarkably successful in describing a wide range of experimental observations. Despite this success, it is widely understood to be incomplete [1]. Notable open questions include:

- The absence of a quantum description of gravity compatible with general relativity [9].
- The lack of a mechanism capable of explaining the observed matter-antimatter asymmetry in the Universe [10].
- The absence of a viable particle candidate to account for astrophysical and cosmological evidence for dark matter [11].
- The inability of the minimal SM to accommodate the experimentally observed non-zero neutrino masses [12].

Addressing these and related questions requires the exploration of physics Beyond the Standard Model (BSM). Such searches motivate collider experiments operating at increasingly high energies and rates of production.

1.2 Particle Accelerators

Particle accelerators are devices designed to increase the energy of charged particles and control their motion using electromagnetic fields. The scale and purpose of accelerator facilities vary greatly, ranging from compact systems used in hospitals to large-scale installations dedicated to fundamental research.

1.2.1 Applications of Particle Accelerators

Particle accelerators are widely used across science, medicine, and industry as sources of charged particle beams. Accelerators enable interactions inaccessible to conventional laboratory techniques. While high-energy colliders are among the most prominent examples of accelerator technology, they represent only one class of accelerator.

In medicine, accelerators are used in radiotherapy to deliver localised radiation doses to tumours and in the production of radioisotopes for diagnostic imaging. In research, accelerator-based light sources such as synchrotron radiation facilities provide intense photon beams for studies of materials, chemistry, and biological systems. Industrial and security applications further illustrate the breadth of accelerator technology.

Accelerator performance is not determined solely by the maximum achievable energy, but by the ability to produce stable, well-controlled beams suited to a specific task.

1.2.2 Collider Architectures

A particle collider is an accelerator facility in which two or more particle beams are brought into controlled collision. This beam-beam configuration contrasts with fixed-target experiments, in which a single beam impacts a stationary target. By comparison, colliders allow for a substantially larger fraction of the beam energy to be made available for particle production in the centre-of-mass frame. Colliders may incorporate one or more interaction points (IPs), at which particle detectors are located and collision data are recorded. One of the primary performance metrics is the luminosity, a measure of the rate at which particles collide in an accelerator, determining how frequently a given process occurs for a fixed cross section.

Collider designs can be classified according to how particle beams are accelerated and brought into collision. In linear colliders, beams are accelerated through linear accelerating structures and brought into collision without long-term storage in a ring. In circular colliders, beams are stored for many turns in a storage ring and brought into collision repeatedly.

These two approaches reflect different trade-offs between achievable energy, repetition rate, and beam utilisation. Linear colliders favour very high beam energies by avoiding synchrotron radiation losses, while circular colliders exploit repeated beam circulation to achieve higher interaction rates, subject to energy-dependent radiation effects for light particles.

Most existing and proposed colliders rely on radio-frequency electromagnetic fields to accelerate particles, using either normal-conducting or superconducting technologies. Alternative acceleration concepts, such as plasma-based schemes, are the subject of active research and hold long-term promise for compact high-gradient acceleration, but are not yet employed in high-energy collider facilities.

1.2.3 Lepton and Hadron Colliders

Particle colliders may also be classified according to the type of particles collided, most commonly leptons or hadrons.

The use of hadrons, which are composite particles, enables access to very high centre-of-mass energies but leads to complex initial states and a large background of secondary interactions. As a result, hadron colliders are well suited to the direct production of heavy particles and to discovery-driven physics at the highest available energies. Hadron colliders define the energy frontier of accelerator-based particle physics, with the highest-energy collider to date being the Large Hadron Collider [13], which operates with proton beam energies of up to 6.8 TeV per beam [14].

Leptons are point-like particles with well-defined initial states, allowing collision energies and quantum numbers to be known with high accuracy. This leads to cleaner experimental signatures and lower levels of experimental background compared to hadron machines [15]. Lepton colliders are therefore well suited to detailed studies of the electroweak interaction and the properties of the Higgs boson, where precision measurements can provide sensitivity to BSM physics. Lepton colliders define the luminosity frontier of accelerator-based particle physics, with the highest-luminosity collider to date being SuperKEKB [16], which holds the world record for instantaneous luminosity of $5.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [17]. This emphasis on luminosity places stringent demands on beam stability and control.

Hadron and lepton colliders play complementary roles within the global collider programme, with design choices driven by the specific physics goals they are intended to address.

1.2.4 Future Colliders

In the near term, the High-Luminosity LHC at CERN [18, 19] will maintain LHC’s energy frontier, extending the physics reach of the Large Hadron Collider through increased luminosity. In parallel, existing lepton facilities such as SuperKEKB continue to deliver record luminosities for precision measurements at lower centre-of-mass energies.

Looking beyond the HL-LHC era, a range of future collider concepts have been proposed to address open questions in particle physics, particularly in the electroweak and Higgs sectors. These proposals are evaluated within international strategic processes, including the European Strategy for Particle Physics [20, 21] and the Particle Physics Project Prioritization Panel [22], which assess the scientific case, technical feasibility, cost, and timescale of candidate facilities.

Proposed options span both lepton and hadron colliders, as well as a variety of accelerator technologies. Circular electron–positron colliders, such as the FCC-ee [23] and CEPC [24, 25], aim to deliver extremely high luminosities for precision studies at

the electroweak and Higgs scales. Linear lepton collider proposals, including CLIC [26] and ILC [27], prioritise higher centre-of-mass energies using single-pass acceleration. At the energy frontier, hadron collider concepts such as the FCC-hh [28] explore the extension of proton–proton collisions to significantly higher energies. More exploratory approaches, including muon colliders [29] and plasma-based acceleration schemes [30], are the subject of ongoing research and development.

Throughout reviews, a theme of balance in energy reach, achievable luminosity, experimental cleanliness, and technological maturity has emerged. High-luminosity circular lepton colliders have emerged as particularly strong candidates for the next major facility following the HL-LHC, offering a combination of precision physics potential and incremental technological development. The performance of such machines is intrinsically linked to beam stability and control, providing strong motivation for the study of feedback systems and beam–beam effects that form the focus of this thesis.

1.3 CERN’s Future Circular Lepton Collider

1.3.1 Overview of the FCC-ee Collider

The Future Circular Electron–Positron Collider (FCC-ee) is a proposed ~ 91 km electron–positron collider to be hosted at CERN as part of the Future Circular Collider (FCC) programme. FCC-ee is envisaged as the first stage of an integrated collider complex, to be followed at a later stage by a hadron collider (FCC-hh) in the same tunnel [31]. The FCC programme was formalised in the 2019 Conceptual Design Report (CDR) [23, 28, 32], building on earlier proposals for TLEP [33] and VHE-LHC [34] developed from 2012. The design was refined in the 2025 Feasibility Study Report (FSR) [35–37].

FCC-ee is designed as a high-luminosity storage ring collider. Electrons and positrons circulate in separate beam pipes in a double-ring configuration, allowing operation with a large number of bunches and high total beam current. A limit of 50 MW per beam of synchrotron radiation power is imposed [23], which constrains the achievable beam current as a function of energy.

To reach the targeted luminosities, FCC-ee employs a nano-beam collision scheme with a large horizontal crossing angle of 30 mrad and crab-waist optics. This collision scheme enables extremely small transverse beam sizes at the IP, particularly in the vertical plane. Collisions are foreseen to occur at four IPs.

Electron and positron beams are generated and conditioned in a dedicated injector complex and subsequently accelerated to the collision energy in a booster synchrotron located in the same tunnel as the collider, enabling continuous top-up injection during operation [38, 39]. Once injected into the collider rings, particles are stored at constant energy and collide repeatedly at the IPs. A schematic layout of the FCC-ee collider is

presented in Fig. 1.2.




Figure 1.2: Schematic layout of FCC-ee from [40].

1.3.2 Physics Programme and Working Points

The physics programme of FCC-ee is centred on precision measurements of the SM through the collection of extremely large data samples with integrated luminosities on the order of inverse atto-barns in clean electron–positron collisions. FCC-ee is designed to operate as a factory for the production of the heaviest known particles, enabling stringent tests of the electroweak, Higgs, and strong interactions, thus offering a high potential for new discoveries [41].

FCC-ee is foreseen to operate at four principal centre-of-mass energies [23], corresponding to the production of Z bosons, W^+W^- pairs, Higgs bosons via the ZH process, and top–antitop quark pairs. Operation at the Z pole enables unprecedented precision in electroweak measurements. Operation near the W^+W^- threshold allows precision measurements of the W boson mass and width, while operation at the ZH threshold provides clean access to Higgs boson properties. At the highest energies, FCC-ee enables detailed studies of the top quark through threshold and continuum production. These huge data samples are predicted to be on the order of 10^{12} Z bosons, 10^8 W^+W^- pairs, 10^6 Higgs bosons and 4×10^5 $t\bar{t}$ pairs per year [42].

Each working point has a distinct combination of beam energy, beam current, radiation damping times and beam–beam interaction strength. Machine performance limitations and stability considerations depend on the working point. The impact of beam dynamics and feedback performance must therefore be assessed across the FCC-ee physics programme.

1.3.3 Precision Requirements and Stability Sensitivity

The collider is designed to operate with vertical beam sizes on the order of tens of nanometres at the IP to reach the target luminosities and thus achieve the physics goals. Due to the small beam sizes, the performance of FCC-ee is highly sensitive to relative beam offset errors at collision over a wide range of spatial and temporal scales. Beam stability is therefore a central consideration for FCC-ee. Beam offsets not only impact luminosity but also beam lifetime, collimation efficiency, and overall operational performance.

Understanding the implications of the sensitivity to offsets is a prerequisite for assessing the stability challenges faced by FCC-ee and motivates the use of active beam-based control systems.

1.3.4 Sources of Perturbation and Stability Challenges

It is foreseen that FCC-ee will be subject to a wide range of both environmental and technical disturbances that can affect beam orbit and collision conditions during operation [36].

Slow perturbations, on the scale of minutes to hours, arise from ground motion, mechanical deformations, and thermal effects, leading to gradual changes in the alignment of accelerator components and in turn orbit drifts. Faster disturbances may originate from ground motion and technical noise sources, such as power-supply ripple or RF-related effects, which can excite beam motion on turn-by-turn or bunch-by-bunch timescales.

When operating with collisions, additional challenges arise from the strong beam-beam force. Offsets between colliding bunches modify the forces experienced. The sensitivity of FCC-ee to such effects varies across the physics programme, depending on beam energy, damping rates, and beam-beam parameters.

Perturbations cannot be eliminated entirely through mechanical design or passive stabilisation at FCC-ee, especially due to the challenging scale of the collider. Maintaining the required beam stability at FCC-ee requires continuous monitoring and active feedback correction of beam motion.

1.3.5 Role of Beam-Based Feedback Systems

The stable operation of FCC-ee will rely on active beam-based feedback systems to maintain the required collision conditions [36].

Feedback systems are needed to address beam motion over a broad range of timescales. Slow feedback systems are used to correct global orbit drifts while faster systems are required to maintain alignment at the IPs. In collision, feedback plays a particularly

important role in maintaining optimal beam overlap and mitigating the impact of beam–beam effects.

This thesis focuses on the modelling, analysis, and optimisation of beam-based feedback systems for FCC-ee, with particular emphasis on beam–beam effects and performance across different working points.

1.4 SuperKEKB

1.4.1 Overview of the SuperKEKB Collider

SuperKEKB is an electron–positron collider located at the High Energy Accelerator Research Organization (KEK) in Tsukuba, Japan [16, 43]. It is the successor to the KEKB collider [44] and represents the current state of the art in high-luminosity circular lepton colliders. SuperKEKB consists of two independent storage rings housed in a common tunnel of circumference approximately 3 km. Electrons are stored in the 7 GeV High Energy Ring (HER), while positrons are stored in the 4 GeV Low Energy Ring (LER). Collisions occur at a single IP hosting the Belle II detector [45].

SuperKEKB is designed to operate at a luminosity of approximately forty times that of KEKB. This is achieved through the adoption of a nano-beam collision scheme, a large horizontal crossing angle of 83 mrad, and crab-waist optics, enabling operation with sub-millimetre vertical beta functions at the IP. First beam circulation occurred in 2016 and first collisions in 2018.

SuperKEKB holds the world record for instantaneous luminosity in an electron–positron collider, reaching $5.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in December 2024 [17]. Beam generation and preparation are provided by an injector complex that also services two photon physics experiments, the Photon Factory (PF) and the Photon Factory Advanced Ring (PF-AR) [46]. The presented work focuses exclusively on the collider rings and interaction region. A schematic layout of SuperKEKB is presented in Fig. 1.3.

1.4.2 Physics Programme and Belle II

SuperKEKB is designed to operate at a centre-of-mass energy of 10.58 GeV, corresponding to the $\Upsilon(4S)$ resonance [45]. At this energy, the $\Upsilon(4S)$ decays almost exclusively into coherent $B\bar{B}$ pairs, providing a clean experimental environment for precision studies of flavour physics. The asymmetric beam energies of SuperKEKB enable time-dependent measurements of CP violation through the boost of the B -meson system in the laboratory frame.

The physics programme is carried out by the Belle II experiment. Belle II aims to record an integrated luminosity of approximately 50 ab^{-1} over the lifetime of the facility [48], corresponding to 55 billion $B\bar{B}$ pairs, 47 billion $\tau^+\tau^-$ pairs, and 65 billion




Figure 1.3: Schematic layout of SuperKEKB from [47].

$c\bar{c}$ states. This dataset enables high-precision tests of the SM and sensitive searches for BSM physics through rare decays, flavour-changing neutral currents, and lepton flavour violation. The physics goals of Belle II place stringent requirements on the delivered luminosity.

The strong solenoidal field of 1.5 T used by the Belle II detector for charged-particle tracking extends over the interaction region and heavily impacts and complicates the beam optics. Dedicated compensation solenoids and skew magnetic elements are therefore an integral part of the SuperKEKB lattice; achievable beam sizes and coupling correction at the IP are tightly constrained by the detector configuration.

1.4.3 Precision Requirements and Collision Sensitivity

The luminosity performance of SuperKEKB relies on operation with IP vertical beam sizes on the tens of nanometres scale; the design vertical beta function is sub-millimetre scale. The vertical beam size thus becomes comparable to typical orbit and alignment tolerances. As a consequence, the luminosity is highly sensitive to small deviations from the nominal collision conditions.

Sensitivities are further enhanced by the strong beam–beam forces, where small perturbations may be amplified through collisions and result in luminosity loss through instabilities. The detector environment also plays a central role in defining the precision requirements at SuperKEKB, due to the strong Belle II solenoidal field. Residual coupling and chromatic coupling through imperfect compensation directly translates into changes in the effective collision conditions.

These factors place tight requirements on the mechanical, magnetic, and beam stability of the SuperKEKB interaction region. Sustained high-luminosity operation requires continuous control of the beam orbit and optics at the IP, necessitating the extensive use of fast beam-based correction systems.

1.4.4 Sources of Perturbation and Stability Challenges

Due to SuperKEKB's adoption of the nano-beam scheme, the collider is sensitive to perturbations in the interaction region that may be negligible in conventional storage ring operation, but become performance-limiting with the nominal vertical beam size and beta function at the IP.

Mechanical motion of accelerator components represents a primary source of perturbation. Ground motion, cultural noise, and vibration of magnets and support structures induce transverse orbit fluctuations at the IP. Although the absolute amplitudes of these motions are small, the system is highly sensitive.

Magnetic field fluctuations provide an additional source of instability. Power supply ripple in magnets leads to time-dependent orbit distortions and optics variations. In the interaction region, these effects are compounded by the Belle II solenoid and its compensation system, where small errors introduce coupling and vertical orbit errors.

At high beam currents, collective effects further challenge stability, ranging from the beam-beam interaction to impedance and wakefield effects. These processes can introduce oscillatory behaviour or drifts that are difficult to correct through static tuning alone.

These perturbations result in a continual disturbance of the collision conditions at SuperKEKB across a range of time scales, from turn-by-turn fluctuations to slow drifts over minutes or hours. Maintaining high luminosity requires active stabilisation of the beam orbit and collision parameters, rather than reliance on passive mechanical or magnetic stability.

1.4.5 Role of Beam-Based Feedback Systems at SuperKEKB

Beam-based feedback systems are an essential component of SuperKEKB operation to actively mitigate the impact of perturbations. In contrast to storage rings with less challenging optics, where static tuning maintains acceptable performance over extended periods, sustained high-luminosity operation at SuperKEKB requires continuous correction of beam orbit and collision parameters.

Orbit feedback systems are used throughout the machine to suppress beam motion arising from ground vibration, magnetic field fluctuations, and collective effects. At the IP, the requirements are particularly stringent: small offsets between the beams can lead to significant luminosity degradation. As a result, fast feedback systems

operating on few-turn time scales are required to stabilise the collision conditions during physics operation. Beam-based feedback plays a dual role, both suppressing external disturbances and damping beam–beam driven motion that would otherwise limit stable operation. This approach, in which feedback systems form an integral part of machine operation, provides a valuable reference for future high-luminosity lepton colliders, where similar stability challenges are expected to arise.

1.4.6 SuperKEKB–FCC-ee Synergies

Although SuperKEKB and FCC-ee differ significantly in scale, energy, and scientific scope, they operate in closely related beam dynamics and stability regimes. In particular, both machines pursue high luminosity through strong focusing at the IP and operation close to beam–beam performance limits.

SuperKEKB provides an operational reference for several concepts that are central to FCC-ee such as the use of a nano-beam collision scheme, operation with a large horizontal crossing angle, and the management of strong beam–beam interactions in the presence of tight interaction-region tolerances. As a result, many of the challenges encountered at SuperKEKB are directly relevant to the design and operation of FCC-ee. Experience in maintaining stable collisions under these conditions informs both modelling approaches and operational strategies for future high-luminosity lepton colliders. The relevance is particularly strong for the Z working point of FCC-ee, with ampere-scale beam currents. Studies of beam-based stabilisation and feedback at SuperKEKB therefore provide valuable insight into the feasibility and performance limits of similar systems envisioned for FCC-ee.

Beyond technical considerations, SuperKEKB also represents an important source of operational experience. As the highest-luminosity e^+e^- collider currently in operation, it offers a unique environment for validating beam–beam models and feedback concepts under realistic conditions. These synergies motivate the use of SuperKEKB as a reference throughout this thesis when discussing beam stability and feedback strategies for FCC-ee. These synergistic studies are strongly supported by international collaboration schemes such as the Europe-America-Japan Accelerator Development Exchange Programme (EAJADE) [49].

1.5 Thesis Outline

This thesis contributes to the design and optimisation of feedback systems for current and future luminosity-frontier circular lepton colliders. An analytic model of IP feedback systems in circular colliders, including the effects of multiple IPs, is developed and presented. The implementation of IP feedback at SuperKEKB is described, with

particular emphasis on the beam–beam deflection based iBump feedback system. Results from machine development studies and full lattice tracking simulations incorporating the iBump system are presented, and their implications for IP feedback at FCC-ee discussed. Finally, dedicated simulations are presented whereby viable IP feedback signals are identified and the required tolerances for FCC-ee are established, culminating in full lattice tracking studies of a proposed beam–beam deflection IP feedback scheme.

Chapter 2 presents the theoretical background relevant to discussion of beam dynamics and control theory for IP feedback systems.

Chapter 3 introduces an analytic model for IP feedback at FCC-ee, presenting studies on the optimisation of the PID control methods for the IP feedback, considering configurations with one, two, and four IPs.

Chapter 4 describes the IP feedback systems implemented at SuperKEKB and presents machine development studies of the iBump feedback as well as an analysis of historical data.

Chapter 5 discusses the modelling of the SuperKEKB optical lattices and presents tracking studies including the full optical lattices and the iBump IP feedback system.

Chapter 6 discusses the potential IP feedback signals at FCC-ee as well as required feedback tolerances. Additionally, simulations of FCC-ee with a proposed IP feedback system are presented, demonstrating effective mitigation of luminosity degradation caused by vibrational errors.

Chapter 7 summarises the work and its conclusions as well as presenting recommendations for future studies.

Chapter 2

Accelerator Physics and Feedback Systems

2.1 Beam Dynamics

2.1.1 Phase Space Coordinates and Reference Trajectory

The motion of charged particles in a circular accelerator is described relative to a reference trajectory, defined as the ideal closed orbit of a particle with nominal mass, charge, and momentum. Particle coordinates are expressed in a Frenet–Serret frame that follows this trajectory. The longitudinal coordinate s denotes the distance along the reference trajectory, while the transverse coordinates x and y represent horizontal and vertical displacements with respect to it.

In this work, single-particle phase-space coordinates are written in the six-dimensional form

$$\mathbf{x} = (x, x', y, y', \zeta, \delta), \quad (2.1)$$

where $x' = dx/ds$ and $y' = dy/ds$ are the transverse slopes with respect to the longitudinal coordinate s . The longitudinal coordinate ζ denotes the displacement of a particle along the reference trajectory relative to the synchronous particle, and the relative momentum deviation is defined as

$$\delta = \frac{p - p_0}{p_0}, \quad (2.2)$$

with p_0 the reference momentum.

Throughout this thesis, two levels of description are employed. In four-dimensional (4D) beam dynamics, only the transverse degrees of freedom (x, x', y, y') are retained, while the longitudinal coordinates (ζ, δ) are neglected. In six-dimensional (6D) beam dynamics, all coordinates are retained, enabling a self-consistent description of longitudinal motion, energy spread, and dispersion-driven transverse effects.

2.1.2 Linear Transverse Motion and Betatron Oscillations

In the absence of strong non-linearities, transverse particle motion in a circular accelerator is well described by linear dynamics about the reference trajectory, with decoupled horizontal and vertical planes.

Beam bending and focusing are provided by magnetic elements distributed around the ring, collectively referred to as the magnetic lattice [50]. Dipole magnets bend the reference trajectory and define the closed orbit, while quadrupole magnets provide linear focusing in one transverse plane and defocusing in the orthogonal plane. Higher-order multipole fields introduce non-linear effects and are neglected at this stage.

Under these assumptions, the transverse displacement $a(s)$ from the reference orbit satisfies Hill's equation,

$$\frac{d^2 a}{ds^2} + K_a(s) a = 0 \quad a \in \{x, y\}, \quad (2.3)$$

where the periodic focusing functions $K_{x,y}(s)$ are determined by the lattice and satisfy $K_{x,y}(s + C) = K_{x,y}(s)$, with C the machine circumference.

The solutions of (2.3) describe oscillatory motion about the reference orbit, known as betatron oscillations. The betatron tunes $Q_{x,y}$ represent the number of transverse oscillations executed per revolution and are defined in terms of the betatron phase $\mu_{x,y}(s)$ as

$$Q_{x,y} = \frac{1}{2\pi} \int_0^C \frac{d\mu_{x,y}}{ds} ds. \quad (2.4)$$

The tunes determine the spectral locations of the natural transverse modes. Perturbations with frequency components near harmonics of $Q_{x,y}$ can resonantly excite transverse motion, making the choice of working point in tune space central to beam stability. This underpins later discussions of resonances, beam-beam effects, and feedback system design.

At this level of description the transverse motion is conservative. Radiation damping and quantum excitation, relevant for lepton colliders, are introduced in Section 2.1.5.

2.1.3 Courant–Snyder Formalism and Optical Functions

The general solution of Hill's equation for linear transverse motion may be expressed using the Courant–Snyder formalism, in which the effect of lattice focusing is encoded by optical functions. The Twiss functions $\beta_a(s)$, $\alpha_a(s)$, and $\gamma_a(s)$ are defined such that

$$\alpha_a(s) = -\frac{1}{2} \frac{d\beta_a}{ds}, \quad \gamma_a(s) = \frac{1 + \alpha_a^2(s)}{\beta_a(s)}, \quad \beta_a \gamma_a - \alpha_a^2 = 1, \quad a \in \{x, y\}. \quad (2.5)$$

These functions are periodic around the ring and are fully determined by the magnetic lattice.

The transverse phase-space coordinates satisfy the Courant–Snyder invariant

$$\varepsilon_a = \gamma_a a^2 + 2\alpha_a a a' + \beta_a a'^2, \quad a \in \{x, y\}, \quad (2.6)$$

which is conserved under linear, conservative motion. Geometrically, this invariant defines an ellipse in transverse phase space with constant area whose shape varies with s , illustrated schematically in Fig. 2.1.




Figure 2.1: Schematic illustration of a Courant–Snyder ellipse representing transverse phase space from [51].

2.1.4 Longitudinal Motion and Synchrotron Oscillations

Longitudinal beam dynamics arise from the combined action of radio-frequency (RF) accelerating cavities and the energy dependence of the closed-orbit length [52]. Particles with different energies follow trajectories of different path length, arriving at the RF cavities at different times. The synchronous particle is defined as the particle that arrives at the RF cavities with the design phase and energy on each turn. Particles with non-zero longitudinal displacement ζ or relative momentum deviation δ experience restoring forces that drive oscillations in longitudinal phase space.

For small amplitudes, the longitudinal equations of motion may be linearised, yielding harmonic oscillations known as synchrotron oscillations,

$$\frac{d^2\zeta}{ds^2} + K_s \zeta = 0, \quad (2.7)$$

where K_s is an effective longitudinal focusing strength determined by the RF voltage, harmonic number, reference energy, and momentum compaction. The momentum

compaction factor α_c characterises the relative change in closed-orbit length with momentum,

$$\alpha_c = \frac{1}{C} \frac{dC}{d\delta}. \quad (2.8)$$

The resulting oscillation frequency defines the synchrotron tune Q_s , corresponding to the number of longitudinal oscillations executed per revolution.

A finite synchrotron tune couples longitudinal and transverse motion through energy-dependent transport effects such as dispersion. As a result, purely transverse perturbations may acquire a longitudinal dependence, and energy deviations can manifest as transverse position and angle offsets.

2.1.5 Synchrotron Radiation, Damping, and Equilibrium

Synchrotron radiation plays a central role in determining the beam dynamics of high-energy lepton colliders [53]. As particles traverse curved trajectories they lose energy by emission of electromagnetic radiation. This energy loss is compensated on average by the RF system, but the emission process introduces two fundamentally important effects: radiation damping and quantum excitation.

Radiation damping acts to reduce the amplitudes of particle oscillations in all degrees of freedom. The horizontal, vertical, and longitudinal motions experience exponential damping characterised by damping times τ_x , τ_y , and τ_s toward the reference trajectory of the synchronous particle. Synchrotron radiation is, however, quantum by nature. Random perturbations are introduced to the particle energy and transverse momenta, leading to quantum excitation. The competition between radiation damping and quantum excitation results in an equilibrium distribution in phase space characterised by finite transverse emittances and a finite energy spread. The resulting equilibrium emittances and energy spread are determined by the lattice optics and radiation integrals.

As a consequence of radiation effects, the Courant–Snyder invariants introduced in Section 2.1.3 are no longer strictly conserved. Instead, particle motion may be viewed as that of a damped linear system subject to stochastic excitation, evolving toward a radiation-defined equilibrium. The resulting beam distribution is approximately Gaussian in all phase-space coordinates.

2.1.6 Linear Dispersion and Energy-Dependent Transport

Particles with different momenta follow different closed orbits. This effect gives rise to dispersion, which couples longitudinal momentum deviations to transverse motion. The transverse particle coordinates may be decomposed into a betatron component and a

dispersive component,

$$a(s) = a_\beta(s) + D_a(s)\delta, \quad a \in x, y \quad (2.9)$$

where a_β describes the betatron motion and $D_a(s)$ is the linear dispersion function. The dispersion functions are periodic lattice properties determined by the lattice.

Horizontal dispersion arises naturally in bending magnets, where off-momentum particles experience a different curvature. Vertical dispersion is typically minimised by design, but in practice is generated by a combination of lattice imperfections and interaction-region effects. These include magnet misalignments, residual coupling and solenoidal fields associated with particle detectors. As a result, non-zero vertical dispersion is generally present at some level in realistic collider lattices.

Dispersion provides a direct mechanism by which longitudinal dynamics influence transverse motion. A finite energy spread, established through synchrotron radiation and quantum excitation, therefore leads to a corresponding spread in transverse position and angle. Energy fluctuations and longitudinal oscillations can manifest as transverse centroid motion, particularly in regions of large dispersion.

In addition to closed-orbit distortions, the momentum dependence of the focusing strength introduces an energy dependence of the betatron tunes. This effect is quantified by the chromaticities $Q'_{x,y}$, defined as

$$Q'_{x,y} = \frac{dQ_{x,y}}{d\delta}. \quad (2.10)$$

In the absence of correction, the natural chromaticity of a focusing lattice is typically negative, resulting in a tune spread across a beam with finite energy spread. Chromatic effects therefore represent a further coupling between longitudinal and transverse dynamics and play an important role in beam stability.

2.2 Interaction Region Optics

The Interaction Region (IR) of a circular lepton collider is highly constrained and specialised. Its role is to bring two stored beams into collision with extremely small transverse dimensions at the Interaction Point (IP). IR optics represent a compromise between luminosity optimisation, detector requirements, and beam dynamics constraints.

2.2.1 Final Focus Systems

The primary function of the final focus system is to minimise the transverse beam size at the IP. This is achieved by strong quadrupole doublets or triplets placed around

the IP, producing very small Twiss beta functions at the IP. The achievable luminosity scales inversely with the transverse beam sizes, motivating aggressive reduction of $\beta_{x,y}^*$.

Strong focusing in the IR introduces several challenges. Large beta functions in the final quadrupoles enhance sensitivity to magnet misalignments and field errors. Chromatic aberrations generated by the final focus must be corrected, typically using dedicated sextupoles arranged to preserve dynamic aperture, where the dynamic aperture refers to the region of transverse phase space over which particle motion remains stable over many turns. These effects place stringent tolerances on orbit control and optical stability in the vicinity of the IP; the small beam size generated by the final focus system defines the scale of observable orbit perturbations and correction effectiveness.

2.2.2 Nano-Beam Scheme

Modern high-luminosity lepton colliders employ the nano-beam scheme, in which extremely small vertical beam sizes are achieved at the IP while allowing comparatively larger horizontal dimensions. The nano-beam collision scheme was first proposed for the SuperB project by Raimondi in 2006 [54]. This approach relies on very small vertical beta functions and the use of a finite crossing angle between the two beams.

A crossing angle is used to suppress parasitic collisions and enables strong vertical focusing without excessive hourglass degradation (see Section 2.3.7). In the nano-beam regime, the overlap region of the colliding bunches is dominated by the vertical beta function, crossing angle, and horizontal beam size, making luminosity particularly sensitive to vertical orbit offsets and angular errors.

The nano-beam configuration enhances the importance of transverse stability at the IP. Nanometre-scale relative offsets can lead to significant luminosity loss and coherent beam–beam deflections, directly motivating the use of fast beam-based feedback systems.

2.2.3 Crab Waist Optics

In the presence of a finite crossing angle, the waist, the longitudinal position of the minimum vertical beta function, plays a critical role in determining the effective beam overlap at the IP. For head-on collisions this waist is naturally at the IP, however in a crossing-angle geometry the region of strongest beam–beam interaction is tilted with respect to the laboratory frame.

The crab waist scheme is employed to restore optimal overlap by rotating the vertical waist of each beam such that it follows the trajectory of the opposing beam [55]. It was first experimentally demonstrated at DAΦNE [56]. This scheme is achieved using dedicated sextupole magnets placed near the IP at locations with appropriate phase advances. The resulting transformation suppresses horizontal-vertical beam–beam resonances as the vertical tune shift no longer varies with the horizontal coordinate,

allowing stronger vertical focusing and higher luminosity. The effect of the crab waist scheme is presented schematically in Fig. 2.2.




Figure 2.2: Schematic of the crab-waist collision scheme with crossing-angle geometry and sextupole locations from [51].

The crab waist transformation is described by the effective Hamiltonian

$$H_{CW} = kxy'^2, \quad (2.11)$$

where x is the horizontal position, y' is the vertical slope, and k parametrises the strength of the crab waist transformation. A perfect crab waist condition is obtained for $k = 1/\theta_c$ where θ_c is the full crossing angle.

2.2.4 Impacts of Experimental Detectors

The IR optics must be compatible with the geometry and performance requirements of the experimental particle detectors.

Magnetic elements are excluded from a region around the IP to allow for detector components and to minimise backgrounds from synchrotron radiation. Bending magnets are avoided close to the IP to suppress synchrotron radiation directed toward the detector. This constraint leads to a largely straight beam trajectory through the central detector region but places additional demands on the surrounding lattice to provide flexibility for chromatic correction and matching to the arcs.

Modern experimental particle detectors typically include a strong detector solenoid field aligned with the average beam axis. The solenoid modifies the local optics and introduces transverse coupling: the horizontal and vertical motions are not strictly

independent in the IR. Without compensation, this coupling can significantly degrade beam size and luminosity. Solenoidal effects are mitigated using compensating solenoids and skew quadrupoles arranged to restore decoupled transverse motion. Residual coupling and orbit distortions remain important considerations, particularly for vertical stability and beam–beam interactions.

2.3 Beam–Beam Effects

The effect of the electromagnetic fields of the opposing beam at the IP plays a central role in determining achievable luminosity, beam stability, and operational limits.

2.3.1 Beam–Beam Force

During the collision of two ultra-relativistic bunches, each bunch experiences the transverse electromagnetic fields generated by the charge distribution of the other. In the paraxial approximation, the interaction is effectively instantaneous and localised at the IP, and to first approximation can be modelled as a transverse momentum kick.

For a bunch with transverse charge density $\rho(x, y)$, the transverse electric field satisfies Poisson’s equation in the transverse plane. For Gaussian beams, the resulting force can be expressed analytically, providing a practical description of the beam–beam interaction for both simulation and analytic modelling.

The transverse kick experienced by a particle at transverse position (x, y) can be written as

$$\Delta x' = -\frac{qe}{p_0 c} E_x(x, y), \quad \Delta y' = -\frac{qe}{p_0 c} E_y(x, y), \quad (2.12)$$

where $E_{x,y}$ are the transverse electric field components of the opposing bunch. The magnetic field contributes an equal force, resulting in an effective doubling of the electric field contribution for ultra-relativistic beams [57, 58].

2.3.2 Bassetti–Erskine Formalism

For beams with a bi-Gaussian transverse distribution, the transverse beam–beam force can be expressed in closed form using the Bassetti–Erskine formalism [59]. This provides an exact solution for the complex transverse electric field generated by a Gaussian charge distribution with unequal horizontal and vertical beam sizes.

Introducing the complex transverse coordinate $z = x + iy$, the complex field $E_x + iE_y$

can be expressed as

$$E_x + iE_y = \frac{Q}{2\pi\epsilon_0} \frac{1}{z} \left[w\left(\frac{z}{\sqrt{2(\sigma_x^2 - \sigma_y^2)}}\right) - \exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right) w\left(\frac{x\sigma_y + iy\sigma_x}{\sqrt{2\sigma_x\sigma_y(\sigma_x^2 - \sigma_y^2)}}\right) \right], \quad (2.13)$$

where Q is the bunch charge, $\sigma_{x,y}$ are the transverse RMS beam sizes, and w denotes the complex error function. The Bassetti–Erskine expression is widely used in beam–beam tracking codes.

2.3.3 Linearised Beam–Beam Force and Tune Shift

For low amplitudes, the beam–beam force may be linearised [58]. In this limit, the beam–beam force is linearly proportional to the offset. As such the interaction acts as an effective thin quadrupole, focusing in both transverse planes.

The linearised kick can be expressed as

$$\Delta a' = \pm \frac{1}{2} k_a \Delta a \quad a \in x, y, \quad (2.14)$$

where the effective focusing strengths k_a depend on the opposing beam parameters.

This quadrupolar nature leads to a shift in the betatron tunes: the beam–beam tune shift. Using the formula for the change of tune due to a quadrupolar error [58], we can quantify the strength of the beam–beam force from a measurable change in tune:

$$\cos(2\pi(Q_0 + \Delta Q)) = \cos(2\pi Q_0) - \beta_0^* k \sin(2\pi Q_0) \quad (2.15)$$

with ΔQ the resulting tune shift with respect to the reference tune Q . In terms of beam parameters, the beam–beam tune shift is given by

$$\xi_a = \frac{r_e N}{2\pi\gamma} \frac{\beta_a^*}{\sigma_a(\sigma_x + \sigma_y)}, \quad a \in x, y \quad (2.16)$$

where r_e is the classical electron radius, N the bunch population, γ the relativistic Lorentz factor, β_a^* the beta function at the IP, and $\sigma_{x,y}$ the RMS transverse beam sizes. The beam–beam tune shift provides a convenient single parameter characterising the strength of the interaction and plays a central role in defining operational limits for circular colliders [60]. A common equivalent representation of the beam–beam force in the linear collider community is the ‘disruption’ parameter [61].

2.3.4 Beam–Beam Deflection

When two colliding beams have a relative transverse offset at the IP, the beam–beam force produces a coherent angular deflection of each beam [57]. To first order, the

deflection is proportional to the relative offset and the beam–beam tune shift.

This beam–beam deflection forms the basis of beam–based IP feedback systems, in which the outgoing beam angle is used as a diagnostic for transverse offsets at the collision point. Unlike incoherent beam–beam effects, which act on individual particles, the beam–beam deflection is a collective effect that directly couples the centroids of the two beams. The sensitivity of the deflection signal depends on the beam sizes, bunch population, and collision geometry, and is strongest in the plane with the largest beam–beam tune shift. For flat beams, this typically corresponds to the vertical plane.

2.3.5 Luminosity

The luminosity is the primary figure of merit for a collider at a given energy and is defined as the interaction rate per unit cross section. For two colliding Gaussian bunches, the instantaneous luminosity per IP is given by

$$\mathcal{L} = \frac{n_b f_{\text{rev}} N^2}{4\pi\sigma_x\sigma_y} F, \quad (2.17)$$

where n_b is the number of bunches per beam, f_{rev} the revolution frequency, N the bunch population, $\sigma_{x,y}$ the transverse RMS beam sizes at the IP, and F a luminosity reduction factor accounting for the crossing angle and the hourglass effect arising from the finite bunch length [62]. Reduction due to the geometric factor may be mitigated using crab-crossing or crab-waist schemes. The luminosity may also be expressed in terms of the beam–beam tune shift. In the flat-beam limit, $\sigma_x \gg \sigma_y$, this yields

$$\mathcal{L} \simeq \frac{n_b f_{\text{rev}} \gamma N}{2r_e \beta_y^*} \xi_y F, \quad (2.18)$$

scaling linearly with beam current, inverse β_y^* , and beam–beam tune shift. For a given optics and beam energy, the achievable luminosity is typically limited by the maximum sustainable beam–beam tune shift. Increasing ξ_y increases luminosity, but enhances non-linear beam–beam effects and radiative processes.

2.3.6 Beamstrahlung

Beamstrahlung is the emission of synchrotron radiation by beam particles as they traverse the electromagnetic field of the opposing bunch during collision [61]. Beamstrahlung is a localised effect occurring only at the IP. As particles pass through the field of the opposing beam, they experience rapid transverse acceleration, resulting in the emission of photons and a corresponding loss of energy. Beamstrahlung radiation contributes to particle loss and lifetime limitations, but also modifies the equilibrium longitudinal beam parameters. In machines with multiple IPs this impact is exacerbated.

The total RMS relative energy spread, including the effect of beamstrahlung radiation, can be expressed as

$$\sigma_{\delta,\text{tot}}^2 = \sigma_{\delta,\text{SR}}^2 + \frac{n_{\text{IP}} \tau_{\text{s,SR}} c}{4C} \langle u^2 \rangle_{\text{BS}}, \quad (2.19)$$

where $\sigma_{\delta,\text{SR}}$ is the equilibrium energy spread due to synchrotron radiation in the arcs, n_{IP} is the number of IPs, $\tau_{\text{s,SR}}$ is the longitudinal damping time, and $\langle u^2 \rangle_{\text{BS}}$ the second moment of the relative photon energy emitted through beamstrahlung. Through the relative photon energy, there is a strong dependence on beam energy and beam–beam field strength.

An increase in the equilibrium energy spread leads to an increase in bunch length: longitudinal emittance growth [61]. As a result, beamstrahlung will play a central role in defining the achievable luminosity and stability margins in high-energy, high-luminosity circular colliders. For FCC-ee operating conditions, beamstrahlung is expected to dominate the equilibrium energy spread, whereas at SuperKEKB it is a negligible effect.

2.3.7 Other Effects

In addition to the effects discussed above, the beam–beam interaction gives rise to several further phenomena that influence collider performance. The strong electromagnetic focusing during collision modifies the effective beta functions at the IP, known as the dynamic beta, altering the local optics and beam sizes [57]. The finite bunch length also leads to the hourglass effect, in which particles collide away from the beta waist where the beta function is larger, reducing the effective luminosity [62]. Non-linear beam–beam forces drive diffusion and resonance excitation, producing dynamic emittance growth that can degrade beam stability and lifetime [57]. The mutual focusing of the two beams during collision, referred to as the pinch effect, increases the local charge density and can enhance the luminosity beyond the geometric expectation while strengthening non-linear and radiative effects [63]. Radiative Bhabha scattering constitutes an important luminosity-related loss mechanism, in which hard photon emission during electron–positron scattering causes particles to fall outside the machine momentum acceptance [64]. Together, these effects define practical limits on achievable luminosity, motivating the use of feedback systems to stabilise collisions and preserve performance.

2.4 Oscillation Modes and Resonances

2.4.1 Sigma and Pi Mode Oscillations

In a circular collider, the beam–beam interaction results in the coupling of the transverse centroid motion of the colliding bunches. Considering the linearized beam–beam force

described in Section 2.3.3, in the case of symmetric Gaussian beams, each beam feels the kick

$$\Delta y'_1 = -\frac{1}{2}k_y(y_1 - y_2) \quad \Delta y'_2 = +\frac{1}{2}k_y(y_1 - y_2), \quad (2.20)$$

which can be re-written into a matrix system as

$$\Delta \mathbf{y}' = -\frac{1}{2}k_y \begin{bmatrix} +1 & -1 \\ -1 & +1 \end{bmatrix} \mathbf{y} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}, \quad \mathbf{y}' = \begin{pmatrix} y'_1 \\ y'_2 \end{pmatrix}, \quad (2.21)$$

and further diagonalised to identify the oscillation modes, with two distinct eigenvectors:

$$\sigma = \begin{pmatrix} +1 \\ +1 \end{pmatrix} \quad \pi = \begin{pmatrix} +1 \\ -1 \end{pmatrix}. \quad (2.22)$$

The σ mode represents the in phase motion and has an eigenvalue of 0. No focussing occurs so the tune is unchanged. The two beams are moving together, such that the separation $(y_1 - y_2)$ remains constant.

The π mode represents the out of phase motion and has an eigenvalue of 1. In the π mode a focussing strength of $-k_y$ is observed. The two beams are moving oppositely such that the separation is double the centroid offset of each beam. Thus each beam sees the full kick, leading to an effective lens strength of k_y , rather than $\frac{1}{2}k_y$.

2.4.2 Sigma and Pi Mode Tunes

The σ -mode tune is unshifted with respect to the bare lattice tune while the π -mode tune is shifted due to the impact of the beam-beam interaction.

The one-turn matrix of the bare lattice can be expressed as a rotation by the phase advance $\mu_0 = 2\pi Q_0$ as

$$R(\mu_0) = \begin{bmatrix} \cos \mu_0 & \sin \mu_0 \\ -\sin \mu_0 & \cos \mu_0 \end{bmatrix}. \quad (2.23)$$

The impact of the beam-beam force can then be expressed as a linear dipole kick proportional to the beams' relative displacement. Based on the eigenvectors described above, symmetric and antisymmetric variables describing the beam centroids are introduced as

$$y_\sigma = \frac{1}{2}(y_a + y_b) \quad p_\sigma = \frac{1}{2}(p_a + p_b) \quad y_\pi = \frac{1}{2}(y_a - y_b) \quad p_\pi = \frac{1}{2}(p_a - p_b), \quad (2.24)$$

where the effect of the beam-beam force is a block diagonal expressed as

$$F(2k) = \begin{bmatrix} 1 & 0 \\ -4\pi\xi & 1 \end{bmatrix}. \quad (2.25)$$

The tunes are then calculated by considering the effect of the combined map; the product of the previous one-turn matrix and the beam-beam interaction:

$$M_\sigma = R(\mu_0), \quad M_\pi = R(\mu_0) F(2k). \quad (2.26)$$

As the tune corresponds to the trace of the one-turn matrix in a 2×2 symplectic system, the σ -mode and π -mode tunes (Q_σ and Q_π) can be calculated as

$$\cos(2\pi Q_\sigma) = \cos(2\pi Q_0) \quad \cos(2\pi Q_\pi) = \cos(2\pi Q_0) - 2\pi\xi \sin(2\pi Q_0). \quad (2.27)$$

For complex cases involving flat beams, crossing angles, or unequal parameters, a geometric (‘Yokoya’) factor [60] modifies the effective coupling. The coupling factor of $2\pi\xi$ is changed to $2C_\pi\pi\xi$, where the Yokoya factor C_π accounts for the complexity.

2.4.3 Pi Mode with Multiple IPs

In the case of N_{IP} identical IPs with constant phase advances separating them in each plane, one considers N_{IP} transfer matrices and N_{IP} beam-beam interactions. For the cases of 1, 2 and 4 IP colliders, the resulting matrices are all squares of the original. The tune calculation can be simplified as for 2×2 matrices the trace of a squared matrix can be calculated using

$$\text{Tr}M^2 = (\text{Tr}M)^2 - 2\det M \quad (2.28)$$

from which, it can be seen that the π -mode tunes in the 1, 2 and 4 IP cases can be expressed as:

$$Q_{\pi, 1\text{IP}} = \frac{1}{2\pi} \arccos[\cos(2\pi Q_0) - 2\pi\xi \sin(2\pi Q_0)], \quad (2.29)$$

$$Q_{\pi, 2\text{IP}} = \frac{1}{2\pi} \arccos[2(\cos(2\pi Q_0) - 2\pi\xi \sin(2\pi Q_0))^2 - 1], \quad (2.30)$$

$$Q_{\pi, 4\text{IP}} = \frac{1}{2\pi} \arccos[2(2[\cos(2\pi Q_0) - 2\pi\xi \sin(2\pi Q_0)]^2 - 1)^2 - 1]. \quad (2.31)$$

These functions are in general complex, where an imaginary tune indicates growth of oscillatory modes.

2.4.4 IP Feedback and the Pi Mode

IP Feedback operates to minimise the offset between colliding bunches at the IP. This relative centroid offset, $y_a - y_b$, is exactly the antisymmetric π -mode described above.

IP feedback is therefore effectively blind to σ -mode oscillations. This oscillation mode must be addressed with an alternative feedback, for example a global orbit feedback, or separate slow control bumps at the IP to ensure that bunches do not move

too far from the nominal IP. Secondly, the stability and optimal phase and gain of the control system depend on the π -mode tune. Knowing Q_π is therefore central to robust feedback design and tuning.

2.4.5 Resonances

In circular accelerators, the transverse motion of a particle can be described by betatron oscillations with horizontal and vertical tunes Q_x and Q_y per revolution. Instabilities arise when the tunes satisfy integer relations [65]. Considering transverse motion only, the requirement for instability is given by:

$$mQ_x + nQ_y = p, \quad m, n, p \in \mathbb{Z}. \quad (2.32)$$

If this condition is met, after a number of turns the phase advance closes on itself such that imperfections or non-linear fields with the same periodicity add coherently to the motion causing resonance.

Resonances can be categorised as:

- **Integer resonances** ($Q = p$): driven by dipole errors and typically avoided by ensuring the working point is never close to an integer.
- **Half-integer resonances** ($2Q = p$): driven by quadrupole errors, especially dangerous as focusing errors change sign each turn and add coherently, producing rapid growth of oscillation amplitude.
- **Higher-order resonances** ($mQ = p$ with $m > 2$): driven by non-linear fields such as sextupoles ($m = 3$), octupoles ($m = 4$), and higher order components. The strength of these components typically decreases with order reducing the impact.
- **Sum and difference resonances** ($mQ_x \pm nQ_y = p$): driven by skew multipoles, either from skew components or misalignment errors. Such resonances couple the horizontal and vertical planes, with the $Q_x \pm Q_y = p$ typically most dangerous.

Meeting a resonance condition does not imply particle loss. The strength of a resonance depends on the presence of the associated Resonance Driving Terms (RDTs), corresponding to the types of field that drive particular resonances. For example, a resonance driving term associated with integer resonances is due to dipolar errors. The strength of the RDTs determines how quickly particles are driven out of the stable region once the working point is close to the corresponding line.

Tune Diagram and Footprint

The resonance condition presented in Eq. (2.32) defines straight lines in the (Q_x, Q_y) plane, known as resonance lines. To visualise these resonances, the tune diagram is typically used. In this way, the nominal tune, as well as the expected tune spread (known as the tune footprint), due to beam–beam interactions, space charge or other effects, can be visualised. The region between strong lines defines the “safe” operating window. For stability, the entire footprint must remain clear of strong resonance lines. An example resonance diagram is presented in Fig. 2.3.




Figure 2.3: Tune diagram showing low-order betatron resonances in the (ν_x, ν_y) plane from [65]. Solid lines indicate integer and half-integer resonances, while dashed lines denote third-order resonances.

For stability with beam–beam collisions both the σ and π -mode tunes must be placed in safe regions of the tune diagram.

2.5 Accelerator Imperfections

Theoretical descriptions of accelerator lattices typically assume ideal magnetic fields and perfect alignment. In practice, accelerators are subject to a wide range of imperfections that perturb the beam trajectory, optics, and collision conditions.

Accelerator errors can be broadly classified according to their physical origin and characteristic timescale. Imperfections arise from mechanical tolerances, magnet and power supply limitations and environmental disturbances. Static or slowly varying errors determine the baseline machine performance and are typically corrected during commissioning. Dynamic errors introduce time-dependent perturbations that cannot be fully eliminated by static correction and motivate the use of fast feedback systems.

2.5.1 Magnet Alignment Errors

Misalignments of magnets represent one of the dominant sources of orbit and optics errors. Offsets of magnets generate unintended steering fields that deflect the beam from its design trajectory. This behaviour can be described generally in terms of multipole feed-down [50, 65]. Writing the transverse magnetic field using complex multipole coefficients C_m , the magnetic field may be expressed as

$$B_y + iB_x = \sum_{m=1}^{\infty} C_m z^{m-1}, \quad (2.33)$$

where $z = x + iy$ is the complex transverse coordinate. If the magnet is displaced by a complex offset $\Delta = \Delta x + i\Delta y$, the field expressed in the shifted coordinate $z' = z - \Delta$ becomes

$$B_y + iB_x = \sum_{n=1}^{\infty} \left(\sum_{m=n}^{\infty} \binom{m-1}{n-1} C_m \Delta^{m-n} \right) z'^{n-1}. \quad (2.34)$$

The effective multipole coefficient of order n observed by the beam is therefore

$$C_n^{\text{eff}} = \sum_{m=n}^{\infty} \binom{m-1}{n-1} C_m \Delta^{m-n}, \quad (2.35)$$

resulting in a non-zero component in the lower order multipoles. For small residual offsets after alignment, it is sufficient to retain only the leading linear term in the displacement. In this limit, a multipole of order m generates an effective component of order $m - 1$,

$$\Delta C_{m-1} = (m-1) C_m \Delta + \mathcal{O}(\Delta^2), \quad (2.36)$$

[65]. A displaced dipole therefore produces no feed-down, while a displaced quadrupole generates an effective dipole field, and a displaced sextupole generates an effective quadrupole component. In integrated strength notation, a quadrupole with strength $k_1 L$ displaced horizontally by Δx produces an effective dipole kick $k_0 L \approx -k_1 L \Delta x$ leading to closed-orbit distortion. In strong-focusing lattices, and particularly in interaction regions where beta functions are large, even micron-scale magnet offsets can therefore result in significant orbit and collision errors.

In addition to orbit distortion, misalignments can drive coupling between the transverse planes and generate vertical dispersion. At the IP, vertical dispersion directly degrades the effective overlap of the colliding beams.

2.5.2 Magnet Strength and Powering Errors

Errors in magnet strengths may arise from limitations in power supply stability, calibration errors, and magnetic hysteresis. Deviations of magnet strengths from their nominal values lead to beta-beating, phase advance errors, and tune shifts. In circular

colliders, such errors can move the working point closer to resonances, reducing dynamic aperture and beam lifetime. Time-dependent fluctuations in magnet currents introduce additional complications. Power supply ripple and noise generate orbit jitter and optics modulation, particularly at frequencies below a few hundred hertz. These effects are often correlated across families of magnets and can excite coherent beam motion. At the IP, even small variations in focusing strength can translate into changes in beam size or collision offset [58].

2.5.3 Ground Motion and Mechanical Vibrations

Ground motion and mechanical vibrations constitute a major source of time-dependent alignment errors in accelerator facilities. These disturbances originate from a combination of natural seismic activity, atmospheric effects, and cultural noise associated with human activity and technical infrastructure. The resulting motion spans a broad frequency range, from sub-hertz drifts to kilohertz-scale vibrations.

At low frequencies, below 1 Hz, ground motion exhibits diffusive behaviour and leads to slow relative displacement of accelerator components. A commonly used phenomenological description of this regime is the ATL law,

$$\langle \Delta x^2 \rangle = A T L, \quad (2.37)$$

where T is the elapsed time, L the separation between components, and A a site-dependent coefficient [66, 67]. Such motion results in gradual orbit drift and requires periodic correction through beam-based alignment and orbit correction systems.

At higher frequencies, ground motion is dominated by local vibrations and cultural noise. These effects are more appropriately described using power spectral density models, which vary strongly with site characteristics such as tunnel depth, local geology, and nearby technical equipment. Magnet supports and girders act as mechanical filters, suppressing some frequency components while amplifying others near structural resonances. As a result, magnet vibration spectra often exhibit narrow resonant peaks superimposed on a broadband background.

In interaction regions, with strong final-focus quadrupoles, ground-motion induced vibration can produce direct beam offset and angle jitter at the IP. This motion is difficult to correct using global orbit feedback and represents a key limitation for luminosity stability in high-performance colliders.

2.5.4 Environmental and External Perturbations

In addition to intrinsic machine imperfections, accelerators are influenced by external environmental factors. Temperature variations cause thermal expansion and contraction

of accelerator components, leading to slow drifts in alignment and magnet strength. These effects typically occur on timescales of minutes to hours and can introduce reproducible orbit changes.

Electromagnetic interference and fluctuations in the electrical power grid can couple into accelerator systems through power supplies and instrumentation. Such perturbations may manifest as coherent noise in beam position monitors or correlated orbit motion across the machine.

2.6 Control Systems

2.6.1 Feedback Loops

Accelerator orbit feedback systems form a closed control loop in which BPM measurements are processed by a controller to generate corrective kicks applied by steering magnets. The resulting beam motion depends on the response of the actuator chain and the beam transport. A typical accelerator feedback loop contains several components, represented using discrete transfer functions, such as:

1. Feedback PID controller $C(z)$,
2. Controller delay $D(z)$,
3. Power supply and magnet dynamics $G_m(z)$,
4. Beam-chamber eddy current lag $G_e(z)$,
5. Monitor (sensor) dynamics and resolution $M(z)$.

The plant transfer function represents the mapping from actuator command to beam-centroid motion, including the response of the corrector magnet power supply, vacuum chamber eddy-currents, and linear beam transport, and can be expressed as

$$P(z) = G_m(z)G_e(z) \cdots . \quad (2.38)$$

Including the controller, delay, plant and measurement response gives the open-loop transfer function

$$L(z) = C(z)D(z)P(z)M(z). \quad (2.39)$$

The performance of the feedback loop is then quantified by the sensitivity function, which characterises how strongly disturbances are transmitted through the closed loop,

$$S(z) = \frac{1}{1 + L(z)}. \quad (2.40)$$

For a disturbance $U(z)$ (e.g. an input magnetic error) entering at the plant input and transferred to the beam by $P_d(z)$, the closed-loop transfer from disturbance to beam motion is

$$\frac{Y(z)}{U(z)} = P_d(z)S(z). \quad (2.41)$$

2.6.2 Feedback and Feedforward Feedback

Feedback and feedforward are two central paradigms of control systems. Feedback corrects based on measurements of the current state, whereas feedforward corrects based on predictions of future disturbances, acting before the error occurs.

Moving Average Feedback

A simple example of a feedback is based on a moving-average over previous feedback cycles. A feedback system based on the last N feedback cycles may be expressed as:

$$u_n = -K \frac{1}{N} \sum_{k=0}^{N-1} e_{n-k}, \quad (2.42)$$

where e_n is the measured error at cycle n , K is the gain and u_n is the correction. This is feedback as the command depends on past and present measured errors only.

To move to frequency space, the unilateral z -transform, the discrete-time analogue of the Laplace transform, is applied [68]. For a discrete-time signal y_n , the z -transform is defined as

$$Y(z) = \sum_{n=0}^{\infty} y_n z^{-n}, \quad (2.43)$$

where n is the discrete time index, y_n denotes the signal in the time domain, $Y(z)$ its representation in the complex z -domain, and z is a complex variable. The transfer function of this system is then expressed as

$$G_{\text{MA}}(z) = -K \frac{1 - z^{-N}}{N(1 - z^{-1})}. \quad (2.44)$$

This function has the form of an N -point moving-average finite impulse response (FIR) filter [69]. The maximum gain is at $\omega = 0$, and higher frequencies are attenuated, therefore a low-pass filter. The primary drawback of this system is the introduction of phase lag at higher frequencies. This system cannot cancel disturbances that are fast relative to loop delay.

Linear Prediction Feedforward

A simple example of a feedforward is linear prediction. An estimate of the error at a future cycle (τ cycles ahead) is calculated from past measurements:

$$u_n = -K\hat{e}_{n+\tau} \quad \hat{e}_{n+\tau} = \sum_{k=0}^{N-1} \alpha_k e_{n-k}, \quad (2.45)$$

where $\hat{e}_{n+\tau}$ is the predicted error at cycle $n + \tau$ and α_k are the prediction coefficients. As the predicted command is applied without waiting for the measured error, the action is feedforward, acting on the predicted disturbance. This enables the system to cancel predictable components before they appear in the measured error. The predictor provides phase lead and can overcome delay-limited rejection. The drawbacks here are the sensitivity to model error and measurement noise. Good predictors require consistent, repeatable disturbance structure or robust estimation of coefficients $\{\alpha_k\}$.

Hybrid Approaches

It is possible to combine the two paradigms of feedback and feedforward control. For example, a predictor embedded inside a closed loop yields a predictive controller, offering the benefits of anticipation with feedback-based robustness. A hybrid approach may yield suppression of repeatable disturbances while retaining stability under uncertainty.

2.6.3 PID Feedback

Control action is commonly split into proportional, integral and derivative parts. Each term targets a distinct control objective and frequency band. PID controls may be described in two ways: discrete time (position form) or incrementally updating (velocity) form [70]. The position form description of a PID system is expressed as

$$u_n = K_P e_n + K_I \sum_{k=0}^n e_k + K_D (e_n - e_{n-1}), \quad (2.46)$$

where e_n is the error at cycle n and u_n the actuator command. The velocity form description of a PID system is expressed as

$$\Delta u_n = u_n - u_{n-1} = K_P (e_n - e_{n-1}) + K_I e_n + K_D (e_n - 2e_{n-1} + e_{n-2}). \quad (2.47)$$

The proportional term gives an immediate correction proportional to the error. The integral term removes steady-state offsets by accumulating persistent error. The derivative term provides phase lead and predictive damping to counteract resonances and delay. Suppression, robustness and noise sensitivity are chosen by adjusting K_P , K_I and K_D (the proportional, integral and derivative gains) independently.

In frequency domain a PID controller is expressed as

$$G_{\text{PID}}(z) = K_P + K_I/z + K_D z. \quad (2.48)$$

From this expression, one can identify the phase contribution of each aspect. Proportional has a flat magnitude (0 dB/dec slope), with 0° phase lag. Integral has a decreasing magnitude $\propto 1/\omega$ (-20 dB/dec slope), with -90° phase lag. Derivative has an increasing magnitude $\propto \omega$ ($+20$ dB/dec slope), with $+90^\circ$ phase lead.

Splitting the control into these three aspects makes design modular. One can tune the low-frequency behaviour (I), the mid-band gain (P), and the phase margin (D) separately. Larger K_P increases suppression but reduces stability margin. Larger K_I improves low-frequency rejection but can introduce slow oscillatory modes if too large. K_D can improve phase margin and reduce resonant peaks, but excessive derivative gain amplifies noise. PID is simple to implement in hardware and software, making it ubiquitous in control systems.

2.6.4 Feedback Stability

For a system described by the open loop transfer function $L(z)$, the closed-loop system is stable if all roots of the characteristic equation lie inside the unit circle in the z -plane [71]. The characteristic equation is given by

$$1 + L(z) = 0. \quad (2.49)$$

For a sampled (e.g. circular accelerator turn-by-turn) system, stability requires every pole z_i to satisfy $|z_i| < 1$.

Instability is commonly caused by delay and phase lag. If the control loop phase approaches -180° at a frequency where the gain of the system is positive $|L(e^{i\omega})| \geq 1$ then the corrective action becomes driving action. The loop is therefore marginal when, for some ω ,

$$L(e^{i\omega}) = -1. \quad (2.50)$$

At this frequency the closed-loop denominator vanishes and oscillations are driven. In discrete-time systems this can happen at frequencies near tune harmonics where the resonant peaks of the plant function occur.

2.6.5 Transfer Functions

A transfer function, H characterises the input to output behaviour of a linear time-invariant system in the frequency domain under zero initial conditions [68]. Zero initial conditions is the constraint that the system's initial stored energy is zero. For a

second-order mechanical oscillator (sinusoidal motion), the initial offset and velocity is zero. This excludes analysis of any ‘residual’ motion of the system.

For a linear, time-invariant system, driven with an input $X(t)$ and yielding an output $Y(t)$, the transfer function is given by

$$H(s) = \frac{Y(s)}{X(s)}, \quad (2.51)$$

where $X(s) = \mathcal{L}\{X(t)\}$ is the Laplace transform applied to the input and output respectively.

The magnitude of $H(s)$ identifies how strongly a given frequency component of a driving force will be transferred or suppressed. As such, the poles of $H(s)$ (denominator roots) correspond to the system’s natural modes and the zeros of $H(s)$ (numerator roots) are frequencies fully suppressed by the filter.

The argument of $H(s)$ relates the phase shift for a given input frequency to the frequency component of the output oscillation

2.6.6 Power Spectral Density

A power spectral density (PSD) describes the distribution of a signal’s power in frequency space [72]. A PSD has units of power per hertz. For example, a vibration PSD has units of m^2/Hz . The integral of the PSD returns the variance of the signal.

For a process described by $X(t)$, the power spectral density is defined as

$$S_{XX}(\omega) = \langle X(\omega) X^*(\omega) \rangle = \int_{-\infty}^{\infty} \langle X(t) X(t + \tau) \rangle e^{i\omega\tau} d\tau \quad (2.52)$$

where $\langle \cdot \rangle$ denotes the expectation value, or equivalently the ensemble average.

Interplay with Transfer Functions

During systems analysis, we may know the PSD of an input signal, but be interested in the PSD of the output motion. The relationship between the output observable $X(\omega)$ and the driving force $F(\omega)$ is given by a transfer function $H(\omega)$ as

$$X(\omega) = H(\omega)F(\omega). \quad (2.53)$$

Substituting this into the definition of the Power Spectral Density:

$$S_{XX}(\omega) = \langle X(\omega) X^*(\omega) \rangle \quad (2.54)$$

$$= \langle H(\omega) F(\omega) [H(\omega) F(\omega)]^* \rangle \quad (2.55)$$

$$= H(\omega) H^*(\omega) \langle F(\omega) F^*(\omega) \rangle \quad (2.56)$$

$$= H(\omega) H^*(\omega) S_{FF}(\omega), \quad (2.57)$$

the overall impact of a transfer function on a power spectral density is given by

$$S_{XX}(\omega) = |H(\omega)|^2 S_{FF}(\omega). \quad (2.58)$$

As the gain function is complex, multiplying $H(\omega)$ by its conjugate removes phase, yielding the power transmission factor. Furthermore, the transfer function acts in amplitude, identifying the gain of the system at a given frequency, where the PSD identifies the power at a given frequency.

2.6.7 Mechanical Resonators

A mechanical oscillator may be modelled as a mass m coupled to an ideal spring of stiffness k with a viscous damping coefficient of b . The equation of motion is then:

$$F(t) = m\ddot{x}(t) + b\dot{x}(t) + kx(t) \quad (2.59)$$

where $F(t)$ is an external driving force. Taking the Fourier transform to move to frequency domain, the transfer function can be identified:

$$X(\omega) = H(\omega) F(\omega) \quad (2.60)$$

$$= \frac{1}{k - m\omega^2 + i b \omega} F(\omega) \quad (2.61)$$

$$= \frac{1}{m(\omega_0^2 - \omega^2 + i\gamma\omega)} F(\omega) \quad (2.62)$$

where the natural (resonant) angular frequency, ω_0 is given by $\sqrt{k/m}$ and the damping rate γ is given by b/m .

Contribution to the Power Spectral Density

The contribution of a mechanical resonator to a PSD is then equal to the product of the force and the square of the magnitude of its transfer function $|H(\omega)|^2$

$$S_{XX}(\omega) = |H(\omega)|^2 S_{FF}(\omega) \propto \frac{1}{(\omega_0^2 - \omega^2)^2 + (\gamma\omega)^2} S_{FF}(\omega) \quad (2.63)$$

For a white noise spectrum, we take the drive spectrum $S_{FF}(\omega) = c$ where c is a constant. This function quickly decays, and so approximating near the resonant peak:

$$\omega \approx \omega_0 \implies \omega_0^2 - \omega^2 = (\omega_0 - \omega)(\omega_0 + \omega) \approx 2\omega_0(\omega_0 - \omega) \quad (2.64)$$

yields the contribution

$$S_{XX}(\omega) \propto \frac{1}{(\omega - \omega_0)^2 + \left(\frac{\gamma}{2}\right)^2}, \quad (2.65)$$

which takes the form of a Lorentzian peak, with a full-width at half-maximum of γ .

2.6.8 Driven Oscillator Transfer Functions

A circular accelerator undergoing a driven error can be modelled as a discrete-time second-order resonator [73]. The transverse motion of a particle beam near a betatron tune is itself a harmonic oscillator sampled once per turn. A further external driver acts as an additional input frequency.

For the IP offset, $(y_1 - y_2)$, we consider the π mode. Starting from the discrete oscillator equation for the π -mode oscillation,

$$\begin{bmatrix} y_{n+1} \\ y'_{n+1} \end{bmatrix} = \mathbf{R}_\pi \begin{bmatrix} y_n \\ y'_n \end{bmatrix}, \quad (2.66)$$

the prime coordinate can be eliminated to yield the second-order recursion relation:

$$y_{n+1} - 2 \cos(2\pi Q_\pi) y_n + y_{n-1} = F_n, \quad (2.67)$$

where F_n is an external driving term, for example a dipolar error. This is the discrete analogue of a driven harmonic oscillator. To move to frequency space, the unilateral z -transform is applied, yielding

$$(1 - 2 \cos(2\pi Q_\pi) z^{-1} + z^{-2}) Y(z) = F(z). \quad (2.68)$$

The transfer function can then be identified as

$$H(z) = \frac{Y(z)}{F(z)} = \frac{1}{1 - 2 \cos(2\pi Q_\pi) z^{-1} + z^{-2}} \quad (2.69)$$

where z is the discrete-time complex variable $z = e^{i\omega}$ and $\omega = 2\pi f/f_{\text{rev}}$ is the discrete-time frequency in radians per turn. This can be simplified using the trigonometric identity $e^{i\omega} + e^{-i\omega} = 2 \cos(\omega)$, reducing the expression to

$$|H(f)| = \frac{1}{2 |\cos(\omega) - \cos(2\pi Q_\pi)|} \quad (2.70)$$

with resonant peaks occurring where $\cos(\omega) = \cos(2\pi Q_\pi)$, or equivalently

$$\omega = \pm 2\pi q_\pi + 2n\pi \quad \Rightarrow \quad f = (n \pm q_\pi)f_{\text{rev}}, \quad (2.71)$$

where q_π is the fractional part of the π -mode tune $q_\pi = Q_\pi \bmod 1$.

2.6.9 The Decibel Scale

The bel (B) is defined as the base-10 logarithm of a power ratio:

$$\text{Bel} = \log_{10} \left(\frac{P}{P_{\text{ref}}} \right), \quad (2.72)$$

where P is the signal power and P_{ref} is a reference power. The decibel (dB) is one tenth of a bel.

For many comparisons, we consider amplitudes. As power scales with the square of amplitude, $P \propto A^2$, the ratio expressed in decibels yields:

$$[\text{dB}] = 10 \log_{10} \left(\frac{P}{P_{\text{ref}}} \right) = 10 \log_{10} \left(\frac{A^2}{A_{\text{ref}}^2} \right) = 20 \log_{10} \left(\frac{A}{A_{\text{ref}}} \right), \quad (2.73)$$

where A is the oscillation amplitude of interest and A_{ref} is the reference driven-error amplitude.

2.7 Accelerator Feedback Systems

Accelerator feedback systems maintain beam stability by measuring deviations from the desired operating point and applying corrections. For a static offset the objective is rapid decay of the error toward zero. For a driven disturbance the objective is attenuation: minimise the residual amplitude when removal of the source is not possible. Correction is limited by sensor resolution, calibration, processing latency, and the finite bandwidth of the correction hardware.

2.7.1 Error Detection

Error detection converts beam motion into signals suitable for control. In circular colliders the primary diagnostic devices for transverse feedback are beam position monitors (BPMs), which provide bunch centroid measurements. Complementary diagnostics include synchrotron-light monitors (beam size and position), beam loss monitors (protection and collimation-related signals), tune and coupling measurements, and dedicated luminosity monitors at the IP. For IP feedback, the relevant ‘error’ is the relative overlap of the colliding beams; consequently, diagnostics that are directly sensitive to collision

conditions (e.g. luminosity, beamstrahlung or outgoing angle/position changes induced by collision) can provide particularly effective error signals. Two features of detection are especially important for feedback design: (1) resolution and noise, and (2) latency and synchronisation.

2.7.2 Signal Processing

Raw diagnostics are rarely used directly as feedback inputs. Signal processing typically performs:

1. Calibration and normalisation.
2. Removal of known systematic components (offsets, gain errors, timing skews).
3. Filtering to remove spectral components outside the actuator bandwidth.

In accelerator applications, filtering is often implemented digitally with Finite-Impulse-Response (FIR) structures [74], designed to achieve sufficient rejection at noise-dominated frequencies while preserving phase margin.

A second role of signal processing is computation of derived quantities that cannot be measured directly. For example, deriving offsets at the IP from multiple BPMs' measurements and an optics model.

2.7.3 Correction Hardware

Correction hardware converts the controller output into an electromagnetic action on the beam. For transverse motion the primary actuators are dipole correctors/kickers. The achievable feedback bandwidth is typically limited by the combined response of the actuator, power supply, and eddy-current effects.

For local feedback, correction is typically applied through closed orbit bumps formed by multiple correctors [75]. Ideally a bump is fully closed in orbit and dispersion so that correction does not leak, though this may not be realised in practice.

2.7.4 Types of Accelerator Feedback

Accelerator feedback systems can be classified by their objective, timescale, and locality. The categories below are not exhaustive, but they capture the dominant feedback families used in modern rings.

Closed Orbit Feedback

Closed Orbit feedback (COFB) stabilises the stored-beam trajectory against drifts and low-frequency disturbances. COFB is typically used to address orbit deviations due

to slow oscillations such as thermal drifts, ground motion, and slow magnet/power-supply variations. COFB is a multi-input multi-output (MIMO) problem: many BPM signals are mapped to many correctors, commonly via response-matrix methods with regularisation to control noise amplification and maintain robustness.

IP Feedback

IP feedback aims to maintain collision by correcting relative beam offsets at the IP. Because luminosity depends strongly on beam overlap when β^* is small, the IP is often the most sensitive location in the machine to orbit jitter and drifts. Unlike COFB, IP feedback is typically local and must respect tight constraints imposed by the complex optics and coupling associated with the interaction region.

Transverse Dampers and Instability Feedback

Transverse bunch-by-bunch feedback (transverse dampers) suppresses coherent instabilities and injection oscillations by applying fast kicks synchronised to the bunch pattern. Their objective is to damp unstable coupled-bunch modes and limit emittance growth from coherent motion. These systems operate at high bandwidth and are strongly constrained by latency, kicker impedance, and digital processing throughput.

RF Feedback

RF feedback regulates the accelerating voltage and phase to stabilise the longitudinal dynamics. Typical objectives include suppression of cavity field perturbations from beam loading, reduction of phase noise, control of coupled-bunch longitudinal instabilities, and stabilisation of bunch RF phase, length and synchrotron tune. RF feedback often combines multiple loops each operating on different timescales.

2.7.5 Interaction Point Feedback

Interaction Point feedback stabilises the collision by correcting offsets and/or angles of one beam with respect to the other at the IP. To maximise luminosity, relative offset must be minimised locally, avoiding perturbing the rest of the machine. This is achieved with dedicated local bumps on one or both beams, together with diagnostics in the IR sensitive to the collision state. IP feedback for future linear colliders has been widely explored, such as for CLIC [76] and for ILC [77].

Two conceptually distinct strategies are widely used:

- Beam–beam deflection based feedback.
- Luminosity optimisation by modulation feedback (dithering).

Beam–Beam Deflection Feedbacks

For small offsets about nominal collision, the beam–beam induced deflection can be approximated as linear in relative displacement as discussed in Section 2.3.3. The relative offset of the beams in collision can be reconstructed from measurements at downstream BPMs and known optics. A full derivation of this reconstruction is presented in Section 2.7.6. A correction bump is then applied to reduce the inferred offset on subsequent turns. Because the measured signal is derived from the collision itself, care is required to manage noise amplification, to handle coupling and dispersion in the interaction region, and to ensure stability in the presence of loop delay. Beam–beam deflection feedback has been successfully implemented at SLC [78] and SuperKEKB [79].

Dithering Feedbacks

Luminosity dithering feedback treats luminosity as the performance metric and applies small, controlled modulations of one or more knobs (e.g. transverse offsets, waist position, coupling knobs) about their nominal values. Since luminosity is a scalar function, a single measurement is insufficient to determine which side of the optimum the operating point is on. Dithering resolves this ambiguity by introducing a known modulation and extracting the response at the modulation frequency, estimating the local gradient of luminosity with respect to the chosen knob. The knob is then steered in the direction that increases luminosity. While dithering feedback is robust to modelling errors and can optimise directly on physics performance, it necessarily introduces modulation that can reduce average luminosity. Luminosity dithering feedback has been successfully implemented at PEP II [80], SLC [81] and SuperKEKB [82].

2.7.6 Linear Prediction Feedback

Linear prediction feedback assumes that future beam offsets can be predicted from previous-turn measurements. Each feedback cycle produces a prediction for a turn some turns ahead, with the prediction horizon advancing by one turn per cycle. In this derivation, prediction for the next feedback cycle is analysed.

For feedback cycle n the offset is measured to be Δy_n . This offset, and offsets for previous turns, are stored in memory. Linear prediction feedback assumes that there is a linear transform relating the measurements of previous turns. For a system storing the results of $M + N$ previous measurements, and predicting based on the most recent

N measurements, this is expressed as

$$\begin{pmatrix} \Delta y_{n-1} & \Delta y_{n-2} & \cdots & \cdots & \Delta y_{n-N} \\ \Delta y_{n-2} & \Delta y_{n-3} & \cdots & \cdots & \Delta y_{n-N-1} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \Delta y_{n-M} & \Delta y_{n-M-1} & \cdots & \cdots & \Delta y_{n-M-N+1} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ \vdots \\ c_N \end{pmatrix} = \begin{pmatrix} \Delta y_n \\ \Delta y_{n-1} \\ \vdots \\ \vdots \\ \Delta y_{n-M+1} \end{pmatrix}, \quad (2.74)$$

where the coefficients c_i encompass the linear relationship. This equation is more compactly expressed by the matrix equation

$$\mathbf{A}\vec{c} = \vec{b}. \quad (2.75)$$

For robustness the system must be overdetermined such that $M > N$. It is in general not possible to solve (2.75). To obtain a working solution it is reformulated as a linear least squares problem, expressed by

$$\min \|\mathbf{A}\vec{c} - \vec{b}\|^2. \quad (2.76)$$

Solving for $\vec{c}$, yields a set of prediction coefficients that are then used to predict the offset at the next feedback cycle, based on the most recent N measurements as

$$\Delta y_{n+1} = \sum_{k=1}^N c_k \Delta y_{n-k+1} \quad (2.77)$$

where Δy_{n+1} is the predicted offset in the next feedback cycle.

Normal Equations Formalism

Least squares problems are often solved by translating the problem into normal equations. From Eq. (2.75) the normal equations can be identified as

$$\mathbf{A}^T \mathbf{A} \vec{c} = \mathbf{A}^T \vec{b}. \quad (2.78)$$

This approach has a number of benefits [83]:

- The problem is reduced to an $N \times N$ linear system, which should be smaller than the original $M \times N$ system.
- By construction, the matrix $A^T A$ is symmetric and positive semidefinite, allowing efficient solvers (Cholesky, conjugate gradient, etc.).
- The formulation clarifies uniqueness: if A has full column rank then $A^T A$ is positive definite and the solution is unique.

Regularisation

When A is ill-conditioned or rank-deficient, $A^T A$ is nearly singular and the normal-equation solution can be unstable. This can occur when data is noisy or the condition that $M > N$ breaks down due to degeneracy in the measurements. To address this, it is useful to add Tikhonov regularisation (ridge regularisation) [84].

$$\left(\mathbf{A}^T \mathbf{A} + \varepsilon \mathbf{I}\right) \vec{c} = \mathbf{A}^T \vec{b}, \quad \varepsilon > 0. \quad (2.79)$$

Where ε is chosen to be small but non-zero.

The benefits of regularisation are:

- The term εI shifts the eigenvalues of $A^T A$ away from zero, improving conditioning and making the system easier and more stable to solve.
- Over-fitting control by penalising large coefficients, which is important when data is noisy or when M is not large relative to N .

2.7.7 Deriving IP Offsets from IP BPM Measurements

During the beam–beam interaction, the force applied by the beam–beam interaction results in a change to the outgoing vertical angle y'^{*b} . The impact of the offsets at the IP can be computed by calculating the expected momentum at the IP from either side and taking the difference.

Assumptions

In this derivation, the offset between colliding bunches is calculated from four BPM measurements local to the IP. The four BPMs are assumed to be located one upstream and one downstream of the IP for each beam. It is assumed that the BPMs provide measurements of the beam centroids in the transverse plane $(\bar{x}, \bar{y})$.

Bunches are assumed to be colliding at the design energy with well corrected longitudinal timing. It is further assumed that there are no elements between the BPMs and the IP that strongly impact the longitudinal plane. Therefore the values of zeta and delta are assumed to be zero at the IP and at the BPMs.

Finally, as this computation is related to feedback it is assumed that this is the regime of small offsets. The beam–beam kick is therefore well approximated by the linear beam–beam force

$$\Delta a' = \pm \frac{1}{2} k_a \Delta a \quad k_a = \frac{4\pi}{\beta_a} \xi_a \quad a \in x, y, \quad (2.80)$$

where ξ_a is the beam–beam tune shift in plane a.

Required Transfer Matrices

Considering a single bunch, the bunch passes through the incoming BPM (denoted ‘A’) to the IP (denoted ‘*’) and out through the outgoing BPM (denoted ‘B’). From simulation of the lattice, the transfer matrices between the IP and the BPMs are obtained. These matrices can then be used to relate the positions at the IP to the positions at the BPMs, expressed as

$$\vec{x}^A = \mathbf{A}\vec{x}^* \quad \vec{x}^B = \mathbf{B}^{-1}\vec{x}^*, \quad (2.81)$$

where the transfer matrix from the IP to BPM ‘A’ is $\mathbf{A}$, the transfer matrix from BPM ‘B’ to the IP is $\mathbf{B}$ and the 6D coordinates of the particles at position a is denoted $\vec{x}^a$.

Reducing to Only BPM Observables

As the BPM observables are only the transverse beam centroids $(\bar{x}, \bar{y})$ the number of computations is reduced. Considering the measurements for a single beam at its ingoing and outgoing BPMs, the computation is the following:

$$x^M = \mathbf{M}_{11}x^* + \mathbf{M}_{12}x'^* + \mathbf{M}_{13}y^* + \mathbf{M}_{14}y'^{*m} + \mathbf{M}_{15}\zeta^* + \mathbf{M}_{16}\delta^* \quad (2.82)$$

$$y^M = \mathbf{M}_{31}x^* + \mathbf{M}_{32}x'^* + \mathbf{M}_{33}y^* + \mathbf{M}_{34}y'^{*m} + \mathbf{M}_{35}\zeta^* + \mathbf{M}_{36}\delta^* \quad M \in A, B. \quad (2.83)$$

Reducing to Solely Transverse Motion

As it is assumed that the variation of zeta and delta between the IP and BPMs are approximately zero, the equations are reduced to

$$x^M = \mathbf{M}_{11}x^* + \mathbf{M}_{12}x'^* + \mathbf{M}_{13}y^* + \mathbf{M}_{14}y'^{*m} \quad (2.84)$$

$$y^M = \mathbf{M}_{31}x^* + \mathbf{M}_{32}x'^* + \mathbf{M}_{33}y^* + \mathbf{M}_{34}y'^{*m} \quad M \in A, B. \quad (2.85)$$

Coordinate Decomposition with Weak Coupling

It is typical, when considering small orbital errors, that the coupling between the vertical and horizontal plane remains small compared to the transfer coefficients within the plane. With this assumption, it is possible to decouple the two transverse planes, such that in the vertical plane

$$y^M = \mathbf{M}_{33}y^* + \mathbf{M}_{34}y'^{*m} \quad M \in A, B. \quad (2.86)$$

Vertical Position Offset Computation

A vertical position offset at the IP induces a beam–beam kick. This kick is given by the difference in y' before and after the collision. The values of y' are expressed as:

$$y'^{*|M} = \frac{1}{\mathbf{M}_{34}} \left(y^M - \mathbf{M}_{33} y'^{*|M} \right) \quad M \in A, B, \quad (2.87)$$

where the values at the IP are labelled distinctly. It is noted that the IP position cannot change during the interaction, while the momentum does, and as such the value of the IP position is denoted equally as

$$y'^{*|A} = \frac{1}{A_{34}} \left(y^A - A_{33} y^* \right) \quad y'^{*|B} = \frac{1}{B_{34}} \left(y^B - B_{33} y^* \right). \quad (2.88)$$

From this, computing the difference at the IP:

$$\Delta y'^* = y'^{*|B} - y'^{*|A} \quad (2.89)$$

$$= \frac{1}{B_{34}} \left(y^B - B_{33} y^* \right) - \frac{1}{A_{34}} \left(y^A - A_{33} y^* \right) \quad (2.90)$$

$$= \left(\frac{y^B}{B_{34}} - \frac{y^A}{A_{34}} \right) + \left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}} \right) y^*. \quad (2.91)$$

If we now consider the motion of bunch 2, with transfer matrix $\mathbf{C}$ denoting the transfer between the IP and bunch 2's incoming BPM (analogous to $\mathbf{A}$), and $\mathbf{D}$ denoting the transfer between the IP and bunch 2's outgoing BPM (analogous to $\mathbf{B}$):

$$\Delta y_1'^* = \left(\frac{y^B}{B_{34}} - \frac{y^A}{A_{34}} \right) + \left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}} \right) y_1^*, \quad (2.92)$$

$$\Delta y_2'^* = \left(\frac{y^D}{D_{34}} - \frac{y^C}{C_{34}} \right) + \left(\frac{C_{33}}{C_{34}} - \frac{D_{33}}{D_{34}} \right) y_2^*. \quad (2.93)$$

From these two pieces, the overall “canonical” kick can now be calculated from the vertical positions across the 4 BPMs as

$$\Delta y_{\text{canonical}}'^* = \frac{\left(\frac{y^A}{A_{34}} - \frac{y^B}{B_{34}} \right)}{\left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}} \right)} - \frac{\left(\frac{y^C}{C_{34}} - \frac{y^D}{D_{34}} \right)}{\left(\frac{C_{33}}{C_{34}} - \frac{D_{33}}{D_{34}} \right)} \quad (2.94)$$

$$= \frac{\Delta y_1'^*}{\left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}} \right)} - \frac{\Delta y_2'^*}{\left(\frac{C_{33}}{C_{34}} - \frac{D_{33}}{D_{34}} \right)} + (y_1^* - y_2^*). \quad (2.95)$$

Now considering the strength of the beam–beam kick in terms of known parameters,

$$\Delta y_1'^* = -k_{1,y} \Delta y^* \quad (2.96)$$

$$\Delta y_2'^* = +k_{2,y} \Delta y^*, \quad (2.97)$$

the canonical kick can now be further simplified

$$\Delta y_{\text{canonical}}'^* = \frac{\Delta y_1'^*}{\left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}}\right)} - \frac{\Delta y_2'^*}{\left(\frac{C_{33}}{C_{34}} - \frac{D_{33}}{D_{34}}\right)} + (y_1^* - y_2^*) \quad (2.98)$$

$$= \frac{-k_{1,y} \Delta y^*}{\left(\frac{A_{33}}{A_{34}} - \frac{B_{33}}{B_{34}}\right)} - \frac{+k_{2,y} \Delta y^*}{\left(\frac{C_{33}}{C_{34}} - \frac{D_{33}}{D_{34}}\right)} + \Delta y^* \quad (2.99)$$

$$= -\left(\frac{2k_{1,y}}{m_{1,y}} + \frac{2k_{2,y}}{m_{2,y}} - 1\right) \Delta y^* \quad (2.100)$$

$$= -K_y \Delta y^*, \quad (2.101)$$

where K_y is fully determined from BPM readings, lattice transfer matrices, and beam parameters.

Vertical Angle Calculation

The vertical angle of interest for correction is the incident vertical angle, $y_1'^{*a}$ and $y_2'^{*a}$. The relationship between the angle before and after collision with the beam–beam force is given by

$$\Delta y_1'^* = y_1'^{*b} - y_1'^{*a}, \quad (2.102)$$

from which one can eliminate the vertical angle after collision, resulting in

$$\frac{y_1^A}{A_{34}} + \frac{y_1^B}{B_{34}} = \left(\frac{A_{33}}{A_{34}} + \frac{B_{33}}{B_{34}}\right) y_1^* + 2y_1'^{*b} + \Delta y_1'^* \quad (2.103)$$

$$\frac{y_2^C}{C_{34}} + \frac{y_2^D}{D_{34}} = \left(\frac{C_{33}}{C_{34}} + \frac{D_{33}}{D_{34}}\right) y_2^* + 2y_2'^{*b} + \Delta y_2'^*, \quad (2.104)$$

and finally an analogous angular offset incident at the IP of

$$\theta_y = \frac{\left(\frac{y_e^A}{m_{34}^A} + \frac{y_e^B}{m_{34}^B}\right)}{\left(\frac{m_{33}^A}{m_{34}^A} + \frac{m_{33}^B}{m_{34}^B}\right)} - \frac{\left(\frac{y_p^C}{m_{34}^C} + \frac{y_p^D}{m_{34}^D}\right)}{\left(\frac{m_{33}^C}{m_{34}^C} + \frac{m_{33}^D}{m_{34}^D}\right)} \quad (2.105)$$

$$= \frac{2y_e'^{*b}}{\left(\frac{m_{33}^A}{m_{34}^A} + \frac{m_{33}^B}{m_{34}^B}\right)} - \frac{2y_p'^{*b}}{\left(\frac{m_{33}^C}{m_{34}^C} + \frac{m_{33}^D}{m_{34}^D}\right)}. \quad (2.106)$$

2.8 Modelling of Particle Accelerators

The design and optimisation of particle accelerators relies heavily on numerical modelling. While analytic descriptions provide essential physical insight, they can be insufficient to capture the combined effects of non-linear lattice dynamics and collective phenomena. Accelerator studies therefore employ a hierarchy of modelling approaches, ranging from simplified transfer maps to full lattice tracking with self-consistent beam–beam forces.

2.8.1 Transfer Maps and Full Lattice Tracking

Beam dynamics in accelerators may be described using symplectic maps that transport a particle’s phase-space coordinates between locations in the machine. In simplified models, the action of many lattice elements or even an entire turn is approximated by a single linear or low-order non-linear transfer map. Such truncated descriptions are computationally efficient and are well suited to analytic studies and early-stage optimisation, particularly when the modelling is intentionally simplified to isolate specific physical effects, such as beam–beam interactions.

In full lattice tracking, particles are instead transported element-by-element through the accelerator, with each element represented by its own symplectic map. This avoids truncation of the underlying dynamics and allows local perturbations to propagate naturally around the ring. Such models are essential for realistic studies of stability, dynamic aperture, and feedback performance.

2.8.2 Beam–Beam Interaction Modelling

The electromagnetic interaction between colliding bunches at an IP is a dominant non-linear effect in high-luminosity lepton colliders. Accurate modelling of the beam–beam force is therefore essential for understanding beam stability, luminosity performance, and feedback operation [85]. Several modelling approaches are commonly employed, differing in physical fidelity and computational cost.

Soft Gaussian Approximation

In many beam–beam models, the interaction at the IP is evaluated assuming a Gaussian transverse charge distribution throughout the collision, known as the soft Gaussian approximation [86]. Within this framework, the beam is represented by a continuous transverse charge density characterised by its RMS beam sizes. The resulting electromagnetic fields, and hence the transverse kicks on particles in the opposing beam, can be computed analytically from Poisson’s equation. The force varies smoothly with particle amplitude and is independent of the microscopic particle distribution.

The soft Gaussian approximation neglects non-Gaussian deformation of the transverse charge distribution during the interaction and therefore does not capture effects associated with strong disruption or the development of non-Gaussian beam profiles. It nevertheless provides an accurate description of incoherent beam–beam forces when the beam distributions remain close to Gaussian and the beam–beam parameters are moderate, as has been tested for the LHC [87]. This approximation is computationally efficient and captures many essential features of beam–beam dynamics, including tune shifts, non-linear focusing, and coherent deflections, making it well suited to long-term tracking studies.

Weak–Strong Method

In the Weak–Strong method, one beam (the strong beam) is assumed to have a fixed, prescribed charge distribution, while the other beam (the weak beam) is tracked through the resulting electromagnetic field. The strong beam is unaffected by the interaction, and the beam–beam force acting on the weak beam is typically computed analytically using the soft-Gaussian approximation.

It is widely used for parameter scans, tolerance studies, and early-stage feedback investigations. However, it neglects the mutual interaction between the two beams and cannot describe coherent beam–beam modes. The weak–strong approximation is widely used in beam–beam studies of circular colliders due to its favourable balance between physical fidelity and computational efficiency [85].

Strong–Strong Method

In the Strong–Strong method, both beams are treated dynamically, and the beam–beam force is computed self-consistently from their evolving distributions. This allows coherent beam–beam oscillation modes and coupled beam dynamics to be studied. The strong–strong model is therefore essential when investigating feedback systems or when coherent effects play a significant role.

Strong–strong simulations are computationally more demanding than weak–strong simulations, particularly when many turns or large particle populations are required. As a result, they are typically reserved for detailed studies once the relevant parameter space has been identified.

Particle-In-Cell Methods

The most general treatment of the beam–beam interaction employs Particle-In-Cell (PIC) methods, in which the electromagnetic fields are computed numerically on a grid from the particle distribution. This approach allows arbitrary beam shapes, crossing

angles, hourglass effects, and strong disruption to be modelled without relying on analytic field expressions.

PIC simulations provide the highest level of physical fidelity but are computationally intensive. They are commonly used for single-collision studies, luminosity optimisation, and beamstrahlung calculations, rather than for long-term multi-turn tracking.

2.8.3 Numerical Tools and Simulation Codes

These modelling approaches are implemented in a variety of accelerator simulation codes. Several established tools are used in this thesis.

Xsuite is a modern, Python-based accelerator tracking framework designed for high-performance simulations and large-scale parameter scans [88]. It supports full lattice tracking, symplectic integration, weak-strong, strong-strong and PIC beam-beam modelling. Its modular structure and native compatibility with contemporary computing environments make it particularly well suited for the studies presented in this thesis. Unless otherwise stated, simulations in the following chapters use Xsuite.

The Strategic Accelerator Design code (SAD) [89] has a long history in the design and operation of electron-positron colliders and has been extensively used for optics design and beam-beam studies. It provides both transfer map capabilities and full lattice tracking, and has played a central role in the development of SuperKEKB.

The GUINEA-PIG code [90, 91] is widely used for detailed modelling of beam-beam interactions at linear and circular colliders, particularly when strong disruption and beamstrahlung effects are important. It employs PIC techniques and is therefore used in this work primarily for single-interaction studies and benchmarking.

Chapter 3

Analytic IP Feedback Modelling

3.1 Model Overview

3.1.1 Model Aims

The performance of an IP feedback system depends on a myriad of factors. To model the response to low frequency errors for accelerators with kilohertz revolution frequencies thousands of turns must be simulated. High-fidelity simulations including lattice tracking, quantum radiation modelling and self-consistent beam-beam are computationally prohibitive.

To facilitate rapid exploration for design, a simplified analytic model has been developed. This enables studies into the optimal type of feedback (PID), feedback gains, and more. The priorities are computational efficiency and conceptual transparency. The primary aim of this model is to evaluate and compare different feedback implementations in multiple IP configurations. The motion of bunch centroids is computed in response to driven errors over ten thousand turns. The IP offset between the bunch centroids is a key figure of merit as this directly impacts the luminosity.

The system is analysed from a control-systems and signal-processing viewpoint. The primary metric is the frequency-domain response (closed-loop transfer function) from driven error to centroid motion from which the integrated (RMS) centroid motion amplitude is also computed to provide a quantitative metric for suppression performance.

3.1.2 Computational Model

Lattice

For computational efficiency, the lattice between two IPs is modelled by a single transfer matrix, as:

$$\vec{x}_{n+1} = \mathbf{M}_{\text{IP} \rightarrow \text{IP}} \vec{x}_n. \quad (3.1)$$

This approach assumes linear optics.

Tracking is performed in the transverse plane only (“4D tracking”) and longitudinal effects are neglected. Modelling the effects of the longitudinal plane requires consideration of non-linear dynamics. Given the linear optics assumption of the transfer matrix approach, this is not feasible in this model.

The IP to IP transfer matrix is constructed from the IP beta functions and the global tunes. With multiple IPs, it is assumed that all arcs have the same phase advance, and that all IPs have the same beta functions and are at waist ($\alpha_{x,y}^* = 0$). The IP to IP transfer matrix in the standard Courant-Snyder form is given by:

$$\mathbf{M}_{\text{IP} \rightarrow \text{IP}} = \begin{pmatrix} \cos \mu_x & \beta_x^* \sin \mu_x & 0 & 0 \\ -\sin \mu_x / \beta_x^* & \cos \mu_x & 0 & 0 \\ 0 & 0 & \cos \mu_y & \beta_y^* \sin \mu_y \\ 0 & 0 & -\sin \mu_y / \beta_y^* & \cos \mu_y \end{pmatrix}, \quad (3.2)$$

where $\beta_{x,y}^*$ are the IP beta functions and $\mu_{x,y}$ are the phase advances IP to IP, given by $\mu_{x,y} = 2\pi Q_{x,y} / N_{\text{IP}}$, where Q is the tune. This IP to IP transfer matrix is symplectic. By construction, this is a decoupled transfer matrix, where the horizontal and vertical planes are treated independently.

Apertures

An aperture check is applied once per turn. Particles are retained only if the magnitude of the transverse orbit is less than one metre. This criterion is introduced to suppress numerical instabilities arising from exponentially growing trajectories.

Errors

To examine the response of the system to a wide range of frequencies, many different single-frequency pure sinusoids are tested. All test frequencies are treated independently.

The driven error in this model is applied as a transverse-angle kick to the beam around the location of the last Final Focuss Quadrupole (FFQ) at one IP. To model this error, particles are drifted back from the IP by a distance L^* , then the dipolar kick is applied, then particles are drifted L^* back to the IP. This approach does not factor in the thickness of the FFQ.

In this model, it is arbitrarily chosen that the error applies to the electron beam-line just prior to IP I. For a bunch arriving with delay Δt (in turns) relative to the reference bunch the kick is

$$\Delta a'(T; \Delta t) = k_{\text{err}} \sin(2\pi f_{\text{err}}(T + \Delta t) / f_{\text{rev}}), \quad a \in \{x, y\}, \quad (3.3)$$

where k_{err} is the kick amplitude, f_{err} the error frequency, T the turn number, and f_{rev} the revolution frequency. For two evenly-spaced bunches, required for simulation of 4 IPs, $\Delta t = 0.5$. The kick amplitude k_{err} is calculated from the integrated strength of the FFQ and the offset as $k_{\text{err}} = k_1 L \Delta y$.

As a linear model, no amplitude-dependent effects are considered here. Further discussion of the error modelling is presented in Section 3.2.

Beam–Beam

In this model the beam–beam force is modelled by a linearised centroid kick $\Delta a' = -1/2k_{\text{bb}}a$, derived in Section 2.3.3. Using the linearised beam–beam force preserves a linear time-invariant system description which is essential for the stability and frequency-domain analyses. The effective beam–beam strength is parameterised by the beam–beam tune shift,

$$k_{\text{bb},a} = \frac{4\pi \xi_a}{\beta_a^*}, \quad (3.4)$$

allowing the tune shifts $\xi_{x,y}$ to be used as inputs to the model.

In stable feedback operation the residual orbit offsets at the IP are expected to be much less than the beam size. In this regime, the beam–beam force is linear to leading order and this model captures the primary contribution of the beam–beam interaction to the centroid motion. Non-linear effects associated with the full Bassetti–Erskine force, including amplitude-dependent tune shifts, tail dynamics, and beam-size evolution, are outside the scope of this analytic treatment. These effects are instead addressed using dedicated tracking simulations.

IP Feedback

In this model, a simplified feedback implementation is used; the positions of both beams are known exactly at the IP. In a real accelerator, offsets are detected with measurement devices such as BPMs, luminosity monitors or beamstrahlung light.

The structures of the Closed Orbit (CO) correction bumps are also neglected. A real accelerator uses a series of corrector magnets with pre determined tuning knobs to specify the CO bump. In this implementation, instantaneous additions or subtractions to the IP position coordinates are made. No correction is made to the angular coordinates at the IP.

Feedback calculations are performed once per turn. When performing predictive feedforward, least squares computations are carried out using normal equation form. A regularisation epsilon of 1×10^{-12} is included, which was found to yield numerical stability across the range of frequencies tested. Further discussion on the use of normal equations and regularisation is presented in Section 2.7.6.

Feedback is performed on a turn-by-turn basis unless otherwise specified. For the revolution frequency of FCC-ee (~ 3.3 kHz) this is deemed reasonable. The details of the different tested feedback approaches are specified in Section 3.4.

3.1.3 Approximations

Transverse Decoupling

As mentioned in Section 3.1.2, the transfer matrices used are without betatron coupling. In a real collider, lattice errors and experimental solenoids at the IPs induce coupling. With good optics corrections, this coupling can be well corrected at the IP. In this model, such correction is assumed and the coupling at each IP is taken to be approximately 0.

Vertical Plane Only

Current and future high-luminosity lepton colliders adopt the nano-beam scheme (Section 2.2.2), in which the vertical beam size is orders of magnitude smaller than the horizontal beam size. As a result, vertical position and momentum offsets constitute the dominant sources of luminosity degradation. Since the model is already formulated as decoupled four-dimensional tracking, a further simplification is made by restricting the dynamics to the vertical plane only. This reduction improves computational efficiency and enhances the clarity of the subsequent analysis.

Centroid Modelling

In the linear regime considered here, the centroid motion of a bunch can be modelled by a single particle. This equivalence is validated by benchmarks presented in Section 3.3. Consequently, only one particle per bunch is tracked, significantly reducing memory usage and enabling dense frequency scans for the computation of transfer functions.

Radiation Exclusion

Synchrotron radiation effects are neglected in this model. While the inclusion of radiation damping and excitation in the presence of orbit offsets is non-trivial, radiation damping times for all FCC-ee working points considered are long compared to the turn-by-turn feedback timescales. Radiation effects are instead studied using dedicated tracking simulations.

3.1.4 Performance Evaluation

A complete feedback loop comprises multiple elements, as discussed in Section 2.6.1. In this analysis, the focus is restricted to the controller dynamics. Effects associated with

magnet ramping, offset detection, and other hardware-related processes are therefore excluded. Given that the FCC-ee hardware design is not yet finalised, any such modelling would rely on poorly constrained assumptions.

Performance comparisons are based on closed-loop transfer functions and on the residual RMS IP offset obtained by propagating a realistic disturbance PSD through the system, as described in Section 2.6.6. The RMS offset provides a physically meaningful performance metric, as feedback effectiveness must be assessed relative to the disturbances it is intended to suppress. Since beam-centroid oscillations span several orders of magnitude, results are presented on a decibel scale to enable compact and transparent sensitivity comparisons.

3.1.5 Adopted Parameters

Parameters used in this model are presented in Table 3.1. Accelerator parameters correspond to the $t\bar{t}$ working point from the FCC-ee Mid-Term Review (MTR) [92]. Feedback parameters are taken from the SuperKEKB iBump feedback system.

Table 3.1: Parameters adopted for the analytic IP feedback modelling simulations discussed in this chapter.

Parameter	Unit	Value	Description
N_{turns}	1	10,000	Number of turns tracked
N_{BI}	1	500	Number of turns excluded (burn in)
β_y	mm	1.6	IP Vertical Beta Function
Q_y	1	398.182	Vertical Tune
f_{rev}	Hz	3306.85	Revolution Frequency
ξ_y	1	0.134	Vertical beam-beam tune shift
A_y	m	1.0	Vertical Aperture
Δ_{driven}	μm	10	Driven Error Magnet Offset
$K1L_{\text{FFQ}}$	m^{-1}	0.5	FFQ Integrated Quadrupole Strength

To remove the impact of transients, the first 500 turns are excluded. This greatly exceeds the memory length of all tested feedback algorithms (up to ~ 50 turns) ensuring that transient behaviour has decayed.

3.1.6 Parameter Issues

FCC-ee parameters remain under active development, with refinements driven by inputs from multiple working groups. For the parameters considered here, a residual issue arises in that the π -mode tune for the $t\bar{t}$ working point lies very close to the half-integer. This proximity poses a significant challenge for feedback schemes based on π -mode oscillations, as discussed in Section 2.4.4. While similar considerations apply to other working points, their lower beam-beam tune shifts provide additional operational

margin. This issue has been noted by the optics team, though revised parameters are yet to be fully developed. The issue is demonstrated in Fig. 3.1. In the 4 IP case the π -mode tune in the vertical plane is very close to half-integer at the nominal working point. Calculations are performed analytically for the π -mode tune of an N_{IP} equally spaced IP system with the same beam-beam tune shift at each IP using the equations presented in Section 2.4.3.




Figure 3.1: π -mode tune variation around the FCC-ee $t\bar{t}$ working point. Real (blue) and imaginary (red) parts of the π -mode tune are shown as functions of the bare vertical tune Q_y (left) and beam-beam parameter ξ_y (right) for different numbers of IPs. Vertical dashed lines indicate nominal parameter values.

To model any stable feedback with four IPs the tune or beam-beam tune shift must be adjusted. From Fig. 3.1, the dependence of the π -mode tune on the vertical tune is much stronger than on the beam-beam tune shift. For this analysis, it is thus chosen to reduce the vertical tune. The approach is discussed further in Section 3.7.1.

3.2 Error Modelling

To maintain luminosity the offsets between colliding bunches must be minimised. Such offsets are orbit errors that result from dipolar errors throughout the lattice. These dipolar errors will be the focus of this analysis. The implementation of errors within the tracking is discussed in Section 3.1.2.

3.2.1 Origins of Dipolar Errors

FCC-ee adopts the nano-beam scheme, described in Section 2.2.2. As such, the strongest magnets in the lattice (excluding experimental solenoids) are the FFQs. As discussed in Section 2.5.1, offsets to quadrupoles result in dipolar errors. Offsets of the FFQs are therefore expected to result in strong dipolar errors. The FFQs are foreseen to be superconducting magnets, and are therefore subject to not only ground motion, but also vibrational motion due to cryocooling. Previous studies have shown that the offsets of these FFQs will be the dominant source of errors [93].

For this study, the dipolar driven error will be applied as a thin dipolar kick located at L^* , the distance from the IP to the face of the last FFQ, with a strength derived from the dipolar kick of an offset thin quadrupole.

3.2.2 Generating Vibrational Errors

For transfer function analysis, white noise is used where all frequencies are driven at the same amplitude. To compute the RMS offsets, a realistic Power Spectral Density (PSD) is approximated from measurements of the PSDs of other collider tunnels and FFQs.

To produce the approximated PSDs the following recipe has been adopted:

1. Obtain reference PSD data (from tunnel/FFQ to be approximated).
2. Generate a logarithmically spaced array of test frequencies up to revolution frequency (N.B. this may go beyond the frequency range of the data).
3. Identify key aspects of the profile from the data: PSD values at boundaries between regions of different gradient, locations and sizes of resonant peaks.
4. Logarithmically fit the spectrum between the different gradient frequency bins.
5. Add resonances as Lorentzian peaks (see Section 2.6.7).
6. Apply decibel banded noise to mimic the high frequency roll-off.

The frequency range must be extended to the revolution frequency for the convergence of the RMS by integration. Frequencies between the upper bound of the reference data and the revolution frequency are modelled as a noise free exponential decay to a noise floor.

Decibel Banded Noise

When approximating PSDs, decibel-banded noise is used to reproduce the frequency-dependent, banded structure observed in measured vibration spectra. As presented in Fig. 3.2, Fig. 3.3, and Fig. 3.5, real-world PSDs are typically not smooth, but

instead exhibit a combination of narrow- and broad-band features, resonances, and jagged variability across logarithmic frequency. Introducing banded noise in the logarithmic domain preserves these characteristic spectral features while allowing controlled perturbations of the underlying spectrum.

For an initial noise-free PSD $S(f)$, the application of decibel-banded noise $N_{\text{dB}}(f)$ results in a perturbed spectrum

$$10 \log_{10} S'(f) = 10 \log_{10} S(f) + N_{\text{dB}}(f). \quad (3.5)$$

Working in decibels ensures that the perturbation scales multiplicatively in power, reflecting how instrumental noise and measurement uncertainty vary across decades of frequency.

3.2.3 Calculation of RMS Motion at the IP

A PSD describes how the variance of the beam motion is distributed across frequency. The total RMS displacement is obtained by integrating the PSD over frequency and taking the square root of the resulting power. For a one-sided PSD $S(f)$, the RMS amplitude is therefore given by

$$A = \sqrt{\int_0^\infty S(f) \, df}. \quad (3.6)$$

To ensure convergence of the RMS, integration is performed up to the revolution frequency.

3.2.4 Approximating SuperKEKB Vibrational Motion

Measurements of vibrational motion in the SuperKEKB tunnel and FFQs are reported in [94]. Based on the observed differences in spectral content, approximate PSD models are constructed for use in the feedback studies presented here.

Tunnel Motion

The SuperKEKB tunnel lies ~ 10 m underground [95], and its vibrational motion is largely governed by the local geology at KEK. Low-frequency horizontal oscillations in the range 0.2 to 0.3 Hz, driven by wind and ocean waves, are observed; their amplitude is weather dependent but shows no significant day–night variation [95]. In contrast, vertical oscillations around ~ 3 Hz, induced by human activity such as traffic, exhibit a clear day–night dependence [95]. This frequency also corresponds to a characteristic resonance of the ‘Kanto loam’ soil in the region [96]. A broader survey of accelerator sites shows that KEK experiences comparatively high levels of ground motion [97].

A comparison between an approximated vibrational PSD and measurement data from [94] for the SuperKEKB tunnel is presented in Fig. 3.2. Only vertical-plane motion is modelled. Small, high-frequency spectral peaks are omitted for simplicity, as they contribute negligibly to the total power. The approximation reproduces the measured spectrum well across the full frequency range of interest.




Figure 3.2: Vertical ground-motion power spectral density in the SuperKEKB tunnel. The published figure from [94] is presented, with an overlaid model estimate of the vertical motion (solid red line).

FFQ Motion

The SuperKEKB FFQs are cryogenically cooled superconducting magnets located within the SuperKEKB tunnel. A comparison between an approximated vibrational PSD and measurement data from [94] for the FFQ QC1RE is presented in Fig. 3.3. QC1RE is the first FFQ on the right-hand side of the IP in the HER (electron) ring. Only vertical-plane motion is modelled. Small, high-frequency peaks are omitted for simplicity, as they contribute negligibly to the total power. The approximation reproduces the measured spectrum well across the frequency range of interest.

A comparison of Fig. 3.3 and Fig. 3.2 shows that the primary differences arise from resonant features. A strong peak at 25 Hz corresponds to a resonant mode of the FFQ, with a weaker additional peak observed around 12 Hz.

Integrated Amplitudes

Although the PSD of QC1RE differs visibly from that of the surrounding tunnel due to the presence of resonance peaks, this difference has only a minor impact on the RMS motion predicted at the IP. While these peaks rise several orders of magnitude above the local spectral background, their contribution is dwarfed by that of the low-frequency motion. A comparison of the integrated amplitudes for the PSDs presented




Figure 3.3: Vertical vibrational power spectral density for SuperKEKB’s QC1RE. The published figure from [94] is presented, with an overlaid model estimate of the vertical motion (solid red line).

in Fig. 3.3 and Fig. 3.2 is presented in Fig. 3.4. Despite the differing spectral structures, the total integrated amplitudes are comparable, with the RMS motion dominated by low-frequency noise rather than by the vibrational peak around 25 Hz.




Figure 3.4: Comparison of integrated vibrational amplitudes for the SuperKEKB tunnel floor and final-focus quadrupole QC1RE. Integrated amplitudes derived from approximate ground-motion spectra are shown as a function of frequency.

3.2.5 Approximating FCC-ee Vibrational Motion

As the FCC tunnel has not yet been constructed, no direct vibration measurements exist. The tunnel motion is therefore approximated using measurements from the LHC tunnel. The LHC tunnel lies at an average depth of 100 m [98], while the FCC tunnel is planned to be deeper, at approximately 200 m [37], and is expected to be more geologically stable. For this analysis, a conservative assumption is adopted: despite the increased depth, the FCC tunnel is assumed to lie within a similar geological environment, and

its motion is therefore approximated using LHC tunnel measurements. A survey of vibrational motion across multiple accelerator sites has shown the LHC tunnel to be particularly stable, exhibiting very low levels of ground motion [97].

Tunnel Motion

A comparison between an approximated vibrational PSD and measurement data from [99] for the LHC tunnel without active vibration control is presented in Fig. 3.5. The approximation reproduces the measured spectrum well across the frequency range of interest.




Figure 3.5: Vertical ground-motion power spectral density in the LHC tunnel. The published figure from [99] is presented, with an overlaid model estimate assuming no active vibration control (solid red line).

FFQ Motion

Based on observations at SuperKEKB, the FCC-ee FFQ vibrational spectrum is modelled using the LHC tunnel ground-motion PSD with additional resonance peaks corresponding to magnet vibrational modes. The reference FCC-ee FFQ PSD adopted for the following analysis is presented in Fig. 3.6. As the FCC-ee FFQs are still under design, no measured vibrational mode data are available; a representative resonance at 20 Hz is therefore included to illustrate the impact of an FFQ vibrational mode.




Figure 3.6: Approximate ground-motion power spectral density adopted for an FCC-ee final-focus quadrupole. The spectrum includes an illustrative resonance representative of final-focus quadrupole motion and is used as the reference PSD for subsequent analysis.

3.3 Benchmarking

Before studying feedback implementation, the tracking model was benchmarked. All benchmarks model collisions at a single IP, using the parameters presented in Table 3.2.

Table 3.2: Parameters adopted for benchmarking the analytic IP feedback model.

Parameter	Unit	Value	Description
$N_{\text{turns}}^{\text{plot}}$	1	200	Number of turns tracked
N_{turns}	1	10,000	Number of turns tracked
N_{p}	1	10,000	Number of particles
ε_y	pm	1.6	Vertical Geometric Emittance
Δ_{static}	nm	100	Static Error Beam Offset
f_{driven}	Hz	50	Driven Error Frequency

3.3.1 Response to an Initial Static Error

Without Beam–Beam

The beam–beam interaction was disabled by setting the beam–beam tune shift to zero. The resulting dynamics are presented in Fig. 3.7. The system energy is conserved. Only beam 1 exhibits oscillatory motion, corresponding to phase-space rotation between y and y' at the IP with the betatron tune.

With Beam–Beam

The beam–beam tune shift was set to its nominal value, and the resulting dynamics are presented in Fig. 3.8. Inclusion of the beam–beam force couples the motion of the




Figure 3.7: Orbit response of the analytic feedback test system to an initial static offset with beam–beam interaction disabled. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 (top), together with the corresponding IP offset (bottom).

two beams, producing a more complex offset pattern. The total system energy remains conserved, and the motion is stable despite the large RMS offset.




Figure 3.8: Orbit response of the analytic feedback test system to an initial static offset with beam–beam interaction enabled. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 (top), together with the corresponding IP offset (bottom).

3.3.2 Response to Driven Error

Without Beam–Beam

The beam–beam interaction was disabled by setting the beam–beam tune shift to zero. The resulting response to a driven error is presented in Fig. 3.9. Only beam 1 exhibits oscillatory motion, containing the frequency of the driven excitation with an additional higher–frequency component at the betatron tune.

To characterise the response over a broad frequency range, transfer functions were constructed by exciting the system with white noise. A comparison between the tracked




Figure 3.9: Orbit response of the analytic feedback test system to a driven error with beam–beam interaction disabled. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 (top), together with the corresponding IP offset (bottom).

transfer function and the analytic expression derived in Section 2.6.8 is presented in Fig. 3.10, with both responses normalised to unity at $f = 1$ Hz. Small discrepancies arise from normalisation effects, as tracking introduces noise near minima, while resonant peaks remain finite due to the discrete set of frequencies evaluated.




Figure 3.10: Transfer function of the analytic feedback test system for a driven error with beam–beam interaction disabled. The tracked response (solid) is compared to the analytic discrete–time second–order resonator model (dashed). Vertical lines indicate the revolution–frequency harmonic nf_{rev} (green) and the upper and lower sidebands $(n \pm q)f_{rev}$ (blue/red).

With Beam–Beam

The beam–beam tune shift was set to its nominal value. The resulting response to a driven error is presented in Fig. 3.11. Inclusion of the beam–beam force couples the motion of the two beams, causing both to oscillate at the drive frequency with an additional higher–frequency component at the betatron tune. A difference in oscillation amplitude leads to a residual offset between the beams.




Figure 3.11: Orbit response of the analytic feedback test system to a driven error with beam–beam interaction enabled. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 (top), together with the corresponding IP offset (bottom).

A comparison between the tracked transfer function and the analytic form derived in Section 2.6.8 is presented in Fig. 3.12, with the response normalised to unity at $f = 1$ Hz. Small discrepancies arise from normalisation effects, as tracking introduces noise near minima, while resonant peaks remain finite due to the discrete frequency sampling. In contrast to the case without beam–beam, the resonant peaks are shifted away from exact revolution–frequency harmonics to the upper and lower betatron sidebands.




Figure 3.12: Transfer function of the analytic feedback test system for a driven error with beam–beam interaction enabled. The tracked response (solid) is compared to the analytic discrete–time second–order resonator model (dashed). Vertical lines indicate the revolution–frequency harmonic nf_{rev} (green) and the upper and lower sidebands $(n \pm q_n)f_{rev}$ (blue/red).

3.3.3 Single vs Multi-particle Modelling

As described in Section 3.1.3, bunch motion is approximated using a single particle representing the centroid. This approximation was benchmarked by comparison with

multi-particle tracking. Owing to the chaotic nature of the system and its sensitivity to initial conditions, the initial centroid position of the 10,000-particle bunch was enforced to match the single-particle position at generation. In all cases, the two approaches show agreement at the level of numerical precision.

An example comparison for a driven error with beam–beam interaction enabled is presented in Fig. 3.13. The residual differences are negligible compared to the offsets of interest in this analysis and lie well below the regularisation scale.




Figure 3.13: Comparison of the orbit response of the analytic feedback test system with single-particle and multi-particle simulation. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 (top) for the multi-particle mean (blue, solid) and single particle (orange, dashed). The corresponding difference is also presented (bottom).

3.4 Feedback Modelling

The objective of this model is to optimise the IP feedback control configuration for FCC-ee. Several control strategies are evaluated, including PID control implemented as pure feedback, predictive feedforward, and hybrid feedback–feedforward schemes.

3.4.1 Stable and Unstable Feedback

If the controller timing or phase is incorrect, the applied correction can instead drive oscillations at that frequency. This manifests as growing beam–centroid motion, leading to rapid luminosity degradation, increased hardware stress, and potentially beam aborts. Further details are given in Section 2.6.4. Avoiding instability is therefore the highest priority in controller design.

In this analysis, stability is enforced through conservative gain limits, the use of phase–lead and predictive feedback approaches, and optimisation of the PID parameters. An example of the centroid motion from cases with no feedback, a weak stable feedback and unstable feedback is presented in Fig. 3.14.




Figure 3.14: Examples of beam centroid response for different feedback configurations. Centroid position at the IP is shown as a function of turn for beam 1 and beam 2 with feedback disabled (top), stable feedback (middle), and unstable feedback (bottom).

3.4.2 Averaged Error Feedback

The simplest feedback implementation considered is based on averaging previous measurements, as described in Section 2.6.2. Transfer functions for different PID control configurations using a five turn averaged feedback scheme are presented in Fig. 3.15.

With purely proportional control, the averaged feedback is stable over a wide range of gains and provides attenuation across the full frequency range. High-frequency suppression is strong, while low-frequency attenuation is modest. Purely integral or derivative control perform poorly. Pure integral control is unstable for all tested gains, leading to particle loss, reflecting its inability to compensate for phase delay. Pure derivative control remains stable but exhibits amplification at some frequencies, with the relative transfer function exceeding unity, indicating that disturbances are driven rather than suppressed. When full PID control is applied with equal gains ($K_P = K_I = K_D$), stable operation with attenuation is achieved for a subset of gains, although performance degrades rapidly at lower gains. This sensitivity is concerning, particularly as it is expected to worsen in multi-IP configurations.

The dependence on the number of turns averaged was also investigated by scanning the averaging window between one and ten turns, using full PID control with fixed gains $K_P = K_I = K_D = 0.01$. This low gain was required to maintain stability for certain values of N_{av} . The results are presented in Fig. 3.16.

For some averaging lengths, the relative transfer function exceeds unity by a large margin, indicating strong amplification associated with the π -mode frequency. The averaging filter introduces frequency-dependent phase lag, and modest shifts in the π -mode tune can place the mode near a loop phase of -180° while the gain remains high, triggering instability. As a result, averaged feedback is operationally fragile: small lattice or beam-beam tune shifts can destabilise the loop, making this approach difficult




Figure 3.15: Dependence of the relative transfer function on feedback gain for averaged feedback. The proportional (P), integral (I), derivative (D), and combined PID contributions are shown as a function of frequency for different feedback gains. Curves corresponding to unstable cases with particle loss are not shown.

to deploy robustly in practice.

3.4.3 Linear Prediction Feedback

The simplest feedforward implementation considered is based on linear prediction using previous measurements, as described in Section 2.6.2. By predicting the future beam position, this approach provides effective phase lead, although it is susceptible to error propagation within the prediction algorithm.

While NumPy [100] provides a native least-squares solver, its interface is not well suited to solving large ensembles of independent systems in a vectorised manner and looping over frequencies was computationally prohibitive. Instead, the implementation uses the normal-equations formulation described in Section 3.1.2. The system is solved using NumPy’s broadcastable linear solver, providing an efficient implementation.

Transfer functions for different PID control configurations using this linear prediction feedback, with a prediction window of shape (48, 6), are presented in Fig. 3.17.




Figure 3.16: Dependence of the relative transfer function on the number of turns averaged for averaged feedback. The relative transfer function is shown as a function of frequency for different values of N_{av} .

With purely proportional control, linear prediction provides stable attenuation across most gains, with strong high-frequency suppression and modest low-frequency performance. In contrast to averaged feedback, the phase lead introduced by linear prediction significantly improves integral control. Pure integral feedback is highly effective at low frequencies and near the betatron tune, although performance degrades at intermediate frequencies. This behaviour is favourable, as the dominant power in the disturbance PSDs lies at low frequencies, as presented in Section 3.2. Pure derivative control remains stable for all tested gains, but exhibits amplification at some frequencies, with the relative transfer function exceeding unity. When full PID control is applied with equal gains ($K_P = K_I = K_D$), stable attenuation is achieved for a subset of gains. Owing to the predictive nature of this approach, sensitivity to modest tune variations is reduced compared to averaged feedback.

3.4.4 Hybrid Linear Prediction Feedback

A hybrid feedback approach is considered in which linear prediction is used for the predictive proportional (P) component, while averaged feedback is applied to the integral and derivative (ID) components.

Transfer functions for different PID control configurations using this hybrid scheme, with a prediction window of shape (48, 6) and averaging length of $N = 6$, are presented in Fig. 3.18.

Although $N = 6$ was found to be a suboptimal choice for averaged feedback alone (see Fig. 3.16), following a scan $N = 6$ was found to be optimal for this hybrid configuration. The hybrid scheme does not provide a clear improvement over the individual approaches. The phase behaviour introduced by averaging degrades the performance of the integral component, while the predictive element improves the behaviour of the derivative




Figure 3.17: Dependence of the relative transfer function on feedback gain for linear prediction feedback. The proportional (P), integral (I), derivative (D), and combined PID contributions are shown as a function of frequency for different feedback gains. Curves corresponding to unstable cases with particle loss are not shown.

term. When full PID control is applied with equal gains ($K_P = K_I = K_D$), stable attenuation is achieved for a subset of gains. However, the increased complexity and limited performance gains make this approach less attractive than pure linear prediction feedback.

3.4.5 SuperKEKB Feedback Implementation

SuperKEKB operates with a vertical deflection system known as the iBump feedback system [79]. The correction applied at each feedback cycle is based on a linear prediction of the beam-centroid motion, such that the orbit bump for the next cycle is given by

$$y_{n+1}^{\text{CO}} = -G \Delta y_{n+1}^{\text{pred}} + y_n^{\text{CO}}, \quad (3.7)$$

where a single gain parameter G is used [101]. Under the convention adopted in Section 2.6.3, this corresponds to a purely integral controller (often described as PI control in the SuperKEKB literature, with the predictive feedforward term interpreted




Figure 3.18: Dependence of the relative transfer function on feedback gain for hybrid linear prediction feedback. The proportional (P), integral (I), derivative (D), and combined PID contributions are shown as a function of frequency for different feedback gains. Curves corresponding to unstable cases with particle loss are not shown.

as the proportional component). Operationally, SuperKEKB uses a prediction window of $M = 48$ samples and an averaging window of $N = 6$ turns.

A key distinction between SuperKEKB’s IP feedback and the schemes considered in this chapter is the feedback update rate. At SuperKEKB, the feedback cycle spans approximately three turns, rather than operating on a turn-by-turn basis. With a revolution frequency of ~ 99.4 kHz [102], around thirty times higher than that of FCC-ee, turn-by-turn feedback is not feasible. The feedback correction is derived from an averaged IP offset, calculated from BPM signals averaged over all bunches over three turns. It in general may not be equivalent to compute offsets on a bunch-by-bunch basis and average them afterwards,

$$f(\bar{y}_A, \bar{y}_B, \bar{y}_C, \bar{y}_D) \neq \frac{1}{N} \sum f(y_A, y_B, y_C, y_D). \quad (3.8)$$

An approximate implementation was simulated by applying a purely integrating linear-prediction feedback with updates every three turns. The results are presented in

Fig. 3.19, using the parameters listed in Table 3.3.




Figure 3.19: Dependence of the transfer function on feedback gain for the SuperKEKB iBump feedback system. The transfer function is shown as a function of frequency for different feedback gains.

Table 3.3: Parameters adopted for the analytic IP feedback modelling simulations of SuperKEKB.

Parameter	Unit	Value	Description
β_y	mm	1.0	IP Vertical Beta Function
Q_y	1	46.580	Vertical Tune
f_{rev}	Hz	99400.0	Revolution Frequency
ξ_y	1	0.035	Vertical beam-beam tune shift
Δ_{driven}	μm	10	Driven Error Magnet Offset
$K1L_{FFQ}$	m^{-1}	0.5	FFQ Integrated Quadrupole Strength

No further analysis is pursued, as the fundamental differences in signal processing prevent a faithful representation of the SuperKEKB feedback within this model.

3.4.6 Selected Approach to Feedback

Based on the preceding analysis, linear predictive feedback is selected for further study. The hybrid scheme provides no performance advantage and is therefore discarded. Although averaged feedback with PID control can achieve strong suppression in isolated cases, its sensitivity to the averaging window and associated phase instabilities makes it operationally fragile. In contrast, linear predictive feedback demonstrates robust performance and benefits from extensive validation at SuperKEKB.

Optimisation

A grid-search optimisation was performed over PID gain combinations for the linear predictive feedback. Performance was evaluated using the residual RMS beam offset

at the IP, computed by propagating the reference FFQ vibration PSD described in Section 3.2.5 through the closed-loop system. The results are presented in Fig. 3.20. For clarity, purely proportional, integral, and derivative configurations are shown separately.




Figure 3.20: Dependence of the cost function on feedback gain for linear predictive feedback. The cost function (RMS beam offset at the IP) is shown as a function of the proportional K_P (left), integral K_I (centre), and derivative K_D (right) gains. Blue points correspond to stable seeds, while red points indicate unstable seeds for which the feedback fails.

Consistent with the preceding studies, aggressive optimisation rapidly leads to instability. Integral control provides the best overall performance, driven by its strong suppression of low-frequency noise. Derivative control remains stable but offers minimal rejection, while proportional control provides limited suppression before becoming unstable. Combined PID control yields only marginal improvements over pure integral control, which remain within the numerical uncertainty of the simulation and do not constitute a meaningful gain in performance.

It should be noted that this optimisation is based on an approximated reference PSD. Given the uncertainty in the underlying noise spectrum, the modest differences observed across the parameter scan are considered insignificant. The purpose of this analytic model is therefore to identify the dominant control mechanisms rather than exact optimal gains. Once more refined inputs become available, such as measured ground and quadrupole vibration spectra or the influence of a global orbit feedback system, it would be appropriate to revisit this optimisation to determine precise controller settings.

Reference Choice

Although increasing gain improves suppression, residual RMS offset was observed to saturate before the instability threshold. Unless stated otherwise, the following configuration is adopted as the reference:

- Feedback Type: Linear Predictive.

- Proportional Gain: 0.00.
- Integral Gain: 0.50.
- Derivative Gain: 0.00.

In addition, the following parameters are fixed as reference values, informed by SuperKEKB operational experience and FCC-ee hardware considerations:

- Feedback memory length $M = 48$, corresponding to the number of previous beam position measurements used to estimate the predictor coefficients.
- Prediction horizon $N = 6$, corresponding to the number of previous measurements used to compute the predicted beam offset.
- Feedback scheme: correction applied to the positron beam, with the electron beam taken as the reference.
- Feedback implementation: turn-by-turn correction, with one feedback update applied per machine turn.

The choice of correcting the positron beam is a modelling convention and does not affect the qualitative conclusions of this study.

3.5 One IP Feedback Model

3.5.1 Model Description

Modelling single-IP feedback was performed using two pairs of bunches. In this configuration, the two bunches do not interact with each other directly; instead, they can be treated as two independent systems separated by half a machine turn in time. The reference vibration spectrum is taken from Section 3.2.5, and the selected feedback scheme is a linear predictive, purely integral controller, as discussed in Section 3.4.6.

A schematic of the simulation setup is presented in Fig. 3.21. The bunches collide at a single IP, denoted IP I, and a driven error is applied to the electron beam at the FFQ upstream of IP I.

Feedback performance is quantified using a relative transfer function, defined as the ratio of the single-IP transfer function with feedback enabled to the corresponding single-IP reference case without feedback.




Figure 3.21: Schematic collision state for analytic feedback modelling with one IP. The positron and electron closed orbits are shown in red and blue respectively; the IP is indicated by a star; and the applied error is shown by a green oscillation.




Figure 3.22: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The lines represent different vertical tunes Q_y (right-hand scale) tested between the integer and half-integer. Test values for which particles were lost are not shown.

3.5.2 Tune Response

A tune scan was performed in steps of 0.025 between $Q_y = 398.025$ and 398.475 , excluding integer and half-integer values. The resulting relative transfer functions are presented in Fig. 3.22.

When the bare lattice tune approaches the half-integer, sharp peaks appear in the relative transfer function. In this regime, the π -mode resonance moves into the frequency range under investigation and the feedback can amplify disturbances, with relative gains exceeding unity. For tunes above $Q_y \approx 398.250$, no stable solutions exist across the full frequency range: the π -mode approaches the half-integer and becomes unstable, leading to exponential growth and particle loss, as discussed in Section 3.1.6.

For the single-IP configuration, a tune range of $398.150 \leq Q_y \leq 398.225$ provides robust feedback performance.

3.5.3 Impact of the Beam–Beam Tune Shift

A beam–beam tune–shift scan was performed in steps of 0.0125 between $\xi_y = 0.000$ and 0.250. The upper bound corresponds to approximately twice the nominal vertical beam–beam tune shift at the $t\bar{t}$ working point. The resulting relative transfer functions are shown in Fig. 3.23.




Figure 3.23: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The lines represent different beam–beam tune shifts ξ_y (right-hand scale) tested between 0 and 0.25. Test values for which particles were lost are not shown.

As the beam–beam tune shift increases, significant changes are observed at higher frequencies. The π –mode resonance moves into the frequency range of interest, leading to amplification of disturbances rather than suppression. This behaviour reflects the strong dependence of the feedback response on the π –mode tune.

Maintaining separation between the π –mode resonance and the dominant noise frequencies is therefore essential for stable feedback operation. If the beam–beam tune shift were to be reduced, raising the fractional lattice tune would help keep the π –mode resonance outside the frequency range populated by the dominant error sources.

3.5.4 Gain Scan

A gain scan was performed using 20 logarithmically spaced values of the integral gain between 1×10^{-4} and 1.00. The resulting relative transfer functions are shown in Fig. 3.24.

Only the highest gain value ($K_I = 1.00$) resulted in instability. For all other tested gains, the system remained stable across the full frequency range with the relative




Figure 3.24: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The lines represent different feedback gains (right-hand scale) tested between 1×10^{-4} and unity, logarithmically spaced. Test values for which particles were lost are not shown.

transfer function remaining below unity at all frequencies, indicating robust attenuation without excitation of resonant modes. As discussed in Section 3.4.6, the single-IP configuration strongly favours high-gain integral feedback. In particular, integral control provides effective suppression of low-frequency disturbances, where the dominant power in the vibration spectrum resides.

3.5.5 Linear Prediction Feedback Parameters Scan

Scans over the linear prediction parameters M and N were performed for the single-IP configuration. The history parameter M was varied between 6 and 48 turns, and the prediction horizon parameter N between 1 and 12 turns, both in integer steps. The resulting relative transfer functions are shown in Fig. 3.25.

Increasing either M or N improves disturbance rejection at low frequencies, where the dominant vibrational power resides. At higher frequencies, the impact of these parameters is less pronounced. For small values of M , the relative transfer function exceeds unity in parts of the spectrum, indicating that insufficient history leads to inaccurate prediction and local instability. This behaviour defines a minimum data window required for robust linear prediction. Larger values of M and N increase memory usage and computational cost due to the growth of the underlying matrix operations. While these considerations are not treated as constraints in the present analytic study, they are expected to be relevant for a real-time hardware implementation.

3.5.6 Detuning Scan

Detuning the two beams has been proposed as a means of suppressing coherent π -mode oscillations by introducing a frequency mismatch between the electron and positron




Figure 3.25: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The lines represent different values of the linear prediction parameter M (right-hand scale, top) tested between 6 and 48, and different values of the parameter N (right-hand scale, bottom) tested between 1 and 12. Test values for which particles were lost are not shown.

centroids, thereby preventing sustained coherent energy exchange through the beam–beam interaction. Such behaviour has previously been observed at PEP–II and in hadron colliders [103, 104].

In this model, detuning is introduced by shifting the tune of one beam, taken here to be the electron beam, relative to the nominal working point. A scan over the detuning parameter was performed between -1×10^{-2} and 1×10^{-2} in steps of 1×10^{-3} . The resulting relative transfer functions are presented in Fig. 3.26.

All simulated cases remained stable, with no particle loss observed across the full detuning range. However, no significant improvement in feedback performance is obtained. Only a marginal reduction at high frequencies is observed for positive detuning, well outside the frequency range containing the dominant vibrational power. Within the framework of this model, beam detuning therefore provides no practical benefit for single–IP feedback performance.




Figure 3.26: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The lines represent different electron beam detuning with respect to the nominal tune (right-hand scale) tested between -1×10^{-2} and 1×10^{-2} . Test values for which particles were lost are not shown.

3.5.7 Turn-by-Turn vs Bunch-by-Bunch Correction

Thus far, feedback has been assumed to operate on a turn-by-turn basis. At FCC-ee, the minimum proposed bunch spacing is 25 ns, comparable to that of the LHC, where bunch-by-bunch transverse feedback systems are routinely employed [105]. This motivates a comparison between turn-by-turn (TxT) and bunch-by-bunch (BxB) correction schemes. A comparison of the relative transfer functions for the two approaches is presented in Fig. 3.27.




Figure 3.27: Relative transfer function as a function of frequency for bunch 1 (left) and bunch 2 (right) with one IP. The turn-by-turn (TxT) and bunch-by-bunch (BxB) correction cases are shown for comparison.

Over the majority of the frequency range, the responses are essentially identical. Differences arise only at high frequencies, where the expected disturbance power is small. In this region, a modest improvement is observed for the BxB scheme, with a larger effect seen for bunch 2 than for bunch 1.

It should be noted that this comparison considers only two bunches separated by half a revolution period. With a larger number of bunches and dominant high-frequency disturbances, BxB correction could become advantageous; however, for practical magnet operation, limited variation between successive corrections is required. In this regime, the effective behaviour approaches the TxT assumption adopted throughout this chapter.

3.6 Two IP Feedback Model

3.6.1 Model Description

Modelling of a two-IP feedback configuration was performed following the same approach as for the single-IP case described in Section 3.5.1. A schematic of the simulation setup is presented in Fig. 3.28. The bunches collide at two IPs, denoted IP I and IP K, and a driven error is applied to the electron beam at the FFQ upstream of IP I.




Figure 3.28: Schematic collision state for analytic feedback modelling with two IPs. The positron and electron closed orbits are shown in red and blue respectively; the IPs are indicated by stars; and the applied error is shown by a green oscillation.

As in the single-IP case, feedback performance is quantified using relative transfer functions, defined as the ratio of the two-IP transfer functions with feedback enabled to the corresponding two-IP reference cases without feedback.

3.6.2 Same vs Alternate Beam Correction

With two IPs, a choice must be made as to which beam receives the fast IP feedback correction; Only one beam is corrected at each IP, while the opposing beam may be adjusted more slowly. Three schemes are considered: always correcting the positron beam, always correcting the electron beam, and alternating the corrected beam between IPs. The resulting centroid motion for these schemes is presented in Fig. 3.29.




Figure 3.29: Comparison of orbit response for different correction schemes with two IPs. Columns correspond to no correction, positron correction, electron correction, and alternating correction, from left to right. Rows show the centroid position at IP I (top), IP J (second), IP K (third) and IP K (bottom) as a function of turn. Blue and red curves denote the electron and positron beams respectively. Panels with a grey background indicate that no beam–beam interaction occurred at that IP.

Clear differences are observed in the individual beam orbits, reflecting differences in the σ -mode response. These differences arise from the modelling choice that the driven error is applied only to the electron beam: correcting the electron beam restores it towards the nominal orbit, while correcting the positron beam causes it to follow the disturbed electron beam. Despite these orbit-level differences, the IP offsets are identical in all cases. This demonstrates that the π -mode correction, which determines luminosity in the present model, is independent of which beam is corrected. The choice of correction scheme therefore does not affect the feedback performance within this analytic framework.

While this distinction is largely irrelevant in the present model, where radiation and emittance growth are neglected, it becomes important once beam quality and damping effects are included. For this reason, the issue is retained here as a reference point for later studies.

3.6.3 Tune Scan

As for the single-IP case, a tune scan was performed in steps of 0.025 between $Q_y = 398.025$ and 398.475 , with integer and half-integer values excluded. The resulting relative transfer functions for the two-IP configuration are presented in Fig. 3.30.




Figure 3.30: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The lines represent different vertical tunes Q_y (right-hand scale) tested between the integer and half-integer. Test values for which particles were lost are not shown.

As in the single-IP case, sharp peaks in the relative transfer function appear as the bare lattice tune approaches the half-integer. These peaks are noticeably reduced in magnitude compared to the one-IP configuration, reflecting the modified π -mode structure with two IPs. Nevertheless, the underlying mechanism remains the same: the π -mode resonance moves into the frequency range of interest and the feedback amplifies disturbances rather than suppressing them. For tunes above $Q_y \approx 398.250$, no stable solutions exist across the full frequency range. As presented in Fig. 3.1, although the form of the π -mode curves differ between the one- and two-IP cases, the half-integer condition is reached at approximately the same tune. Consequently, the range of tunes providing robust feedback performance remains similar, with $398.150 \leq Q_y \leq 398.225$ optimal for both the one-IP and two-IP configurations.

3.6.4 Beam-Beam Tune Shift Scan

As in the single-IP case, a beam-beam tune-shift scan was performed in steps of 0.0125 between $\xi_y = 0.000$ and 0.250 . The resulting relative transfer functions for the two-IP

configuration are presented in Fig. 3.31.




Figure 3.31: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The lines represent different beam–beam tune shifts ξ_y (right-hand scale) tested between 0 and 0.25. Test values for which particles were lost are not shown.

At IP I, the response closely resembles that of the single-IP case (Fig. 3.23), with the effect of the beam–beam tune shift becoming significant only at higher frequencies. In contrast, at IP K, differences are observed across the full frequency range. Very small tune shifts lead to instability at high frequencies, while moderate reductions relative to the nominal value provide improved attenuation at low frequencies. Within the vicinity of the nominal beam–beam tune shift, operation remains stable, and feedback performance is not expected to be sensitive to small variations in ξ_y .

3.6.5 Gain Scan

As in the single-IP case, a gain scan was performed using 20 logarithmically spaced values of the integral gain between 1×10^{-4} and 1.00. The resulting relative transfer functions for the two-IP configuration are presented in Fig. 3.32.

Only the highest gain value ($K_I = 1.00$) resulted in instability. All other tested gains were stable across the full frequency range, with the relative transfer function remaining below unity. The overall behaviour closely matches that observed in the single-IP case (Fig. 3.24), indicating that the introduction of a second IP does not significantly degrade feedback stability. Although the choice to focus on integral control was motivated by the single-IP model, these results demonstrate that high-gain integral




Figure 3.32: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The lines represent different feedback gains (right-hand scale) tested between 1×10^{-4} and unity, logarithmically spaced. Test values for which particles were lost are not shown.

feedback remains a robust and effective strategy when extended to multiple IPs, with feedback calculations performed independently at each IP.

3.6.6 Linear Prediction Feedback Parameters Scan

As in the single-IP case, scans over the linear prediction parameters M and N were performed. The long-history parameter M was varied between 6 and 48 turns, and the prediction horizon parameter N between 1 and 12 turns, both in integer steps. The resulting relative transfer functions are presented in Fig. 3.33.

The overall trends observed closely follow those of the single-IP case (Fig. 3.25), with improved low-frequency rejection obtained by increasing either M or N . However, the two-IP configuration exhibits reduced stability margins at high frequencies. For insufficient values of M , the relative transfer function approaches or exceeds unity, indicating that limited history degrades prediction accuracy. Even for larger values of M , only a small stability margin remains. The influence of the prediction horizon N is particularly pronounced at IP K, where insufficient N leads to noticeably poorer rejection. These results indicate that larger values of M and N improve both stability and performance when multiple IPs are corrected.




Figure 3.33: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The lines represent different values of the linear prediction parameter M (right-hand scale, top) tested between 6 and 48, and different values of the parameter N (right-hand scale, bottom) tested between 1 and 12. Test values for which particles were lost are not shown.

3.6.7 Detuning Scan

Following the single-IP study, detuning of the electron beam tune was also tested for the two-IP configuration. A scan over the detuning parameter was performed between -1×10^{-2} and 1×10^{-2} in steps of 1×10^{-3} . The resulting relative transfer functions

are shown in Fig. 3.34.




Figure 3.34: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The lines represent different electron beam detuning with respect to the nominal tune (right-hand scale) tested between -1×10^{-2} and 1×10^{-2} . Test values for which particles were lost are not shown.

As in the single-IP case, all simulated configurations remained stable across the full detuning range. However, no meaningful improvement in feedback performance is observed. The detuning primarily affects high frequencies, well outside the range containing the dominant vibrational power, and does not reduce the residual IP motion. Beam detuning therefore provides no practical benefit for feedback performance in the two-IP configuration and is not considered further in this study.

3.6.8 Turn-by-Turn vs Bunch-by-Bunch Correction

As in the single-IP case, both TxT and BxB correction schemes were tested for the two-IP configuration, with feedback applied independently at each IP. A comparison of the relative transfer functions for the two approaches is shown in Fig. 3.35.

Over most of the frequency range, the responses of the two schemes are essentially identical, consistent with the behaviour observed for one IP (Fig. 3.27). Differences arise only at high frequencies, as the revolution frequency is approached. In contrast to the single-IP case, no clear preference for either correction scheme emerges: the relative performance alternates between bunches and between IPs.

The largest relative transfer function is observed for the BxB correction of bunch 1 at IP I; however, the deviation from the TxT case is small. Within the accuracy of this




Figure 3.35: Relative transfer function as a function of frequency for bunch 1 (left) and bunch 2 (right) with two IPs. The turn-by-turn (TxT) and bunch-by-bunch (BxB) correction cases are shown for comparison.

model, no systematic advantage of BxB correction is therefore identified for the two-IP configuration.

3.7 Four IP Feedback Model

3.7.1 Model Description

Modelling of a four-IP feedback configuration was performed following the same approach as for the single- and two-IP case described in Section 3.6.1. In this model, two bunches represent the minimum configuration required to observe cross-talk between the four IPs. A schematic of the simulation setup is presented in Fig. 3.36. The bunches collide at four IPs, denoted IP I, IP J, IP K and IP L, and a driven error is applied to the electron beam at the FFQ upstream of IP I.

As in the single- and two-IP cases, feedback performance is quantified using a relative transfer functions, defined as the ratio of the four-IP transfer functions with feedback enabled to the corresponding four-IP reference cases without feedback.

3.7.2 Tune Issue

As discussed in Section 3.1.6, the current parameter set leads to a π -mode tune very close to the half-integer for the four-IP configuration. To illustrate the instability,




Figure 3.36: Schematic collision state for analytic feedback modelling with four IPs. The positron and electron closed orbits are shown in red and blue respectively; the IPs are indicated by stars; and the applied error is shown by a green oscillation.

tracking results with driven errors at 1 Hz and 500 Hz are presented in Fig. 3.37 for both the nominal tune and an adjusted tune. The adjusted tune is lowered by 0.1, from 398.182 to 398.082.

Figure 3.37 shows that the instability manifests only at higher frequencies, where the reference PSD contains little power. Nevertheless, the instability leads to particle loss and is therefore unacceptable. One possible mitigation is to reduce the feedback gain, weakening the resonance drive. However, low gain prevents adequate suppression of the dominant low-frequency noise and limits luminosity recovery. An alternative would be to introduce frequency filtering to exclude the unstable high-frequency region from the feedback action. Such filtering is not naturally compatible with the white-noise excitation and transfer-function framework adopted here. Stable operation in the presence of feedback therefore requires adjustments to the underlying optics to move the π -mode away from the half-integer. Such optics optimisation is beyond the scope of this thesis.




Figure 3.37: Instability of the π -mode for the four-IP configuration of FCC-ee at the $\bar{t}t$ working point. The left block shows the response to a 1 Hz dipolar error and the right block to a 500 Hz dipolar error. For each block, columns correspond to the nominal tune (left) and the adjusted tune (right). Rows show the centroid position at the four IPs and the corresponding IP offsets as functions of turn. Blue and red curves denote the electron and positron beams respectively.

3.7.3 Same vs Alternate Beam Correction

With four IPs, the question of which beam is corrected at each IP remains relevant. As in the two-IP case, discussed in Section 3.7.3, three schemes are considered: correcting the positron beam, correcting the electron beam, and alternating the corrected beam between IPs. The resulting centroid motion is presented in Fig. 3.38.

As observed previously for two IPs, all correction schemes yield identical IP offsets at all IPs, indicating equivalent correction of the π -mode oscillations that determine luminosity. Differences between the schemes appear only in the individual beam orbits, reflecting variations in the σ -mode response. Since fast IP feedback targets the π -mode, while σ -mode oscillations are addressed by slower orbit corrections or static IP bumps, the choice of beam-correction scheme does not affect the feedback performance in the present model. This confirms that the conclusion drawn for two IPs remains valid for the four-IP configuration.

3.7.4 Tune Scan

As for the one-IP and two-IP cases, a tune scan was performed in steps of 0.025 between $Q_y = 398.025$ and 398.475, with integer and half-integer values excluded. The resulting relative transfer functions for the four-IP configuration are presented in Fig. 3.39.

The tune response differs substantially from the one-IP and two-IP cases presented in Fig. 3.22 and Fig. 3.30. As before, peaks appear in the relative transfer function




Figure 3.38: Comparison of orbit response for different correction schemes with four IPs. Columns correspond to no correction, positron correction, electron correction, and alternating correction, from left to right. Rows show the centroid position at IP I (top), IP J (second), IP K (third) and IP K (bottom) as a function of turn. Blue and red curves denote the electron and positron beams respectively.

as the bare lattice tune approaches the integer; however, their form and frequency dependence are markedly different. This reflects the reduced reference tune adopted for the four-IP case (see Section 3.7.2). This shift moves the resonance to lower frequencies and modifies the shape of the response.

The most significant difference is the strong reduction in the range of viable tunes. While stable operation was possible up to $Q_y \approx 398.250$ for one and two IPs, no stable solutions exist for tunes above $Q_y \approx 398.150$ in the four-IP case. This reflects the increased difficulty of remaining simultaneously far from the integer resonance of the σ -mode and the half-integer resonance of the π -mode as the number of IPs increases. As discussed in Fig. 3.1, the modified π -mode structure causes instability to be reached at lower bare lattice tune values than in the one- and two-IP cases.




Figure 3.39: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. The lines represent different vertical tunes Q_y (right-hand scale) tested between the integer and half-integer. Test values for which particles were lost are not shown.

3.7.5 Beam–Beam Tune Shift Scan

As for the one- and two-IP cases, a beam–beam tune–shift scan was performed in steps of 0.0125 between $\xi_y = 0.000$ and 0.250. The resulting relative transfer functions for the four-IP configuration are presented in Fig. 3.40.

The response differs markedly from the one- and two-IP cases presented in Fig. 3.23 and Fig. 3.31. This difference arises from the lower reference vertical tune adopted for the four-IP configuration. At low beam–beam tune shift, the π -mode tune approaches the σ -mode tune and enters the frequency range populated by vibrational errors, leading to degraded feedback performance.

For lower lattice tunes, maintaining adequate separation between the π -mode resonance and the dominant noise frequencies requires a sufficiently large beam–beam tune shift. In this regime, higher values of ξ_y can be tolerated from a feedback perspective, provided the resonant peak remains outside the noise band.




Figure 3.40: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. The lines represent different beam-beam tune shifts ξ_y (right-hand scale) tested between 0 and 0.25. Test values for which particles were lost are not shown.

3.7.6 Gain Scan

As in the one- and two-IP cases, a gain scan was performed using 20 logarithmically spaced values of the integral gain between 1×10^{-4} and 1.00. The resulting relative transfer functions for the four-IP configuration with the adjusted vertical tune are presented in Fig. 3.41.

Unlike the one-IP and two-IP cases, the highest gain value ($K_I = 1.00$) remains stable for the four-IP configuration when the adjusted tune is used. In this configuration, integral feedback continues to provide effective attenuation across the full frequency range. However, as demonstrated in Fig. 3.37, an incorrectly chosen tune leads to strong excitation of resonant modes, causing the feedback to drive errors rather than suppress them.

An alternative mitigation strategy is to retain the nominal tune and reduce the feedback gain. To assess this case, a corresponding gain scan using the nominal tune is presented in Fig. 3.42. While stability can be recovered in this regime, the allowed gain




Figure 3.41: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. The lines represent different feedback gains (right-hand scale) tested between 1×10^{-4} and unity, logarithmically spaced. Test values for which particles were lost are not shown.

range is significantly reduced, resulting in substantially weaker noise suppression.

The impact of reduced gain is illustrated in Fig. 3.43, which shows the RMS IP offset obtained by integrating the corrected spectrum using the reference FCC-ee PSD from Section 3.2.5. As the gain is reduced, the residual RMS offset increases rapidly, reflecting the diminished suppression of low-frequency motion.

Although operation with the nominal tune and reduced gain may be technically stable, the resulting residual IP motion is unlikely to meet the stringent RMS orbit requirements imposed by the FCC-ee energy and polarisation programme [106].

3.7.7 Linear Prediction Feedback Parameters Scan

As in the one- and two-IP cases, scans over the linear prediction parameters M and N were performed for the four-IP configuration. The long-history parameter M was varied between 6 and 48 turns, and the prediction horizon parameter N between 1 and 12 turns, both in integer steps. The resulting relative transfer functions are presented




Figure 3.42: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and nominal vertical tune. The lines represent different feedback gains (right-hand scale) tested between 1×10^{-4} and unity, logarithmically spaced. Test values for which particles were lost are not shown.




Figure 3.43: Integrated RMS IP orbit amplitude (left-hand scale), integrated over frequency, as a function of feedback gain for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. Lines correspond to the individual IPs as indicated in the legend. Gains were tested between 1×10^{-4} and 1.00, logarithmically spaced. Test values for which particles were lost are not shown.

in Fig. 3.44.




Figure 3.44: Relative transfer function (left-hand scale) as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. The lines represent different values of the linear prediction parameter M (right-hand scale, left) tested between 6 and 48, and different values of the parameter N (right-hand scale, right) tested between 1 and 12. Test values for which particles were lost are not shown.

As in the lower-IP configurations, increasing either M or N improves stability and low-frequency disturbance rejection. In the four-IP case, there was a marked increase in computational cost. This reflects the serial implementation of the analytic model rather than a fundamental limitation of the feedback scheme. In practice, feedback systems would operate locally and independently at each IP, and this scaling does not constrain the physical performance of the system.

3.7.8 Turn-by-Turn vs Bunch-by-Bunch Correction

As in the single-IP and two-IP cases, TxT and BxB correction schemes were tested for the four-IP configuration. A comparison of the relative transfer functions is shown in Fig. 3.45.

As observed for two IPs, differences between the two correction schemes arise only at high frequencies, close to the revolution frequency. In this configuration, the BxB scheme performs slightly worse for bunch 1 at IP I, while the remaining bunch-IP combinations show no clear preference between the two approaches. No systematic trend with IP or bunch number is identified. In the low-frequency region, where the dominant vibrational power is concentrated, the two correction schemes produce nearly identical responses. Within the framework of this model, TxT and BxB correction therefore provide equivalent performance for the four-IP configuration.




Figure 3.45: Relative transfer function as a function of frequency for bunch 1 (left) and bunch 2 (right) with four IPs and adjusted vertical tune. The turn-by-turn (TxT) and bunch-by-bunch (BxB) correction cases are shown for comparison.

3.8 Summary

An analytic feedback model has been developed to compare interaction-point (IP) feedback strategies for FCC-ee configurations with multiple IPs. The model is formulated in the frequency domain and is designed to prioritise clarity and computational efficiency, enabling rapid exploration of wide frequency ranges over many thousands of turns. This framework allows systematic comparison of feedback architectures, control parameters, and machine configurations while retaining the essential beam-beam coupling relevant for IP stability.

Feedback performance is quantified using both relative transfer functions and the resulting RMS residual motion at the IP. To evaluate RMS residuals, representative models of an FCC-ee final-focus-quadrupole vibration spectrum were constructed and used as input excitations. These spectra were benchmarked against SuperKEKB measurements and were found to provide a reasonable representation of observed ground and component motion for the purposes of this study.

Among the feedback strategies tested, linear predictive feedback was identified as

the most robust and effective approach for IP stabilisation at FCC-ee. Performance is dominated by the integral gain K_I , which provides strong suppression of low-frequency disturbances where the majority of vibrational power resides. Pure proportional control offers limited benefit at high frequencies, while derivative control was found to amplify disturbances and is therefore disfavoured. These conclusions are consistent across all IP configurations studied.

All studies were performed using the FCC-ee Mid-Term Review optics for the $t\bar{t}$ working point. This configuration was selected as a sensitive test case due to its relatively large beam-beam tune shift and, in the four-IP configuration, the proximity of the π -mode tune to the half-integer resonance. The aim of this study is to assess feedback sensitivity to such conditions rather than to optimise the collider working point, which lies outside the scope of this thesis.

With the adjusted tune adopted for the four-IP configuration, stable feedback operation was demonstrated for the one-, two-, and four-IP cases. The principal conclusions are:

- Feedback performance is highly sensitive to the bare lattice tune. Stability requires that the π -mode tune remains sufficiently far from the half-integer resonance.
- The beam-beam tune shift has a comparatively weaker influence on stability. Moderate variations in ξ_y do not significantly degrade feedback performance provided the lattice tune remains well positioned.
- High-gain integral control remains effective across all IP configurations when the tune condition is satisfied. Instability arises primarily from resonance proximity rather than from feedback gain alone.
- Increasing the number of IPs reduces stability margins but does not qualitatively change the preferred feedback architecture.

While no difference is observed in centroid offsets between correction schemes within this model, the choice of which beam is corrected may have implications for emittance growth when radiation and non-linear effects are included. Correcting only one beam concentrates fast orbit corrections in that beam, whereas alternating corrections would distribute such effects more evenly. These effects are not captured by the present centroid-based model but become relevant when extending the analysis to include emittance evolution. The optimal hardware configuration would provide independent correction capability on both beam lines at each IP, allowing flexible or alternating operation.

The analysis in this chapter assumes a linear beam-beam force. For flat beams such as FCC-ee, geometric reduction factors (e.g. the Yokoya factor) slightly reduce the effective beam-beam tune shift. However, scans over ξ_y indicate that feedback

stability is significantly more sensitive to the lattice tune than to moderate variations in beam–beam strength. The linear approximation is therefore sufficient for the conclusions drawn here.

The present model does not capture averaging over the full distribution of bunch phases around the ring. Extending the analysis to thousands of bunches would be computationally prohibitive for the parameter scans performed here. Previous optics studies have identified additional resonance constraints that motivated the current working–point choice. The feedback–related constraints identified in this work have been communicated to the optics designers; however, identifying a globally optimal tune configuration for all FCC-ee working points lies beyond the scope of this thesis.

In summary, this analytic study demonstrates that robust IP feedback for FCC-ee is achievable across multiple IPs provided that lattice tunes are chosen to maintain adequate separation from coherent beam–beam resonances. Linear predictive feedback with integral action emerges as a stable and effective architecture, offering strong suppression of low–frequency disturbances that dominate realistic vibration spectra. These results establish a clear set of constraints for feedback design and optics choice, and provide a foundation for subsequent studies incorporating non-linear beam dynamics, radiation effects, and detailed hardware implementation.

Chapter 4

Operational Performance of IP Feedback at SuperKEKB

4.1 SuperKEKB IP Feedbacks

4.1.1 IP Feedback Requirements

SuperKEKB’s IP feedback systems were designed based on a combination of SuperKEKB simulations and operational experience from KEKB, with the design primarily driven by ground-motion-induced vibrations of the final-focus quadrupoles (FFQs).

Ground Motion

The SuperKEKB tunnel is located approximately 10 m underground [95]. The vibrational motion of the tunnel is governed by the geological and environmental characteristics of the KEK site. Low-frequency horizontal oscillations in the range 0.2 Hz to 0.3 Hz arise primarily from wind and ocean-wave effects. The amplitude of these oscillations is weather dependent but does not exhibit a significant day–night variation [95].

Higher-frequency vertical oscillations are observed at around 3 Hz. This frequency corresponds to a characteristic resonance of the local Kanto loam soil [96]. In contrast to the low-frequency horizontal motion, the magnitude of the 3 Hz vertical vibration shows a pronounced day–night dependence, attributed to variations in human activity such as regional traffic [95].

A comparative vibration study across multiple accelerator sites found that KEK exhibits relatively high levels of ground motion [97]. The results are summarised in Table 4.1, which reports the maximum peak-to-peak displacement (P_k - $P_{k_{\text{Max}}}$), full width at half maximum (FWHM), root-mean-square (RMS), and standard deviation (SD) of the measured vertical motion.

Table 4.1: Vertical ground-motion characteristics measured at selected accelerator sites, reproduced from [97]. Values are sorted according to the peak-to-peak displacement.

Site Location	Pk-Pk _{Max} [nm]	FWHM [nm]	RMS [nm]	SD [nm]	Quietest RMS [nm]	Noisiest RMS [nm]
Seismic Station Moxa	7	17	0.6	0.1	0.5	0.9
CERN Geneva	21	53	1.8	0.9	0.9	2.8
LAPP Annecy	35	59	3.3	1.6	1.9	7.0
SLAC Menlo Park	60	105	4.8	1.2	4.1	7.4
DESY HERA Hamburg	170	200	51.8	18.9	34.8	77.0
KEK Tsukuba	170	210	78.0	36.0	38.0	125.1
DESY XFEL Schenefeld	180	245	38.7	16.6	35.1	70.0

Final-Focus Quadrupole Vibration

Operational data from KEKB indicated that orbit motion at the IP driven by FFQ vibrations could reach approximately four to five times the vertical beam size [79]. For SuperKEKB, previous studies therefore performed a modal analysis of the new machine–detector interface (MDI) region using three-dimensional finite-element models implemented in ANSYS [107].

From these simulations, the vertical eigenfrequencies of the first FFQs (QC1s) on the right-hand side of the IP, QC1RP and QC1RE, were reported [79]. The reported vibrational modes, together with the resulting RMS orbit differences at the IP and the corresponding luminosity degradation, are summarised in Table 4.2. These results indicate that a feedback system operating up to 100 Hz should provide sufficient suppression of the induced orbit motion.

Table 4.2: Simulated vibrational eigenmodes of the SuperKEKB final-focus quadrupoles QC1RP and QC1RE (right side), obtained using ANSYS [79]. The resulting RMS orbit difference at the IP and the corresponding luminosity degradation are calculated assuming coherent motion of the quadrupoles.

Mode Number	Frequency [Hz]	RMS Orbit Difference [nm]	Luminosity Degradation [%]
1	25	18.6	4.6
2	38	1.7	0.2
3	69	8.3	0.3
4	100	3.1	0.3

The 25 Hz vibration mode was identified as the most critical due to its large impact on luminosity. The SuperKEKB IP feedback systems were designed to recover the associated luminosity loss. Coherent motion between QC1RP and QC1RE was assumed in these studies, as the two magnets share a common cryostat. If the vibrations are assumed to be incoherent, orbit differences of several times the beam size can occur,

consistent with KEKB experience. ANSYS studies indicate good coherence below 50 Hz [108] and the corresponding luminosity loss from the left-hand side quadrupoles was reported to be comparable.

4.1.2 The iBump Feedback System

The iBump feedback system is a beam–beam deflection–based IP feedback system [79]. The system was first implemented at KEKB, where it was used to correct both horizontal and vertical orbit offsets [109]. At SuperKEKB, iBump is capable of correcting horizontal position, vertical position, and vertical angle offsets; in routine operation, however, it is used primarily for vertical position correction. Horizontal deflection feedback is not routinely required, as the larger horizontal beam size, smaller beam–beam tune shift, and large crossing angle reduce the effectiveness of this approach. A schematic overview is presented in Fig. 4.1.




Figure 4.1: Schematic overview of the iBump Feedback System from [79]

Feedback correctors are installed only in the HER (electron ring). Closed-orbit (CO) bumps are generated on the HER beam, which is steered into collision with the LER beam; in this configuration, the HER beam effectively “chases” the LER beam to maintain collision.

The amplitude of the CO bumps is limited by the current limits of the iBump corrector magnets. When these limits are approached, static orbit bumps are applied to re-centre the correctors within their operating range, while also preventing excessive displacement of the IP from the centre of the Belle II vertex detector.

Offset Calculation

The vertical position and vertical angle offsets at the IP between colliding bunches are calculated from measurements at four IP BPMs, comprising one incoming and one outgoing transverse centroid position ($\bar{x}, \bar{y}$) measurement per bunch.

The calculation additionally requires the linear transfer matrices between the IP and the BPM locations, which are obtained from the SAD model of the SuperKEKB

interaction region, described in Section 5.1. By propagating the BPM measurements to the IP using these transfer matrices, the vertical position and angle offsets at the IP are reconstructed.

A detailed derivation of the offset calculation from BPM measurements is given in Section 2.7.6.

Feedback Method

The iBump feedback system employs a linear prediction method proposed by K. Oide. At each feedback cycle, the height of the orbit bump to be applied is proportional to the negative of the predicted vertical offset,

$$\Delta y_{n+1}^{\text{bump}} = -G \Delta y_{n+1}, \quad (4.1)$$

where G is the feedback gain and Δy_{n+1} is the predicted canonical kick at the IP for the next feedback cycle.

The predicted canonical kick is obtained using a linear least-squares prediction based on a finite history of previous measurements,

$$\Delta y_{n+1} = \sum_{k=1}^N c_k \Delta y_{n-k+1}, \quad (4.2)$$

with the coefficients c_k determined from past data. In SuperKEKB operation, the prediction parameters are set to $N = 6$ and $M = 48$. A detailed description of the linear prediction formalism and coefficient determination is given in Section 2.7.6.

The canonical kick on each cycle is defined relative to a reference offset,

$$\Delta y_n = y_n - y_{\text{target}}, \quad (4.3)$$

where y_{target} is the iBump Feedback Vertical Offset Target (iBFBVOT). This parameter defines the reference vertical collision point and plays a central role in iBump operation and optimisation.

The iBump Feedback Vertical Offset Target

The calculation of the vertical offset at the IP depends on BPM measurements, transfer matrices, and the assumed reference orbit; in practice, all of these inputs are subject to finite accuracy. The iBump Feedback Vertical Offset Target (iBFBVOT) is introduced to absorb systematic errors in the reconstructed IP offset arising from these limitations.

A dominant contribution to such errors originates from random and systematic BPM offsets due to misalignments. In operation, the iBFBVOT is optimised empirically

by scanning its value and maximising the measured luminosity. This optimisation procedure is described in Section 4.2.2.

Hardware Implementation

Beam Position Monitors The BPMs used for the iBump feedback are button BPMs located approximately 0.5 m from the IP on either side. They are mechanically coupled to the Belle II detector structure, reducing sensitivity to relative vibrations. At the design current of 3.6 A, the nominal BPM resolution is approximately $0.2\text{ }\mu\text{m}$ [79]. The BPMs are read out at 32 kHz, corresponding to one measurement approximately every three turns. Measured centroids represent averages over all bunches during this interval.

Corrector Magnets The iBump system employs twelve dedicated corrector magnets on the HER beam line: four horizontal and eight vertical, distributed symmetrically around the IP over a total length of approximately 105 m on each side. Most correctors are located within $\sim 30\text{ m}$ of the IP, with additional vertical correctors placed further into the interaction region. The layout of the corrector magnets is presented in Fig. 4.2. The correctors are stand-alone iron-yoked magnets constructed by the Budker Institute for Nuclear Physics, with a maximum kick angle of $100\text{ }\mu\text{rad}$ [79].




Figure 4.2: iBump feedback corrector locations around the IP as shown in an operational control panel from [79]. An example vertical orbit correction bump is shown in the upper panel. The displayed unit label is believed to be erroneous, although the intended unit is not confirmed. The locations of the horizontal and vertical correctors in the interaction region are indicated schematically.

Signal Processing

The original KEKB iBump system, and the initial SuperKEKB implementation (2018–2021), operated as a slow feedback system. Feedback corrections were computed on a CPU using a SAD script reading data from EPICS process variables, and applied to the correctors at a rate of approximately 0.5 Hz.

Since Autumn 2021, the iBump system has been upgraded to a fast feedback architecture operating at 32 kHz, corresponding to one correction approximately every three turns. This performance increase is enabled by a dedicated hardware implementation based on a Virtex-5 FXT FPGA, which performs the feedback calculations in real time.

A schematic overview of the signal processing chain is presented in Fig. 4.3. Signals from the IP BPMs are transmitted to a local control room, processed by the fast feedback controller, and applied to the corrector magnets via the power-supply controllers.




Figure 4.3: Schematic overview of the iBump feedback system signal processing chain from [101].

4.1.3 The SLAC Dither System

The SLAC dither feedback system is an IP feedback system based on the dithering technique [82]. A schematic overview of the hardware layout is shown in Fig. 4.4. The feedback correctors are installed only on the positron (LER) beam line.




Figure 4.4: Hardware layout of the SLAC dither feedback system from [79]. The locations of the dither correctors around the interaction region are indicated.

Although the system was successfully commissioned, it was not operational during the period covered by this study and is therefore not considered further.

4.2 SuperKEKB Machine Studies Overview

Multiple Machine Development (MD) studies were performed during secondments to KEK under the EAJADE scheme [49]. These MD periods were essential for understanding the operational behaviour of the iBump system and evaluating its performance. The studies were carried out in collaboration with R. Ueki, M. Masuzawa and Y. Funakoshi and published in [110]. This section summarises the machine setup, data acquisition, and data availability relevant to these studies.

4.2.1 SuperKEKB Stable Operation Beam Setup

To perform MDs on the iBump feedback system, the beams were first established in a stable colliding configuration. This discussion concerns the operational setup within a shift; details of the commissioning process are presented in Section 4.3.

No optics changes were applied for these MDs, and global optics corrections had already been performed. At the start of each fill, standard luminosity optimisation procedures were performed, including optimisation of the IR coupling, waist position, and IP dispersion. An initial iBump feedback target scan was then carried out to ensure optimal collision conditions before commencing the MD.

4.2.2 iBump Feedback Target Scans

The iBump Feedback Vertical Offset Target (iBFBVOT) is a free operational parameter, described in Section 4.1. Collision conditions are optimised by scanning the iBFBVOT and maximising the measured luminosity. Two independent luminosity monitors, LUMIBELLE2 and CsI, are used for this optimisation [45]. A representative optimisation curve is presented in Fig. 4.5 showing consistent optimal values of the iBFBVOT are obtained from both luminosity monitors coinciding with the minimum vertical beam size at the IP for both the LER and HER beams.

An iBFBVOT scan is performed by stepping the offset target and measuring the corresponding luminosity with a typical step size of 1 μm . At each step, the system is allowed to equilibrate and then multiple luminosity measurements are recorded, from which a mean and standard deviation are computed. The resulting data are fitted with a quadratic function to determine the optimal iBFBVOT.

When the scan is performed symmetrically about the optimum and within a sufficiently small range, a quadratic model provides a good approximation to the luminosity profile. For scans extending further from the optimum, the quadratic approximation




Figure 4.5: Example iBump feedback vertical offset target optimisation. The upper panel shows the measured vertical beam size at the IP in nano-metres as a function of the iBump vertical offset target in millimetres for the LER and HER (top left-hand scale). The lower panel shows the corresponding variation of the specific luminosity measured by two independent luminosity monitors (bottom left and right-hand scales). The optimal vertical offset target is found at approximately -0.293 mm.

may degrade. The fitted optimum is not applied automatically; the operator manually updates the iFBVOT. Consequently, occasional fit imperfections do not directly impact machine performance.

4.2.3 SuperKEKB Control and Data Storage

The SuperKEKB control system is a distributed system based on EPICS, in which machine parameters are managed and logged as EPICS Process Variables (PVs) [111]. EPICS provides a standardised framework for real-time control, monitoring, and data acquisition across large-scale scientific facilities [112]. The SuperKEKB Control system contains 200 input/output controllers (IOCs) which are categorised as: Magnet (36), RF (17), Beam Monitor (53), Beam Transport (39), Control (20), Vacuum (23) and Safety (5) [111]. The operation interface is prepared with Python, SAD scripts and Control System Studio (CSS) BOY [113]. These PVs provide the primary data source for both machine operation and offline analysis.

Operational data are recorded by the KEKBLog archiver system, originally developed for KEKB, which continuously logs 130,000 PVs for later retrieval and analysis [114]. Additional data access for short-term studies is provided via the CSS archiver, and selected time-series data can also be viewed using Grafana [115].

In addition to digitally archived PVs, shift leaders maintain operational logs containing written notes and selected control-panel screenshots, which provide contextual information for machine studies.

4.2.4 Data for MD Analysis

SuperKEKB produces a wide range of operational data, from which a subset of variables was selected for analysis in collaboration with experts at KEK. The selection was guided by the dominant physics effects relevant to offsets at the IP, with an emphasis on both input error sources and observable beam responses. The variables selected for analysis include:

- iBump Vertical Offset Target.
- Calculated iBump vertical position and angle offsets.
- Luminosity.
- Calculated canonical kick and angle at the IP.
- Beam current.
- Beam emittance.
- IR temperatures.
- Final-focus quadrupole gap sensors.

Multiple EPICS PVs were evaluated within each category.

For offline analysis, EPICS PV data were extracted from the KEKBLog archiver and resampled at 1 s intervals. As many measurements are not acquired at exactly 1 Hz, interpolation was applied. The 1 s sampling interval represents a compromise between time resolution, data volume, and data availability.

4.2.5 Correlation Testing

To quantify correlations between parameters, the Pearson product-moment correlation coefficient was used [116]. This coefficient measures the strength of linear correlation between two variables, with values in the range $[-1, 1]$; it does not capture non-linear relationships.

For paired data $\{(x_1, y_1), \dots, (x_n, y_n)\}$, the Pearson correlation coefficient is defined as

$$r_{xy} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}, \quad \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \quad (4.4)$$

where n is the sample size. An absolute value of $r_{xy} = 1$ corresponds to perfect linear correlation, $r_{xy} = 0$ to no linear correlation, with the sign indicating positive or negative correlation.

The correlation coefficients were computed using NumPy’s “np.corrcoef” function [100].

4.3 SuperKEKB Collision Commissioning

4.3.1 Overview

This section summarises the commissioning procedure for the 2025c SuperKEKB operation period, culminating in the establishment of first collisions and the activation of the IP feedback system on 18/11/2025.

4.3.2 Initial Optics

Following the summer shutdown, the LER and HER rings were recommissioned using highly detuned optics with IP beta functions of (1000, 40) mm to achieve first turns. These optics provide a substantially larger dynamic aperture than the nominal configuration, improving beam lifetime and injection efficiency during early recovery.

Once beam circulation was established, global optics corrections were performed, including orbit correction and the reduction of beta-beating, coupling, and spurious dispersion to recover the design optics.

A major focus during this period was injection optimisation. Synchrotron injection was trialled in the HER, requiring modifications to the injection-region optics, BTe transfer line magnet strengths, and LINAC configuration. To maintain sufficient dynamic aperture for off-energy injection, the sextupole configuration was re-optimised.

Vacuum conditions had degraded during the shutdown due to extensive interventions. Vacuum scrubbing was therefore performed by operating at the highest stable beam currents for extended periods, promoting outgassing and improving beam-pipe vacuum quality for subsequent operation.

4.3.3 Optics Squeezing

In preparation for collisions, the machine optics were progressively squeezed to reduce the IP beta functions and hence the beam size at collision. The squeezing proceeded through a sequence of optics with horizontal and vertical IP beta functions of (1000, 40), (200, 8), (100, 3), and (60, 1) mm (2025c nominal optics) in each ring.

At each stage, global optics corrections were applied, including orbit correction and the reduction of beta-beating, coupling, and dispersion. As the optics were squeezed, injection became increasingly challenging due to the corresponding reduction in dynamic aperture, and significant effort was therefore devoted to injection optimisation.

4.3.4 Collision Tuning Setup

In preparation for collision tuning, a bunch filling pattern of 97 bunches per ring was selected, comprising 96 colliding bunches and one pilot bunch. The HER filling pattern

used during collision tuning is presented in Fig. 4.6.




Figure 4.6: HER bunch filling pattern recorded during 2025c first collision tuning. Each panel shows the measured bunch current as a function of RF bucket number. The monitor reports 99 populated buckets, while the nominal fill pattern corresponds to 97 bunches.

The pattern exhibits clear overspill into adjacent RF buckets, arising from incomplete capture during synchrotron injection, which was still being optimised at this stage. As a result, populated buckets appear as dominant peaks accompanied by smaller satellite peaks in neighbouring buckets.

This overspill complicated the initial stage of collision tuning, which involves aligning the RF buckets of the two rings. In the presence of overspill, it was not always straightforward to distinguish collisions involving the intended HER bunch from those involving neighbouring overspill, necessitating additional checks during RF phase alignment.

Collision tuning was therefore performed twice on the same day. An initial working point was identified in the early optimisation, but due to an RF phase issue the procedure was repeated later. For clarity, only the results of the second optimisation are presented here.

4.3.5 RF Collision Tuning

With the filling pattern established, the first stage of collision tuning was RF phase optimisation. During this procedure, the LER RF phase (ROOMPHASE) was varied while the RF frequency was held fixed, as the LER and HER share a common RF timing system. Adjusting the RF phase modifies the relative arrival time of the positron

bunches at the IP. Due to the large horizontal crossing angle, timing shifts affect both the longitudinal and horizontal collision conditions.

At this stage, luminosity cannot be used as a reliable figure of merit. Given the nanometre-scale vertical beam size, any residual vertical offset prevents effective collision even when the timing is correct. Instead, the RF phase was optimised by maximising the horizontal beam–beam deflection, which is observable from BPM data as a horizontal beam–beam kick on the HER electron bunches. The RF phase was scanned over a wide range (initially -100° to 100° , and subsequently -30° to 50°). The resulting beam–beam kick measurements are presented in Fig. 4.7.




Figure 4.7: Beam–beam kick measured during collision commissioning RF phase sweeps. The horizontal (left) and vertical (right) beam–beam kicks are shown as a function of the RF room phase setting. Results are shown separately for electrons (top row) and positrons (bottom row).

The observed dependence follows the Bassetti–Erskine form described in Section 2.3.2, with a consistent zero crossing observed for both beams. A vertical beam–beam kick is also observed near the optimal RF phase, indicating correct horizontal overlap. Localised luminosity is detected at this point, as limited vertical tuning had already been performed. A fit to the horizontal beam–beam kick curve was used to determine the optimal LER RF phase. Once vertical tuning is sufficiently established, luminosity can also be used as a supplementary optimisation observable.

4.3.6 Vertical Offset Collision Tuning

Slow Orbit Bump Tuning

With RF tuning complete, the beams were aligned in timing but remained vertically offset. As collisions were being recommissioned, it was not guaranteed that the vertical

offset lay within the operating range of the iBump feedback system. A wide scan using static IP vertical orbit bumps was therefore performed to locate the collision region. The results of a representative scan are presented in Fig. 4.8.




Figure 4.8: SuperKEKB iBump commissioning slow IP vertical orbit bump scan. The upper panel shows the vertical fast bump height (left-hand scale, red) and the corresponding slow vertical bump height (right-hand scale, blue) as a function of time. The lower panel shows the measured ECL specific luminosity (left-hand scale, green) and the corresponding specific luminosity (right-hand scale, red) during the scan.

During this procedure the fast iBump feedback was disabled, and only the static orbit bumps were adjusted. Although the luminosity remained small, a preliminary vertical working point was identified.

iBump Feedback Orbit Bump Tuning

With the slow orbit bump optimised, further vertical alignment was performed using the iBump feedback corrector magnets. During this stage the feedback was not enabled; instead, static orbit bumps were generated using the iBump correctors. An extended-range iBFBVOT scan was performed to identify the approximate collision region. The results are presented in Fig. 4.9.

A clear kink is observed in the vertical beam-beam kick, coinciding with the peak in luminosity and indicating a coarse optimum iBFBVOT value. Owing to the extended scan range, the quadratic fit quality is poor, but the scan provides a robust estimate of the collision region.

Following this coarse optimisation and prior to enabling feedback, a finer iBFBVOT




Figure 4.9: Vertical beam–beam kick and luminosity measured during an iBump vertical offset scan. The left panel shows the measured vertical beam–beam kick in units of milliradians as a function of the applied iBump height. The right panel shows the corresponding luminosity in units of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ as a function of the iBump vertical offset, with the measured data overlaid by a fitted curve.

scan was performed. This finer scan range improves sensitivity near the optimum, although the fitted maximum does not exactly coincide with the luminosity peak. The results are presented in Fig. 4.10.




Figure 4.10: iBump feedback vertical offset scan prior to feedback activation. The upper panel shows the measured vertical beam size at the IP as a function of the iBump vertical offset target for the LER (red) and HER (blue) (left-hand scale). The lower panel shows the corresponding variation of the specific luminosity measured by two independent luminosity monitors (left- and right-hand scales). Measured data are shown together with fitted curves.

4.3.7 iBump Feedback Activation

During early commissioning with detuned optics, vertical alignment could be maintained using static orbit bumps. As the optics were further squeezed, active feedback became necessary to preserve luminosity. With the beams in collision, the iBump feedback was therefore enabled. The impact of feedback activation is presented in Fig. 4.11.




Figure 4.11: Activation of the iBump feedback system following collision commissioning. The upper panel shows the vertical fast bump height (left-hand scale, red) and the corresponding vertical bump height (right-hand scale, blue) as a function of time. The lower panel shows the specific luminosity measured by two independent luminosity monitors (left- and right-hand scales).

Upon activation, a transient reduction in luminosity is observed while the correction converges. The system then stabilises at a higher and more stable luminosity level than prior to feedback activation.

BPM Noise Issue

During commissioning, the iBump feedback was briefly activated before optimal collision conditions had been fully established. Although initial RF phase and vertical offset scans had been performed, the applied RF phase setting was incorrect, and the beams were not properly in collision. Under these conditions, the computed canonical kick exhibited large fluctuations, resulting in erratic fast orbit corrections, as presented in Fig. 4.12. This behaviour illustrates the necessity of establishing a well-optimised collision configuration prior to engaging the feedback system.




Figure 4.12: iBump feedback response during non-optimal collision conditions. The upper panel shows the vertical fast bump height (left-hand scale, red) and the corresponding vertical bump height (right-hand scale, blue) as a function of time. The lower panel shows the vertical canonical kick (left-hand scale, red) together with the iBFBVOT target value (blue).

4.4 Varied Current MD

4.4.1 MD Overview

The iBump Feedback Varied Current MD (VCMD) was performed on 4th June 2024, between 14:00 and 18:00 JST. The objective of this study was to investigate the dependence of the iBump Feedback Vertical Offset Target (iBFBVOT) on beam current. Optimal iBFBVOT values were determined by performing iBFBVOT scans, as described in Section 4.2.

A single beam abort occurred during the MD due to losses in the LER, which triggered a global beam abort. Losses on the CLAWS D05V1 collimator led to a large dose on the Belle II diamond abort trigger. The beams were recovered shortly thereafter, and stable operation was maintained for the remainder of the study.

4.4.2 iBFBVOT Scans

During this MD, the iBFBVOT was scanned eleven times. Ten optimised values were obtained at seven distinct beam currents. One scan was abandoned due to an incorrect scan direction, and one scan was interrupted by a beam abort. During each scan, the system was allowed to equilibrate for 5 s at each iBFBVOT step, after which 40

luminosity measurements were acquired. An overview of the MD, showing the LER and HER beam currents and the iBFBVOT as a function of time, is presented in Fig. 4.13.




Figure 4.13: Summary of the iBump feedback varied-current machine development study. The iBump vertical offset target iBFBVOT (black, left-hand scale) and the LER (red) and HER (blue) beam currents (right-hand scale) are shown as a function of time. The grey shaded region indicates a period without beam. Vertical dashed lines mark the start (green) and end (red) of each scan period.

A total of seven beam current settings were investigated. Currents were selected in the LER, with the HER current adjusted according to standard operational ratios. A minimum LER current of 400 mA was chosen to remain well above the noise floor, while the maximum current was limited to 1000 mA to avoid beam-stability issues.

The measurement times and corresponding beam currents for each scan are summarised in Table 4.3. Selected current settings were deliberately repeated before and after the beam abort to assess measurement reproducibility.

Table 4.3: Measurement times and corresponding LER and HER beam currents for the iBump feedback varied-current machine development study.

Scan Number	Start Time [JST]	LER Current [mA]	HER Current [mA]
1	14:38	400	320
2	14:56	500	400
3	15:08	500	400
4	15:22	600	480
5	15:39	700	560
6	15:56	800	640
7	16:25	400	320
8	16:43	600	480
9	16:56	800	640
10	17:11	900	720
11	17:25	1000	800

4.4.3 Erroneous Scans

Opposite Direction Scan

To ensure consistency between measurements, all iFBVOT scans were performed with the offset target increasing monotonically with time. The effect of scan direction on reproducibility was later investigated during the Constant Current MD (see Section 4.5).

One scan was inadvertently performed in the opposite direction, as presented in Fig. 4.14. This scan was abandoned prior to optimisation and subsequently repeated at the same beam current using the predefined scan direction. As no optimised value was obtained, the reverse-direction scan is excluded from further analysis.




Figure 4.14: Scan directions for a subset of scans during the iBump feedback varied-current machine development study. The iBump vertical offset target iBFBVOT (black, left-hand scale) and the LER beam current (red, right-hand scale) are shown as a function of time. Forward scans, in which the iBFBVOT is increased, are indicated by green arrows, while reverse scans are indicated by magenta arrows. Vertical dotted lines mark the start (green) and end (red) of each scan period.

Beam Abort Scan

The beam abort at 16:02 JST occurred during an iBFBVOT scan and interrupted data acquisition. Although the scan was incomplete, sufficient data had been collected to determine an optimal iBFBVOT. In the control-system logging, the recorded value corresponds to the setting at the time of the abort, as no adjustment was applied following beam loss.

To recover the optimal value, the raw luminosity data (interpolated at 1 Hz) were grouped by iBFBVOT step and used to compute per-step mean luminosity values and standard deviations. These values were then replotted as a function of iBFBVOT, and

a quadratic fit was applied to determine the optimal setting. The fitted optimum was used in subsequent analysis rather than the logged abort value.

The time evolution of the luminosity and iBFBVOT during the abort-interrupted scan is presented in Fig. 4.15.




Figure 4.15: Luminosity variation with time during an abort-interrupted iBump vertical offset target scan during the iBump feedback varied-current machine development study. The iBump vertical offset target iBFBVOT is shown in black (left-hand scale), while the luminosity measured by LUMIBELLE2 is shown in red (right-hand scale). The red markers with error bars indicate the mean luminosity and standard deviation computed within each iBFBVOT step. The green dashed horizontal line indicates the iBFBVOT value at which the beam abort occurred, and the grey shaded region marks the beam-abort period.

4.4.4 iBFBVOT Variation with Current

At the end of each scan, the applied iBFBVOT value was taken as the optimal setting for analysis. This value may differ slightly from the fitted optimum for two reasons: first, the applied value is quantised at the $1\text{ }\mu\text{m}$ level, whereas the fit reports arbitrary precision; second, in cases where the fit quality was visibly poor, a value providing better luminosity optimisation was selected. Following the exclusion of scans affected by an incorrect scan direction and a beam abort, the dependence of iBFBVOT on LER beam current is presented in Fig. 4.16. Measurements obtained before and after the beam abort are shown separately. A clear discrepancy is observed between repeated current points, indicating hysteresis associated with the abort.

This behaviour indicates that the optimal vertical collision offset depends sensitively on beam current and that changes in operating conditions, such as a beam abort, can shift the effective collision configuration. Linear fits are applied independently to the pre-abort and post-abort datasets. The resulting Pearson correlation coefficients of -0.99 and 1.00 indicate a strong linear dependence, with increasing LER current




Figure 4.16: Processed iBFBVOT as a function of LER current during the iBump feedback varied-current machine development study. Data obtained before the beam abort are shown in blue, and data obtained after the beam abort are shown in orange. Linear fits are applied separately to the pre-abort and post-abort datasets, with the Pearson product-moment correlation coefficients indicated.

corresponding to decreasing iBFBVOT.

4.4.5 Correlations with Other Parameters

Although the primary objective of the MD was to study the dependence of the iBFBVOT on beam current, additional analyses were performed to test for correlations with other machine parameters. Correlations were evaluated between each parameter and the optimised iBFBVOT values obtained at the different current settings.

Correlation strength was quantified using the Pearson product-moment correlation coefficient. A detailed description of this metric is provided in Section 4.2.5. For the purposes of this analysis, a strong correlation was defined as having a correlation coefficient magnitude greater than 0.95.

Correlations with Means

Correlations were evaluated between the mean value of each parameter during a scan and the corresponding optimised iBFBVOT. Parameters exhibiting strong correlations are presented in Table 4.4. Strong correlations are observed across multiple PV categories, many of which have clear causal interpretations.

With increasing beam current, the bunch charge increases, leading to a larger beam-beam parameter. This directly explains several correlations. The luminosity, being proportional to the beam-beam parameter, exhibits strong correlation with iBFBVOT. Similarly, the IP feedback canonical kick scales with the beam-beam interaction strength

Table 4.4: Parameters exhibiting strong correlation between their mean values and the optimised iFBVOT during the iBump feedback varied-current MD study. Correlations are quantified using the Pearson product-moment correlation coefficient, evaluated separately for pre-abort data, post-abort data, and the full data set.

EPICS PV	Pre-Abort Correlation	Post-Abort Correlation	All Data Correlation
<i>Beam Current</i>			
BMHDCCT:CURRENT	-0.9898	-0.9986	-0.9624
BMLDCCT:CURRENT	-0.9899	-0.9986	-0.9625
<i>Luminosity</i>			
B2_LUMI:LUMIBELLE2:HER:LUMI	-0.9863	-0.9985	-0.9618
B2_LUMI:LUMIBELLE2:LER:LUMI	-0.9696	-0.9887	-0.9407
<i>IP Beam Size</i>			
BMLXRM:BEAM:SIGMAXatIP	-0.9643	-0.9506	-0.8903
<i>IP Feedback Kick</i>			
BM_IPFB:YCanonicalKick	+0.9653	+0.9965	+0.9897
<i>IR Temperature</i>			
B2_IR:paraffin:BP_BWD:T2	-0.9835	-0.9947	-0.9531
B2_IR:paraffin:BP_BWD:T3	-0.9694	-0.9934	-0.9495
B2_IR:paraffin:BP_BWD:T4	-0.9809	-0.9947	-0.9513
B2_IR:paraffin:BP_BWD:T5	-0.9896	-0.9973	-0.9559
B2_IR:paraffin:BP_BWD:T7	-0.9579	-0.9834	-0.9315
B2_IR:paraffin:BP_BWD:T8	-0.9563	-0.9953	-0.9518
B2_IR:paraffin:BP_FWD:T10	-0.9738	-0.9923	-0.9578
B2_IR:paraffin:BP_FWD:T2	-0.9821	-0.9957	-0.9548
B2_IR:paraffin:BP_FWD:T3	-0.9839	-0.9961	-0.9566
B2_IR:paraffin:BP_FWD:T4	-0.9848	-0.9962	-0.9562
B2_IR:paraffin:BP_FWD:T5	-0.9805	-0.9970	-0.9610
B2_IR:paraffin:BP_FWD:T7	-0.9894	-0.9949	-0.9669
B2_IR:paraffin:BP_FWD:T8	-0.9865	-0.9978	-0.9701
<i>FFQ Gap Sensors</i>			
MG_QCS:GAPL:LOWER-OUT	+0.9656	+0.9633	+0.6081
MG_QCS:GAPR:LOWER-OUT	+0.9858	+0.9863	+0.6843
MG_QCS:GAPR:OutIn	-0.9787	-0.9966	-0.6342
MG_QCS:GAPR:UPPER-IN	-0.9845	-0.9900	-0.8560
MG_QCS:GAPR:UpDown	+0.9529	+0.9793	+0.8827

and is therefore strongly correlated. Although less direct, the horizontal beam size also shows correlation; due to the large crossing angle, increased beam-beam interaction leads to horizontal emittance growth, observed in both simulation and measurement.

In addition, higher beam current increases beam-pipe heating through impedance effects, producing the strong correlations observed in IR temperature PVs.

Of particular significance is the strong correlation observed in the FFQ gap sensor

readings (“MG_QCS” PVs). This is consistent with thermally driven FFQ motion as the IR is observed to heat and cool with beam current. Such motion represents a significant challenge for stable machine tuning.

Correlations with Standard Deviations

Correlations were also evaluated between the standard deviation of each parameter during a scan and the corresponding optimised iFBVOT. The strongest correlations, with magnitudes between 0.8 and 0.95, are presented in Table 4.5.

Table 4.5: Parameters exhibiting strong correlation between their standard deviation values and the optimised iFBVOT during the iBump feedback varied-current MD study. Correlations are quantified using the Pearson product-moment correlation coefficient, evaluated separately for pre-abort data, post-abort data, and the full data set.

EPICS PV	Pre-Abort Correlation	Post-Abort Correlation	All Data Correlation
<i>IP Feedback BPMs</i>			
BM_IPFB:DSP:LW_BPM2_Y	0.8381	0.8602	0.8157
BM_IPFB:DSP:LW_BPM3_Y	0.8178	0.9867	0.8828
BM_IPFB:DSP:LW_BPM4_Y	0.9190	0.9537	0.8705
<i>Specific Luminosity</i>			
CG_OPR:SpecificLuminosity	0.8028	0.9419	0.7468

The standard deviations of the IP feedback BPM readings and the specific luminosity increase with increasing optimised iFBVOT, and therefore decrease with increasing beam current. This behaviour is expected. The BPM signal amplitude scales with bunch charge, so lower beam current leads to reduced signal strength and consequently larger relative fluctuations.

Similarly, luminosity measurements are derived from beam-background processes (predominantly radiative Bhabha scattering). At higher luminosity, increased event statistics reduce statistical uncertainty, leading to smaller observed standard deviations.

These correlations emphasise the importance of high beam current in reducing measurement uncertainty in the feedback-relevant observables.

4.4.6 QCS Magnet Variation

While most parameters exhibit strong correlations across the full data set, the QCS gap sensor readings show strong correlations with iFBVOT when considered separately before and after the beam abort, but not when the full data set is combined. This indicates a significant hysteretic effect introduced by the abort. The strongest examples are the PVs:

- “MG_QCS:GAPL:LOWER-OUT”,

- “MG_QCS:GAPR:LOWER-OUT”,
- “MG_QCS:GAPR:OutIn”,

which measure relative gaps within the QCS magnet assemblies. Linear fits to these parameters as a function of iBFBVOT are presented in Fig. 4.17, with pre-abort and post-abort data treated separately. The distinct offsets between the two regimes are evident.




Figure 4.17: Variation of selected QCS gap parameters as a function of the iBFBVOT during the iBump feedback varied-current machine development study. Measurements obtained before the beam abort are shown in blue, and measurements obtained after the beam abort are shown in orange. Linear fits are applied separately to the pre-abort and post-abort datasets in each panel. The corresponding Pearson product-moment correlation coefficients are indicated.

To further illustrate the origin of this behaviour, the time evolution of the “MG_QCS:GAPR:OutIn” gap parameter is presented in Fig. 4.18. A clear step-like change is observed coincident with the beam abort. This behaviour is not reproduced during normal operation.

A plausible explanation is that the abrupt loss of beam current during the abort leads to a rapid change in local heating, producing a sudden thermally driven mechanical shift of the final-focus quadrupoles. In contrast, the gradual temperature evolution during normal operation leads only to slow motion. This observation demonstrates that QCS magnet positions are both measurably mobile and sensitive to transient thermal conditions, posing a significant challenge for reproducible collision tuning.




Figure 4.18: Variation of a QCS gap parameter as a function of time during the iBump feedback varied-current machine development study. The QCS gap value is shown in black (left-hand axis). The LER beam current is shown in red and the HER beam current in blue (right-hand axis). The grey shaded region indicates the period corresponding to the beam abort.

4.5 Constant Current MD

4.5.1 MD Overview

The iBump Feedback Constant Current MD (CCMD) was performed on 11 December 2024 between 19:20 and 01:00 JST. The aim of this MD was to analyse the stability of the iBump Feedback Vertical Offset Target (iBFBVOT) during constant-current operation. This constant-current operation took place during physics running. The optimal iBFBVOT was identified approximately every 20 minutes over this five-hour period. The optimal iBFBVOT was evaluated by performing an iBFBVOT scan, the details of which are discussed in Section 4.2.

Beam conditions were generally good, with the LER beam current held constant at 1200 mA. Although the HER beam current was nominally held at 960 mA, intermittent current drops occurred due to minor RF issues. Measurements taken during periods affected by RF instabilities were rejected.

4.5.2 iBFBVOT Scans

During this MD, the iBFBVOT was scanned sixteen times. Two scans were rejected due to RF issues occurring during the measurements. A further scan was deliberately continued during an RF issue to assess the impact of a reduced beam current. The final scan was used to test the reproducibility of the iBFBVOT scan procedure when performed in both forward and reverse directions. In total, twelve scans were retained for assessing the stability of the iBFBVOT over the study period. During each scan, 30

luminosity measurements were recorded per iBFBVOT step following a 1 s equilibration period. A summary of the LER and HER beam currents together with the iBFBVOT over the MD period is presented in Fig. 4.19.




Figure 4.19: Summary of the iBump feedback constant-current machine development study. The iBump vertical offset target iBFBVOT (black, left-hand scale) and the LER (red) and HER (blue) beam currents (right-hand scale) are shown as a function of time. Vertical dashed lines mark the start (green) and end (red) of each scan period.

The scan performed during the RF issue was used to cross-check the results of the iBump Feedback Varied Current MD (VCMD, Section 4.4). Although this scan was not as controlled, since the beam current was not constant during the scan and the LER and HER currents were not maintained at operator-defined ratios, it nevertheless confirmed the earlier observation that a reduction in beam current requires an increased iBFBVOT. This is visible in Fig. 4.19, where a pronounced peak in the optimal iBFBVOT coincides with a significant HER beam-current drop, reproducing the VCMD result.

4.5.3 iBFBVOT Stability

During periods of stable machine configuration and constant beam current, the iBFBVOT is observed to remain stable over several hours. As presented in Fig. 4.19, the black trace corresponds to the instantaneous iBFBVOT during each scan, including the deliberate variation applied while stepping the offset target. The relevant quantity for assessing stability is the plateau value reached at the end of each scan, once the luminosity optimisation is complete.

Considering the 12 stable scans, the optimised iBFBVOT values between successive scans show variations at the sub-micron level. These variations are smaller than, or comparable to, the uncertainties associated with the quadratic fits applied to the luminosity data during each scan. No systematic drift of the optimised iBFBVOT with

time is observed. Under constant-current operation and stable machine conditions, the iBFBVOT can therefore be regarded as constant within errors and operationally stable.

4.5.4 iBFBVOT Scan Reproducibility

To test the reproducibility of the iBFBVOT scan procedure, an extended scan was performed. The iBFBVOT was scanned forward, then in reverse, and then forward again, before extending to larger offsets from the optimal value. The variation of luminosity and iBFBVOT during this scan is presented in Fig. 4.20.




Figure 4.20: Variation of the iBFBVOT and relative luminosity as a function of time during a reproducibility test scan during the iBump feedback constant-current machine development study. The iBFBVOT is shown in black (left-hand axis), and the relative luminosity measured by LUMIBELLE2 is shown in red (right-hand axis). The scan consists of a forward scan, a reverse scan, and a subsequent forward scan, followed by an extension to larger iBFBVOT values.

Consistent optimal iBFBVOT values were obtained from quadratic fits to the forward, reverse, and second forward scan segments. Although the fitted optimum was reproducible, the absolute luminosity level was reduced during the reverse scan. The origin of this reduction is not understood.

At the end of the scan, larger iBFBVOT excursions were explored, resulting in substantial luminosity degradation. In this regime, the luminosity profile becomes strongly non-quadratic and the fit quality deteriorates, illustrating the importance of restricting scans to a narrow range around the optimum during optimisation.

Notably, the luminosity remains consistent within errors over a finite range of iBFBVOT values around the optimum. This indicates that small iBFBVOT fluctuations (such as those observed during this MD) are unlikely to result in measurable luminosity degradation.

4.5.5 QCS Magnet Stability

Following the QCS motion observed during the VCMD, the stability of the QCS magnets was also examined during the CCMD. The time evolution of the EPICS PV “MG_QCS:GAPR:OutIn”, corresponding to the out-in separation of the right-side QCS magnets, is presented in Fig. 4.21. This is presented for direct comparison with the VCMD results shown in Fig. 4.18.




Figure 4.21: Variation of a selected QCS gap parameter as a function of time during the iBump feedback constant-current machine development study. The final focus quadrupole gap (right, out-in) is shown in black (left-hand axis), while the LER and HER beam currents are shown in red and blue respectively (right-hand axis).

During the CCMD, this gap parameter varied between $257.5\mu\text{m}$ and $258.75\mu\text{m}$ over a five-hour period. By contrast, during the VCMD the same parameter varied between $336.5\mu\text{m}$ and $340.0\mu\text{m}$ over approximately three hours. The substantially reduced motion observed during constant-current operation indicates that variations in beam current play a dominant role in driving QCS displacement. Consistent trends were also observed in the IR temperature sensors, supporting the conclusion that the QCS motion is primarily thermally driven.

4.6 Historical iBump Data

In addition to the dedicated MD periods, historical operational data stored in the KEKBlog (see Section 4.2.3) were analysed over extended periods to investigate whether the trends observed during the MDs persist during regular operation.

4.6.1 Extracting iFBVOT Scans

Over long periods, extracting individual iFBVOT scans is challenging. The variation of the LER and HER beam currents, and iFBVOT during 2022 operation are presented in Fig. 4.22. Scan boundaries are not readily identifiable in the time-domain signals.




Figure 4.22: Long-term variation of the iFBVOT (left-hand axis) together with the LER and HER beam currents (right-hand axis) during 2022 operation.

During regular operation, the iFBVOT is held fixed between optimisation scans and is only varied while a scan is actively being performed. Scan periods can therefore be identified as intervals in which the iFBVOT deviates from an otherwise constant baseline value. For each identified scan, the optimised iFBVOT is taken as the new constant value reached after the scan has completed. Other parameters used for comparison are characterised by their mean values over the duration of the scan.

Depending on machine conditions, scans may be abandoned or coincide with significant changes in machine configuration. This behaviour was observed during both the VCMD and CCMD, where scans were interrupted by beam aborts and RF issues. Even during these controlled MD periods, manual data cleaning was required. Extending such manual cleaning to data sets spanning several months is not feasible within the scope of this thesis. Furthermore, when different machine configurations are used, results from separate scans are not necessarily comparable.

Despite these limitations, a coarse analysis of historical data was performed to assess whether the trends observed during the MDs are also visible over longer timescales.

4.6.2 iFBVOT Variation with Current

To test the observations of the VCMD and CCMD (see Section 4.4 and Section 4.5), that the iFBVOT remains constant under constant-current conditions and varies linearly with beam current, the correlation between the iFBVOT and LER beam

current was evaluated using historical data from the 2022 and 2024 operation periods. The results are presented in Fig. 4.23.




Figure 4.23: Variation of the optimised iBFBVOT with LER beam current for the 2022 and 2024 SuperKEKB operation periods. Each point corresponds to an individual optimisation scan and is coloured according to the time at which the scan was performed.

In contrast to the VCMD results, no clear dependence of the optimised iBFBVOT on beam current is observed. Furthermore, unlike the CCMD results, a large spread in optimised iBFBVOT values is present at fixed beam current. This indicates that over extended operation periods the optimal iBFBVOT is influenced by additional factors beyond beam current alone.

The iBFBVOT compensates for errors arising from multiple sources, including BPM alignment, FFQ position, and transfer matrix uncertainties. Over long time periods, these quantities evolve independently, particularly following optics changes and beam aborts. As demonstrated during the VCMD, aborts can introduce hysteretic changes in machine state. Consequently, over extended operation periods the optimal iBFBVOT cannot be reliably predicted from beam current alone, and collision tuning is required following each abort or optics modification.

4.6.3 Time Dependence

The variation of the iBFBVOT with time was also examined, and results are presented in Fig. 4.24. Some structure is visible in the data, but it is neither continuous nor stable. In particular, even within visually clustered regions, most notably during 2022 operation, the iBFBVOT values exhibit large scatter. This indicates that no single, well-defined optimal offset persists over extended periods.




Figure 4.24: Variation of the optimised iBFBVOT with time for the 2022 and 2024 SuperKEKB operation periods. Each point corresponds to an individual optimisation scan.

Furthermore, the mean optimised iBFBVOT differs between the 2022 and 2024 operation periods by nearly 0.5 mm. This corresponds to a vertical displacement significantly larger than the vertical beam size, highlighting the strong dependence of the optimal collision configuration on machine state over long timescales.

4.7 Summary

Machine Development studies of the SuperKEKB iBump IP feedback system were performed during the 2024ab and 2024c operational periods. In addition, the collision tuning and commissioning process for the 2025c run was examined. Together, these

studies provide detailed insight into the operational behaviour of a fast IP feedback system under realistic machine conditions.

A central focus of this work was the iBump Feedback Vertical Offset Target (iBFBVOT), which compensates for discrepancies between the idealised feedback response predicted by simulation and the optimal collision conditions observed in operation. During the Varied Current MD, a strong linear dependence of the optimal iBFBVOT on beam current was observed. Correlated variations were also identified in interaction-region temperatures and in the positions of the QCS final-focus quadrupoles, indicating that thermal and mechanical effects play a significant role in determining the effective collision offset.

Beam aborts were found to induce hysteretic changes in the machine state. These effects were particularly evident in the QCS magnet positions and in the relationship between iBFBVOT and beam current before and after an abort. As a result, over extended operation periods the optimal iBFBVOT cannot be reliably predicted from previous measurements alone. Collision tuning is therefore required following each beam abort or optics modification.

The stability of the iBFBVOT under steady operating conditions was investigated during the Constant Current MD. Over a five-hour period of stable physics running, the optimised iBFBVOT remained constant within measurement uncertainties. Small variations were observed during periods of RF-induced current fluctuations, reinforcing the current dependence identified in the Varied Current MD. Scan reproducibility tests demonstrated that, despite a small forward–reverse asymmetry, the optimal iBFBVOT could be consistently recovered.

The sensitivity of luminosity to small deviations in iBFBVOT was found to be weak under the SuperKEKB operating conditions studied here. Near the optimum, the luminosity profile is shallow, and variations in the optimised iBFBVOT at the micron level did not result in observable luminosity degradation. However, all measurements were performed far from nominal design beam sizes, and increased sensitivity is expected as tighter collision conditions are approached.

These results have direct implications for the design and operation of IP feedback systems at FCC-ee. In particular, the strong and potentially hazardous dependence of feedback behaviour on beam current is of concern given the planned bootstrap injection scheme, in which collisions occur during filling and beam current varies continuously by design. The SuperKEKB experience demonstrates that feedback performance cannot be assumed to be uniform across operating conditions and that current-dependent and hysteretic effects must be explicitly accounted for. With multiple IPs foreseen at FCC-ee, each requiring independent feedback systems, robust strategies for feedback optimisation, re-tuning, and coordination will be essential to ensure stable and reliable operation.

Chapter 5

Modelling SuperKEKB and the iBump Feedback System

5.1 The SuperKEKB Optical Lattices

5.1.1 The Low Energy Ring (LER)

The Low Energy Ring (LER) of the SuperKEKB collider is a storage ring, for positrons at 4 GeV. The LER is formed of four arc sections and four technical straight sections. The four arcs are each composed of four standard arc cells, and one arc mid cell. The mid cell has a different periodic structure, due to tunnel constraints (inherited from KEKB and Tristan [117]). The arcs have two subtly different optics, with the arcs between the Tsukuba and Nikko straight sections and the arc between the Nikko and Fuji straight sections sharing one optics while the arcs between the Fuji and Oho straight sections and the arc between the Oho and Tsukuba straight sections share another optics. The LER ring is traversed counter-clockwise as seen from above. In order of particle traversal from the positron beam transport line, BTP, the beam encounters the Fuji Straight section, the Oho Straight section, the Tsukuba straight section and finally the Nikko straight section. The Fuji straight section is used for injection and extraction and also features normal conducting RF cavities. The Oho straight section is used for emittance control with a wiggler and also features the Non-Linear Collimator for improved momentum collimation. The Tsukuba straight section is the experimental straight section, where the Belle II experiment is located. Finally the Nikko straight section is used for emittance control with a wiggler and also features normal conducting RF cavities.

The linear optics as well as magnet strengths for the LER (2024ab operation, during the VCMD period described in Section 4.4) is presented in Fig. 5.1.




Figure 5.1: SuperKEKB Low Energy Ring Twiss optical functions and magnet strengths as a function of longitudinal position s . The square root of the horizontal and vertical beta functions, $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$ (left-hand scale) and the horizontal and vertical dispersions, D_x and D_y (right-hand scale) are shown. Coloured bars indicate the longitudinal locations and relative strengths of lattice elements, as identified in the legend.

5.1.2 The High Energy Ring (HER)

The High Energy Ring (HER) of the SuperKEKB collider is a storage ring, for electrons at 7 GeV. The HER is formed of four arc sections and four technical straight sections. The four arcs are each composed of four standard arc cells, and one arc mid cell. The arcs have two subtly different optics, with the arcs between the Tsukuba and Oho straight sections and the arc between the Oho and Fuji straight sections sharing one optics while the arcs between the Fuji and Nikko straight sections and the arc between the Nikko and Tsukuba straight sections share another optics. The HER ring is traversed clockwise as seen from above. In order of particle traversal from the electron beam transport line, BTe, the beam encounters the Fuji Straight section, the Nikko Straight section, the Tsukuba straight section and finally the Oho straight section. The Fuji straight section is used for injection and extraction. The Nikko straight section is used for superconducting RF cavities. The Tsukuba straight section is the experimental straight section, where the Belle II experiment is located. Finally, the Oho straight section is used for emittance control with a wiggler and also features normal conducting RF cavities.

The linear optics as well as magnet strengths for the HER (2024ab operation, during the VCMD period described in Section 4.4) is presented in Fig. 5.2.




Figure 5.2: SuperKEKB High Energy Ring Twiss optical functions and magnet strengths as a function of longitudinal position s . The square root of the horizontal and vertical beta functions, $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$ (left-hand scale) and the horizontal and vertical dispersions, D_x and D_y (right-hand scale) are shown. Coloured bars indicate the longitudinal locations and relative strengths of lattice elements, as identified in the legend.

5.1.3 The SuperKEKB Machine Detector Interface (MDI)

The Machine Detector Interface (MDI) region of SuperKEKB is one of the most challenging aspects of both operation and modelling. Situated in the Tsukuba straight section, this region features a challenging layout with many different key hardware elements. A plan view of the Tsukuba straight section local to the Interaction region is presented in Fig. 5.3. This region is highly complex, with significant bending, especially




Figure 5.3: Plan view of the Tsukuba straight section local to the IP. The horizontal axis shows the coordinate Z in metres while the vertical axis shows the coordinate X in metres. The Low Energy Ring (LER) is shown in red and the High Energy Ring (HER) is shown in blue.

in the LER to accommodate all required elements.

The MDI region features strong local chromatic correction sextupoles for both the horizontal and vertical planes (SLX and SLY type sextupoles respectively). The vertical chromatic correction sextupoles are also used to generate the crab waist. SuperKEKB utilises a virtual crab waist scheme [118]. The chromatic correction sextupoles on each side of the IP would typically be of identical strength. For the virtual crab waist there is a deliberate asymmetry in the sextupole pairs. The imperfect cancellation leaves a residual sextupolar component that acts as the crab waist sextupole.

The Interaction Region (IR) is particularly complex. To reduce the distance between the pole tip of the Final Focus Quadrupoles (FFQs) and the IP (a distance denoted L^*), the FFQs are overlapped with the experimental solenoid of Belle II. The Belle II solenoid is a superconducting solenoid producing a peak field of 1.5 T [45]. This causes strong coupling between the horizontal and vertical planes. To correct for this, a strong anti-solenoid is included each side of the IP and the FFQs feature strong dipolar, quadrupolar, sextupolar and octupolar correctors and the FFQs are offset and rotated [16]. Modelling this complex configuration and the interplay of these different magnetic fields required self-consistent 3D magnetostatics modelling. Simulations were performed using the 3D magnetostatics solver in the code OPERA [119] and some example field maps are presented in Fig. 5.4.




Figure 5.4: Three-dimensional magnetic field map of the SuperKEKB interaction region, originally reported in [120]. The Cartesian coordinate axes x , y , and z are shown in metres. Coloured contours represent the magnitude of the magnetic field (left-hand scale).

To correct for the beam orbit, coupling and dispersion due to the strong solenoidal field, a number of skew quadrupoles (five per side in the LER, three per side in the HER) and vertical bends are used for correction.

5.2 Modelling SuperKEKB Optics with Xsuite

5.2.1 The SAD Model of SuperKEKB

Modelling SuperKEKB has historically been performed using SAD [89]. The development of SAD and the development of the model of SuperKEKB have been intrinsically linked. Many code developments are a direct result of SuperKEKB's simulation needs. As a compiled FORTRAN code, modelling SuperKEKB with SAD has been powerful and efficient.

A majority of elements in the SAD model of the SuperKEKB lattice are standard for a range of tracking and optics codes. Some elements are noted to be more troublesome to convert [121]. These features will be discussed further here.

Reference Frame Choices

The models of the HER and LER feature a combination of unphysical reference frame representations and physical reference frame choices that differ from the conventions adopted at CERN. The choices made do not directly impact single beam performance in tracking or optics, but present challenges when performing tracking with self-consistent beam-beam.

There is a difference in standard choice of global reference system between the accelerator models used at KEK and those at CERN. At KEK, the convention has been adopted that the longitudinal coordinate s proceeds along the beam direction and the horizontal coordinate x is towards the outside of the ring. The vertical coordinate y is chosen such that a right-handed coordinate system is produced. This is different from the CERN conventions, where s is similarly along the beam direction, but y is chosen to point skywards. The horizontal plane x is chosen such that a right-handed coordinate system is produced.

Further to the overall reference frame conventions, there are specific subtleties to the modelling of these rings. The LER is modelled with the bends in the wrong direction, such that the ring is clockwise. Computing the Twiss optical functions in this configuration results in positive dispersion, where the physical dispersion would be negative.

The HER is modelled with the order of the elements reversed and the charge reversed acting as a positron ring going clockwise.

MDI Model

As described in Section 5.1.3 the MDI region of the SuperKEKB is highly complex. To simulate this region full self-consistent magneto-static models are utilised. Such 3D field-maps are not possible to read directly into SAD. As such, magnetic field-maps from

ANSYS simulations [107] (equivalent to the aforementioned OPERA simulations) were expanded into multipolar components along the beam-line axis [122]. The multipolar expansion was performed in 1 cm slices from -4 m to 4 m from the IP for each beam-line along its respective crossing angle through the IP. Thin kicks are also included alongside the thick solenoid/multipole slices. This makes the MDI region highly computationally expensive due to the fine expansion of the fields.

The solenoid at SuperKEKB is a defining bend in the reference particle's trajectory. With a large half crossing angle of 41.5 mrad, the beams traverse the solenoid off centre and are strongly bent. This is accounted for by using reference frame transformations at the boundaries of the sliced region. A combination of transverse and longitudinal shifts, as well as rotations are utilised to account for this.

Fringes

The modelling of magnetic fringes in SAD is distinct from the approach adopted by many other common modern accelerator codes. The parameter F1 indicates the longitudinal extent over which the fringe field is applied. A diagram from the SAD manual is presented in Fig. 5.5.




Figure 5.5: Schematic definition of the SAD fringe parameter F1. The vertical axis shows the vertical magnetic field component B_y as a function of the longitudinal coordinate s .

Such fringes are not easily convertible between optics codes and the optics are strongly dependent on these fringes. Some magnets are highly fringed, for example, the LER wiggler magnet “BW0NLP” (situated in the Nikko straight section wiggler) is parameterised by a length of 216.2 mm, an integrated strength of $k_0l = 10.7248$ mrad and has the fringe parameters $FB1 = 203.6$ mm and $FB2 = 146.7$ mm where $FB1$ and $FB2$ are equivalent to $F1$ demonstrated in Fig. 5.5, but independently specifying the entry and exit lengths. A diagram of this setup is presented in Fig. 5.6. The field of the wiggler magnet is dominated by the effect of the fringes, and the approximation of a thick constant field bulk with thin fringes is lost.




Figure 5.6: Schematic representation of the fringe extent of the LER wiggler BW0NLP in the SAD model of SuperKEKB. The longitudinal coordinate is shown along the horizontal axis. The central region of length L corresponds to the main field region with integrated dipole strength $k_0 L$. The parameters FB1 and FB2 indicate the upstream and downstream fringe-field extents, respectively.

5.2.2 Conversion Motivation

The work of converting the optics of SuperKEKB from SAD to Xsuite has been a major endeavour. The major motivation of using Xsuite for SuperKEKB modelling is as a test for FCC-ee simulation. The tool of choice for FCC-ee particle tracking simulations is Xsuite. The Xsuite ecosystem provides a wide range of tools natively and is also powerful in its ability to interface with other state of the art codes, such as MAD-NG [123]. While the tool has been heavily benchmarked against the machines of the CERN accelerator complex, modelling SuperKEKB as the current state of the art luminosity frontier high energy lepton collider is a vital step towards validating Xsuite for simulation of FCC-ee.

A range of users have already adopted the converted Xsuite SuperKEKB lattices. This is enabling rich and varied studies ranging from collimation to non-linear optics. Comparisons of Xsuite SuperKEKB simulations with measured experimental data have been performed. The simulations show excellent agreement with measurements, supporting that the Xsuite SuperKEKB lattices are a good representation of the real machine. Examples of such studies, analysing detector dose rates and injection efficiency while varying collimators, are presented in Fig. 5.7.

5.2.3 Conversion Process

The process of converting the lattices from SAD to Xsuite is mature [126]. Due to simulation differences between the two codes the models will never be exactly identical. The most notable differences stem from the impact of the fringe fields and in the Hamiltonian used for simulating solenoids with overlapping multipolar components.




Figure 5.7: Results from collimation studies using SuperKEKB Xsuite lattices. Left: relative dose rate as a function of the D06V1 aperture from [124]. Measured experimental data are shown in blue while simulation results are shown in red. Right: normalised dose rate as a function of the D02V1 lower jaw offset from [125]. Experimental data are shown in red while simulation results are shown in blue.

While perfect agreement with SAD would initially be an idealised path, the most critical aspect is the agreement with measured data, which is demonstrated in Section 5.2.2. Benchmarks with SAD remain useful metrics as the SAD model has been benchmarked directly with measurements. Overall, the lattices are deemed performant for studies.

The conversion process occurs in two stages: initial conversion using the SAD2XS converter followed by optics re-matching. An example Twiss of the HER after the initial import and after the optics re-matching process is presented in Fig. 5.8. Dispersion leakage into the straight sections is visible prior to re-matching. Furthermore the uniformity of the periodic structures is broken in both the dispersion and the beta functions. After re-matching a fully periodic set of Twiss parameters are recovered.

The SAD2XS Converter

A SAD to Xsuite converter (SAD2XS) has been developed [127]. The project is publicly available on GitHub and PyPI and has been tested for use with the SuperKEKB LER and HER lattices, SuperKEKB BTeV and BTeV transport lines, the J-PARC Main Ring and many FCC-ee lattice versions. As is visible in Fig. 5.8 the automated conversion cannot account for the modelling differences in the case that fringes are a dominant effect on the optics.

Physical transformations are performed at this import stage. As discussed in Section 5.2.1 the lattices are stored using reference frames that make beam-beam simulations more challenging. At this stage in the conversion process, lines are reversed, bend directions switched and reference particle charges are corrected such that lattices




Figure 5.8: Horizontal and vertical Twiss optical functions for the SuperKEKB High Energy Ring as a function of longitudinal position s . The top panel shows the optics immediately after lattice import, and the bottom panel shows the optics after optics re-matching. The square root of the horizontal and vertical beta functions, $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$ (left-hand scale), and the horizontal and vertical dispersions, D_x and D_y (right-hand scale), are shown. Coloured bars indicate the longitudinal locations and relative strengths of lattice elements.

have physically meaningful element orders, charges and surveys.

The Optics Re-Matching Process

Optics re-matching is required to account for the impact of the SAD fringes and to correct reference frame transformations around the solenoid due to subtle differences in the orbit. The development of optics re-matching recipes for the SuperKEKB LER and HER lattices was a major endeavour. While respecting the powering circuits of the magnets in the SuperKEKB lattices, it is not possible to match all Twiss parameters at

all points; there are insufficient degrees of freedom to meet these constraints. Certain aspects are therefore prioritised in the matching process. The overall structure of the optics re-matching process is the following:

1. Rematch the strength of the wiggler kicker magnets to match the energy loss through the wigglers.
2. Rematch phase advances of the arc cell (overall and sextupole pairs).
3. Rematch the arc mid cell to periodically match the arc cell.
4. Rematch each of the Oho, Fuji and Nikko straight sections:
 - Respecting the key Twiss parameters at the injection point.
 - Respecting the periodicity of structures such as FODO cells.
 - Respecting the phase advances of each straight section and all sextupole pairs.
5. Rematch the Tsukuba straight section:
 - Rematch the solenoid entry and exit reference frame transforms.
 - Rematch the IR Coupling using the skew quadrupoles.
 - Rematch the IP optics (beta, alpha, dispersion) and sextupole phases.
6. Rematch the global tune using quadrupoles in the Fuji straight section.

In addition to the above, a final step of adding Xtrack “ZetaShift” elements is performed. These zeta shifts are equivalent to performing a frequency trim to the RF Cavities (performed in SAD using the `fshift` parameter). This is performed to ensure that computation of closed-ring 6D linear optics, including synchrotron motion, is possible.

Through the re-matching process no machine bends, those which define the reference path of the machine, are altered. The wigglers are the only bending elements which are rematched. As discussed in Section 5.2.1, the effect of the fringes on the wiggler bends is very strong. Significant re-matching is required to match the energy loss through the wiggler. Matching the energy loss results in the creation of a significant orbit discrepancy with respect to SAD within the wiggler. Matching the energy loss has been deemed a greater priority than matching the orbit in this region to ensure agreement on the equilibrium emittances of the lattice, for which the energy loss in the wiggler is defining. The impact on the orbit in the wigglers is presented in Fig. 5.9. A notable reduction in the orbit is observed in both the LER and HER to ensure the energy loss over the wiggler is comparable.

Through the re-matching process all quadrupoles are used for optics re-matching with the exception of the sliced final focus quadrupoles. Strength limits are imposed on




Figure 5.9: Comparison of horizontal orbit in the SuperKEKB wiggler region as a function of longitudinal position s . The top panel shows the horizontal orbit for the LER, x_{LER} , and the bottom panel shows the horizontal orbit for the HER, x_{HER} . Results from SAD are shown by solid red lines, and results from Xsuite are shown by blue dashed lines.

the quadrupoles during optics re-matching to ensure that the new strengths are not far from the values observed in SAD. Histograms showing the strength changes between the import and rematched optics are presented in Fig. 5.10. Strength changes are small, on the order of 10^{-3} or less. During re-matching, strength changes are limited to 1% of the initial strength. Magnets that were initially unpowered remain unpowered. At the IR the initial strengths of the skew quadrupoles are so low that 1% strength changes are too small to be appreciable. In this case the tolerances are relaxed. The quadrupoles with the largest strength changes are the FODO quadrupoles in the straight sections where the wiggler magnets have been adjusted. These quadrupoles must be adjusted to account for the difference in phase advance from the different fringe model and bend strengths.




Figure 5.10: Distribution of quadrupole strength changes for SuperKEKB between imported and rematched Xsuite optics. The left panel shows the LER and the right panel shows the HER. The horizontal axis shows the change in quadrupole strength, and the vertical axis shows the number of elements.

No sextupoles are adjusted in the optics re-matching process. As the beta functions throughout the lattice are in good agreement, and no sextupoles are strongly fringed in SAD, no re-matching is required.

5.2.4 Linear Optics Benchmarking

As discussed in Section 5.2.3, the most important test of the lattice performance is benchmarking against real machine measurements. Due to limitations on beam time at SuperKEKB a direct benchmarking campaign was not possible. Benchmarking against the SAD model is an effective proxy as this model has itself been benchmarked against measurements.

The first benchmark is on the orbit in the machine. While differences in the orbit are known at the wigglers, and demonstrated in Fig. 5.9, at all other points in the lattices the orbit should be in agreement. The most critical region is the IR, where the orbit is highly complex and the beam dynamics are highly dependent on the orbit through the solenoid and offset final focus quadrupoles. A comparison of the orbits computed by the Xsuite and SAD Twiss functions is presented in Fig. 5.11. The orbits in this region are




Figure 5.11: Comparison of horizontal and vertical interaction region orbit for SuperKEKB between SAD and Xsuite as a function of longitudinal coordinate s . The top two panels show the horizontal orbits x_{LER} and x_{HER} , and the bottom two panels show the vertical orbits y_{LER} and y_{HER} . Results from SAD are shown by solid red lines, and results from Xsuite are shown by blue dashed lines.

highly complex due to the solenoid and strong offset final focus quadrupoles. Vertical jumps in the orbit correspond to the reference shifts at the entry and exit of the sliced

solenoid region. The orbits in this region show excellent agreement. The re-matching of the reference frame transforms was below the scale of nanometre differences, and the resulting orbits are fully closed.

Beyond the orbit, the linear optics have been extensively benchmarked, in particular the beta functions, alpha functions, dispersions and momentum dispersions. A comparison of the beta functions between SAD and Xsuite is presented in Fig. 5.12. All functions are observed to be in excellent agreement. Following these benchmarks, it




Figure 5.12: Comparison of horizontal and vertical interaction region beta functions for SuperKEKB between SAD and Xsuite as a function of longitudinal coordinate s . The top two panels show the horizontal beta functions $\beta_{x,LER}$ and $\beta_{x,HER}$, and the bottom two panels show the vertical beta functions $\beta_{y,LER}$ and $\beta_{y,HER}$. Results from SAD are shown by solid red lines, and results from Xsuite are shown by blue dashed lines.

is deemed that the linear optics is well matched.

5.2.5 Non-Linear Optics Benchmarking

Following the linear optics benchmarking, the non-linear optics were tested. While no sextupoles have been re-matched, the beta functions in the ring are subtly different post the re-matching of the quadrupoles. As such the chromatic properties of the lattice are expected to show minor differences; the same strength of sextupole now exists with a slightly different beta function. The first test to compare this is by considering the chromatic functions around the lattices. A comparison of the chromatic functions computed using SAD and Xsuite is presented in Fig. 5.13. The chromatic functions are computed natively by the Xsuite Twiss computations, while in SAD the chromatic




Figure 5.13: Comparison of horizontal and vertical chromatic functions for SuperKEKB between SAD and Xsuite as a function of longitudinal coordinate s . The top two panels show the horizontal chromatic functions $W_{x,LER}$ and $W_{x,HER}$, and the bottom two panels show the vertical chromatic functions $W_{y,LER}$ and $W_{y,HER}$. Results from SAD are shown by solid red lines, and results from Xsuite are shown by blue dashed lines.

function was computed from three Twiss computations (one on-momentum and two off-momentum). The agreement is good, though small discrepancies exist, and this impacts the non-linear behaviour.

To assess the impact of the discrepancies observed in the chromatic functions, the chromaticity is considered. A scan of the fractional tune vs the momentum deviation delta was performed to see the non-linear behaviour. The results of these scans are presented in Fig. 5.14. For some high values of delta Xsuite was unable to compute the tune where SAD was. These values approached the limits of the dynamic aperture of the machine and the Xsuite Twiss, based on particle tracking, becomes unstable in these regions. While discrepancies are observed at high values of delta, the curvature close to 0 is in excellent agreement.




Figure 5.14: Comparison of horizontal and vertical tunes for SuperKEKB between SAD and Xsuite as a function of relative momentum deviation δ . The top two panels show the horizontal fractional tunes $q_{x,\text{LER}}$ and $q_{x,\text{HER}}$, and the bottom two panels show the vertical fractional tunes $q_{y,\text{LER}}$ and $q_{y,\text{HER}}$. Results from SAD are shown by solid red lines, and results from Xsuite are shown by blue dashed lines.

5.3 Radiation and Beam–Beam Modelling for Multi-Turn Tracking

5.3.1 Synchrotron Radiation Modelling

To simulate equilibrium beam dynamics at SuperKEKB using particle tracking, synchrotron radiation must be modelled. Energy loss by synchrotron radiation determines the equilibrium emittances and hence the strength of the beam–beam interaction and the achievable luminosity. After re-matching the wiggler energy loss, the total synchrotron radiation loss is expected to be approximately equivalent in the SAD and Xsuite models.

SuperKEKB does not employ magnet tapering (adjusting magnet strengths according to the local equilibrium momentum deviation). As a result, when synchrotron radiation is included, the equilibrium closed orbit differs from the reference orbit obtained without radiation. With RF cavities operated at the zero-crossing phase (0° in SAD or 180° in Xsuite), the reference particle acquires a non-zero longitudinal offset at the IP. For beam–beam performance, the LER and HER bunches must arrive at the IP simultaneously. In these simulations, this condition is achieved by re-matching the RF phase in Xsuite such that the equilibrium longitudinal coordinate satisfies $\zeta = 0$ at the IP.

The turn-by-turn evolution of the longitudinal phase space in SAD and Xsuite with

both nominal RF settings and the RF phase re-matched is presented in Fig. 5.15. In all cases, 1000 particles are initialised at $(0, 0, 0, 0, 0, 0)$ and tracked for 20,000 turns (approximately five damping times) to ensure convergence to equilibrium.




Figure 5.15: Turn-by-turn evolution of the longitudinal coordinate ζ (top row) and relative momentum deviation δ (bottom row). The left column shows the LER and the right column shows the HER. Results from SAD are shown in red. Results from Xsuite with nominal RF settings are shown in blue, and with adjusted RF settings are shown in purple.

During the RF re-matching procedure, a single common RF phase offset is applied uniformly to all cavities, with the matching constrained solely by the longitudinal coordinate. This introduces a small shift in the equilibrium relative momentum deviation δ , which is negligible compared with the natural turn-by-turn and azimuthal variations of δ around the ring and does not affect the overall longitudinal dynamics.

5.3.2 Equilibrium Emittances with Synchrotron Radiation

Both SAD and Xsuite are capable of computing equilibrium emittances from radiation integrals with mean radiation; in Xsuite through a computation of the Twiss optical functions and in SAD using the EMIT module. While this is excellent for a quick calculation, such an estimate does not take into account non-linearities and resonances. A better approach is to use a quantum synchrotron radiation model and particle tracking. In these simulations, 1000 particles were initialised at $(0, 0, 0, 0, 0, 0)$ and tracked for 20,000 turns (approximately five damping times) to allow the longitudinal dynamics to reach equilibrium.

A comparison of the emittance evolution between SAD with nominal RF, Xsuite with nominal RF and Xsuite with RF rematched for zero zeta at the IP is presented in Fig. 5.16. Agreement is observed in both the radiation integral emittance estimations




Figure 5.16: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for SuperKEKB between SAD and Xsuite. The left column shows the LER and the right column shows the HER. Results from SAD are shown in red. Results from Xsuite with nominal RF settings are shown in blue, and with adjusted RF settings are shown in purple. Dashed lines indicate the corresponding equilibrium emittance values obtained from Twiss calculations.

and in the tracking emittance computations. A discrepancy is observed in the evolution of the vertical emittance with the Adjusted RF. As was presented in Fig. 5.15 for the SAD and Xsuite Nominal RF the particles are initialised far from equilibrium. As such the curves presented in Fig. 5.16 are not just the emittance damping rates, but a convolution with the rate of damping to equilibrium orbits. Ultimately the same equilibrium emittances are achieved and as such the lattices are deemed performant with synchrotron radiation.

5.3.3 Generating Vertical Emittance

While finite values of emittances have been obtained, these values correspond to the ideal lattices. Both the longitudinal and horizontal emittances are in agreement with measurements, however the vertical emittance is much lower than measured in the real machine. The horizontal and longitudinal emittances are primarily a result of the energy lost in the main bends as the particles are bent in approximately a circle over the ~ 3 km ring. On the other hand, the only sources of vertical emittance in these model lattices is the IR (solenoid and vertical orbit bumps) and the beam avoidance crossing in the Fuji straight section where the LER and HER beams cross over. In the real machine, the vertical emittance is a result of errors and offsets of the elements. Measurements at SuperKEKB may observe a minimum vertical emittance of approximately 25 pm, but a

more typically observed vertical emittance is on the order of 40 pm [128].

With the lattice only vertical emittance, the beam sizes at equilibrium are too small. The beam–beam tune shift therefore becomes too high and results in unstable beam–beam performance. To obtain a sensible beam–beam tune shift, the vertical emittance must be tuned to achieve a more stable setup.

For this study, it was chosen to generate vertical emittance by offsetting sextupole pairs. To mitigate the negative impact on the machine optics, each sextupole pair is offset oppositely; the first sextupole is offset upwards and the second sextupole downwards with equal magnitudes. This creates closed coupling bumps which result in vertical emittance. For simplicity, all sextupoles in each of the LER and HER were offset by the same magnitude. By scanning across the offset values, it is possible to optimise for a vertical emittance of 40 pm. This scan was performed initialising 100 macroparticles as a point beam at $(0,0,0,0,0,0)$ and allowed to evolve for 20,000 turns under the effect of quantum synchrotron radiation. This smaller number of macroparticles was chosen for computational reasons. The results of this offset scan are presented in Fig. 5.17.




Figure 5.17: Horizontal, vertical, and longitudinal emittances as a function of sextupole offset for SuperKEKB. The left column shows the LER and the right column shows the HER. The top row shows the horizontal emittance ϵ_x , the middle row shows the vertical emittance ϵ_y , and the bottom row shows the longitudinal emittance ϵ_z . The horizontal axis shows the sextupole offset in metres. The red dashed horizontal line indicates the reference vertical emittance value of 40 pm.

Optimised sextupole offsets of 285 μm for the LER and 67.5 μm for the HER were identified. Using these offsets, a final test was performed using now 1000 macroparticles

to ensure good agreement with the target vertical emittance. The results are presented in Fig. 5.18.




Figure 5.18: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for SuperKEKB with optimised sextupole offsets. The left column shows the LER and the right column shows the HER. Dashed horizontal red lines indicate target vertical emittance values.

The resulting equilibrium emittances under synchrotron radiation with sextupole offsets are presented in Table 5.1.

Table 5.1: Equilibrium horizontal, vertical and longitudinal emittances for SuperKEKB obtained from particle tracking including quantum synchrotron radiation, shown for the LER and HER with and without optimised sextupole offsets.

Parameter	Unit	LER		HER	
		W/O Offsets	W/ Offsets	W/O Offsets	W/ Offsets
ε_x	nm	3.19	4.26	3.49	4.71
ε_y	pm	0.755	41.7	0.624	42.0
ε_z	μm	2.95	3.37	2.73	3.22

Impact on Machine Optics

The optimised sextupole offsets of $285\text{ }\mu\text{m}$ for the LER and $67.5\text{ }\mu\text{m}$ for the HER are significant. While in principle offsetting each pair of sextupoles equal and opposite should produce closed coupling bumps, in the case of large offsets or in the presence of synchrotron radiation without tapering this assumption breaks down. In fact, due to

the lack of tapering optics distortions occur with radiation present. In the full machine, optics corrections are performed to correct for this as much as possible. Implementing full optics correction knobs in the presence of radiation was beyond the scope of this thesis. As such, these residual errors remain and are taken to be approximations of residual errors that may be present in the real machine. The impacts of radiation and sextupole offsets on selected key Twiss parameters are presented in Table 5.2.

Table 5.2: IP optical functions and tunes for the LER and HER computed with and without synchrotron radiation and with and without optimised sextupole offsets.

Parameter	Unit	W/O Radiation W/O Offsets	W/ Radiation W/O Offsets	W/O Radiation W/ Offsets	W/ Radiation W/ Offsets
<i>LER</i>					
β_x^*	mm	80.0	75.4	80.0	75.4
β_y^*	mm	1.00	0.996	1.00	0.995
D_x^*	μm	-0.011	86.2	36.8	123
D_y^*	μm	-0.002	0.131	-205	-208
$D_{p_x}^*$	μrad	-0.143	-823	361	-471
$D_{p_y}^*$	μrad	-0.294	298	-63500	-61800
Q_x	1	44.526	44.526	44.526	44.526
Q_y	1	46.585	46.582	46.585	46.582
<i>HER</i>					
β_x^*	mm	60.0	56.0	60.0	56.0
β_y^*	mm	1.00	0.973	1.00	0.973
D_x^*	μm	0.0231	-28.3	-0.0328	-28.4
D_y^*	μm	-0.0165	-0.0127	204	206
$D_{p_x}^*$	μrad	0.0528	-255	0.194	-254
$D_{p_y}^*$	μrad	-3.52	-4.50	25900	26500
Q_x	1	45.533	45.532	45.534	45.532
Q_y	1	43.575	43.574	43.575	43.574

Bunch Generation

In simulations including beam-beam forces, bunches must be generated close to equilibrium, as excessively small initial beam sizes lead to unrealistically large beam-beam tune shifts and immediate beam loss. For the SuperKEKB lattices considered here, the large equilibrium momentum deviation and longitudinal offset therefore require additional care during bunch generation.

A matched Gaussian distribution was constructed by first generating a longitudinal distribution in ζ and δ matched to the RF bucket, then shifting it to the radiation equilibrium orbit obtained from optics calculations. The transverse distribution was subsequently generated using the specified geometric emittances.

The procedure was validated by tracking 100,000 macroparticles for 1000 turns,

during which the emittances remained stable with only small synchrotron oscillations. This bunch generation method is used for all subsequent studies.

5.3.4 Beam–Beam Modelling

To model the impact of the beam–beam force and compute luminosity the Xfields “xfields.BeamBeamBiGaussian3D” element is used. A single beam–beam element is included at the IP. This element can be configured for weak–strong, quasi–strong–strong and strong–strong simulations. A Particle-In-Cell (PIC) solver is also available in Xsuite but was not used in this analysis.

When testing the beam–beam tune shift, the weak–strong model is used. For all other simulations around the iBump Feedback system, strong–strong simulation is used for self-consistency and is deemed a requirement. Firstly, strong–strong simulation is used to find the equilibrium, as no reference parameters are available with this particular lattice configuration. Furthermore, when simulating the beam–beam deflection for IP feedback, both beams must deflect in a self-consistent way. There are 3 plausible options: quasi–strong–strong, strong–strong and PIC solvers. The quasi–strong–strong option was deemed insufficient, as the variation within the bunch would not be well modelled, and may be important for the deflections. The PIC option is the ultimate in performance for beam–beam simulation, but it is accordingly costly; this computational load is unfeasible when performing simulations over thousands of turns to simulate multiple damping times. As such, the natural choice has been to use strong–strong beam–beam simulation. The parameters used for the strong–strong Xsuite simulations are presented in Table 5.3.

Table 5.3: Parameters adopted for the beam–beam simulations used in the SuperKEKB iBump feedback simulation studies.

Parameter	Unit	Value
Beam–beam Mode	/	Strong–Strong
LER Beam Current	mA	1300
HER Beam Current	mA	1000
Number of Bunches	1	2300
Number of Macroparticles	1	100,000
Number of Slices	1	101
Update Frequency	Turn	1
Slicer Mode	/	‘Shatilov’
Beamstrahlung Model	/	Quantum
Bhabha Model	/	None

The choice of 100,000 macroparticles and 101 slices is on the lower end for strong–strong beam–beam simulation with long lepton bunches [129]. The performance of the beam–beam element was therefore tested with a range of macroparticle and slicing

configurations. From testing with the weak–strong beam–beam the beam–beam tune shift was well converged with these simulation parameters. For more subtle beam–beam related effects these parameters may be lower than ideal, however to use more would be prohibitively slow.

5.3.5 Equilibrium with Beam–Beam

By including beam–beam elements in the lattice, the bunches experience the beam–beam force at the IP. In its simplest form, the beam–beam kick at the IP acts similarly to an additional thin quadrupole causing a beam–beam tune shift. This quadrupole is not accounted for in the optics, and as such they are greatly distorted. The phase advances between elements become disrupted and previously periodic structures are distorted such that beta-beating and dispersion leakage are observed. The combination of these effects results in the particles experiencing a very different lattice and as such new equilibrium emittances are found.

To analyse the impact of offsets and feedback fairly, future simulations must be start from an equilibrium including the beam–beam force. This allows the effect of the nominal beam–beam force and the offsets to be decoupled. To identify the new equilibrium parameters particle tracking was performed. Matched Gaussian bunches were initialised as described in Section 5.3.3 with the sextupole offset equilibrium emittances presented in Table 5.1. Beam–beam was included in strong–strong mode with the parameters presented in Table 5.3 and the bunches were tracked for 20,000 turns. The emittance evolution is presented in Fig. 5.19. The resulting equilibrium emittances under synchrotron radiation with strong–strong beam–beam and sextupole offsets are presented in Table 5.4.

Table 5.4: Equilibrium horizontal, vertical and longitudinal emittances for SuperKEKB obtained from particle tracking including quantum synchrotron radiation, optimised sextupoles offsets and strong–strong beam–beam shown for the LER and HER.

Parameter	Unit	LER	HER
ε_x	nm	7.62	7.58
ε_y	pm	53.7	68.2
ε_z	μm	3.40	3.23

In addition to emittances, the equilibrium luminosity was identified. The evolution of the luminosity is presented in Fig. 5.20. The luminosity is seen to decay with turns. This decay was not due to losses, rather a fast luminosity degradation as the beams increase in size reaching new equilibrium emittances.




Figure 5.19: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for SuperKEKB with strong-strong beam-beam and optimised sextupole offsets. Emittances for the LER are shown in red while emittances for the HER are shown in blue. Dashed horizontal lines indicate the relevant initial emittance values.




Figure 5.20: Turn-by-turn evolution of luminosity for SuperKEKB with strong-strong beam-beam and optimised sextupole offsets. Results for the LER are shown in red and results for the HER are shown in blue.

5.3.6 Beam-Beam Tune Shifts

The iBump feedback system is based on beam-beam deflection. In the feedback regime, where offsets are small compared to the beam sizes, this deflection is proportional to the product of the offset and the strength of the beam-beam force. It is therefore important to characterise the strength of the beam-beam tune shift.

The beam–beam tune shifts have been identified in two ways. Firstly, using the analytic formula presented in Section 2.3.3. Secondly, by computing the optical functions including a weak–strong beam–beam element. For the Twiss to remain stable, the beam–beam strength must be linearly rescaled, such that the beam–beam element in the line is one tenth the nominal strength. The resulting beam–beam tune shift is then rescaled up by a corresponding factor of ten. The beam–beam tune shifts for the three identified sets of emittances (bare lattice, including offset sextupoles and with strong–strong beam–beam) are presented in Table 5.5.

Table 5.5: Horizontal and vertical beam–beam tune shifts for the SuperKEKB LER and HER evaluated for different emittance configurations. The emittance configurations correspond to those values presented in Table 5.1 and Table 5.4. The beam–beam simulation settings follow Table 5.3.

Configuration	Mode	$\xi_{x,\text{LER}}$ [10 ⁻³]	$\xi_{y,\text{LER}}$ [10 ⁻³]	$\xi_{x,\text{HER}}$ [10 ⁻³]	$\xi_{y,\text{HER}}$ [10 ⁻³]
Offset Sextupoles +	Analytic	2.814	36.19	1.954	30.09
Synchrotron Radiation	Weak–strong	2.812	36.15	1.565	27.17
Offset Sextupoles +	Analytic	2.794	28.29	1.922	26.30
Synchrotron Radiation + Beam–beam	Weak–strong	2.796	28.73	1.552	24.31

There is good agreement between the analytic and rescaled Weak-Strong Twiss computation. This indicates that the beam–beam elements are correctly installed with appropriate crossing angles, bunch charges and beam sizes.

5.4 iBump Feedback Modelling

5.4.1 iBump Feedback Correction Knobs

During operation, the strengths of the iBump feedback corrector magnets are continuously updated to steer the beam. With the new 32 kHz feedback, calculating the optimal strengths at every feedback cycle from scratch is not feasible. As such, “knobs” are created prior to operation, and from the computed offsets and the pre-computed knobs the corrector strengths are calculated. The knob is a set of corrector strength scaling values that are matched to be linear in the targeted offset. For the iBump feedback system, offset correction knobs are produced using the SAD model of the SuperKEKB HER, based on a series of matching conditions. There are three feedback knobs available for the iBump feedback system: horizontal position, vertical position and vertical angle.

Horizontal Position Knob

The horizontal position knob is used for horizontal position correction at the IP. To create the knob, the following conditions are used in SAD:

1. $x = x_{\text{target}}$, $p_x = 0$, $y = 0$, and $p_y = 0$ at the IP,
2. $x = 0$, $p_x = 0$, $y = 0$, $p_y = 0$, D_y , and $D_{p_y} = 0$ in the arc.

This feedback knob is not typically used in operation.

Vertical Position Knob

The vertical position knob is used for vertical position correction at the IP. To create the knob, the following conditions are used in SAD:

1. $x = 0$, $p_x = 0$, $y = y_{\text{target}}$, and $p_y = 0$ at the IP,
2. $x = 0$, $p_x = 0$, $y = 0$, $p_y = 0$, D_y , and $D_{p_y} = 0$ in the arc.

This feedback knob is typically used in operation.

Vertical Angle Knob

The vertical angle knob is used for vertical angle correction at the IP. To create the knob, the following conditions are used in SAD:

1. $x = 0$, $p_x = 0$, $y = 0$, and $p_y = p_{y,\text{target}}$ at the IP,
2. $x = 0$, $p_x = 0$, $y = 0$, $p_y = 0$, D_y , and $D_{p_y} = 0$ in the arc.

This feedback knob is not typically used in operation.

5.4.2 Xsuite iBump Feedback Knob Implementation

In operation, typically only the vertical position knob is used and this will be the focus of my analysis. The iBump vertical position knob was implemented in Xsuite using the Xtrack “line.match_knob” function. An additional match constraint was added, attempting to minimise D_y and D_{p_y} at the IP. A maximum kick angle per corrector of $100\text{ }\mu\text{rad}$ was enforced. The knob was successfully matched and the shape of the orbit due to the knob is presented in Fig. 5.21. Five orbits are shown, varying linearly from $+37\text{ }\mu\text{m}$ to $-37\text{ }\mu\text{m}$ of offset at the IP. The knob is highly asymmetric in its orbit. This is expected as the IR orbit is in general highly asymmetric.

During matching, the target vertical position at the IP was taken to be $25\text{ }\mu\text{m}$. While this is below the maximum value of $37\text{ }\mu\text{m}$, the match process is different between Xsuite and SAD. In Xsuite, for the knob to be produced all constraints must be strictly met




Figure 5.21: Horizontal and vertical orbit in the interaction region for the iBump Feedback vertical position knob as a function of longitudinal coordinate s . Vertical orbits are shown in red and horizontal orbits are shown in black. The top panel shows the full longitudinal extent of the knob region. The bottom panel shows a zoomed view of the sliced solenoid region around the IP.

within tolerances. In the SAD matching routine, values are minimised, but a solution may not be within expected tolerances. To test the performance of the feedback knob, the linearity of the knob was assessed. The results of this test are presented in Fig. 5.22. The response of the orbit position at the IP varies linearly with the applied iBump vertical offset knob over the range considered. Less than 0.1 nm of horizontal orbit residual and a maximum of 1 nrad of angle residual are observed with the maximum value of the knob. In operation, static orbit bumps are implemented when the maximum value of the knob is approached, and so in operation the working region is typically far from these extents.

In addition to the offset at the IP, implementing such an orbit bump may be detrimental to the machine optics. To test this, the Twiss parameters were evaluated with the knob set to its maximum value in both positive and negative. The results of this test are presented in Fig. 5.23. The beta beating in the lattice is negligible. In operation, the target beta-beating to correct to is on the order of a few percent, far above the beta-beating introduced by the use of the iBump Feedback knob. The additional horizontal dispersion is also negligible, as the horizontal dispersion in the arcs is several orders of magnitude greater. The vertical dispersion is of greater interest, however the dispersion at the IP itself is matched to be below 1 μm . With the current beam parameters this is negligible, however as the optics become more squeezed this




Figure 5.22: Orbit positions and angles at the IP as a function of the iBump Feedback vertical offset knob. The top panel shows the horizontal orbit position (x) and vertical orbit position (y) at the IP. The bottom panel shows the horizontal orbit angle (p_x) and vertical orbit angle (p_y) at the IP. Horizontal quantities are shown in black and vertical quantities are shown in red.




Figure 5.23: Relative beta beating and dispersion differences for two iBump Feedback vertical position knob settings as a function of longitudinal coordinate s . The top panel shows the relative horizontal beta beating ($\Delta\beta_x/\beta_x$), the second panel shows the relative vertical beta beating ($\Delta\beta_y/\beta_y$), the third panel shows the horizontal dispersion difference (ΔD_x), and the bottom panel shows the vertical dispersion difference (ΔD_y). Results for an iBump vertical offset of $+37 \mu\text{m}$ are shown in red, and results for an offset of $-37 \mu\text{m}$ are shown in blue.

will be a greater relative contribution to the beam size.

5.4.3 BPM Implementation

The IP Feedback BPMs are modelled using Xtrack “BeamPositionMonitor” elements. Four BPMs are included, inserted ± 0.5 m from the IP in each of the LER and HER. These monitors return average centroid positions in the transverse plane and counts of the number of macroparticles passing through the monitor within an observation period. The frequency of observations is set at 32 kHz, corresponding to the frequency of the FPGA based fast iBump feedback. To simulate the resolution of the real BPMs, post processing is applied. A resolution of $0.2\text{ }\mu\text{m}$ is used, which is the quoted resolution at nominal beam current of 3.6 A. While this resolution is likely not achieved with lower bunch currents, in the real machine there are many bunches to average over, while for computational reasons only one electron and one positron bunch are simulated here.

5.4.4 Offset Calculation

For a realistic computation of the IP offset from measurable signals, many pieces of information are required. Measurements of the centroid positions at each of the four BPMs are calculated as described in Section 5.4.3. To transport these measurements to the IP, the reference orbit at the BPMs and the transfer matrices between the BPMs and the IP are required. These are computed from closed-ring 6D linear optics computations including synchrotron motion and synchrotron radiation effects in Xsuite. To calculate the canonical kick experienced during the beam–beam collision, the beam–beam tune shifts are required. These are computed analytically from IP beta functions, bunch charges, emittances and the crossing angle. The bunch charges are computed from the BPM counts. Emittances are monitored using an Xcoll “EmittanceMonitor” element, effectively simulating the measurements of the synchrotron radiation monitors in the ring. The beta functions are taken from a computation of the Twiss optical functions with mean radiation. From this information, the offset at the IP is calculated following the method described in Section 2.7.6.

5.4.5 Feedback Implementation

The feedback implementation uses a linear predictive algorithm. A full description of the iBump feedback system is given in Section 4.1.2; only the elements relevant to the simulation are summarised here. The feedback operates at 32 kHz, corresponding to the fast FPGA-based iBump system.

The prediction matrix parameters are set to $N = 6$ and $M = 48$, matching the operational iBump configuration, such that 48 previous cycles are used to compute the

prediction weights and 6 cycles are used for the forward prediction. At each feedback cycle, a new target vertical position is computed from the BPM data, scaled by a gain factor, and cumulatively applied to the existing iBump vertical offset. No magnet ramp times or eddy-current effects are included; changes in magnet strength between feedback cycles are small.

To demonstrate the performance of the prediction algorithm, a simplified test simulation was performed in which the orbit offset was constructed from a linear combination of pure sinusoids. The predictive algorithm was applied directly to this signal, without beam-beam effects, in order to isolate the impact of the feedback gain. The results are presented in Fig. 5.24. A strong dependence on the gain is observed, with high gains required to achieve effective suppression of the residual offset.




Figure 5.24: Orbit correction signals and residual offsets as a function of time for different feedback gains for the iBump Feedback system’s linear prediction algorithm as a function of time. The top panel shows the computed correction signal. The bottom panel shows the residual orbit offset. The reference orbit offset is shown in black. Coloured dashed lines show correction signals obtained with feedback gains of $G = 0.25$ (blue), $G = 0.50$ (orange), $G = 0.75$ (green), and $G = 0.95$ (red).

5.5 Multi-Turn Particle Tracking Simulations

5.5.1 Error Implementation

To model the iBump feedback system in a way that truly represents the operational system and is comparable with measurements, it would be required to realistically model the errors experienced by the machine in operation. In particular, as this system considers the orbit at the IP, this would include magnet powering errors, in particular of the dipoles, and the vibrational errors, most notably of the FFQs. The impact of other orbit correction, such as from the global orbit feedback system would have to be considered. Additionally, errors in the BPMs, synchrotron radiation monitors and more

would be considered. Finally, the simulated lattice would be configured with optics corrections performed realistically with radiation and pre-defined correction knobs. The implementation of all of these errors and corrections is beyond the scope of this thesis.

The impact of some of these may be small in particular simulations. The global orbit correction system is slow, operating on the order of 10 s. The 20,000 turn simulations performed to simulate the iBump feedback system correspond to 0.2 s. As such the global orbit feedback can be considered static on the timescales of these simulations.

To provide an orbit error for the iBump feedback to correct, a dipole corrector is driven sinusoidally in the LER. The corrector, “ZVQR3P”, is in the Fuji straight section of the LER. It is driven with a normalised dipole strength of $k_0 = 1 \times 10^{-5}$. The frequency depends on the test of interest. While not representative of the oscillations of the real system, this allows for systematic testing of the feedback’s response to different error frequencies. It was chosen to oscillate the LER beam as the HER beam is steered to chase the LER beam.

5.5.2 Tracking with Driven Sinusoidal Errors: iBump Feedback Disabled

With driven errors the quality of the beams is expected to degrade. Tracking without feedback is performed to establish a reference for comparison with the cases where feedback is enabled. To illustrate the impact of the driven error a 10 Hz driven error is assessed with the feedback disabled. Beams were initialised with the equilibrium emittances obtained from tracking with strong–strong beam–beam, optimised sextupole offsets and quantum synchrotron radiation presented in Table 5.4.

The evolution of the beam centroids and beam sizes without feedback for a 10 Hz driven error is presented in Fig. 5.25. The longitudinal beam size stays approximately constant and centred on zero. Horizontal beam-size growth is observed in both the LER and the HER beams but centroids remain centred on zero. In the vertical plane, some beam-size growth is observed, but primarily the motion of the centroid is visible, with clear oscillation of the LER beam due to the driven error. From this, the evolution of the IP vertical offset without feedback is found and presented in Fig. 5.26. A maximum offset of ~ 200 nm is observed, on the order of the vertical beam size. A clear sinusoidal form is observed, with greatest noise at the peaks where the beams have greatest separation.

To assess the impact on the performance, the offset is not the only parameter. The other key driver to luminosity degradation is emittance growth, and the evolution of the emittances without feedback is presented in Fig. 5.27. An initial decrease in horizontal emittance is observed before a rise and oscillations in both beams. The vertical emittances for both the LER and HER are observed to grow, and a periodic oscillation is




Figure 5.25: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for SuperKEKB with a single frequency error and no IP feedback. The top panel shows the horizontal centroid position x , the middle panel shows the vertical centroid position y , and the bottom panel shows the longitudinal centroid position z . The Low Energy Ring (LER) is shown in red and the High Energy Ring (HER) is shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.




Figure 5.26: Turn-by-turn evolution of vertical position offset at the IP for SuperKEKB with a single frequency error and no IP feedback.

observed corresponding to the frequency of the driven error. The longitudinal emittances remain consistent with the initial values throughout.

Coupling all of these effects, the evolution of the luminosity without feedback for a 10 Hz driven error is presented in Fig. 5.28. The 10 Hz error is clearly visible in the measured luminosity, and some degradation is observed. Due to the large vertical emittance of the considered case, the sensitivity to the offsets is not so critical as in the




Figure 5.27: Turn-by-turn evolution of beam emittances for SuperKEKB with a single frequency error and no IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . Emittances for the LER are shown in red while emittances for the HER are shown in blue. Dashed horizontal lines indicate the relevant initial emittance values.




Figure 5.28: Turn-by-turn evolution of luminosity for SuperKEKB with a single frequency error and no IP feedback. Luminosity computed for the positron beam is shown in red, while luminosity computed for the electron beam is shown in blue.

design optics of SuperKEKB, so the impact of the offset is reduced.

5.5.3 Tracking with Driven Sinusoidal Errors: iBump Feedback Enabled

The iBump Feedback system is now included in the simulation. To illustrate the impact of the feedback, the same 10 Hz driven error is assessed with a feedback gain of 0.80. Beams were initialised with the equilibrium emittances obtained from tracking with strong-strong beam-beam, optimised sextupole offsets and quantum synchrotron radiation presented in Table 5.4.

The evolution of the beam centroids and beam sizes with feedback for a 10 Hz driven error is presented in Fig. 5.29. The longitudinal beam size stays approximately




Figure 5.29: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for SuperKEKB with a single frequency error and IP feedback. The top panel shows the horizontal centroid position x , the middle panel shows the vertical centroid position y , and the bottom panel shows the longitudinal centroid position z . The Low Energy Ring (LER) is shown in red and the High Energy Ring (HER) is shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.

constant and centred on zero. Horizontal beam-size growth is observed in both the LER and the HER beams but centroids remain centred on zero. In the vertical plane, the motion of the centroids is initially different. The LER beam shows a clear sinusoidal motion throughout. The HER beam, until approximately turn 150, does not move as the feedback has not yet buffered. Following this, the HER beam is driven such that the centroid position matches the centroid of the driven LER beam. From this, the evolution of the IP vertical offset with feedback is found and presented in Fig. 5.30. The offset initially rises to ~ 50 nm prior to feedback engaging, but is then maintained on




Figure 5.30: Turn-by-turn evolution of vertical position offset at the IP for SuperKEKB with a single frequency error and IP feedback.

the order of ~ 10 nm, much below the vertical beam size. The feedback has effectively mitigated the offset due to the driven error.

In comparison with the no feedback case, the evolution of the emittances with feedback is presented in Fig. 5.31. The emittance response of the system with feedback




Figure 5.31: Turn-by-turn evolution of beam emittances for SuperKEKB with a single frequency error and IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . Emittances for the LER are shown in red while emittances for the HER are shown in blue. Dashed horizontal lines indicate the relevant initial emittance values.

is very similar to the case without feedback. The greatest discrepancy is in the form of the vertical emittance for the HER. As the correction at the IP is now applied in the

HER beam-line, the beams are being kicked in order to correct the oscillation. This additional vertical kick causes an increase in the vertical emittance of the HER beam.

The key metric for the performance of this IP feedback is the luminosity performance. The evolution of the luminosity with feedback for a 10 Hz driven error is presented in Fig. 5.32. An initial luminosity degradation is observed as the LER beam is driven




Figure 5.32: Turn-by-turn evolution of luminosity for SuperKEKB with a single frequency error and IP feedback. Luminosity computed for the positron beam is shown in red, while luminosity computed for the electron beam is shown in blue.

and moves away from collision. This luminosity degradation is rapidly slowed when the iBump feedback system is then enabled. While some luminosity is recovered, the system overall shows degraded luminosity due to the emittance growth in both beams from the errors and corrections.

As was presented in Fig. 5.24, there is a strong dependence on the gain, and better residuals might be expected if the gain was increased.

5.5.4 Tracking with Driven Sinusoidal Errors: Frequency Dependence

Following the successful implementation of the iBump feedback system, a scan over a range of frequencies was performed to assess the system's response to different frequencies. Due to the computational expense of this study, only a limited range of frequencies, being 1, 5, 10, 50, 100, 500, and 1000 Hz were assessed.

The first area assessed was the equilibrium RMS IP offset. The variation of the RMS IP Vertical offset with frequency in cases with and without feedback is presented in Fig. 5.33. With feedback disabled, the RMS becomes purely related to the magnitude of the driven error and remains broadly constant across frequency. For low frequencies, this does not strictly hold as insufficient turns were performed to fully capture the form of the oscillation. With feedback enabled, the RMS IP vertical offset is attenuated at




Figure 5.33: RMS vertical offset at the IP as a function of driven error frequency for SuperKEKB. The horizontal axis shows the error frequency in hertz, and the vertical axis shows the RMS vertical offset in metres. Results without IP feedback are shown in blue, and results with IP feedback are shown in orange. The reference case with no applied error is shown by the black dashed line.

all frequencies, but most strongly at low frequencies, pushing towards the floor of the no error case. To more clearly demonstrate the frequency dependence, the transfer function was computed, and the variation of the feedback suppression with frequency is presented in Fig. 5.34. The system is clearly most effective at low frequencies, as




Figure 5.34: Feedback suppression of vertical offset at the IP as a function of driven error frequency for SuperKEKB. The horizontal axis shows the error frequency in hertz. The vertical axis shows the feedback suppression of the vertical offset at the IP in decibels.

expected of the integrating feedback. Results at 1 and 5 Hz may not be fully converged, as only 20,000 turns could be simulated.

Finally, the effect of the feedback on the luminosity is considered. The variation of the luminosity with frequency in cases with and without feedback is presented in Fig. 5.35. The luminosity is observed to be equivalent or better with feedback enabled in all cases. The luminosity gain with feedback is low in this case, as the dominant




Figure 5.35: Luminosity as a function of driven error frequency for SuperKEKB. The horizontal axis shows the error frequency in hertz, and the vertical axis shows the luminosity. Results without IP feedback are shown in blue, and results with IP feedback are shown in orange. The reference case with no applied error is shown by the black dashed line.

source of luminosity loss is the emittance growth due to the driven errors, not the reduction due to IP offsets. For a more squeezed set of parameters, it is expected that the feedback will show greater impacts.

5.6 Summary

Self-consistent 6D tracking simulations of the iBump Feedback system have been developed.

A new model of the SuperKEKB optical lattices has been developed in Xsuite. This involved the development of a general-purpose SAD to Xsuite converter tool, the SAD2XS converter, which is publicly available and has also been used for J-PARC, KEK BT and FCC-ee studies. Due to the strong fringes in the element descriptions of the SuperKEKB lattices, a custom re-matching approach was required and developed. These rematched lattices have then been strictly benchmarked against the SAD model. The lattices are shown to be in strong agreement in linear and non-linear optics cases. Even more promisingly, strong agreement with measured data is observed.

In addition to the optics benchmarks, the behaviour of the lattices with radiation was also tested. The equilibrium emittances obtained from the lattices are in excellent agreement with SAD and in turn in good agreement with measurements from SuperKEKB. The vertical emittance with the bare lattice is too small by comparison to measurements, as in the real machine the vertical emittance largely originates from errors and offsets of elements. To achieve a more realistic vertical emittance, a scheme where sextupole pairs were oppositely offset to create local vertical dispersion bumps around the lattice was utilised. The sextupole offsets were optimised to achieve a

realistic single beam emittance of 40 pm. A strong-strong beam-beam element was then included and new equilibrium emittances including all of the effects of beam-beam, quantum synchrotron radiation and offsets were calculated.

The key elements of the iBump feedback system were also included. Knobs for the iBump Feedback system horizontal position, vertical position and vertical angle were developed in Xsuite. The vertical position knob was the focus of this study as the knob used typically in operation. The linearity of this knob and its impact on beta functions and dispersions were considered. The BPMs, offset calculation and feedback algorithm were also implemented.

Tracking simulations were then performed to identify the performance of the feedback system. First reference simulations were performed with driven errors to identify the impact on the emittances and luminosity with no feedback enabled. It was beyond the scope of this thesis to develop a full realistic driven error model of the accelerator, and as such a single driven corrector was used for these studies. Following this, the iBump feedback system was simulated. The system was observed to work, and the HER beam was observed to track the LER beam and reduce the residual offset and maintain the luminosity. The system performance was assessed by scanning over a range of single-error frequencies, with excellent suppression of the RMS IP offset observed at low frequencies. Luminosity degradation remained due to the emittance growth caused by driving errors. This growth may be pessimistic, due to the residual optics distortion from using offset sextupoles to generate emittance.

Chapter 6

Modelling IP Feedback for FCC-ee

6.1 The FCC-ee Lattice

6.1.1 Lattice Development Status

The optics for FCC-ee have been in development for over a decade. There are currently two competing choices for the optics. An optics comparison process is ongoing as of February 2026. The study aims to compare the two lattice choices on the primary categories of:

1. Design and Parameters.
2. Hardware Requirements and Technical Infrastructure.
3. Performance and Operational Robustness.

A formal comparison report has been produced [130]. The report is designed to inform the recommendation of an independent evaluation panel. Following the panel's input, a decision on the new baseline optics is expected for March 2026.

Throughout the course of these doctoral studies all proposed optical lattices have been in constant development. This has resulted in a highly variable set of optics and beam parameters throughout the course of these studies. Studies were performed with the most up to date optics and beam parameters available at the time of each study. Due to the computational intensity of the studies, studies have not been repeated when newer bunch parameters and optics have become available. Parameters are specified where relevant.

6.1.2 FCC-ee Lattices

The two optics under consideration for FCC-ee are: the Global Hybrid Correction (GHC) Optics [118] and the Local Chromaticity Correction (LCC) Optics [131]. Both

designs are observed to fulfil the performance objectives of the FCC-ee at the Z working point, however both designs are currently observed to have significant limitations at the $t\bar{t}$ working point [130]. The other working points are not addressed in the comparison.

Other optics also exist, either based on the GHC and LCC designs or designed independently. Some examples include:

- Ballistic optics for commissioning [132].
- Monochromatization optics [133].
- Combined function magnet optics [134].

As this thesis considers FCC-ee’s IP feedback system, the nominal tuned lattice configuration will be considered.

The Global Hybrid Correction Lattice (GHC)

The GHC lattice has been designed primarily by K. Oide [118]. The lattice has been in development since prior to the Conceptual Design Report. The design concept of the GHC lattice is based on the use of multi-family non-interleaved sextupole pairs [135], which was applied in both rings of KEKB and later at SuperKEKB. Local chromaticity correction is included around the IPs [136], although only in the vertical plane. The arc cell is a simple FODO design, with $90^\circ/90^\circ$ phase advances for all energies. The lattice has a super-periodicity of four featuring four IPs. A single universal straight section is adopted.

The Local Chromatic Correction Lattice (LCC)

The LCC lattice has been designed primarily by P. Raimondi [131]. The lattice has been in development since 2022. The design concept of the LCC lattice is based on the use of a hybrid FODO cell structure with a sextupole complement to achieve near cancellation of chromatic and geometric aberrations up to fourth order. Local chromaticity correction is included around the IPs in both planes. The optics assumes an identical layout of the magnetic lattice at all operation energies. The lattice has dedicated technical and experimental insertions, with transparency conditions to mitigate the impact on global beam dynamics. The final-focus design also includes additional specific sextupoles (Brinkmann sextupoles [137]), octupoles and decapoles to minimise the non-linearities intrinsic to the section, with strengths determined using multi-objective algorithms to balance achromaticity and anharmonicity of the entire ring.

6.1.3 The FCC-ee Machine Detector Interface (MDI)

The machine detector interface is a key challenge of FCC-ee. The MDI region is common across both lattice models. The IR features a full crossing angle of 30 mrad. The experimental detectors are foreseen to feature strong solenoid magnets. The exact schemes have not yet been determined with both local and non-local compensation schemes under study. No solenoid is included in the current lattice models. A section view of the MDI region of FCC-ee is presented in Fig. 6.1.




Figure 6.1: Section view of the FCC-ee interaction region from the IP to the first final-focus quadrupole (QC1) from [40]. The figure shows the beam pipe, cryostat shell, compensation solenoid, screening solenoid, final-focus quadrupoles, and associated vacuum and cooling components. The longitudinal extent shown is approximately 8.4 m from the IP.

Several key pieces of hardware relevant to IP feedback are situated in this region. Luminosity calorimeter monitors are foreseen to be situated at 1.1 m from the IP (denoted LumiCal in Fig. 6.1), on either side, with a target absolute measurement precision of 1×10^{-4} [138]. A button Beam Position Monitor (BPM) is foreseen to be fitted on the shared elliptical beam pipe, attached to each luminosity calorimeter [139] and therefore strongly mechanically linked with the IP. Further BPMs are foreseen to be located outside of the final focus quadrupole system between the first and second quadrupole (QC1 and QC2). A beamstrahlung (BS) dump is foreseen to be located ~ 500 m downstream of the IP, based on a liquid lead absorber [140]. It is proposed to have a BS monitor along this photon line.

6.2 FCC-ee IP Feedback Signals

6.2.1 IP Feedback Signal Requirements

To correct for IP offsets, the offset must first be calculated. The offset at the IP can be identified by a number of methods, using a variety of measured signals. The following questions are key to determining which approach is preferred:

- What is the sensitivity of this signal to offsets?

- Is this a direct measurement of the offset, or is that a secondary computation?
- Is it possible to perform this measurement with currently available hardware?

Based on the current hardware design for FCC-ee and learning from IP feedback systems at other accelerators, the leading choices for identifying IP offsets are: IP Beam Position Monitor (BPM) signals, luminosity signals, and beamstrahlung (BS) radiation signals [141].

6.2.2 Simulation Setup

To study the sensitivity of the observables to offsets, simulation studies were performed using the Particle-In-Cell (PIC) solver GUINEA-PIG [90]. For this study, the C++ version was used [91]. GUINEA-PIG was chosen as a highly benchmarked and reliable PIC solver. Using a PIC solver offers the most self-consistent approach to simulating the beam-beam interaction during a bunch crossing. Due to the high computational load of using PIC solvers, single bunch crossings rather than multi-turn tracking studies were performed.

Scans of both horizontal and vertical position offset at the IP were performed in the interval of $[0, 1]$ in units of nominal beam sigma for each plane at each working point. For each offset, 50 single bunch crossings were simulated with different random seeds used to generate the input bunches. Quoted errors correspond to the one sigma statistical error calculated across the 50 seeds.

Bunch Parameters

This study was performed using beam parameters from the FCC-ee Feasibility Study Mid Term Review (MTR). These parameters remain similar to the current nominal parameters so the trends are expected to be representative of current parameters. The bunch parameters used for this study are presented in Table 6.1.

Bunch Initialisation

Gaussian bunches may be generated natively in GUINEA-PIG. However, it was chosen to generate the initial particle distributions using python, and to pass them to GUINEA-PIG. As the crab waist collision scheme is critical to the performance of FCC-ee, directly impacting luminosity and the rate of background processes, it was chosen to include it in this study.

Based on the parameters presented in Table 6.1, Gaussian bunches were generated with a random seed. Bunches were generated using equilibrium emittances including beamstrahlung. Following this, an analytic crab waist transform was applied, according

Table 6.1: Bunch parameters used for FCC-ee IP feedback signal studies performed with GUINEA-PIG. Parameter values are taken from the FCC-ee Feasibility Study Mid-Term Review [92].

Working Point	Unit	Z	WW	ZH	$t\bar{t}$
Beam energy	GeV	45.6	80	120	182.5
Bunch population	10^{11}	2.14	1.45	1.15	1.55
Bunches per beam	1	11200	1780	440	60
Horizontal geometric emittance	nm	0.71	2.17	0.71	1.59
Vertical geometric emittance	pm	1.9	2.2	1.4	1.6
Horizontal IP beta	mm	110	220	240	1000
Vertical IP beta	mm	0.7	1	1	1.6
Bunch Length (SR)	mm	5.60	3.47	3.40	1.81
Bunch Length (BS)	mm	15.5	5.41	4.70	2.17
Horizontal Beam–Beam Tune Shift	10^{-3}	2.2	13	10	73
Vertical Beam–Beam Tune Shift	10^{-3}	97.3	128	88	134
Crab waist Ratio	%	70	55	50	40
Luminosity /IP	$10^{34} \text{ cm}^{-2} \text{ s}^{-1}$	141	20	5.0	1.25

to the Hamiltonian:

$$H_{CW} = kxy'^2, \quad (6.1)$$

where x and y' are horizontal position and vertical slope at the IP, and k is the strength of the crab waist. Further details are presented in Section 2.2.3. This process is performed for the electron and positron bunch generation independently.

While it would be preferred to use an equilibrium tracked beam, it was not feasible to perform such a simulation given the computational load of the PIC solver and the long damping times of FCC-ee.

GUINEA-PIG Configuration

The simulation parameters used for the GUINEA-PIG simulations are presented in Table 6.2. The grid requirements at the $t\bar{t}$ working point are reduced compared to the other working points due to the much lower Piwinski angle.

The choice of 100,000 macroparticles was a compromise. To perform simulations for all working points of FCC-ee, across a range of both vertical and horizontal position offsets, with a sufficient quantity of repetitions to obtain statistical errors, it was preferred to reduce the number of macroparticles used as much as possible. On the other hand, for the performance of the PIC computations, avoiding errors due to low populations in the tails, a large number of macroparticles were required. Furthermore, the quality of the beamstrahlung radiation observables was identified as being highly sensitive to the number of macroparticles. Beamstrahlung spectra for the four nominal working points of FCC-ee in nominal collision computed with GUINEA-PIG are presented in

Table 6.2: Numerical grid and macroparticle parameters used in the GUINEA-PIG simulations for FCC-ee.

Parameter	Z, WW, ZH	$t\bar{t}$	Description
cut_x	$150\sigma_x$	$12\sigma_x$	Horizontal transverse grid cut
cut_y	$60\sigma_y$	$10\sigma_y$	Vertical transverse grid cut
cut_z	$3\sigma_z$	$10\sigma_z$	Longitudinal grid cut
n_x	256	128	Number of horizontal grid cells
n_y	256	128	Number of vertical grid cells
n_z	512	256	Number of longitudinal grid cells
n_t	1	1	Number of time steps per cell
n_m	100,000	100,000	Number of macroparticles

Fig. 6.2. The spectra show a log-concave form spanning several orders of magnitude. To sufficiently sample this spectrum, a large number of macroparticles was required.




Figure 6.2: Photon energy spectra of beamstrahlung radiation for the four nominal FCC-ee working points. The horizontal axis shows the photon energy and the vertical axis shows the photon count. Spectra corresponding to the Z, WW, ZH, and $t\bar{t}$ working points are shown. Vertical dashed lines indicate the corresponding average photon energies for each working point.

6.2.3 Luminosity

The variation of luminosity with both horizontal and vertical position offset was studied. GUINEA-PIG features multiple definitions of luminosity and for this analysis the ‘lumi_ee’ variable was studied. A comparison of the luminosities computed in GUINEA-PIG with nominal collision and the luminosities in the parameter table is presented in Table 6.3. Excellent agreement is observed for the Z, WW and ZH working points, but a large discrepancy is observed for the $t\bar{t}$ working point. No issue was detected in the input file. The discrepancy is between weak-strong and PIC estimates and may be due to the pinch effect.

Table 6.3: Nominal and GUINEA-PIG-computed collision luminosities per IP for the four FCC-ee working points. GUINEA-PIG luminosities are evaluated for nominal collisions. Quoted uncertainties correspond to statistical errors from 100 simulation repeats (50 horizontal and 50 vertical position offset scans).

Luminosity per IP [$10^{34} \text{ cm}^{-2} \text{ s}^{-1}$]	Z	WW	ZH	$t\bar{t}$
FCC-ee MTR	141	20	5.0	1.25
GUINEA-PIG Nominal	141.49 ± 0.09	20.47 ± 0.00	5.04 ± 0.00	1.48 ± 0.00

Variation with Offset

The variation of relative luminosity with horizontal position offset is presented in Fig. 6.3, while the variation with vertical position offset is presented in Fig. 6.4. The use of relative luminosity and offset in units of beam sigma is chosen for comparability between the different working points.




Figure 6.3: Relative luminosity as a function of horizontal IP position offset. The horizontal axis shows the horizontal position offset normalised to the horizontal beam size, and the vertical axis shows the relative luminosity. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

Clear degradation of luminosity is observed with horizontal position offset at the IP. The sensitivity of luminosity to horizontal position offset is strongly correlated with the Piwinski angle. At the Z working point ($\phi_{pw} \sim 26$), the collision occurs in a large- ϕ_{pw} regime. In this case the effective interaction region is longitudinally extended and luminosity is accumulated from many longitudinal slices. This longitudinal extension results in a reduced sensitivity to transverse offsets. As the Piwinski angle decreases through the WW and ZH working points ($\phi_{pw} \sim 3.7$ and $\phi_{pw} \sim 5.4$ respectively), the overlap region becomes progressively more compact and the luminosity degrades more rapidly with offset. At the $t\bar{t}$ working point ($\phi_{pw} \sim 0.8$), the collision is effectively near head-on. In this case the horizontal position offsets coherently suppress the overlap, leading to the steepest luminosity loss.




Figure 6.4: Relative luminosity as a function of vertical IP position offset. The horizontal axis shows the vertical position offset normalised to the vertical beam size, and the vertical axis shows the relative luminosity. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

Clear degradation of luminosity is also observed with vertical position offset at the IP. The Z, ZH and $t\bar{t}$ working points are highly consistent in the form of their luminosity degradation. The WW working point exhibits a reduced impact of the vertical position offsets. The sensitivity of luminosity to vertical position offset is primarily governed by the effective vertical overlap. This is modified by the hourglass effect and the impact of non-linear beam-dynamics. As such, there is no single parameter to identify why the WW working point shows greater offset tolerances.

To summarise the sensitivity to offset for each of the four working points, the offsets corresponding to a particular luminosity reduction were identified. These offsets were computed using an order 4 polynomial fit to the data. The offsets are presented in Table 6.4. The horizontal position offsets required for a 1% reduction in luminosity

Table 6.4: Horizontal and vertical beam position offsets required to reduce the luminosity to 99% and 95% of its nominal value for the four FCC-ee working points.

Offset Plane	Luminosity Factor	Z	WW	ZH	$t\bar{t}$
Horizontal Position Offset μm	99%	2.18	3.83	2.90	4.72
	95%	4.97	8.70	6.73	10.3
Vertical Position Offset nm	99%	4.96	7.62	5.03	6.97
	95%	11.2	17.3	11.4	16.0

are all on the micron scale and as such are not deemed a major concern. The vertical position offsets are much smaller. These are the tolerances for a single bunch crossing however, and with multiple IP crossings and emittance blow-up the true tolerances over multiple turns are likely tighter.

Viability for FCC-ee Feedback

Luminosity measurements will be vital to the operation of FCC-ee. Multiple measurements of luminosity are foreseen at each IP, through both the detector luminosity calorimeters, and downstream fast luminosity calorimeters. It is foreseen that the detector luminosity calorimeter will be highly accurate, potentially with an integration time long with respect to the feedback timescale. The downstream luminosity monitors will likely be faster but less accurate. The fast luminosity monitors are foreseen to measure radiative Bhabha losses which could allow for bunch-by-bunch relative luminosity measurements. Such monitors are already successfully operated at SuperKEKB [142]. As such, the required hardware is already available and expected to be performant.

The major challenge of using luminosity as a method of offset identification is that luminosity is a scalar signal. As such, there is no direct information on which direction to drive the beam for correction. It is possible to overcome this limitation but adopting a dithering strategy, however this approach entails sacrificing luminosity by driving the beam.

6.2.4 Beam Deflection

The variation of outgoing beam angle with both horizontal and vertical position offset was studied. From GUINEA-PIG outputs, the outgoing angles for both beams are available. As offsets are made symmetrically, for this analysis the mean outgoing deflection angles were studied.

Variation with Offset

The variation of the horizontal and vertical outgoing beam angles with horizontal position offset is presented in Fig. 6.5, while the variation with vertical position offset is presented in Fig. 6.6.

The horizontal deflection angle is observed to increase with horizontal IP position offset. In the region of small offsets the deflection can be well approximated as linear. As this plot is scaled by the nominal beam size, the axis is weighted by the square root of the beta function, but also by the emittance which affects the gradient observed. It is therefore not observed that the ordering directly follows the horizontal beam-beam tune shift, however in this case the gradient does follow the ordering presented in Table 6.1. With only horizontal position offset the vertical deflection remains consistent with zero within statistical uncertainty.

The vertical deflection angle is observed to increase with vertical IP position offset. In the region of small offsets the deflection can be well approximated as linear. In the linearised beam-beam deflection, the outgoing beam angle is proportional to the offset and the beam-beam tune shift, and inversely proportional to the IP beta function. As




Figure 6.5: Outgoing beam angles as a function of horizontal IP position offset normalised to the horizontal beam size. The upper panel shows the horizontal outgoing angle and the lower panel shows the vertical outgoing angle. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

this plot is scaled by the nominal beam size, the axis is weighted by the square root of the beta function, but also by the emittance which affects the gradient observed. It is therefore not observed that the ordering directly follows the vertical beam–beam tune shift. With only vertical position offset the horizontal deflection remains consistent with zero within statistical uncertainty.

Viability for FCC-ee Feedback

The preferred BPMs for IP feedback are foreseen to be located approximately 1.1 m from the IP, within the solenoid field prior to the first FFQ. For successful feedback it is targeted to have offsets on the order of 2% of the RMS beam size. Such offsets, on the order of 1 nm, correspond to deflection angles on the order of half a micro-radian. Excluding the effect of the solenoidal field, the position offset at a BPM 1.1 m downstream would be less than one micron. This implies a requirement for sub-micron resolution for these BPMs. This is a strong technical challenge when also taking into account impedance constraints. Using BPMs further downstream relaxes the resolution requirement, but introduces errors from the offsets of the FFQs in the transfer matrix for the offset computation. Position information may be improved by averaging across multiple bunches.

As a vector signal, the deflection angle is strongly correlated with the offsets and is




Figure 6.6: Outgoing beam angles as a function of vertical IP position offset normalised to the vertical beam size. The upper panel shows the horizontal outgoing angle and the lower panel shows the vertical outgoing angle. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

a highly promising signal. To use the four BPM readings to compute the offset, a good model of the beam motion in the IR is required, including the model of the solenoid field.

6.2.5 Beamstrahlung Signals

The variation of beamstrahlung radiation power and angle with both horizontal and vertical position offset was studied. For this study, beamstrahlung photons with “positive energies” (forward firing) were analysed.

Variation with Offset

The variation of relative beamstrahlung radiation power with horizontal position offset is presented in Fig. 6.7, while the variation with vertical position offset is presented in Fig. 6.8. The variation of horizontal and vertical beamstrahlung radiation angle with horizontal position offset is presented in Fig. 6.9, while the variation with vertical position offset is presented in Fig. 6.10. The use of relative beamstrahlung power and offset in units of beam sigma is chosen for comparability between the different working points.

The beamstrahlung power remains consistent with the nominal value for all horizontal IP position offsets at the WW and ZH working points. At the Z working point a small




Figure 6.7: Relative beamstrahlung power as a function of horizontal IP position offset. The horizontal axis shows the horizontal position offset normalised to the horizontal beam size, and the vertical axis shows the relative beamstrahlung power. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

increase is observed. In contrast, at the $t\bar{t}$ working point the beamstrahlung power decreases strongly with horizontal position offset. This is consistent with its much larger horizontal beam–beam tune shift. In this case the high-field core interactions are rapidly suppressed as the beams are transversely separated. While this broadly follows the same overlap/Piwinski-angle logic as the degradation of the luminosity, beamstrahlung is more strongly weighted to the high-density core of the collision. In the strong beam–beam regime it therefore exhibits a much steeper reduction with offset than luminosity.




Figure 6.8: Relative beamstrahlung power as a function of vertical IP position offset. The horizontal axis shows the vertical position offset normalised to the vertical beam size, and the vertical axis shows the relative beamstrahlung power. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

For vertical IP position offsets, the beamstrahlung power increases with separation for

all working points. In the flat-beam collision scheme of FCC-ee, a vertical displacement moves particles into regions of stronger transverse electric field of the opposing beam. The stronger transverse field enhances the radiation even as the luminosity is reduced.




Figure 6.9: Outgoing beamstrahlung angles as a function of horizontal IP position offset normalised to the horizontal beam size. The upper panel shows the horizontal outgoing angle and the lower panel shows the vertical outgoing angle. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

In the horizontal plane, the crab-crossing collision scheme introduces an energy dependence in the nominal emission angle. The observed non-zero mean horizontal beamstrahlung photon angles with zero transverse offset reflects the coherent beam–beam deflection at the IP. Beamstrahlung photons are emitted approximately along the instantaneous direction of the emitting particles, and the photon angular centroid therefore follows the mean transverse beam angle during the collision.

The dependence on working point is consistent with the differing beam–beam tune shifts and the associated variations in energy loss and beam rigidity. At nominal collision, the mean horizontal photon emission angles are 36.9, 18.8, 14.2, and 5.1 μrad for the Z, WW, ZH, and $t\bar{t}$ working points, respectively.

Additional horizontal deflection is observed with increasing horizontal position offset. In contrast, for purely horizontal position offsets, the vertical photon emission angle remains consistent with zero within statistical uncertainty.

With vertical position offset, strong variation of the beamstrahlung vertical angle is observed. With only vertical position offset the horizontal deflection remains consistent with the nominal horizontal angles. It is noted that the deflection angles of the




Figure 6.10: Outgoing beamstrahlung angles as a function of vertical IP position offset normalised to the vertical beam size. The upper panel shows the horizontal outgoing angle and the lower panel shows the vertical outgoing angle. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

beamstrahlung radiation are reduced by $\sim 30\%$ compared to the deflection angles of the outgoing beam centroids presented in Fig. 6.6.

Viability for FCC-ee Feedback

The use of beamstrahlung radiation is an enticing option. Both power and deflection angle are measurables that are well correlated with the IP offset. While power is a scalar signal, the deflection is a vector signal. Furthermore, the signal is a direct measure of the opposing beam, not a convolution of the separation of both beams.

The measurement of the beamstrahlung radiation is a major technical challenge. The FCC-ee beam instrumentation team are actively researching the topic. The major issue is the intensity of the radiation power. A summary of the radiation power per IP, per direction is presented in Table 6.5. Especially at the Z working point, this power may make beamstrahlung feedback measurements impossible. Furthermore, the beamstrahlung measurement is foreseen near the beamstrahlung dump at approximately 500 m downstream from the IP. This distance makes detection of the angular signals easier, as there is a long propagation distance for these small angle deviations to manifest over. The distance however incurs additional signal transmission time which may be impactful for a fast feedback system. This signal is ultimately the most challenging and cannot be assumed to be feasible.

Table 6.5: GUINEA-PIG-computed beamstrahlung radiation power for the four FCC-ee working points. Results are shown for nominal collisions and for one beam-sigma offsets in the horizontal and vertical planes. Quoted uncertainties correspond to statistical errors from 50 simulation repeats.

Beamstrahlung Power [kW]	Z	WW	ZH	$t\bar{t}$
Nominal collision	229 ± 5	95.5 ± 1.8	63.6 ± 1.1	40.2 ± 0.7
1σ x-offset	239 ± 6	94 ± 2	63.4 ± 1.2	29.0 ± 0.6
1σ y-offset	299 ± 7	120.2 ± 1.8	84.5 ± 1.8	54.3 ± 0.8

6.3 FCC-ee IP Feedback Tolerances

6.3.1 Contributing Factors

Machine Stability

FCC-ee is a double-ring storage ring collider, in which beams are stored for extended periods. Consequently, stable beam–beam collisions are a fundamental requirement for operation. With offsets at the IP a number of detrimental effects occur. Due to the beam–beam force, offsets result in beam–beam deflection which may cause further orbit errors. This can in turn cause tune errors due to the strong sextupoles at FCC-ee. IP offsets also result in increased exposure to the strong transverse electric fields of the opposing beam which increases emitted beamstrahlung radiation power and results in further bunch lengthening. If this process occurs asymmetrically a 3D flip–flop instability may occur [143]. The combination of these many effects results in the excitation of resonances and a corresponding increase of tune diffusion. This in turn results in emittance growth leading to degradation of luminosity and particle losses.

This is not just an issue for the loss of luminosity, but losses are a major issue for the collimation of a machine such as FCC-ee, as well as the background rate the detectors are exposed to. With the high stored beam energies of FCC-ee, this is a particular challenge.

Physics Performance

FCC-ee is designed to be a precision physics machine. To achieve this, systematic errors must be minimised. One key systematic error is the measurement of the beam energy. The beam energy measurement precision is influenced by offsets at the IP. For beams with a vertical dispersion difference of $1\text{ }\mu\text{m}$ and a vertical position offset of 1 nm , the average collision energy shifts by 100 keV , which is the target energy precision [144]. Thus, the maximum vertical position offset should not exceed 1 nm . Offsets also directly degrade luminosity, reducing the volume of data for physics analysis.

Error Sources

FCC-ee will be subject to a wide range of errors. The primary sources of error leading to offsets at the IP are ground motion, mechanical vibration and power converter noise. To achieve the small beta functions at the IP, the FFQs at FCC-ee are the strongest magnets in the collider. For small offsets, the quadrupole therefore applies a strong dipolar kick. Previous studies have indicated that the motion of the Final Focus Quadrupoles will be a dominant error source [93]. Studies at SuperKEKB support these conclusions [145].

The dynamics of the FCC-ee MDI region, including the final focus system, are currently under study. The amplitude and frequency of the FFQ resonant modes are expected to dictate the required frequency of the IP feedback, and the maximum strengths of the correctors. This information is not available at the time of writing.

Interplay with Other Feedback Systems

The IP feedback system will not exist in isolation. In particular, the interplay with the global Closed Orbit Feedback (COFB) is a vital consideration. It is expected that large, low frequency oscillations will be addressed by the global feedback over the full ring, reducing the required closed orbit bumps at the IP. Furthermore, feedback systems such as the transverse damper will have interplay with regards to the bunch length. The design of these other feedback systems is also currently under study [36].

The current estimates for the performance of the COFB are that suppression remains good on average at 2 Hz, with a maximum of 10 Hz based on simulations by S. Kostoglou (private communication, January 2026).

6.3.2 Studying Machine Tolerance

While the tightest IP offset tolerance is expected to be the 1 nm tolerance required for the beam energy measurement, the machine tolerances remain an accelerator issue. It must first be ensured that the beam can remain stable in the presence of vibrations before pushing to achieve the requirements of other groups. To study machine stability with IP offsets tracking simulations were performed. Due to the nano-beam optics at the IP and the results observed in Section 6.2, the most impactful errors on the beam are expected to be in the vertical plane: vertical position offset and vertical momentum offset. As such, simulations have been focussed on these tolerances and are published in [146].

6.3.3 Simulation Setup

Multi-turn tracking simulations were performed using Xsuite. These simulations were performed using the GHC lattice version 24.3. Simulations were performed using simplified lattices. These lattices were composed of:

1. A weak–strong beam–beam element.
2. A transfer map IP to outgoing Crab Sextupole.
3. A thin Crab Sextupole.
4. A transfer map from outgoing Crab Sextupole to incoming Crab Sextupole.
5. A thin Crab Sextupole.
6. A transfer map from incoming Crab Sextupole to IP.
7. A longitudinal aperture.
8. A transverse aperture.

The longitudinal aperture is used to model the momentum acceptance of the lattice while the transverse aperture is used to model the dynamic aperture of the lattice. Effective synchrotron radiation damping and noise is applied in the ‘arc’ transfer map (outgoing crab sextupole to incoming crab sextupole). This simplified lattice represents one super-period (IP to IP).

The GHC lattice uses an imperfect cancellation of the Local Chromatic Correction sextupole pair close to the IP to produce the crab waist effect. In this model, a single thin multipole each side of the IP is used to represent the crab waist. These multipoles have phase advances from the IP of $(1.0, 1.25) [2\pi]$ in the horizontal and vertical plane respectively. The integrated sextupolar strengths, $\pm k_2 l_{CW}$, are calculated according to

$$k_2 l_{CW} = k_{CW} \frac{1}{2\theta_c \beta_y^* \beta_y^{CW}} \sqrt{\frac{\beta_x^*}{\beta_x^{CW}}}, \quad (6.2)$$

where θ_c is the half crossing angle and β is the Twiss beta function, taken from the full lattice description.

Beam–beam simulations are performed in ‘weak–strong’ mode. The ‘strong’ beam is held offset and the distribution of the ‘weak’ beam is then analysed. The strong beam is identically offset at each IP due to the choice to only model one super-period. Monitoring occurs after the outgoing crab sextupole to remove the impact of the non-linear crab waist transform.

The choice of weak–strong simulation was made to facilitate tracking of thousands of particle turns. While weak–strong beam–beam simulations are not strictly self-consistent, these studies give a good indication for a fraction of the computational time of a strong–strong simulation.

Scans of both vertical position and vertical angle offset at the IP were performed in the interval of $[0, 1]$ nominal beam sigma for each offset at each working point.

Bunch Parameters

This study was performed using beam parameters for the GHC lattice version 24.3. These parameters remain similar to the current nominal parameters so the trends are expected to be representative of current parameters. The bunch parameters used for this study are presented in Table 6.6.

Table 6.6: Bunch parameters used for FCC-ee IP feedback tolerance studies performed with Xsuite. Parameter values are taken from the GHC v24.3 lattice.

Working Point	Unit	Z	WW	ZH	$t\bar{t}$
Beam energy	GeV	45.6	80	120	182.5
Bunch population	10^{11}	2.16	1.38	1.69	1.48
Bunches per beam	1	11200	1852	300	64
Horizontal geometric emittance	nm	0.70	2.16	0.66	1.51
Vertical geometric emittance	pm	1.9	2.0	1.0	1.36
Horizontal IP beta	mm	110	220	240	900
Vertical IP beta	mm	0.7	1	1	1.4
Bunch Length (SR)	mm	5.57	3.46	3.26	1.91
Bunch Length (BS)	mm	15.6	5.28	5.59	2.32
Horizontal Beam–Beam Tune Shift	10^{-3}	2.2	13	10.8	65
Vertical Beam–Beam Tune Shift	10^{-3}	97.7	129	130	136
Crab waist Ratio	%	70	55	50	40
Luminosity /IP	$10^{34} \text{ cm}^{-2} \text{ s}^{-1}$	143	20	7.5	1.38

Gaussian Beam Tracking Studies

To study the dynamics of realistic bunches, bunches were initialised as Gaussians with equilibrium collision parameters including beamstrahlung radiation. Luminosity, quantum synchrotron radiation, quantum beamstrahlung radiation and quantum Bhabha scattering were enabled.

Particles were tracked for ~ 8 longitudinal damping times, taken to be 1171, 218, 65.4 and 19.4 turns for each of the Z, WW, ZH and $t\bar{t}$ working points respectively. At the WW, ZH and $t\bar{t}$ working points 10,000 macroparticles were tracked. Due to the long damping time at the Z working point, only 1,000 macroparticles were tracked. This reduced number of macroparticles resulted in greater noise at the Z working point.

Frequency Map Analysis Studies

To study the resonances experienced by particles with different action amplitudes grids of particles were tracked. Particle grids of shape $[161 \times 161]$ were initialised for square root action amplitude $\sqrt{2J}$ in the interval 0 to $8\sqrt{\varepsilon_{RMS}}$ across four test phases $[0, \pi/4, \pi/2, 3\pi/4]$. Luminosity, synchrotron radiation, beamstrahlung radiation and Bhabha scattering were disabled. This allows analysis of the impact of the beam offset, without stochastic processes obscuring the results. Particles were tracked for 1024 turns.

6.3.4 Previous Studies and Benchmarking

Previous studies were performed using a linearised Conceptual Design Report lattice at the Z working point. The study analysed a linearised lattice with 4 IPs but the same phase advance per super-period as the CDR baseline 2IP lattice. This lattice was not optimised for operation with four IPs. These simulation studies were performed using the tracking code Lifetrac [147]. A vertical position offset tolerance of $0.05\sigma_y$ was identified [148]. While this tolerance is greater than the 1 nm tolerance for energy calibration, this tolerance is of considerable concern as it implies the beam becomes unstable with only small offsets.

The previous Lifetrac studies were reproduced in Xsuite. For comparison, the study of analysing the action space density of a tracked beam was performed in weak-strong mode, where the strong beam had a $0.2\sigma_y$ vertical position offset. Contours of the resulting normalised action space density from both Lifetrac and Xsuite are presented in Fig. 6.11. Good agreement is observed between the two codes, in particular for the action space density distribution of the beam core. Some discrepancies are observed on the halos of the beam. Some of the visual discrepancy may be due to the plotting and binning process.

6.3.5 Tracking to Equilibrium

To study equilibrium bunch parameters and distributions, the bunches must be tracked for many damping times to equilibrate. Bunches were tracked for ~ 8 damping times so it is expected that a new equilibrium is reached. With only 1,000 – 10,000 macroparticles, the emittances, luminosity and beamstrahlung radiation power are noisy on a turn-by-turn basis. The mean and standard deviation over many turns is therefore used. The final 7500, 1500, 375 and 150 turns are analysed for each of the Z, WW, ZH and $t\bar{t}$ working points respectively. An example showing how the equilibrium parameters are extracted from tracking data is presented in Fig. 6.12.

For this study, all bunches are initialised with the parameters in the parameter table. From Fig. 6.12, in the nominal collision case while the horizontal and longitudinal




Figure 6.11: Action-space density contour plots for the FCC-ee Z working point. Results from Lifetrac are shown in the left panel and results from Xsuite are shown in the right panel. The horizontal axis shows the horizontal action coordinate ($\sqrt{2J_x}$) and the vertical axis shows the vertical action coordinate ($\sqrt{2J_y}$).

emittances are in agreement with the nominal values, the vertical emittance of the parameter table is not reached due to the use of the simplified lattice rather than the




Figure 6.12: Turn-by-turn normalised beam-size evolution in the horizontal (top panel), vertical (middle panel), and longitudinal (bottom panel) planes. Solid blue curves show results with zero IP position offset, and solid orange curves show results with a vertical IP position offset of $1.0\sigma_y$. Dashed blue and orange horizontal lines indicate the corresponding mean equilibrium beam sizes. The vertical dashed black line indicates the turn number from which equilibrium is assumed.

full lattice. The equilibrium emittances quoted for the GHC lattices are a result of tracking with a vertical emittance wiggler in the full FCC-ee lattice in SAD. Without the non-linearities of the full lattice, a reduced vertical emittance is observed. For comparison, the reference values will be the equilibrium values identified by tracking in the no offset configuration.

6.3.6 Emittance Variation with Offset

The variation of horizontal, vertical and longitudinal relative emittance with vertical position offset is presented in Fig. 6.13, while the variation with vertical momentum offset is presented in Fig. 6.14. The use of relative emittance and offset in units of beam sigma is chosen for comparability between the different working points.




Figure 6.13: Relative emittance growth as a function of vertical IP position offset normalised to the vertical beam size. The top panel shows the horizontal emittance growth, the middle panel shows the vertical emittance growth, and the bottom panel shows the longitudinal emittance growth. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

For the WW, ZH and $t\bar{t}$ working points at all vertical position offsets the horizontal emittance appears consistent with the nominal case within errors. For the Z working point, the errors are significantly greater, but consistency remains. Vertical emittance is observed to increase with vertical position offset. This is most clear for the $t\bar{t}$ working point, which has the largest vertical beam-beam parameter. Longitudinal emittance is also observed to increase with vertical position offset. With vertical position offset, the bunches are exposed to higher transverse fields and so emit more

beamstrahlung radiation. This increases the energy spread of the bunch and results in a larger longitudinal emittance. This is most impactful at the Z working point, which has the lowest beam energy, and the bunch length is therefore most sensitive to the increase in beamstrahlung radiation.




Figure 6.14: Relative emittance growth as a function of vertical IP momentum offset normalised to the vertical momentum spread. The top panel shows the horizontal emittance growth, the middle panel shows the vertical emittance growth, and the bottom panel shows the longitudinal emittance growth. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown. Error bars indicate the statistical uncertainty across 50 random seeds.

For all working points at all vertical momentum offsets the horizontal emittance appears consistent with the nominal case. Vertical emittance is observed to increase with vertical momentum offset. This effect is strongest for the WW working point. Furthermore, the effect of vertical momentum offset appears to be stronger than the effect of vertical position offset. Longitudinal emittance is also observed to increase with vertical momentum offset. This effect is noticeably weaker for the Z working point compared to the vertical position offset case.

6.3.7 Luminosity Variation with Offset

The variation of relative luminosity with vertical position offset is presented in Fig. 6.15, while the variation with vertical momentum offset is presented in Fig. 6.16. The use of relative luminosity and offset in units of beam sigma is chosen for comparability between the different working points.

Significant luminosity degradation is observed with vertical position offset. The




Figure 6.15: Relative luminosity as a function of vertical IP position offset. The horizontal axis shows the vertical position offset normalised to the vertical beam size, and the vertical axis shows the relative luminosity. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown, as indicated in the legend. Error bars indicate the statistical uncertainty across 50 random seeds.

form of the degradation is similar for all working points, especially for small offsets. For large offsets, the impact is worst at the Z working point. This aligns with the emittance observation presented in Fig. 6.13, where strong longitudinal emittance blow-up was observed with large offsets at the Z working point. The degradation here seems reduced with respect to Fig. 6.4, likely due to the combination of different parameters and a different beam-beam model.




Figure 6.16: Relative luminosity as a function of vertical IP momentum offset. The horizontal axis shows the vertical momentum offset normalised to the vertical momentum spread, and the vertical axis shows the relative luminosity. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown, as indicated in the legend. Error bars indicate the statistical uncertainty across 50 random seeds.

Significant luminosity degradation is observed with vertical momentum offset. The form of the degradation is different across all working points, especially at large offsets. For large offsets, the impact is worst at the WW working point. This aligns with the

emittance observation presented in Fig. 6.14, where strong vertical emittance blow-up was observed with large offsets at the WW working point.

6.3.8 Beamstrahlung Variation with Offset

To successfully model beamstrahlung radiation a large number of beamstrahlung photons must be simulated to converge to the log-concave spectrum presented in Fig. 6.2. With the ‘low’ number of macroparticles used here, especially for the Z working point, there was significant turn-by-turn variation in the beamstrahlung emission. As each turn only samples part of the beamstrahlung spectrum, this is considered a numerical issue. These significant variations in beamstrahlung power between turns results in extremely large errors on the observed equilibrium beamstrahlung power. Despite the resulting lack of statistical significance, it is possible to observe the overall trends and analyse the underlying physics processes.

The variation of relative beamstrahlung power with vertical position offset is presented in Fig. 6.17, while the variation with vertical momentum offset is presented in Fig. 6.18. The use of relative beamstrahlung power and offset in units of beam sigma is chosen for comparability between the different working points.




Figure 6.17: Relative beamstrahlung power as a function of vertical IP position offset. The horizontal axis shows the vertical position offset normalised to the vertical beam size, and the vertical axis shows the relative beamstrahlung power. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown, as indicated in the legend. Error bars indicate the statistical uncertainty across 50 random seeds.

Relative beamstrahlung power is observed to increase with vertical position offset. The errors at the Z working point greatly obscure any observations as a result of the reduced number of macroparticles used in simulation. With vertical position offset, the bunches are exposed to higher transverse fields and so emit more beamstrahlung radiation. The effect is most significant at the $t\bar{t}$ working point. It would naively be expected that this would result in the largest longitudinal emittance increase, however

the increase in energy spread also depends on the beam energy. As such, the longitudinal emittance of the Z working point is observed to increase most as presented in Fig. 6.13.




Figure 6.18: Relative beamstrahlung power as a function of vertical IP momentum offset. The horizontal axis shows the vertical momentum offset normalised to the vertical momentum spread, and the vertical axis shows the relative beamstrahlung power. Results for the Z, WW, ZH, and $t\bar{t}$ working points are shown, as indicated in the legend. Error bars indicate the statistical uncertainty across 50 random seeds.

Relative beamstrahlung power is also observed to increase with vertical momentum offset. The errors at the Z working point once again greatly obscure any observations as a result of the reduced number of macroparticles used in simulation.

6.3.9 Action Space Density Variation with Offset

To further analyse the observations of Section 6.3.6, Section 6.3.7 and Section 6.3.8, the bunch distributions were analysed in action space. At each turn of tracking, particles were binned based on their action amplitude, computed as:

$$J_a = \frac{a^2 + p_a^2}{2}, \quad a \in x, y, \quad (6.3)$$

where a and p_a are the position and momentum coordinate in a given plane. Analysing the action space distributions of the beams enables insight into the cause of the observed emittance changes: to distinguish between emittance growth due to distortion of the beam core or significant population of the beam halo. The first ~ 2.5 longitudinal damping times are rejected as the bunch equilibrates.

A comparison of the action space density distributions for the previously studied CDR 4IP lattice at the Z working point and the GHC version 24.3 lattice at the Z working point with vertical position offsets is presented in Fig. 6.19, while the variation with vertical momentum offset is presented in Fig. 6.20. This comparison was chosen to understand if the previously identified strict requirements (presented in Section 6.3.4)

remain with more current versions of the lattice.




Figure 6.19: Action-space density distributions in the transverse action space ($\sqrt{2J_x}$, $\sqrt{2J_y}$) for increasing vertical IP position offsets. Columns correspond to nominal collision, a $0.05\sigma_y$ offset, and a $0.20\sigma_y$ offset, respectively. The top row shows results for the FCC-ee CDR lattice with four IPs, while the bottom row shows results for the FCC-ee GHC v24.3 lattice. Colour contours represent the logarithm of the normalised particle density.

For the CDR 4IP lattice, it is clear that the beam is strongly distorted in action space with vertical position offset at the IP. The beam core becomes highly stretched and depopulated. This leads to significant luminosity reductions. For the GHC v24.3 lattice, this strong beam core depopulation is not observed, and only minor changes are observed to the halo of the beam. This implies the GHC v24.3 lattice shows much greater tolerance to IP offsets.

For the CDR 4IP lattice, it is clear that the beam is strongly distorted in action space with vertical momentum offset at the IP. The beam core becomes highly stretched and depopulated. It appears that the impact of the vertical momentum offset is reduced compared to the vertical position offset. For the GHC v24.3 lattice, this strong beam core depopulation is not observed, and now some significant changes can be observed to only the halo of the beam. Once again this implies the GHC v24.3 lattice shows much greater tolerance to IP offsets.

The overall improved tolerances with the GHC v24.3 compared to the CDR 4IP lattice is likely due to the more optimised tune working point. This implies from a machine point of view, the tolerances on IP offset are relaxed.




Figure 6.20: Action-space density distributions in the transverse action space ($\sqrt{2J_x}$, $\sqrt{2J_y}$) for increasing vertical IP momentum offsets. Columns correspond to nominal collision, a $0.05\sigma_{p_y}$ offset, and a $0.20\sigma_{p_y}$ offset, respectively. The top row shows results for the FCC-ee CDR lattice with four IPs, while the bottom row shows results for the FCC-ee GHC v24.3 lattice. Colour contours represent the logarithm of the normalised particle density.

6.3.10 Frequency Map Analysis Studies

To further analyse the observations of Section 6.3.9 Frequency Map Analysis was performed. Frequency map analysis (FMA) is a numerical technique used to characterise the non-linear dynamics of particle motion in accelerators by analysing the evolution of fundamental oscillation frequencies over time [149]. In this method, particle trajectories are tracked for a large number of turns, and the dominant betatron frequencies are extracted from the phase-space coordinates using high-resolution spectral analysis techniques, such as the Numerical Analysis of Fundamental Frequencies (NAFF) [150]. The frequencies are then computed over two or more distinct time windows and compared to quantify their temporal stability. For regular motion, the extracted frequencies remain constant, whereas chaotic trajectories exhibit measurable frequency drift, reflecting sensitivity to initial conditions. By evaluating the frequency diffusion and mapping the resulting tune footprints in frequency space, FMA provides a powerful diagnostic for identifying non-linear resonances, delineating stable and unstable regions of phase space, and assessing the impact of non-linear elements such as beam-beam interactions, magnetic errors, and radiation effects on long-term beam stability.

The tune diffusion, a parameter encapsulating the stability of the tune of a given

particle, was studied. The tune diffusion is defined as

$$D = \frac{1}{2} \log_{10} (\sigma_{Q_x}^2 + \sigma_{Q_y}^2). \quad (6.4)$$

The tune diffusion was computed for each tracked particle. For each particle, the horizontal and vertical tunes were calculated ten times in windows of 512 turns (half of the tracking duration) using the NAFF algorithm. From the ten computed tunes per particle per plane, the standard deviation was calculated. Particles with high tune diffusions are more likely to cross resonances and eventually be lost.

A comparison of the tune diffusion coefficient, presented in action space, for the previously studied CDR 4IP lattice at the Z working point and the GHC version 24.3 lattice at the Z working point with vertical position offsets is presented in Fig. 6.21. The same comparison, presented in tune space, is presented in Fig. 6.22.




Figure 6.21: Tune diffusion coefficient distributions displayed in transverse action space ($\sqrt{2J_x}$, $\sqrt{2J_y}$) for increasing vertical IP position offsets. Columns correspond to nominal collision, a $0.05\sigma_y$ offset, and a $0.20\sigma_y$ offset, respectively. The top row shows results for the FCC-ee CDR lattice with four IPs, while the bottom row shows results for the FCC-ee GHC v24.3 lattice. Colour indicates the logarithm of the tune diffusion coefficient.

Due to the very different lattice configurations, the tune footprints look very different between the CDR 4IP lattice and the GHC v24.3 lattice. The particles closest to

$(0,0)$ are not presented as these experience the strongest beam–beam force and the computation of the tunes fails. With vertical position offset, additional tune diffusion is observed in action space. For the CDR lattice there is a strong diffusion appearing close to $(0,0)$. This is of particular concern as the greatest beam density exists in this region. For the GHC lattice, some tune diffusion also appears close to $(0,0)$, but more weakly.




Figure 6.22: Tune diffusion coefficient distributions displayed in tune space (Q_x , Q_y) for increasing vertical IP position offsets. Columns correspond to the nominal case, a $0.05\sigma_y$ offset, and a $0.20\sigma_y$ offset, respectively. The top row shows results for the FCC-ee CDR lattice with four IPs, while the bottom row shows results for the FCC-ee GHC v24.3 lattice. Colour indicates the logarithm of the tune diffusion coefficient.

Analysing the tune diffusion in tune space is more instructive as the resonance lines can be more clearly identified. The CDR tune footprint is clearly heavily distorted when vertical position offsets are introduced at the IP. It appears that the $3Q_y = 2$ resonance is highly driven by the beam–beam force with offset. On the other hand, while a strong resonance is observed in the GHC v24.3 lattice, this is present even in the nominal case so no major degradation is observed with offset; no further resonances are driven in the GHC v24.3 lattice with vertical position offset at the IP.

This explains the improved stability of the GHC v24.3 lattice in the presence of IP offsets with beam–beam.

6.4 Radiation and Beam–Beam Modelling for Multi-Turn Tracking

To simulate IP Feedback for FCC-ee, a lattice must be prepared for multi-turn tracking simulations. As discussed in Section 6.1, at the time of writing there are not only four working points to consider, but also two competing optics.

There are three key factors influencing the number of turns tracked for feedback simulation: the frequencies of the errors, the revolution frequency and the radiation damping time. FCC-ee has a revolution frequency of ~ 3.3 kHz making it possible to observe relevant error frequencies (those above the response of a Closed Orbit Feedback ~ 10 Hz) in order hundreds of turns. To reduce the number of turns required to track to equilibrium, using a lattice with a low radiation damping time is preferable.

The working point with the lowest radiation damping time is the $t\bar{t}$ working point. However, the $t\bar{t}$ lattices may still be considered an upgrade to the baseline. As such, the preferable choice is the ZH working point. This is also preferable considering the physics programme of FCC-ee: aiming to fulfil the ESPP’s requirement for a Higg’s factory. For this working point, no current LCC lattice version was available. As such, the choice was made to use a GHC ZH lattice for the feedback simulations. No solenoid is included in this lattice.

6.4.1 Lattice Reversals

During optics development a single FCC-ee lattice is developed per working point for each lattice scheme, typically the positron lattice. For studies involving offsets at the IP, it is vital to have two distinct lattices where the particles are bent in the correct direction. If this were not the case, the oscillations in the lattices would not be self-consistent. To achieve this, a series of transformations on the bending magnets powering and reference frame transforms were performed to create a distinct electron lattice. The lattices are aligned spatially at the RF. No reversals of the technical insertions were performed.

6.4.2 Synchrotron Radiation Modelling

As discussed in Section 5.3.1, synchrotron radiation modelling is required to simulate equilibrium with particle tracking in high energy lepton colliders. Unlike the SuperKEKB case, FCC-ee is foreseen to utilise tapering: magnet strengths will be adjusted to the expected local relative momentum deviation, δ . The full tapering scheme has not been fully developed. As such, idealised tapering is used, where each magnet is individually tapered. In the real machine, this level of granularity may not be possible due to the cost of the additional power converters.

The variation of delta with longitudinal coordinate for the electron and positron beams from a closed-ring 6D optics calculation including synchrotron motion, synchrotron radiation effects and idealised tapering is presented in Fig. 6.23.




Figure 6.23: Variation of the local relative momentum deviation δ as a function of longitudinal coordinate s along the FCC-ee ring at the ZH working point. The positron beam is shown in red and the electron beam in blue. Vertical dashed black lines indicate the longitudinal locations of the four IPs.

The use of tapering greatly reduces the effect of synchrotron radiation energy loss on the beam optics. The process is imperfect so some residual optics distortion is present. To identify the strength of the optics distortion, the relative change of the horizontal and vertical beta functions between the reference no radiation case and the radiation plus tapering case is presented in Fig. 6.24. A non-negligible optics distortion remains




Figure 6.24: Relative beta-beating $\Delta\beta/\beta$ of the horizontal beta function β_x (red) and vertical beta function β_y (blue) as a function of longitudinal coordinate s for FCC-ee at the ZH working point. The optics include synchrotron radiation and lattice tapering.

present.

6.4.3 Generating Vertical Emittance

Current FCC-ee lattices do not feature solenoids, nor any vertical orbit by design. There are therefore no sources of vertical dispersion and the vertical emittance is zero. In the real machine, vertical emittance will be generated by coupling throughout the beam-lines, as well as the experimental solenoids. In lattice parameter tables, vertical emittances are assumed based on expected minimal coupling after optics correction.

To generate vertical emittance, two approaches were tested. The approach that has been standard for FCC-ee vertical emittance generation has been the use of a rotated wiggler. The approach of offsetting sextupole pairs, as presented in Section 5.3.3 was also tested.

Wiggler Emittance Generation

To generate vertical emittance with a wiggler in FCC-ee, the polarisation wiggler is rotated to be vertical. This wiggler is initially unpowered in the lattice. The wiggler powering is parameterised by a single strength parameter. The vertical emittance is then matched to the target value using Xsuite’s match function by varying this strength parameter. The match function uses Xsuite’s ‘line.twiss’ method to compute the optics parameters and calculate the emittances from a normal mode decomposition. Some discrepancies are observed between tracked and optics computation emittances. Therefore, the strength parameter was then scanned around the previously optimised value and tracking to equilibrium was performed. This scan was performed by initialising 100 macroparticles as a point beam at $(0, 0, 0, 0, 0, 0)$ then allowing the system to evolve for 1,000 turns under the effect of quantum synchrotron radiation. From this scan an optimised strength was identified. The results of the tracking scan over the wiggler strength parameter are presented in Fig. 6.25 and the results of tracking with the optimised wiggler are presented in Fig. 6.26.

During the wiggler scans, the horizontal and longitudinal emittances are seen to be consistent within errors throughout. This indicates the wiggler is not causing significant changes in the damping partitions. The vertical emittance is seen to vary strongly over the restricted range of the wiggler strength parameter considered.

From tracking, both the horizontal and longitudinal emittances are seen to slightly exceed the values initially predicted using Xsuite’s optics computation. Most importantly, the vertical emittance is now observed to be in excellent agreement with the target.

Offset Sextupole Emittance Generation

To generate vertical emittance with offset sextupoles for FCC-ee the approach was broadly the same as described in Section 5.3.3. To mitigate the negative impact on




Figure 6.25: Horizontal, vertical, and longitudinal emittances as a function of wiggler parameter offset for FCC-ee at the ZH working point. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The horizontal axis shows the wiggler parameter. The red dashed line indicates the target vertical emittance.




Figure 6.26: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for FCC-ee at the ZH working point with optimised wiggler parameter. The left column shows the positron ring and the right column shows the electron ring. Blue lines indicate the tracked emittances. Black horizontal lines indicate the Twiss emittances. Magenta lines indicate the observed equilibrium values.

the machine optics, each sextupole pair is offset oppositely; the first sextupole is offset upwards and the second sextupole downwards with equal magnitudes. In the case of the GHC lattice, the sextupoles are split. As such, the first two sextupole parts are offset upwards, and the second two downwards. This ensures that the two offset parts have a phase advance of π between them. This ideally creates closed coupling bumps which result in vertical emittance. All sextupoles in both the positron and electron ring were offset by the same magnitude. By scanning the offset, it was possible to optimise for the target vertical emittance. This scan was performed by initialising 100 macroparticles as a point beam at $(0, 0, 0, 0, 0, 0)$ then allowing the system to evolve for 1,000 turns under the effect of quantum synchrotron radiation. The results of the tracking scan over sextupole offset are presented in Fig. 6.27 and the results of tracking with the optimised sextupole offset are presented in Fig. 6.28.




Figure 6.27: Horizontal, vertical, and longitudinal emittances as a function of sextupole offset for FCC-ee at the ZH working point. The top row shows the horizontal emittance ϵ_x , the middle row shows the vertical emittance ϵ_y , and the bottom row shows the longitudinal emittance ϵ_z . The horizontal axis shows the sextupole offset in metres. The red dashed line indicates the target vertical emittance.

During the offset sextupole scans, the horizontal and longitudinal emittances are seen to be consistent within errors throughout. This indicates the offset sextupole scheme is not causing significant changes in the damping partitions. The vertical emittance is seen to vary over two orders of magnitude over the one order of magnitude of sextupole offset scanned. This is of concern for the real machine, highlighting the sensitivity of FCC-ee’s current optics to machine errors.

From tracking, the vertical and longitudinal emittances are now observed to be in




Figure 6.28: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for FCC-ee at the ZH working point with optimised sextupole offsets. The left column shows the positron ring and the right column shows the electron ring. Blue lines indicate the tracked emittances. Black horizontal lines indicate the Twiss emittances. Magenta lines indicate the observed equilibrium values.

excellent agreement with the target. The horizontal emittances are seen to slightly exceed the value initially predicted using Xsuite’s optics computation. This is likely a result of the residual optics distortion in the lattice.

Emittance Generation Choice

There are some notable issues with both approaches. The vertical emittance generation by use of the wiggler is a highly localised effect. When beam–beam is included in simulations, strong beta-beating occurs in the lattice due to the beam–beam tune shift. It was recently noted that the beta-beating at the location of the vertical emittance wiggler is very high [151]. As such, the vertical emittance greatly increases in the presence of beam–beam, in a way that would not be expected with a more distributed error. On the other hand, when using the offset sextupoles, residual optics distortions are observed. In the SuperKEKB case, these optics distortions were observed to be small as presented in Table 5.2; These optics distortions resulted in unstable beams during beam–beam simulation. Optics correction knobs for FCC-ee have not yet been standardised. It was not within the scope of this thesis to create a full optics correction in order to use the offset sextupole scheme. As such, for future simulations, the wiggler has been adopted to generate vertical emittance. The resulting synchrotron radiation

vertical emittances are presented in Table 6.7. Small increases in the horizontal and

Table 6.7: Equilibrium horizontal, vertical and longitudinal emittances for FCC-ee obtained from particle tracking including quantum synchrotron radiation, with and without the optimised vertical emittance wiggler.

Parameter	Unit	W/O Wiggler	W/ Wiggler
ε_x	nm	0.69	0.72
ε_y	pm	0.00	0.76
ε_z	μm	3.33	3.64

longitudinal emittances are observed due to the additional energy loss of the wiggler.

6.4.4 Beam–Beam Modelling

The beam–beam elements were included as described in Section 5.3.4. As was discussed in Section 3.7.1, to see the cross talk between the bunches at different IPs, a minimum of two electron bunches and two positron bunches must be simulated. This results in eight beam–beam interactions per turn. Lattice tracking for four bunches is also required. This makes simulating self-consistent beam–beam interactions at FCC-ee highly computationally expensive.

Because of computational constraints, the number of macroparticles per bunch was chosen to be 10,000, fewer than ideal for such a strong–strong simulation and increasing the noise of the simulation. In tests over 325 turns, the observed behaviour appeared roughly comparable between 10,000 and 100,000 macroparticles. The parameters used for the strong–strong Xsuite simulations are presented in Table 6.8.

Table 6.8: Parameters adopted for the beam–beam simulations used in the FCC-ee IP feedback simulation studies.

Parameter	Unit	Value
Beam–beam Mode	/	Strong–strong
Beam Current	mA	26.8
Number of Bunches	1	300
Number of Macroparticles	1	10,000
Number of Slices	1	51
Update Frequency	Turn	1
Slicer Mode	/	‘Shatilov’
Beamstrahlung Model	/	Quantum
Bhabha Model	/	None

6.4.5 Equilibrium with Beam–Beam

As discussed in Section 5.3.5, the inclusion of beam–beam elements in the lattice greatly distorts the optics. For FCC-ee, due to the large beam–beam tune shifts the optics are greatly distorted.

To analyse the impact of offsets and feedback fairly, future simulations must start from an equilibrium including the beam–beam force. This allows the effect of the nominal beam–beam force and the offsets and feedback to be decoupled. To identify the new equilibrium parameters particle tracking was performed. Gaussian bunches were initialised as described in Section 5.3.3 with the emittances presented in Table 6.7 for the horizontal and vertical planes. For the longitudinal plane, the emittance was scaled by a factor of 2.5 to account for the impact of beamstrahlung radiation. This is done to avoid instabilities that occur in collisions where the bunch length was much shorter than equilibrium. Beam–beam was included in strong–strong mode with the parameters presented in Table 6.8 and the bunches were tracked for 2,000 turns. The emittance evolution is presented in Fig. 6.29.




Figure 6.29: Turn-by-turn evolution of the horizontal emittance ε_x (top row), vertical emittance ε_y (middle row), and longitudinal emittance ε_z (bottom row) for FCC-ee at the ZH working point with strong–strong beam–beam and optimised wiggler parameter. Dashed horizontal lines indicate the relevant synchrotron radiation equilibrium emittance values.

An issue was identified with the lattice at this stage. When beam–beam was included an XZ instability was observed. The lattice currently features a total horizontal beam–beam tune shift that exceeds the synchrotron tune. This results in the tune space crossing Q_s resonances and an XZ instability. This is visible in Fig. 6.29 in the strong

oscillations in the horizontal and longitudinal emittances. Despite the oscillations of the emittances, the system reaches a quasi stable state, and appears suitable for tracking studies. Using this lattice remains preferable to the long damping times of the WW and Z working points. Furthermore, the luminosity performance was not observed to suffer.

The resulting equilibrium emittances under synchrotron radiation with strong-strong beam-beam and optimised wiggler strength are presented in Table 6.9.

Table 6.9: Equilibrium horizontal, vertical and longitudinal emittances for FCC-ee obtained from particle tracking including quantum synchrotron radiation, optimised sextupoles offsets and strong-strong beam-beam.

Parameter	Unit	Value
ε_x	nm	1.76
ε_y	pm	1.36
ε_z	μm	8.83

In addition to emittances, the equilibrium luminosity was identified. The evolution of the luminosity is presented in Fig. 6.30. The luminosity is seen to initially sharply




Figure 6.30: Turn-by-turn evolution of luminosity for FCC-ee at the ZH working point with strong-strong beam-beam and optimised wiggler parameter. Results for the positrons are shown in red and results for the electrons are shown in blue.

decay as the beams increase in size reaching new equilibrium emittances. Oscillations can be observed due to the XZ instability though these oscillations are small.

6.4.6 Beam-Beam Tune Shifts

The proposed FCC-ee feedback system is based on beam-beam deflection. In the feedback regime, where offsets are small compared to the beam sizes, this deflection is proportional to the product of the offset and the strength of the beam-beam force. It is therefore important to characterise the strength of the beam-beam force using the beam-beam tune shift.

The beam–beam tune shifts have been identified in two ways. Firstly, using the analytic formula presented in Section 2.3.3. Secondly, by performing a closed-ring 6D linear optics calculation including synchrotron motion, synchrotron radiation effects, tapering and weak–strong beam–beam elements at each IP. For the Twiss computation to remain stable, the beam–beam strength must be linearly rescaled. When using the equilibrium emittance computed including beam–beam, the beam–beam element in the line is one half the nominal strength. When using the equilibrium emittances computed including synchrotron radiation only it is one tenth of the nominal strength. The computed beam–beam tune shifts are then rescaled by the corresponding factor of two or ten. The computed beam–beam tune shifts are presented in Table 6.10. There is

Table 6.10: Horizontal and vertical beam–beam tune shifts for the FCC-ee evaluated for different emittance configurations. The emittance configurations correspond to those values presented in Table 6.7 and Table 6.9. The beam–beam simulation settings follow Table 6.8.

Configuration	Mode	$\xi_x [10^{-3}]$	$\xi_y [10^{-3}]$
Wiggler + Synchrotron	Analytic	108.9	879.2
Radiation	Weak–strong	106.0	846.4
Wiggler + Synchrotron	Analytic	42.3	435.0
Radiation + Beam–Beam	Weak–strong	44.9	422.5

excellent agreement between the analytic and rescaled optics computation tune shifts. This indicates that the beam–beam elements are correctly installed with appropriate crossing angles, bunch charges and beam sizes.

6.5 Feedback Modelling

6.5.1 FCC-ee Feedback Correction Knobs

During operation, FCC-ee feedback corrector magnets must be continuously updated to steer the beam. Given SuperKEKB’s demonstrated operation at 32 kHz, it is foreseen that FCC-ee will operate with turn-by-turn feedback. At such rates, it remains impractical to compute optimal corrector strengths from scratch at each cycle. Instead, correction “knobs” are pre-computed, consisting of linear scaling factors for the corrector strengths matched to a target offset. Corrector strengths are then obtained from the measured offsets and the pre-computed knobs.

For the proposed FCC-ee IP feedback, offset correction knobs were generated using the Xsuite FCC-ee model. Knobs were constructed on the electron beam-line at two IPs and on the positron beam-line at two IPs. In operation, feedback hardware is foreseen at all IPs for both beam-lines to provide redundancy.

In the present FCC-ee lattice models, no solenoid is included, significantly reducing the complexity of the feedback knobs. Following recommendations by K. Oide (private

communication, October 2025), the primary correctors are taken to be dipole correction coils on the first quadrupoles downstream of the final focus system. Two fast-ramping correctors are used on each side of the IP (foreseen to be installed on ceramic beam pipe to mitigate eddy currents). Hardware feasibility considerations are beyond the scope of this thesis.

In this study, only vertical position correction is considered. The knob is matched to produce the desired offset at the IP while closing the orbit outside the interaction region. Dispersion matching is not explicitly imposed, as only two correctors per side are used; dispersion leakage is instead evaluated a posteriori.

The vertical position correction knob is constructed in Xsuite by imposing the following conditions:

1. $x = 0$, $p_x = 0$, $y = y_{\text{target}}$, and $p_y = 0$ at the IP,
2. $x = 0$, $p_x = 0$, $y = 0$, and $p_y = 0$ in the arc.

Although this corresponds to eight constraints with four degrees of freedom, the match is possible here due to full transverse decoupling. With a solenoid included, this matching may no longer be feasible.

6.5.2 Xsuite FCC-ee Feedback Knob Implementation

The FCC-ee IP feedback vertical position knob was implemented in Xsuite using the Xtrack “line.match_knob” function. A maximum kick angle of $100\text{ }\mu\text{rad}$ per corrector was enforced. Five closed orbit bumps, varying linearly from $+20\text{ }\mu\text{m}$ to $-20\text{ }\mu\text{m}$ of IP offset, are presented in Fig. 6.31. The knob is approximately symmetric, with small asymmetries arising from unequal magnet spacing upstream and downstream of the IP.

During matching, the target vertical position offset at the IP was taken to be $20\text{ }\mu\text{m}$. The linearity of the knob was assessed by evaluating orbit positions and angles at the IP as a function of the applied offset, as shown in Fig. 6.32. Excellent performance is observed, with horizontal orbit residuals below 0.1 nm and angle residuals below 0.1 nrad across the full knob range.

In addition to producing the desired IP offset, orbit bumps may perturb the machine optics. To quantify this effect, closed-ring 6D Twiss calculations including synchrotron motion, synchrotron radiation, and tapering were performed with the knob applied. While the beta beating in the lattice is far below the target value of 3% set in previous studies [152], at large offsets ($\pm 20\text{ }\mu\text{m}$), the induced vertical dispersion reaches several millimetres throughout the lattice, which is undesirable for routine operation.

In practice, FCC-ee operation is foreseen to rely on slow orbit bumps when approaching the limits of the fast feedback knob. The effective working range of the fast




Figure 6.31: Horizontal and vertical orbit in the interaction region for the FCC-ee IP Feedback vertical position knob as a function of longitudinal coordinate s . Vertical orbits are shown in red and horizontal orbits are shown in black.




Figure 6.32: Orbit positions and angles at the IP as a function of the FCC-ee IP Feedback vertical offset knob. The top panel shows the horizontal orbit position (x) and vertical orbit position (y) at the IP. The bottom panel shows the horizontal orbit angle (p_x) and vertical orbit angle (p_y) at the IP. Horizontal quantities are shown in black and vertical quantities are shown in red.

knob is therefore expected to be significantly smaller. To assess optics perturbations in this regime, the knob was evaluated over a reduced range of $\pm 1 \mu\text{m}$. The resulting beta beating and dispersion are presented in Fig. 6.33. In this case, vertical dispersion remains below $200 \mu\text{m}$ throughout the lattice, with a maximum of $0.1 \mu\text{m}$ at the IP. This corresponds to an operational range of approximately 40 beam sigma based on the nominal emittances, and is considered acceptable for this feedback simulation.




Figure 6.33: Relative beta beating and dispersion differences for two reduced FCC-ee IP Feedback vertical position knob settings as a function of longitudinal coordinate s . The top panel shows the relative horizontal beta beating ($\Delta\beta_x/\beta_x$), the second panel shows the relative vertical beta beating ($\Delta\beta_y/\beta_y$), the third panel shows the horizontal dispersion difference (ΔD_x), and the bottom panel shows the vertical dispersion difference (ΔD_y). Results for an FCC-ee IP vertical offset of $+20\ \mu\text{m}$ are shown in red, and results for an offset of $-20\ \mu\text{m}$ are shown in blue.

6.5.3 BPM Implementation

The IP Feedback BPMs are modelled using Xtrack “BeamPositionMonitor” elements. Four BPMs are included at each IP, inserted $\pm 1.1\text{ m}$ from the IP. These monitors return average centroid positions in the transverse plane and counts of the number of macroparticles passing through the monitor within an observation period. The frequency of observations is set at 3.3 kHz , corresponding to the revolution frequency of the machine. To simulate the resolution of a real BPM, post processing is applied. A resolution of $0.2\ \mu\text{m}$ is used. While no resolution for these BPMs has been confirmed by the FCC-ee Beam Instrumentation team, this is the same resolution used for SuperKEKB simulations. It is therefore taken as a feasible value for FCC-ee.

6.5.4 Offset Calculation

For a realistic computation of the IP offset from measurable signals, many pieces of information are required. Measurements of the centroid positions at each of the four BPMs per IP are calculated as described in Section 6.5.3. To transport these measurements to the IP, the reference orbit at the BPMs and the transfer matrices

between the BPMs and the IP are required. These are calculated by performing a closed-ring 6D linear optics calculation including synchrotron motion, synchrotron radiation and tapering. To calculate the canonical kick experienced during the beam–beam collision, the beam–beam tune shifts are required. These are computed analytically from IP beta functions, bunch charges, emittances and the crossing angle. The bunch charges are computed from the BPM counts. Emittances are monitored using an Xcoll “EmittanceMonitor” element, simulating measurements of synchrotron radiation monitors foreseen in the collider. The Twiss beta functions are taken from a closed-ring 6D linear optics calculation including synchrotron motion, synchrotron radiation and tapering. From this information, the offset at the IP is calculated following the method described in Section 2.7.6.

6.5.5 Feedback Implementation

The feedback implementation is based on the results of the analytic modelling and operational experience from the SuperKEKB iBump system (described in Section 4.1.2).

Feedback is implemented independently at each IP. IPs A and C are corrected using the positron beam-line, while IPs B and D are corrected using the electron beam-line. The feedback knob calculations are performed independently, consistent with fast local FPGA-based systems operating at each IP. Based on the results presented in Section 3, turn-by-turn feedback is employed.

The linear prediction algorithm uses history matrix parameters $N = 5$ and $M = 20$ corresponding to 20 previous cycles for weight computation and 5 cycles for forward prediction. A feedback gain of 0.9 is adopted based on initial testing, without further optimisation. Magnet ramp times and eddy-current effects are not included, as turn-to-turn variations in magnet strength are small. A demonstration of the prediction algorithm using pure sinusoidal inputs is shown in Fig. 5.24.

During testing, large transients were observed at feedback initialisation. To mitigate this, the gain is smoothly ramped using a cosine profile over 65 turns, corresponding to approximately one damping time, which was found to suppress these transients effectively.

6.6 Multi-Turn Particle Tracking Simulations

6.6.1 Error Implementation

To model IP feedback at FCC-ee realistically would require a realistic model of the errors experienced by the machine in operation. This would include magnet powering errors, vibrational errors, the impact of the global Closed Orbit Feedback (COFB), errors in the BPMs, errors in the synchrotron radiation monitors and more. Finally, the

machine’s initial state, and the vertical emittance, would be generated by errors and optics correction performed in a realistic way with radiation and pre-defined correction knobs. The implementation of all of these errors and corrections is beyond the scope of this thesis.

Two error models have been tested. The first error model uses driven sinusoidal errors. As was performed in the analysis of SuperKEKB’s iBump feedback system, described in Section 5.5.1, the errors were driven by a corrector magnet. The corrector chosen was one of the vertical wiggler magnets, “bwigm.0”, and the error was driven in the positron beam-lines. The corrector was driven with an integrated dipole strength of $k_0L = 1 \times 10^{-7}$ for a range of frequencies. This approach allowed for systematic testing of the feedback’s response to different frequencies of error. The choice of oscillating the positron beam was arbitrary.

The second error model was realistic vibration of the Final Focus Quadrupole magnets. In Section 3.2.5 an approximate model of the vibration of the FCC-ee tunnel was estimated based on LHC tunnel measurements. An estimated Power Spectral Density for the vibrational motion of the FFQs was then presented in Fig. 3.6 and will be used for these studies. From the PSD, vibration time series were generated, and the final focus quadrupoles were shifted vertically according to this time series. The time series was generated with a random seed and all offsets started at zero, as if an ideal orbit correction had just been performed. This allows for testing of the feedback’s response to more realistic errors.

FCC-ee Closed Orbit Feedback

As discussed in Section 6.3.1, the global Closed Orbit Feedback (COFB) at FCC-ee is still under development. FCC-ee is a highly challenging machine with sensitive optics. Initial tests with vibration time series generated using the full estimated PSD resulted in significant beam losses. As such, the PSD considered will be a PSD adjusted for the inclusion of the COFB.

Assuming that the COFB uses PI control proportional to frequency with a crossover frequency of 5 Hz (somewhat pessimistically below the quoted 10 Hz), a new PSD is generated. The resulting PSD and a comparison to the original estimate is presented in Fig. 6.34. A clear suppression of the noise is seen for low frequencies.

Correlation of FFQ Offsets

As discussed in Section 4.1.1, strong correlations were observed in the offsets of the Final Focus Quadrupoles within the same cryostat at SuperKEKB. The correlations were primarily in the low frequency oscillations.

For this study, it has been assumed that all three sub-components of each QC1




Figure 6.34: Comparison of the estimated vibrational motion power spectral density of a final focus quadrupole at FCC-ee with and without global closed orbit feedback. The blue curve shows the original estimate of the power spectral density, while the orange curve shows the corresponding spectrum after accounting for the suppression provided by closed orbit feedback.

magnet oscillate identically. Therefore one vibration time series is generated per side of the IP per beam-line. Correlations have been ignored between the positron and electron beam-line vibrations. The strong correlations observed at SuperKEKB were for the low frequency motion which is highly suppressed due to the inclusion of the COFB. Testing with perfect correlations resulted in beams tracking each other in a way that seemed unrealistically positive.

Examples of generated vibrational time series with the original and post COFB PSDs are presented in Fig. 6.35. Vibrational time series are initialised at zero offset.

6.6.2 Tracking with Driven Sinusoidal Errors: IP Feedback Disabled

With driven errors the quality of the beams is expected to degrade. As the beams orbit off axis, they traverse magnetic elements off-centre leading to emittance increases. Tracking without feedback is performed to establish a reference for comparison with the cases where feedback is enabled. Beams were initialised with the equilibrium emittances obtained from tracking with strong-strong beam-beam, optimised wiggler and quantum synchrotron radiation presented in Table 6.9.

The evolution of the beam centroids and beam sizes at IP A without feedback for a driven error frequency of 10 Hz is presented in Fig. 6.36.

The evolution of the vertical position offset at each IP with turns for a driven error frequency of 10 Hz is presented in Fig. 6.37. With this low frequency of driven error, no discrepancy is observed between the first and second bunch crossing, as the driven error frequency is much below the revolution frequency. The vertical position offset




Figure 6.35: Comparison of approximated vibrational motion time series for the final focus quadrupole triplets at FCC-ee, generated from PSD estimates without closed orbit feedback (left column) and with closed orbit feedback (right column). Blue curves show the vertical motion of the quadrupole triplets on the left-hand side of each IP, while orange curves show the vertical motion of the quadrupole triplets on the right-hand side. Results correspond to single realisations generated using one random seed.

differs between the IPs as the error is driven in one arc only. The vertical position offset at a given IP therefore depends on the phase advance between the error and the IP. Significant offsets, on the order of the nominal beam size were driven at the IPs.

In addition to the IP vertical position offset, two quantities of particular interest are the emittances and luminosity. The evolution of the horizontal, vertical and longitudinal emittances with turns for a range of driven errors is presented in Fig. 6.38. In the presence of the driven error the vertical emittance is observed to sharply increase for the positrons. The error source is on the positron beam-line so this is expected. The vertical emittance initially increases also on the electron line, but damps back down. In the longitudinal plane, the positron emittances are seen to increase while the electron emittances decrease. This appears as a 3D flip-flop instability, where the increased vertical emittance of the positron bunches reduces the amount of beamstrahlung emitted by the electron bunches. This in turn causes the electron bunches to shrink and the positron bunches in turn emit more beamstrahlung and further blow up. A new equilibrium appears to be reached.

In addition to the emittance evolution, the evolution of the luminosity with turns for a range of driven errors is presented in Fig. 6.39. Some minor discrepancies are observed in the computed luminosities for the electron and positron bunches due to the




Figure 6.36: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for FCC-ee with a single frequency error and no IP feedback. The top row shows the horizontal centroid position x , the middle row shows the vertical centroid position y , and the bottom row shows the longitudinal centroid position z . The left column corresponds to the first bunch crossing and the right column corresponds to the second bunch crossing. Positrons are shown in red and electrons are shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.

soft Gaussian approximation used in the beam-beam element. A clear oscillation of the drive frequency is observed in the luminosity. The luminosity is degraded with respect to the nominal value, however the system appears to have reached a stable state.

6.6.3 Tracking with Driven Sinusoidal Errors: IP Feedback Enabled

By including IP feedback, the offsets at the IP are expected to be reduced and thus some of the luminosity loss should be recovered. Some degradation is still expected: the beams are still being driven in the arcs, and now further driven by the IP feedback correctors at the IPs. The aim is that the reduced IP offsets still result in an increase to luminosity. Beams were initialised with the equilibrium emittances obtained from tracking with strong-strong beam-beam, optimised wiggler and quantum synchrotron radiation presented in Table 6.9.

The evolution of the beam centroids and beam sizes at IP A with feedback for a driven error frequency of 10 Hz is presented in Fig. 6.40.

The evolution of the vertical position offset at each IP with turns for a driven error frequency of 10 Hz is presented in Fig. 6.41.




Figure 6.37: Turn-by-turn evolution of IP vertical position offsets for FCC-ee with a single frequency error and no IP feedback. The top row shows the offset at IP A, the second row shows the offset at IP B, the third row shows the offset at IP C, and the bottom row shows the offset at IP D. Offsets during the first bunch crossing are shown in blue, and offsets during the second bunch crossing are shown in orange.

With this low frequency of driven error, no discrepancy is observed between the first and second bunch crossing, as the driven error frequency is much below the revolution frequency. An initial growth of offset is seed, and this is rapidly damped away as the feedback is enabled after turn 25. The damping is not instant as the feedback ramps up over the first 65 turns.

The evolution of the emittance with turns for a range of driven errors is presented in Fig. 6.42. As in the case without feedback, in the presence of the driven error the vertical emittance is observed to sharply increase for the positrons.

In addition to the emittance evolution, the evolution of the luminosity with turns for a range of driven errors is presented in Fig. 6.43. Some minor discrepancies are observed in the computed luminosities for the electron and positron bunches due to the soft Gaussian approximation used in the beam-beam element. A clear oscillation of the drive frequency remains observable in the luminosity. The luminosity remains degraded with respect to the nominal value, however, by comparison to Fig. 6.39, the luminosity degradation has been visibly mitigated.




Figure 6.38: Turn-by-turn evolution of beam emittances for FCC-ee with a single frequency error and no IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The first column shows the first positron bunch, the second column the second positron bunch, the third column the first electron bunch and the fourth column the second electron bunch. Emittances are shown in blue, and the corresponding initial values are indicated by red dashed lines.




Figure 6.39: Turn-by-turn evolution of luminosity for FCC-ee with a single frequency error and no IP feedback. Luminosity computed for the positron bunches are shown in red, while luminosity computed for the electron bunches are shown in blue.

6.6.4 Tracking with Driven Sinusoidal Errors: Frequency Dependence

Following the demonstrations presented in Section 6.6.2 and Section 6.6.3, a scan across a range of error frequencies was undertaken. This scan tested driven error frequencies




Figure 6.40: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for FCC-ee with a single frequency error and IP feedback. The top row shows the horizontal centroid position x , the middle row shows the vertical centroid position y , and the bottom row shows the longitudinal centroid position z . The left column corresponds to the first bunch crossing and the right column corresponds to the second bunch crossing. Positrons are shown in red and electrons are shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.

spanning DC to 1 kHz, using approximately logarithmic spacing with IP feedback enabled and disabled. The dependence of the RMS vertical position offset at each IP with driven error frequency is presented in Fig. 6.44. With no feedback, the RMS vertical position offset is stable at frequencies up to 500 Hz. At 1000 Hz the RMS offsets are greatly increased indicating that some resonance has been hit. With feedback on, the RMS vertical position offset is reduced with respect to the no feedback case throughout. The impact of the feedback is greater for lower frequencies, as expected of this integral feedback and the transfer function is presented in Fig. 6.45.

Further to the IP offsets, the dependence of the horizontal, vertical and longitudinal emittances on the frequency of the driven error was studied and results are presented in Fig. 6.46. The form of the emittance dependence on frequency is complex. In general, the emittances of the positron bunches are higher than the emittances of the electron bunches. This is expected as the error source is situated on the positron beam-line. A decrease in all emittances for the positrons and a corresponding increase is observed for the electrons between 100 Hz and 500 Hz.

Finally, considering both the IP offset and emittance information, the luminosity was considered. The dependence of the per IP luminosity on driven error frequency is




Figure 6.41: Turn-by-turn evolution of IP vertical position offsets for FCC-ee with a single frequency error and IP feedback. The top row shows the offset at IP A, the second row shows the offset at IP B, the third row shows the offset at IP C, and the bottom row shows the offset at IP D. Offsets during the first bunch crossing are shown in blue, and offsets during the second bunch crossing are shown in orange.

presented in Fig. 6.47. The luminosity is observed to degrade with frequency of driven error. A sharp decrease in luminosity is observed above 500 Hz. This decrease is clearly explained by the significant increase in vertical emittance for the positron bunches as presented in Fig. 6.46. With IP feedback enabled, the luminosity is improved, though this effect is diminishing with frequency. Overall, this comparison clearly shows the positive impact of the IP feedback system.

6.6.5 Tracking with Realistic Final Focus Quadrupole Vibrations: IP Feedback Disabled

Now a vibration time series is pre-generated for each set of FFQs either side of the IP and the FFQs are oscillated. Offsets of these FFQs results in strong dipolar kicks and residual vertical dispersion which is expected to degrade the emittance. Tracking without feedback is performed to establish a reference for comparison with the cases where feedback is enabled. Beams were initialised with the equilibrium emittances obtained from tracking with strong-strong beam-beam, optimised wiggler and quantum synchrotron radiation presented in Table 6.9.

The evolution of the beam centroids and beam sizes at IP A without feedback for one FFQ vibration time series random seed is presented in Fig. 6.48.




Figure 6.42: Turn-by-turn evolution of beam emittances for FCC-ee with a single frequency error and IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The first column shows the first positron bunch, the second column the second positron bunch, the third column the first electron bunch and the fourth column the second electron bunch. Emittances are shown in blue, and the corresponding initial values are indicated by red dashed lines.




Figure 6.43: Turn-by-turn evolution of luminosity for FCC-ee with a single frequency error and IP feedback. Luminosity computed for the positron bunches are shown in red, while luminosity computed for the electron bunches are shown in blue.

The evolution of the vertical position offset at each IP with turns for one FFQ vibration time series random seed is presented in Fig. 6.49. Small discrepancies are observed between the first and second bunch crossing. The high frequency noise varies on the scale of the revolution frequency so is experienced differently between the two




Figure 6.44: RMS vertical offset at the IP as a function of driven error frequency for FCC-ee. The horizontal axis shows the error frequency in hertz, and the vertical axis shows the RMS vertical offset across all IPs in metres. Results without IP feedback are shown in blue, and results with IP feedback are shown in orange. The reference case with no applied error is shown by the black dashed line.




Figure 6.45: Feedback suppression of vertical offset at the IP as a function of driven error frequency for FCC-ee. The horizontal axis shows the error frequency in hertz. The vertical axis shows the feedback suppression of the mean vertical offset across the IPs in decibels.

bunch crossings. Sharp spikes are observed due to particle losses. Significant offsets, on the order of the nominal beam size are observed at the IPs even using time series generated using a PSD including the effect of the COFB.

The evolution of the emittance with turns for one FFQ vibration time series random seed is presented in Fig. 6.50. In the presence of the error the vertical emittance is observed to sharply increase for all bunches. As this is not a single sinusoidal oscillation, the variation does not appear periodic. Over this time scale, it is unclear if these emittances will return to equilibrium or if a 3D flip-flop process is occurring towards the end, indicated by the shrinking longitudinal emittances of both the positron bunches and the increasing longitudinal emittances of both electron bunches.

In addition to the emittance evolution, the evolution of the luminosity with turns for




Figure 6.46: Emittance as a function of driven error frequency for FCC-ee. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The left column shows the mean of the positron bunches and the right column shows the mean of the electron bunches. Results without IP feedback are shown in blue, and results with IP feedback are shown in orange. The reference case with no applied error is shown by the black dashed line.




Figure 6.47: Luminosity as a function of driven error frequency for FCC-ee. The horizontal axis shows the error frequency in hertz, and the vertical axis shows the luminosity per IP. Results without IP feedback are shown in blue, and results with IP feedback are shown in orange. The reference case with no applied error is shown by the black dashed line.

one FFQ vibration time series random seed is presented in Fig. 6.51. Some significant discrepancies are observed in the computed luminosities for the electron and positron bunches due to the soft Gaussian approximation used in the beam-beam element. The emittances of the two beams were observed to differ greatly during the simulation.




Figure 6.48: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for FCC-ee with vibrated FFQs and no IP feedback. The top row shows the horizontal centroid position x , the middle row shows the vertical centroid position y , and the bottom row shows the longitudinal centroid position z . The left column corresponds to the first bunch crossing and the right column corresponds to the second bunch crossing. Positrons are shown in red and electrons are shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.

Furthermore, the assumption of a Gaussian profile may be lost as beam halos are populated with the FFQ vibrational kicks and residual vertical emittance and induced coupling. The luminosity is greatly degraded with respect to the nominal value.

6.6.6 Tracking with Realistic Final Focus Quadrupole Vibrations: IP Feedback Enabled

With the inclusion of IP feedback it is expected that the offsets at the IP are reduced and thus some of the luminosity loss should be recovered. Some degradation is still expected due to the emittance increase expected when machine errors and vibrations are included. Beams were initialised with the equilibrium emittances obtained from tracking with strong-strong beam-beam, optimised wiggler and quantum synchrotron radiation presented in Table 6.9.

The evolution of the beam centroids and beam sizes at IP A with feedback for one FFQ vibration time series random seed is presented in Fig. 6.40.

The evolution of the vertical position offset at each IP with turns for the same FFQ vibration time series random seed used in Section 6.6.5 is presented in Fig. 6.53. Small discrepancies are observed between the first and second bunch crossing. This is especially




Figure 6.49: Turn-by-turn evolution of IP vertical position offsets for FCC-ee with vibrated FFQs and no IP feedback. The top row shows the offset at IP A, the second row shows the offset at IP B, the third row shows the offset at IP C, and the bottom row shows the offset at IP D. Offsets during the first bunch crossing are shown in blue, and offsets during the second bunch crossing are shown in orange.

visible now with the feedback, as the lower frequency noise is highly suppressed and the high frequency noise may not be in phase between the two bunch crossings. Sharp spikes are observed due to particle losses. An initial large offset is visible, at IP C and IP D for this seed. This is quickly damped once the feedback system is enabled after 25 turns. The damping is not instant, as the feedback is ramped over a further 65 turns to mitigate transients. The offsets at the IP have been significantly reduced from an RMS offset of 28.2 nm without IP feedback (presented in Fig. 6.49) to an RMS offset of 10.6 nm with IP feedback. This remains far above the energy calibration requirement of 1 nm. Due to the computational cost of these simulations, fine optimisation of the feedback has not been possible.

The evolution of the emittance with turns for the same FFQ vibration time series random seed used in Section 6.6.5 is presented in Fig. 6.54. The form of the emittances with IP feedback enabled is very similar to the case with no IP feedback presented in Fig. 6.50. The horizontal and longitudinal emittances are particularly similar. The vertical emittance is reduced with IP feedback on compared to with IP feedback off. The emittances are still expected to increase as a result of the driven error and resulting residual vertical dispersion throughout the collider storage ring. FCC-ee is clearly highly sensitive to the tested errors.

The evolution of the luminosity with turns for the same FFQ vibration time series




Figure 6.50: Turn-by-turn evolution of beam emittances for FCC-ee with vibrated FFQs and no IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The first column shows the first positron bunch, the second column the second positron bunch, the third column the first electron bunch and the fourth column the second electron bunch. Emittances are shown in blue, and the corresponding initial values are indicated by red dashed lines.




Figure 6.51: Turn-by-turn evolution of luminosity for FCC-ee with vibrated FFQs and no IP feedback. Luminosity computed for the positron bunches are shown in red, while luminosity computed for the electron bunches are shown in blue.

random seed used in Section 6.6.5 is presented in Fig. 6.51. Discrepancies are still observed in the computed luminosities for the electron and positron bunches due to the soft Gaussian approximation used in the beam-beam element. The emittances of the two beams were observed to differ greatly during the simulation. Furthermore, the




Figure 6.52: Turn-by-turn evolution of beam centroid positions and beam-size envelopes at the IP for FCC-ee with vibrated FFQs and IP feedback. The top row shows the horizontal centroid position x , the middle row shows the vertical centroid position y , and the bottom row shows the longitudinal centroid position z . The left column corresponds to the first bunch crossing and the right column corresponds to the second bunch crossing. Positrons are shown in red and electrons are shown in blue. Solid lines indicate the beam centroids, and the shaded regions indicate the corresponding beam-size envelopes.

assumption of a Gaussian profile may be lost as beam halos are populated with the FFQ vibrational kicks and residual vertical emittance and induced coupling. The luminosity remains greatly degraded compared to the no error case, however the luminosity has been recovered with respect to the no feedback case.

A comparison of the luminosity with and without IP feedback for multiple FFQ vibrational error seeds is presented in Fig. 6.56. Here the luminosities are the mean across the four computed values (for two positron and two electron bunches). The feedback system is clearly effective in mitigating luminosity degradation in some cases, and far less effective in others. To test the reason for the failure, the IP offsets were considered. A comparison of the vertical position offset at each IP with and without IP feedback for multiple FFQ vibrational error seeds is presented in Fig. 6.57. It is clear that the feedback system is highly effective at reducing the IP offset across all tested seeds. This indicates that the collision alignment is not the issue in the reduction of the luminosity. Other factors are clearly involved, such as the horizontal alignment/longitudinal timing, the vertical angle offset or the emittance growth. Given the complex form of the emittances, it seems likely that the vertical emittance blow-up and 3D flip-flop effect are having significant detrimental effects.




Figure 6.53: Turn-by-turn evolution of IP vertical position offsets for FCC-ee with vibrated FFQs and IP feedback. The top row shows the offset at IP A, the second row shows the offset at IP B, the third row shows the offset at IP C, and the bottom row shows the offset at IP D. Offsets during the first bunch crossing are shown in blue, and offsets during the second bunch crossing are shown in orange.




Figure 6.54: Turn-by-turn evolution of beam emittances for FCC-ee with vibrated FFQs and IP feedback. The top row shows the horizontal emittance ε_x , the middle row shows the vertical emittance ε_y , and the bottom row shows the longitudinal emittance ε_z . The first column shows the first positron bunch, the second column the second positron bunch, the third column the first electron bunch and the fourth column the second electron bunch. Emittances are shown in blue, and the corresponding initial values are indicated by red dashed lines.




Figure 6.55: Turn-by-turn evolution of luminosity for FCC-ee with vibrated FFQs and IP feedback. Luminosity computed for the positron bunches are shown in red, while luminosity computed for the electron bunches are shown in blue.




Figure 6.56: Turn-by-turn evolution of luminosity for FCC-ee with vibrated FFQs and IP feedback across many random seeds with and without IP feedback. Luminosities with IP feedback disabled are shown in red, while luminosities with IP feedback enabled are shown in blue.




Figure 6.57: Turn-by-turn evolution of IP A vertical position offset for FCC-ee with vibrated FFQs and IP feedback across many random seeds with and without IP feedback. IP A vertical position offset with IP feedback disabled are shown in red, while offsets with IP feedback enabled are shown in blue.

6.7 Summary

A number of vital aspects of the FCC-ee IP feedback system have been identified and tested. An IP feedback system to mitigate beam offsets at the IPs has been proposed for FCC-ee.

A variety of IP feedback input signals have been tested. Beam–beam deflection, luminosity and beamstrahlung radiation signals are the leading candidates. These would be measured at the IP BPMs, the luminosity calorimeters and by a beamstrahlung monitor. The deflection measurements at the IP BPMs appear the most promising as a vector signal with well understood detection hardware. The luminosity signal is also feasible, though remains a scalar signal so luminosity dithering would have to be employed. Beamstrahlung is a very exciting signal, however measuring radiation power on the order of several hundred kilowatts is a major technical challenge.

The requirements of the IP feedback system for FCC-ee have also been outlined. The stability of the beams in the presence of machine offsets was studied. Compared to the CDR lattice with four IPs, it appears that the beam remains stable even with significant IP offset. Further to a stable beam, for physics performance strict tolerances are enforced. The strictest tolerance comes from the measurement of the beam energy for energy calibration. This tolerance is currently set at a 1 nm offset, which is highly challenging.

Self-consistent 6D tracking simulations of the proposed FCC-ee feedback system have been developed. A model of FCC-ee, using a GHC lattice at the ZH working point, was used in Xsuite. The behaviour of the lattice with radiation and two distinct methods to generate vertical emittance for simulations were tested. The sensitivity of the FCC-ee lattice to sextupole offsets was noted, and the use of a rotated wiggler was chosen to generate vertical emittance. Strong–strong beam–beam elements were then included at each IP for two positron and two electron bunches. Equilibrium emittances including all of the effects of beam–beam, quantum synchrotron radiation and optimised wiggler were calculated.

The key elements of the proposed FCC-ee feedback system were included. Knobs for the IP feedback vertical position offset correction were developed in Xsuite based on discussions with the optics designers. The linearity of this knob and its impact on beta functions and dispersions were considered. With the current knob, vertical dispersion leakage may be an issue, and it would be worth considering the use of more correctors on each side of the IP, especially once the impact of the experimental solenoid is included. The BPMs, offset calculation and feedback algorithm were also implemented.

Tracking simulations were then performed to identify the performance of the feedback system. Simulations were first performed with driven sinusoidal errors. The impact on the IP vertical position offsets, emittances and luminosity with error frequency was

studied with and without IP feedback enabled. The system was observed to be effective at mitigating the IP vertical position offset and luminosity degradation. Simulations were then also performed with vibrational errors applied to the final focus quadrupoles. Here, the effect of the feedback was less clear. The offsets at the IP were well corrected, though in some tested seeds, the luminosity remained highly degraded. This appeared to be an emittance issue, where due to the errors in the lattice some seeds exhibited the 3D flip-flop instability.

There are a number of key limitations of this study. This study was performed on a GHC ZH FCC-ee lattice, which was negatively impacted by an XZ instability. For performance reasons, lattices with longer damping times were not chosen. Furthermore, no solenoid exists in the current lattice models, and as such the feedback knobs are trivial compared to the real machine. A highly simplified model of the Closed Orbit feedback system was implemented, and future simulations should include a more realistic form of the feedback. The choice of the wiggler is also far from optimal; the polarisation wiggler is located at a point in the lattice that is subject to very high beta beating when beam-beam is included. As such, significant vertical emittance blow-up is observed with beam-beam. With a more realistic distributed source of errors this is not expected. The aim was therefore to use the offset sextupoles as a distributed source of errors, however correction knobs were not available to correct the optics distortions and it was beyond the scope of this thesis to develop such corrections.

Finally, it was not feasible to optimise the system due to the heavy computational cost of these studies. All simulations would benefit from an increase in the number of macroparticles and number of slices in the beam-beam element. This would be too costly to perform simulations on CPU. GPU simulations offer the opportunity for significant speed-up, however the current synchrotron radiation model is not performant on GPU, so at the time of writing this is not possible. It is hoped that this model may be implemented in the near future for further development of this system.

Chapter 7

Conclusions

7.1 Summary

CERN’s Future Circular Lepton Collider, FCC-ee, is a proposed circular electron–positron collider designed to study the heaviest known Standard Model particles at the highest achievable luminosities. To deliver unprecedented luminosity and precision, strict demands are placed on beam stability at the interaction points (IPs). With IP beam sizes on the nanometre scale, operational performance is not limited by optics design alone, but by the ability to control beam motion in the presence of strong beam–beam forces and time-dependent disturbances. In this regime, passive stability and static correction are insufficient, and fast beam-based IP feedback becomes a fundamental requirement for achieving and maintaining design performance, motivating a detailed study of IP feedback for FCC-ee as an integral component of the collider design.

KEK’s SuperKEKB represents the current state of the art luminosity frontier lepton collider operating in the beam–beam-limited regime. SuperKEKB routinely collides beams with nanometre-scale vertical beam sizes and relies on IP feedback as an essential component of stable luminosity production. Operational experience at SuperKEKB demonstrates not only the feasibility but the necessity of feedback in collision, while exposing the complexity of real machine conditions. As such, SuperKEKB provides a unique experimental benchmark against which models of IP feedback can be validated, and a critical reference for assessing the feasibility of feedback systems for future colliders such as FCC-ee, where no comparable operational experience yet exists.

This thesis develops and applies a consistent framework for the study of IP feedback in current and future circular lepton colliders. Analytic modelling, experimental observations from SuperKEKB, and realistic end-to-end tracking simulations are combined to assess feedback performance under conditions representative of both present and future machines. This work provides an approach to modelling IP feedback for FCC-ee,

enabling quantitative studies of performance, limitations, and sensitivity to machine parameters.

An analytic framework has been developed to study the performance of IP feedback in a form that makes the underlying beam dynamics explicit while remaining computationally lightweight, designed to enable efficient simulation over thousands of turns to enable modelling of Hz scale errors at FCC-ee. This approach has been used to investigate the relative performance of different feedback strategies, including proportional, integral, and derivative components. A purely integral controller has been identified as the most effective choice in the operating regime relevant to FCC-ee. On this basis, the beam dynamics in collision have been analysed across a range of studies involving beam-beam strength, detuning, feedback gain, delay, and turn-by-turn vs bunch-by-bunch feedback, allowing systematic exploration of their respective roles. Analysis of configurations with up to four IPs enables direct application to FCC-ee and provides a clear reference for interpreting the trends observed in fully self-consistent numerical simulations.

Experimental data from studies of SuperKEKB’s iBump IP feedback system provide direct insight into the operational behaviour of beam-beam deflection based IP feedback in collision. Analysis of the iBump vertical offset target, a free correction parameter, showed systematic variations with time, beam current, and machine conditions, with a direct impact on feedback performance in operation. The data demonstrates the presence of slow temporal drifts and current-dependent behaviour that must be accommodated in any realistic feedback design. These operational observations inform the choice and tuning of feedback parameters, and contribute directly to understanding how IP feedback systems should be operated in future colliders such as FCC-ee.

Analysis of operational data from SuperKEKB machine studies revealed a potentially hazardous dependence of feedback behaviour on beam current. This behaviour has clear implications for FCC-ee operation under the bootstrap injection scheme, where beams must remain in collision during filling. The beam current will therefore vary while collisions are maintained, making current-dependent feedback effects a direct operational concern. While stable operation was observed to be achievable over finite time intervals, the results demonstrate that feedback performance cannot be assumed to be uniform across operating conditions, and that current-dependent effects must be explicitly accounted for in future design and optimisation studies.

A major outcome of this work is the development of a simulation framework for SuperKEKB in Xsuite. The successful conversion of the SuperKEKB lattices provides a flexible platform that has supported a range of simulation studies beyond the scope of this thesis. This capability enabled detailed studies of SuperKEKB’s iBump IP feedback system in self-consistent simulations using full lattice tracking. The iBump feedback system has been implemented together with the full detector solenoidal field and its

correction, radiation effects, beam–beam interaction, and distributed sources of vertical emittance. This framework enabled systematic investigation of feedback performance in response to disturbances across a range of frequencies, and it was observed that the iBump system is effective at suppressing the low-frequency noise dominant in the vibrations of the final-focus quadrupoles and the resulting beam orbit oscillations.

Candidate IP feedback signals for FCC-ee have been identified through simulations of FCC-ee, experience from previous lepton colliders and operational insights from SuperKEKB. The studies focus on observables that are directly sensitive to beam–beam offsets while remaining compatible with the measurement constraints imposed by the FCC-ee machine-detector interface. Signal sensitivity has been evaluated using both single-collision particle-in-cell simulations with GUINEA-PIG and multi-turn tracking simulations with Xsuite. In parallel, the tolerances that FCC-ee’s IP feedback system must meet have been collated. Using Xsuite tracking simulations, it was identified that current FCC-ee lattice designs featured relaxed stability tolerances compared to those at the time of the Conceptual Design Report. Together, these studies demonstrate that useful feedback information can be extracted and that the signal and tolerance requirements are broadly compatible with existing and foreseeable instrumentation performance.

The central result of this thesis is the proposal and simulation of an IP feedback system for FCC-ee, including realistic lattice configurations, and beam–beam effects. Two complementary studies have been performed to assess feedback performance: controlled sinusoidal excitation of beam offsets over a range of frequencies to characterise the frequency response of the system, and simulations with realistic final-focus quadrupole motion using vibration spectra representative of FCC-ee operating conditions. The driven studies show that IP feedback can effectively suppress beam–beam-induced offsets over a broad frequency range, whereas simulations with realistic vibrations exhibit significantly poorer performance, in particular due to the onset of 3D flip–flop behaviour. While current FCC-ee lattice designs offer improved robustness compared to earlier lattices, these results demonstrate that achieving the stringent one-nanometre offset tolerance required for precision energy calibration remains extremely challenging under realistic operating conditions.

Taken together, this work provides an integrated and experimentally informed study of IP feedback for FCC-ee. By combining analytic modelling, operational insight from SuperKEKB, and realistic end-to-end simulation, a coherent framework has been developed for assessing feedback behaviour, limitations, and design sensitivity in circular lepton colliders. This thesis therefore establishes a quantitative foundation on which further optimisation, design refinement, and operational studies of IP feedback systems for FCC-ee and future machines can be built.

7.2 Further Work

The principal limitation of the studies presented in this thesis arises from the level of realism achievable in large-scale beam–beam and feedback simulations. While the models employed capture the essential dynamics of IP feedback, both SuperKEKB and FCC-ee studies are restricted by simplified lattice descriptions, idealised correction schemes, and limited macroparticle statistics. Addressing these limitations represents the most important direction for further work.

For SuperKEKB, the simulations performed rely on lattices that reproduce the nominal optics and beam–beam parameters, but do not incorporate realistic machine errors and operational correction procedures. In practice, orbit feedback systems, magnet misalignments, residual coupling, and time-dependent imperfections play a significant role in determining the beam conditions presented at the IP. In addition, the number of macroparticles used in beam–beam simulations is constrained by computational cost, limiting the fidelity. Extending the present studies to lattices with realistic alignment errors, beam-based correction algorithms, and increased macroparticle statistics would allow a more direct comparison between simulation and operational performance.

For FCC-ee, limitations are more pronounced. At present, no lattice exists simultaneously including a realistic interaction region solenoid, error distributions, and correction strategies for long-term tracking with feedback. The absence of such a lattice constrains the ability to design and evaluate physically realistic IP feedback bumps in the solenoidal field around the IP. In addition, the requirement to model the full collision pattern, involving eight beam–beam interactions per turn, places particularly stringent constraints on the usable number of macroparticles. The development of an FCC-ee lattice incorporating a solenoid, with consistent optics correction is therefore a prerequisite for fully quantitative IP feedback studies.

The computational cost associated with such simulations remains substantial, and many of the studies presented rely on reduced models or limited statistics. Future improvements in the performance of tracking with quantum synchrotron radiation on graphical processing units (GPUs) are expected to significantly alleviate these constraints. Once high-fidelity radiation tracking becomes routinely available on GPUs, large-scale studies combining realistic lattices, feedback systems, and beam–beam effects will become feasible, enabling systematic optimisation of IP feedback performance for FCC-ee.

The conclusions of this thesis may therefore be viewed as establishing the conceptual and modelling framework for IP feedback, rather than providing a final quantitative design. The extension to fully realistic lattices and computationally intensive simulations represents a natural and necessary next step toward the deployment of IP feedback as an integral component of future circular lepton colliders.

Bibliography

- [1] R. L. Workman et al. ‘Review of Particle Physics’. In: *Progress of Theoretical and Experimental Physics* 2022.8 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [2] Cush. *Standard Model of Elementary Particles Anti.svg*. Wikimedia Commons. An SVG diagram of the Standard Model including anti-particles. 2018. URL: commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles_Anti.svg.
- [3] S. Weinberg. ‘A Model of Leptons’. In: *Physical Review Letters* 19 (1967), pp. 1264–1266. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [4] D. J. Gross and F. Wilczek. ‘Ultraviolet Behavior of Non-Abelian Gauge Theories’. In: *Physical Review Letters* 30 (1973), pp. 1343–1346. DOI: [10.1103/PhysRevLett.30.1343](https://doi.org/10.1103/PhysRevLett.30.1343).
- [5] C. S. Wu et al. ‘Experimental Test of Parity Conservation in Beta Decay’. In: *Physical Review* 105 (1957), pp. 1413–1415. DOI: [10.1103/PhysRev.105.1413](https://doi.org/10.1103/PhysRev.105.1413).
- [6] P. W. Higgs. ‘Broken Symmetries and the Masses of Gauge Bosons’. In: *Physical Review Letters* 13 (1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [7] F. Englert and R. Brout. ‘Broken Symmetry and the Mass of Gauge Vector Mesons’. In: *Physical Review Letters* 13 (1964), pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [8] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble. ‘Global Conservation Laws and Massless Particles’. In: *Physical Review Letters* 13 (1964), pp. 585–587. DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585).
- [9] S. Weinberg. ‘The Search for Unity: Notes for a History of Quantum Field Theory’. In: *Daedalus* 106 (1977), pp. 17–35. URL: www.jstor.org/stable/20024506.
- [10] A. D. Sakharov. ‘Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe’. In: *JETP Letters* 5 (1967), pp. 24–27. URL: jetpletters.ru/ps/1643/article_25089.pdf.
- [11] G. Bertone, D. Hooper and J. Silk. ‘Particle dark matter: Evidence, candidates and constraints’. In: *Physics Reports* 405 (2005), pp. 279–390. DOI: [10.1016/j.physrep.2004.08.031](https://doi.org/10.1016/j.physrep.2004.08.031).
- [12] Y. Fukuda et al. ‘Evidence for oscillation of atmospheric neutrinos’. In: *Physical Review Letters* 81 (1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).
- [13] L. Evans and P. Bryant. ‘LHC Machine’. In: *Journal of Instrumentation* 3.08 (2008), S08001. DOI: [10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001).

- [14] R. Steerenberg. *Accelerator Report: Crescendo at the LHC Following the First Stable Beams at 6.8 TeV*. CERN News. 2023. URL: home.cern/news/news/accelerators/accelerator-report-crescendo-lhc-following-first-stable-beams-68-tev.
- [15] M. Dam. *Experimental Physics at Lepton Colliders*. Lecture presented at the CERN Summer Student Programme. 2019. URL: indico.cern.ch/event/934633/.
- [16] K. Akai, K. Furukawa, H. Koiso et al. ‘SuperKEKB Collider’. In: *Nuclear Instruments and Methods in Physics Research Section A* 907 (2018), pp. 188–199. DOI: [10.1016/j.nima.2018.08.017](https://doi.org/10.1016/j.nima.2018.08.017).
- [17] T. Feder. ‘Japan accelerator pursues nanobeams to boost luminosity’. In: *Physics Today* (2025). DOI: [10.1063/pt.imsn.ildq](https://doi.org/10.1063/pt.imsn.ildq).
- [18] O. Brüning and L. Rossi. ‘The High Luminosity Large Hadron Collider–HL-LHC’. In: *The High Luminosity Large Hadron Collider: New Machine for Illuminating the Mysteries of the Universe*. World Scientific, 2024, pp. 1–53. DOI: [10.1142/9789811278952_0001](https://doi.org/10.1142/9789811278952_0001).
- [19] G. Apollinari et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report*. Tech. rep. CERN-2017-007-M. Geneva: CERN, 2017. DOI: [10.23731/CYRM-2017-004](https://doi.org/10.23731/CYRM-2017-004).
- [20] K. Meier et al. *2020 Update of the European Strategy for Particle Physics*. CERN Report CERN-ESU-015. Geneva: CERN, 2020. DOI: [10.17181/CERN.JSC6.W89E](https://doi.org/10.17181/CERN.JSC6.W89E).
- [21] K. Meier et al. *European Strategy for Particle Physics: 2026 Update — Recommendations by the European Strategy Group*. CERN Report CERN-ESU-2025-002. Ascona: CERN, 2025. DOI: [10.17181/CERN.423R.S20Z](https://doi.org/10.17181/CERN.423R.S20Z).
- [22] S. Ritz et al. *Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context*. P5 Report. Washington, DC: High Energy Physics Advisory Panel, 2014. URL: science.osti.gov/-/media/hep/hepap/pdf/May-2014/FINAL_P5_Report_053014.pdf.
- [23] A. Abada et al. ‘FCC-ee: The Lepton Collider’. In: *European Physical Journal Special Topics* 228.2 (2019), pp. 261–623. DOI: [10.1140/epjst/e2019-900045-4](https://doi.org/10.1140/epjst/e2019-900045-4).
- [24] J. Gao et al. *CEPC Conceptual Design Report: Volume 1 — Accelerator*. CEPC Conceptual Design Report. Beijing: Institute of High Energy Physics, 2018. URL: cepc.ihep.ac.cn/preCDR/volume.html.
- [25] J. Gao et al. ‘CEPC Technical Design Report: Volume 1 — Accelerator’. In: *Radiation Detection Technology and Methods* 8 (2024), pp. 1–1105. DOI: [10.1007/s41605-024-00463-y](https://doi.org/10.1007/s41605-024-00463-y).
- [26] P. N. Burrows et al. ‘The Compact Linear Collider (CLIC): 2018 Summary Report’. In: *CERN Yellow Reports: Monographs* 2018 (2018), pp. 1–98. DOI: [10.23731/CYRM-2018-002](https://doi.org/10.23731/CYRM-2018-002).
- [27] T. Behnke et al. *The International Linear Collider Technical Design Report: Volume 3 — Accelerator*. ILC Technical Design Report. Tsukuba: KEK, 2013. URL: www.linearcollider.org/ILC/Publications/Technical-Design-Report.

- [28] A. Abada et al. ‘FCC-hh: The Hadron Collider’. In: *European Physical Journal Special Topics* 228.4 (2019), pp. 755–1107. DOI: [10.1140/epjst/e2019-900087-0](https://doi.org/10.1140/epjst/e2019-900087-0).
- [29] J. Delahaye et al. ‘Muon Colliders’. In: *Reviews of Accelerator Science and Technology* 10 (2019), pp. 137–159. DOI: [10.1142/9789814651493_0007](https://doi.org/10.1142/9789814651493_0007).
- [30] E. Esarey, C. B. Schroeder and W. P. Leemans. ‘Physics of Laser-Driven Plasma-Based Electron Accelerators’. In: *Reviews of Modern Physics* 81 (2009), pp. 1229–1285. DOI: [10.1103/RevModPhys.81.1229](https://doi.org/10.1103/RevModPhys.81.1229).
- [31] M. Benedikt and F. Zimmermann. ‘Future Circular Collider: Integrated Programme and Feasibility Study’. In: *Frontiers in Physics* 10 (2022), p. 888078. DOI: [10.3389/fphy.2022.888078](https://doi.org/10.3389/fphy.2022.888078).
- [32] A. Abada et al. ‘FCC Physics Opportunities’. In: *European Physical Journal C* 79.6 (2019), p. 474. DOI: [10.1140/epjc/s10052-019-6904-3](https://doi.org/10.1140/epjc/s10052-019-6904-3).
- [33] M. Bicer et al. ‘First Look at the Physics Case of TLEP’. In: *JHEP* 01 (2014), p. 164. DOI: [10.1007/JHEP01\(2014\)164](https://doi.org/10.1007/JHEP01(2014)164). arXiv: [1308.6176](https://arxiv.org/abs/1308.6176) [hep-ex].
- [34] O. Domínguez and F. Zimmermann. ‘Beam Parameters and Luminosity Time Evolution for an 80-km VHE-LHC’. In: *Proceedings of the 4th International Particle Accelerator Conference (IPAC 2013)*. Paper TUPFI042. Shanghai, China: JACoW Publishing, 2013, pp. 1442–1444. URL: proceedings.jacow.org/IPAC2013/papers/tupfi042.pdf.
- [35] M. Benedikt et al. ‘Future Circular Collider Feasibility Study Report: Volume 1, Physics, Experiments, Detectors’. In: *European Physical Journal C* 85.12 (2025), p. 1468. DOI: [10.1140/epjc/s10052-025-15077-x](https://doi.org/10.1140/epjc/s10052-025-15077-x).
- [36] M. Benedikt et al. ‘Future Circular Collider Feasibility Study Report: Volume 2, Accelerators, Technical Infrastructure and Safety’. In: *European Physical Journal Special Topics* 234.19 (2025), pp. 5713–6197. DOI: [10.1140/epjs/s11734-025-01967-4](https://doi.org/10.1140/epjs/s11734-025-01967-4).
- [37] M. Benedikt et al. ‘Future Circular Collider Feasibility Study Report: Volume 3, Civil Engineering, Implementation and Sustainability’. In: *European Physical Journal Special Topics* 234.17 (2025). See also erratum: [10.1140/epjs/s11734-025-02043-7](https://doi.org/10.1140/epjs/s11734-025-02043-7), pp. 5113–5383. DOI: [10.1140/epjs/s11734-025-01958-5](https://doi.org/10.1140/epjs/s11734-025-01958-5).
- [38] P. Craievich et al. ‘The FCC-ee Pre-injector Complex’. In: *Proceedings of the 13th International Particle Accelerator Conference (IPAC 2022)*. Paper WEPOPT063. Bangkok, Thailand: JACoW Publishing, 2022, pp. 2007–2010. DOI: [10.18429/JACoW-IPAC2022-WEPOPT063](https://doi.org/10.18429/JACoW-IPAC2022-WEPOPT063).
- [39] B. Dalena et al. ‘Status of the High Energy Booster of the Lepton Option of the Future Circular Collider’. In: *Proceedings of Science (PoS)*. Vol. 414. ICHEP2022. 2022, p. 042. DOI: [10.22323/1.414.0042](https://doi.org/10.22323/1.414.0042).
- [40] M. Boscolo et al. ‘Mechanical Model for the FCC-ee Interaction Region’. In: *EPJ Techniques and Instrumentation* 10 (2023), p. 16. DOI: [10.1140/epjti/s40485-023-00103-7](https://doi.org/10.1140/epjti/s40485-023-00103-7).
- [41] A. Blondel and P. Janot. ‘FCC-ee Overview: New Opportunities Create New Challenges’. In: *European Physical Journal Plus* 137.1 (2022), p. 92. DOI: [10.1140/epjp/s13360-021-02154-9](https://doi.org/10.1140/epjp/s13360-021-02154-9).

- [42] D. d’Enterria. ‘Physics at the FCC-ee’. In: *Particle Physics at the Year of Light: Proceedings of the Seventeenth Lomonosov Conference on Elementary Particle Physics*. Singapore: World Scientific, 2017, pp. 182–191. DOI: [10.1142/9789813224568_0028](https://doi.org/10.1142/9789813224568_0028).
- [43] Y. Ohnishi et al. ‘Accelerator Design at SuperKEKB’. In: *Progress of Theoretical and Experimental Physics* 2013.3 (2013), 03A011. DOI: [10.1093/ptep/pts083](https://doi.org/10.1093/ptep/pts083).
- [44] T. Abe et al. ‘Achievements of KEKB’. In: *Progress of Theoretical and Experimental Physics* 2013.3 (2013), 03A001. DOI: [10.1093/ptep/pts102](https://doi.org/10.1093/ptep/pts102).
- [45] T. Abe et al. *Belle II Technical Design Report*. KEK Report 2010-1. Tsukuba: KEK, 2010. arXiv: [1011.0352](https://arxiv.org/abs/1011.0352) [[physics.ins-det](#)].
- [46] T. Honda et al. ‘Operational Status of Photon Factory Light Sources’. In: *Proceedings of the 12th International Particle Accelerator Conference (IPAC 2021)*. Paper MOPAB093. Campinas, Brazil: JACoW Publishing, 2021, pp. 350–353. DOI: [0.18429/JACoW-IPAC2021-MOPAB093](https://doi.org/0.18429/JACoW-IPAC2021-MOPAB093).
- [47] Y. Suetsugu et al. ‘Mitigating the Electron Cloud Effect in the SuperKEKB Positron Ring’. In: *Physical Review Accelerators and Beams* 22 (2019), p. 023201. DOI: [10.1103/PhysRevAccelBeams.22.023201](https://doi.org/10.1103/PhysRevAccelBeams.22.023201).
- [48] B. Wang. *The Belle II Experiment and SuperKEKB Upgrade*. 2015. arXiv: [1511.09434](https://arxiv.org/abs/1511.09434) [[hep-ex](#)].
- [49] European Commission. *Europe-America-Japan Accelerator Development and Exchange Programme (EAJADE)*. CORDIS EU Research Results. 2023. DOI: [10.3030/101086276](https://doi.org/10.3030/101086276).
- [50] E. J. N. Wilson. *An Introduction to Particle Accelerators*. Oxford University Press, 2001.
- [51] E. Bravin. *Transverse Emittance*. 2020. arXiv: [2005.05770](https://arxiv.org/abs/2005.05770) [[physics.acc-ph](#)]. URL: cds.cern.ch/record/2723977.
- [52] D. A. Edwards and M. J. Syphers. *An Introduction to the Physics of High Energy Accelerators*. Wiley, 1993.
- [53] M. Sands. *The Physics of Electron Storage Rings: An Introduction*. Tech. rep. SLAC-121. Stanford, CA: Stanford Linear Accelerator Center, 1970. URL: www.osti.gov/servlets/purl/1453884.
- [54] P. Raimondi, D. N. Shatilov and M. Zobov. *Beam–Beam Issues for Colliding Schemes with Large Piwinski Angle and Crabbed Waist*. Tech. rep. INFN Laboratori Nazionali di Frascati, 2007. DOI: [10.48550/arXiv.physics/0702033](https://doi.org/10.48550/arXiv.physics/0702033). arXiv: [physics/0702033](https://arxiv.org/abs/physics/0702033) [[physics.acc-ph](#)].
- [55] M. Zobov et al. *Crab Waist Collision Studies for e^+e^- Factories*. Tech. rep. INFN Laboratori Nazionali di Frascati, 2008. DOI: [10.48550/arXiv.0802.2667](https://doi.org/10.48550/arXiv.0802.2667). arXiv: [0802.2667](https://arxiv.org/abs/0802.2667) [[physics.acc-ph](#)].
- [56] M. Zobov et al. ‘Test of Crab-Waist Collisions at the DAΦNE Phi Factory’. In: *Physical Review Letters* 104 (2010), p. 174801. DOI: [10.1103/PhysRevLett.104.174801](https://doi.org/10.1103/PhysRevLett.104.174801).
- [57] W. Herr and T. Pieloni. *Beam–Beam Effects*. 2016. arXiv: [1601.05235](https://arxiv.org/abs/1601.05235) [[physics.acc-ph](#)].

- [58] A. W. Chao et al., eds. *Handbook of Accelerator Physics and Engineering*. World Scientific, 2023.
- [59] M. Bassetti and G. A. Erskine. *Closed Expression for the Electrical Field of a Two-Dimensional Gaussian Charge*. Tech. rep. CERN-ISR-TH/80-06. CERN, 1980. URL: cds.cern.ch/record/122227/.
- [60] K. Yokoya and H. Koiso. ‘Tune Shift of Coherent Beam–Beam Oscillations’. In: *Particle Accelerators* 27 (1990), pp. 181–186.
- [61] K. Yokoya and P. Chen. ‘Beam–Beam Phenomena in Linear Colliders’. In: *Frontiers of Particle Beams: Intensity Limitations*. Vol. 400. Lecture Notes in Physics. Springer, 1992, pp. 415–445. DOI: [10.1007/3-540-55250-2_37](https://doi.org/10.1007/3-540-55250-2_37).
- [62] W. Herr and B. Muratori. ‘Concept of Luminosity’. In: *CERN Accelerator School: Intermediate Accelerator Physics*. CERN Yellow Reports: School Proceedings. Geneva: CERN, 2006, pp. 361–376. DOI: [10.5170/CERN-2006-002.361](https://doi.org/10.5170/CERN-2006-002.361).
- [63] F. Zimmermann. *1998 SLC Luminosity and Pinch Enhancement*. Tech. rep. SLAC-CN-418. SLAC, 1998.
- [64] H. Burkhardt and R. Kleiss. ‘Monte Carlo Simulation of Radiative Bhabha Scattering in the Very Forward Direction’. In: *Comput. Phys. Commun.* 81 (1994), pp. 372–380. DOI: [10.1016/0010-4655\(94\)90085-X](https://doi.org/10.1016/0010-4655(94)90085-X). arXiv: [hep-ph/9401333](https://arxiv.org/abs/hep-ph/9401333).
- [65] H. Wiedemann. ‘Resonances’. In: *Particle Accelerator Physics*. 4th ed. Cham: Springer, 2015, pp. 539–564. DOI: [10.1007/978-3-319-18317-6_16](https://doi.org/10.1007/978-3-319-18317-6_16).
- [66] V. Shiltsev. ‘Space–Time Ground Diffusion: The ATL Law for Accelerators’. In: *Proceedings of the 4th International Workshop on Accelerator Alignment (IWAA ’95)*. Tsukuba, Japan, 1995, pp. 352–381. URL: www-group.slac.stanford.edu/met/IWAA/TOC_S/Papers/VShil95.pdf.
- [67] A. Seryi. *Ground Motion and Vibration Issues for Accelerators*. Tech. rep. SLAC-PUB-10138. SLAC, 2003.
- [68] A. V. Oppenheim, A. S. Willsky and S. H. Nawab. *Signals and Systems*. 2nd ed. Upper Saddle River, NJ: Prentice Hall, 1997.
- [69] J. G. Proakis and D. G. Manolakis. *Digital Signal Processing: Principles, Algorithms, and Applications*. 4th ed. Upper Saddle River, NJ: Prentice Hall, 2006.
- [70] K. J. Åström and T. Häggglund. *PID Controllers: Theory, Design, and Tuning*. 2nd ed. Research Triangle Park, NC: Instrument Society of America, 1995.
- [71] K. Ogata. *Modern Control Engineering*. 5th ed. Prentice Hall, 2010.
- [72] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. 4th ed. Wiley, 2011.
- [73] G. F. Franklin, J. D. Powell and M. L. Workman. *Digital Control of Dynamic Systems*. 4th ed. Pearson, 2015.
- [74] T. Nakamura. ‘Transverse and Longitudinal Bunch-by-Bunch Feedback for Storage Rings’. In: *Proceedings of the 9th International Particle Accelerator Conference (IPAC’18)*. Paper TUZGBD2. Vancouver, BC, Canada, 2018, pp. 3911–3914. DOI: [10.18429/JACoW-IPAC2018-TUZGBD2](https://doi.org/10.18429/JACoW-IPAC2018-TUZGBD2).

- [75] D. Dinev. ‘Closed-Orbit Correction in Synchrotrons’. In: *Physics of Particles and Nuclei* 28.4 (1997), pp. 398–417.
- [76] J. Resta-López, P. N. Burrows and G. Christian. ‘Luminosity performance studies of the Compact Linear Collider with intra-train feedback system at the interaction point’. In: *Journal of Instrumentation* 5.09 (2010), P09007. DOI: [10.1088/1748-0221/5/09/P09007](https://doi.org/10.1088/1748-0221/5/09/P09007).
- [77] R. J. Apsimon et al. ‘Design and operation of a prototype interaction point beam collision feedback system for the International Linear Collider’. In: *Phys. Rev. Accel. Beams* 21.12 (2018), p. 122802. DOI: [10.1103/PhysRevAccelBeams.21.122802](https://doi.org/10.1103/PhysRevAccelBeams.21.122802).
- [78] W. Koska et al. ‘Beam–Beam Deflection as a Beam Tuning Tool at the SLAC Linear Collider’. In: *Nuclear Instruments and Methods in Physics Research A* 286 (1990), pp. 32–37. DOI: [10.1016/0168-9002\(90\)90203-I](https://doi.org/10.1016/0168-9002(90)90203-I).
- [79] Y. Funakoshi et al. ‘Interaction Point Orbit Feedback System at SuperKEKB’. In: *Proceedings of the 6th International Particle Accelerator Conference (IPAC’15)*. Paper MOPHA054. Richmond, VA, USA, 2015, pp. 921–923. DOI: [10.18429/JACoW-IPAC2015-MOPHA054](https://doi.org/10.18429/JACoW-IPAC2015-MOPHA054).
- [80] A. S. Fisher et al. ‘Commissioning the Fast Luminosity Dither for PEP-II’. In: *Proceedings of the 2007 Particle Accelerator Conference (PAC’07)*. Albuquerque, New Mexico, USA, 2007, pp. 4165–4167. DOI: [10.1109/PAC.2007.4440192](https://doi.org/10.1109/PAC.2007.4440192).
- [81] L. Hendrickson et al. ‘Feedback Systems for Linear Colliders’. In: *Proceedings of the 1999 Particle Accelerator Conference (PAC’99)*. Vol. 1. 1999, pp. 338–342.
- [82] M. Masuzawa et al. ‘Early Commissioning of the Luminosity Dither Feedback for SuperKEKB’. In: *Proceedings of the 7th International Beam Instrumentation Conference (IBIC’18)*. Paper TUPC13. Shanghai, China, 2018, pp. 328–331. DOI: [10.18429/JACoW-IBIC2018-TUPC13](https://doi.org/10.18429/JACoW-IBIC2018-TUPC13).
- [83] Å. Björck. *Numerical Methods for Least Squares Problems*. Philadelphia, PA: SIAM, 1996. DOI: [10.1137/1.9781611971484](https://doi.org/10.1137/1.9781611971484).
- [84] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Washington, DC: W. H. Winston & Sons, 1977.
- [85] P. Kicsiny. ‘Beam–Beam Effects in Future High-Energy Circular Lepton Colliders’. PhD thesis. Lausanne: EPFL, 2024.
- [86] K. Hirata. ‘Analysis of Beam–Beam Interaction with a Large Crossing Angle’. In: *Physical Review Letters* 74.12 (1995), pp. 2228–2231. DOI: [10.1103/PhysRevLett.74.2228](https://doi.org/10.1103/PhysRevLett.74.2228).
- [87] M. P. Zorzano and F. Zimmermann. ‘Simulations of coherent beam–beam modes at the Large Hadron Collider’. In: *Phys. Rev. ST Accel. Beams* 3.4 (2000), p. 044401. DOI: [10.1103/PhysRevSTAB.3.044401](https://doi.org/10.1103/PhysRevSTAB.3.044401).
- [88] G. Iadarola, R. De Maria, S. Lopaciuk et al. *Xsuite: An Integrated Beam Physics Simulation Framework*. 2023. arXiv: [2310.00317](https://arxiv.org/abs/2310.00317) [[physics.acc-ph](https://arxiv.org/archive/physics)].
- [89] K. Hirata. ‘An Introduction to SAD’. In: *Nonlinear Dynamics and the Beam–Beam Interaction*. CERN Yellow Reports. Geneva: CERN, 1988, pp. 62–68. DOI: [10.5170/CERN-1988-004.62](https://doi.org/10.5170/CERN-1988-004.62).

- [90] D. Schulte. ‘Study of Electromagnetic and Hadronic Background in the Interaction Region of the TESLA Collider’. PhD thesis. University of Hamburg, 1997.
- [91] G. Le Meur et al. *Description of GUINEA-PIG++: The C++ Upgraded Version of the GUINEA-PIG Beam–Beam Simulation Program*. Tech. rep. EUROTeV, 2009.
- [92] FCC Collaboration. *FCC Feasibility Study Mid-Term Review: Full Report. 215th Session of Council*. Tech. rep. Geneva: CERN, 2024. URL: cds.cern.ch/record/2888566.
- [93] K. Oide. *Vibrational Orbit Effects*. Presentation at the FCC-EIC Joint & MDI Workshop. CERN, Switzerland, 2022. URL: indico.cern.ch/event/1186798/contributions/5062667/.
- [94] M. Masuzawa, R. Sugahara and H. Yamaoka. *Vibration Issues for SuperKEKB*. 2010. URL: www.slac.stanford.edu/econf/C1009137/PDF/masuzawavibsubmit.pdf.
- [95] Y. Nakayama et al. ‘Characteristics of Ground Motion at KEK and SPring-8’. In: *Proceedings of the 9th European Particle Accelerator Conference (EPAC 2004)*. Paper WEPKF048. Lucerne, Switzerland, 2004, pp. 1711–1713. URL: proceedings.jacow.org/e04/PAPERS/WEPKF048.PDF.
- [96] A. Tada et al. ‘On the Physical Properties of the Kanto Loam, Volcanic Ash Soil’. In: *Transactions of the Japanese Society of Irrigation, Drainage and Reclamation Engineering* 14 (1965), pp. 67–70. DOI: [10.11408/jsidre1965.1965.14_67](https://doi.org/10.11408/jsidre1965.1965.14_67).
- [97] R. Amirikas et al. *Ground Motion and Comparison of Various Sites*. EUROTeV Report EUROTeV-Report-2005-023. Hamburg: DESY, 2005.
- [98] O. Brüning et al. *LHC Design Report*. CERN Report CERN-2004-003. CERN, 2004.
- [99] M. Serluca et al. *Vibration Analysis and Control in Particle Accelerators*. Presentation at the FCC-ee Mechanical Optimisation Workshop. 2019. URL: indico.cern.ch/event/694811/contributions/2863859/.
- [100] C. R. Harris et al. ‘Array Programming with NumPy’. In: *Nature* 585.7825 (2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
- [101] H. Fukuma. *Fast Orbit Feedback System*. Slides included in the presentation by Y. Funakoshi, “Experience of Collision Control Using Beam–Beam Deflection at KEKB/SuperKEKB,” September 2022. 2022.
- [102] K. Akai. ‘Stability Analysis of RF Accelerating Mode with Feedback Loops Under Heavy Beam Loading in SuperKEKB’. In: *Physical Review Accelerators and Beams* 25 (2022), p. 102002. DOI: [10.1103/PhysRevAccelBeams.25.102002](https://doi.org/10.1103/PhysRevAccelBeams.25.102002).
- [103] A. W. Weidemann et al. ‘Multi-Turn Strong–Strong Beam–Beam Simulations and Measurements at PEP-II’. In: *Proceedings of the 2005 Particle Accelerator Conference (PAC’05)*. Paper MOPC006. Knoxville, Tennessee, USA, 2005, pp. 348–350.
- [104] R. Moore and V. Shiltsev. *Coherent Beam–Beam Modes and Tune–Shift Suppression in Hadron Colliders*. Tech. rep. FERMILAB-CONF-13-053-APC. Batavia, IL: Fermi National Accelerator Laboratory, 2013.

- [105] W. Höfle et al. ‘Beam Tests of the LHC Transverse Feedback System’. In: *Proceedings of the 22nd Russian Particle Accelerator Conference (RuPAC 2010)*. Paper THCHX01. Protvino, Russia: JACoW Publishing, 2010, pp. 275–279. URL: proceedings.jacow.org/r10/papers/thchx01.pdf.
- [106] J. Keintzel et al. ‘FCC-ee Energy Calibration and Polarization: Status and Outlook’. In: *Proceedings of the International Conference on High Energy Physics (ICHEP2022)*. PoS (ICHEP2022). 2022, p. 048. DOI: [10.22323/1.414.0048](https://doi.org/10.22323/1.414.0048).
- [107] P. C. Kohnke. ‘ANSYS’. In: *Finite Element Systems: A Handbook*. Springer, 1982, pp. 19–25.
- [108] M. Masuzawa and H. Yamaoka. *Our Experiences with Vibration and Damping Material*. 2018. URL: www.slac.stanford.edu/econf/C1810085/Presentations/321.pdf.
- [109] M. Masuzawa et al. ‘IP Orbital Feedback for Collision Tuning at KEKB’. In: *Proceedings of the 7th European Particle Accelerator Conference (EPAC 2000)*. Vienna, Austria, 2000, pp. 1211–1213.
- [110] J. Salvesen et al. ‘Operational challenges of the SuperKEKB iBump feedback system’. In: *Proceedings of the 16th International Particle Accelerator Conference (IPAC’25)*. Paper THPS011. Taipei, Taiwan, 2025, pp. 2968–2971. DOI: [10.18429/JACoW-IPAC2025-THPS011](https://doi.org/10.18429/JACoW-IPAC2025-THPS011).
- [111] H. Kaji et al. ‘Control System of SuperKEKB’. In: *Proceedings of the 17th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS’19)*. Paper MOAPP01. New York, NY, USA, 2019, pp. 1–7. DOI: doi.org/10.18429/JACoW-ICALEPCS2019-MOAPP01.
- [112] L. R. Dalesio, A. J. Kozubal and M. R. Kraimer. ‘EPICS Architecture’. In: *Proceedings of the International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS’91)*. Tsukuba, Japan, 1991, pp. 278–282.
- [113] Control System Studio Collaboration. *Control System Studio (CSS) — BOY (Best OPI, Yet)*. EPICS Controls Community, 2026. URL: www.controlsystemstudio.org/.
- [114] T. T. Nakamura, N. Yamamoto and T. Katoh. ‘Data Archiving System in KEKB Accelerators Control System’. In: *Proceedings of the 10th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS’05)*. Geneva, Switzerland, 2005, pp. 204–206.
- [115] Grafana Labs. *Grafana: Open Source Analytics and Monitoring Platform*. 2024. URL: grafana.com/.
- [116] K. Pearson. ‘Mathematical Contributions to the Theory of Evolution. III. Regression, Heredity, and Panmixia’. In: *Philosophical Transactions of the Royal Society of London A* 187 (1896), pp. 253–318. DOI: [10.1098/rsta.1896.0007](https://doi.org/10.1098/rsta.1896.0007).
- [117] Y. Kimura and H. Baba. ‘TRISTAN Electron-Positron Collider’. In: *IEEE Transactions on Nuclear Science* 30.4 (1983), pp. 2231–2233.
- [118] K. Oide et al. ‘Design of beam optics for the future circular collider e^+e^- collider rings’. In: *Phys. Rev. Accel. Beams* 19 (11 2016), p. 111005. DOI: [10.1103/PhysRevAccelBeams.19.111005](https://doi.org/10.1103/PhysRevAccelBeams.19.111005).

- [119] Dassault Systèmes. *OPERA Electromagnetic and Multiphysics Simulation Software*. Version 2024, www.3ds.com/products-services/simulia/products/opera/. 2024.
- [120] Y. Arimoto et al. ‘Three-Dimensional Field Analysis for Final Focus Magnet System at SuperKEKB’. In: *Proceedings of the 5th International Particle Accelerator Conference (IPAC’14)*. Paper WEPRI086. Dresden, Germany, 2014, pp. 2595–2597. DOI: [10.18429/JACoW-IPAC2014-WEPRI086](https://doi.org/10.18429/JACoW-IPAC2014-WEPRI086).
- [121] J. P. T. Salvesen et al. ‘Consistent Representation of Lattices Between Optics Codes for FCC-ee, SuperKEKB, and More’. In: *Proceedings of the 70th ICFA Advanced Beam Dynamics Workshop on High Luminosity Circular e^+e^- Colliders (eeFACT 2025)*. Tsukuba, Japan, 2025.
- [122] H. Sugimoto et al. ‘Design Study of the SuperKEKB Interaction Region Optics’. In: *Proceedings of the 5th International Particle Accelerator Conference (IPAC’14)*. Paper TUOAB02. Dresden, Germany, 2014, pp. 950–952. DOI: [10.18429/JACoW-IPAC2014-TUOAB02](https://doi.org/10.18429/JACoW-IPAC2014-TUOAB02).
- [123] L. Deniau. ‘MAD-NG, a standalone multiplatform tool for linear and non-linear optics design and optimisation’. In: *Proceedings of the 14th International Computational Accelerator Physics Conference (ICAP 2024)*. Seeheim, Germany, 2024. URL: arxiv.org/abs/2412.16006.
- [124] G. Broggi et al. ‘Comparison of Xsuite Simulations with Measured Backgrounds at SuperKEKB’. In: *Proceedings of the 16th International Particle Accelerator Conference (IPAC’25)*. Paper MOPM035. Taipei, Taiwan, 2025, pp. 1098–1101. DOI: [10.18429/JACoW-IPAC2025-MOPM035](https://doi.org/10.18429/JACoW-IPAC2025-MOPM035).
- [125] G. Nigrelli et al. *Simulations and Measurements of Injection Backgrounds at SuperKEKB*. Presented at the 17th International Particle Accelerator Conference (IPAC’26). Accepted contribution. Deauville, France, 2026.
- [126] J. P. T. Salvesen et al. ‘Modelling Optics and Beam–Beam Effects of SuperKEKB with Xsuite’. In: *Proceedings of the 16th International Particle Accelerator Conference (IPAC’25)*. Paper MOPM034. Taipei, Taiwan, 2025, pp. 382–385. DOI: [10.18429/JACoW-IPAC2025-MOPM034](https://doi.org/10.18429/JACoW-IPAC2025-MOPM034).
- [127] J. P. T. Salvesen. *SAD2XS: Conversion of SAD Lattices to Xsuite*. Version v0.2.0. 2026. DOI: [10.5281/zenodo.18985397](https://doi.org/10.5281/zenodo.18985397).
- [128] F. Zimmermann. *FCC-ee Lessons from SuperKEKB*. 2025. DOI: [10.48550/arXiv.2509.07448](https://doi.org/10.48550/arXiv.2509.07448). arXiv: [2509.07448](https://arxiv.org/abs/2509.07448) [physics.acc-ph].
- [129] P. Kicsiny et al. ‘Towards Beam–Beam Simulations for FCC-ee’. In: *Proceedings of the eeFACT 2022 Workshop*. JACoW, 2023, pp. 165–170. DOI: [10.18429/JACoW-eeFACT2022-WEZAT0102](https://doi.org/10.18429/JACoW-eeFACT2022-WEZAT0102).
- [130] G. Roy and S. Kostoglou. *FCC-ee Optics Comparison Report*. Tech. rep. To appear on CERN CDS. CERN, 2026.
- [131] P. Raimondi et al. ‘Local chromatic correction optics for Future Circular Collider e^+e^- ’. In: *Phys. Rev. Accel. Beams* 28 (2 2025), p. 021002. DOI: [10.1103/PhysRevAccelBeams.28.021002](https://doi.org/10.1103/PhysRevAccelBeams.28.021002).
- [132] C. M. García Jaimes. ‘Optics Design of the Next Generation Future Lepton Circular Collider FCC-ee’. PhD thesis. École Polytechnique Fédérale de Lausanne (EPFL), 2025.

- [133] Z. Zhang et al. ‘Optimized Physics Performance Evaluation of Monochromatization Interaction Region Optics for Direct s -Channel Higgs Production at FCC-ee’. In: *Proceedings of the 16th International Particle Accelerator Conference (IPAC’25)*. Paper MOPM040. Taipei, Taiwan, 2025, pp. 406–409. DOI: [10.18429/JACoW-IPAC2025-MOPM040](https://doi.org/10.18429/JACoW-IPAC2025-MOPM040).
- [134] C. García-Jaimes et al. ‘Exploring FCC-ee Optics Designs with Combined Function Magnets’. In: *Journal of Physics: Conference Series* 2687.2 (2024), p. 022007. DOI: [10.1088/1742-6596/2687/2/022007](https://doi.org/10.1088/1742-6596/2687/2/022007).
- [135] R. V. Servranckx and K. L. Brown. ‘Circular Machine Design Techniques and Tools’. In: *AIP Conference Proceedings* 153 (1987), pp. 121–149. DOI: [10.1063/1.36364](https://doi.org/10.1063/1.36364).
- [136] K. Oide and H. Koiso. ‘Dynamic Aperture of Electron Storage Rings with Noninterleaved Sextupoles’. In: *Physical Review E* 47 (3 1993), pp. 2010–2018. DOI: [10.1103/PhysRevE.47.2010](https://doi.org/10.1103/PhysRevE.47.2010).
- [137] A. Seryi et al. ‘Experimental and Theoretical Progress of Linear Collider Final Focus Design and ATF2 Facility’. In: *Nuclear Instruments and Methods in Physics Research Section A* 740 (2014), pp. 2–5. DOI: [10.1016/j.nima.2013.12.029](https://doi.org/10.1016/j.nima.2013.12.029).
- [138] M. Dam. ‘Challenges for FCC-ee Luminosity Monitor Design’. In: *European Physical Journal Plus* 137.1 (2022), p. 81. DOI: [10.1140/epjp/s13360-021-02265-3](https://doi.org/10.1140/epjp/s13360-021-02265-3).
- [139] S. Mazzoni et al. ‘CERN’s Beam Instrumentation R&D Study for FCC-ee’. In: *Proceedings of the 14th International Particle Accelerator Conference (IPAC’23)*. Paper THPL088. Venice, Italy, 2023, pp. 3688–3691. DOI: [10.18429/JACoW-IPAC2023-THPL088](https://doi.org/10.18429/JACoW-IPAC2023-THPL088).
- [140] M. Boscolo et al. ‘Status of the FCC-ee Interaction Region Design’. In: *EPJ Techniques and Instrumentation* 12.1 (2025), p. 4. DOI: [10.1140/epjti/s40485-025-00117-3](https://doi.org/10.1140/epjti/s40485-025-00117-3).
- [141] J. P. T. Salvesen, F. Zimmermann and P. Burrows. ‘First studies on error mitigation by interaction point fast feedback systems for FCC-ee’. In: *Proceedings of the 15th International Particle Accelerator Conference (IPAC’24)*. Paper THPG31. Nashville, U.S.A., 2024, pp. 3322–3325. DOI: [10.18429/JACoW-IPAC2024-THPG31](https://doi.org/10.18429/JACoW-IPAC2024-THPG31).
- [142] C. G. Pang et al. ‘A Fast Luminosity Monitor Based on Diamond Detectors for the SuperKEKB Collider’. In: *Nuclear Instruments and Methods in Physics Research Section A* 931 (2019), pp. 225–235. DOI: [10.1016/j.nima.2019.03.071](https://doi.org/10.1016/j.nima.2019.03.071).
- [143] D. Shatilov. ‘FCC-ee Parameter Optimization’. In: *ICFA Beam Dynamics Newsletter* 72 (2017), pp. 30–41.
- [144] R. Aßmann et al. *Accelerator and Physics Requirements for the Calibration of the Collision Energy*. Tech. rep. CERN, 2023. DOI: [10.17181/zb0qc-acj59](https://doi.org/10.17181/zb0qc-acj59).
- [145] M. Serluca et al. ‘Vibration and Luminosity Frequency Analysis of the SuperKEKB Collider’. In: *Nuclear Instruments and Methods in Physics Research Section A* 1025 (2022), p. 166123. DOI: [10.1016/j.nima.2021.166123](https://doi.org/10.1016/j.nima.2021.166123).

- [146] J. P. T. Salvesen, F. Zimmermann and P. Burrows. ‘Interaction point beam offset tolerances for luminosity performance at FCC-ee’. In: *Proceedings of the 16th International Particle Accelerator Conference (IPAC’25)*. Paper THPS012. Taipei, Taiwan, 2025, pp. 2972–2975. DOI: [10.18429/JACoW-IPAC2025-THPS012](https://doi.org/10.18429/JACoW-IPAC2025-THPS012).
- [147] D. Shatilov et al. ‘Lifetrac Code for the Weak-Strong Simulation of Beam–Beam Effects in Tevatron’. In: *Proceedings of the 2005 Particle Accelerator Conference (PAC’05)*. Knoxville, TN, USA, 2005, pp. 4138–4140. DOI: [10.1109/PAC.2005.1591743](https://doi.org/10.1109/PAC.2005.1591743).
- [148] D. Shatilov. *Large Footprint with 4 IP (Half-Integer Crossing), Discussion and Mitigation*. Presentation at the FCC-ee Optics Design Meeting. 2019. URL: indico.cern.ch/event/835526/contributions/3502567/.
- [149] J. Laskar and D. Robin. ‘Application of Frequency Map Analysis to the Dynamics of Particles in Accelerators’. In: *Particle Accelerators* 54 (1996), pp. 183–204.
- [150] J. Laskar. ‘Application of Frequency Map Analysis in Galactic Dynamics’. In: *The Chaotic Universe*. Ed. by S. Gurri et al. World Scientific, 2000, pp. 115–126. DOI: [10.1142/9789812793621_0006](https://doi.org/10.1142/9789812793621_0006).
- [151] M. A. A. Jebramcik. *Artificial Coupling Bumps for Beam–Beam*. Presentation at the FCC-ee Beam–Beam Meeting. CERN, Geneva, Switzerland, 2026. URL: indico.cern.ch/event/1634250/contributions/6877171/.
- [152] T. K. Charles et al. ‘Alignment and Stability Challenges for FCC-ee’. In: *EPJ Techniques and Instrumentation* 10.1 (2023), p. 8. DOI: [10.1140/epjti/s40485-023-00096-3](https://doi.org/10.1140/epjti/s40485-023-00096-3).