

Diploma Thesis

High-Precision Magnetic Field Prediction and Control in Synchrotrons

submitted in satisfaction of the requirements for the degree
Diplom-Ingenieur
of the TU Wien, Faculty of Physics

Diplomarbeit

Hochpräzise Magnetfeldvorhersage und -steuerung in Synchrotronen

ausgefuehrt zum Zwecke der Erlangung des akademischen Grads
Diplom-Ingenieur
eingereicht an der TU Wien, Fakultät für Physik

Anton Lu, M.Sc.Eng. M.Sc.

Matr.Nr.: 12231206

Betreuung: Prof. Dipl.-Ing. Dr.techn. **Michael Benedikt**
Atominstitut
Technische Universität Wien
Stadionallee 2, 1020 Wien, Österreich

Wien, im Date undefined

Kurzfassung

Abstract

Acknowledgements

Contents

Acknowledgements	6
1 Introduction	19
1.1 Motivation and Problem Statement	19
1.2 Research Objectives	19
1.3 Thesis Structure and Contributions	19
2 Background and Literature Review	20
2.1 Accelerator Physics Fundamentals	20
2.1.1 Transverse Beam Dynamics as Quasi-Harmonic Oscillators	20
2.1.2 Nonlinear Transverse Beam Dynamics and Momentum Coupling	21
2.1.3 Principal Magnet Types in Synchrotrons	22
2.1.4 Betatron tune and beta function	24
2.1.5 Dipole Field Errors and Closed Orbit Distortion	24
2.1.6 Magnet errors and resonances	25
2.1.7 Longitudinal Beam Dynamics	26
2.1.8 Beam Injection	27
2.1.9 Beam Extraction	29
2.2 The CERN Accelerator Complex	32
2.3 The CERN Super Proton Synchrotron	32
2.3.1 Principal Machine Design	33
2.3.2 Principal Magnet Families	35
2.3.3 Beam Instrumentation and Diagnostics	36
2.4 Normal Conducting Accelerator Magnets	37
2.4.1 Magnetic Field Generation in Normal Conducting Magnets	38
2.4.2 Physical Mechanisms of Hysteresis	39
2.4.3 Sources of Field Errors and Mitigation	41
2.5 Hysteresis Models	41
2.5.1 Operator-type Models	42
2.5.2 ODE-type Models	47
2.5.3 Machine Learning Approaches	52
2.6 Eddy Current Decay	52
2.6.1 Eddy Currents in Conducting Materials	53
2.6.2 Eddy Currents in Synchrotron Components	54
2.6.3 Beam Impact of Eddy Current Decay in Fast Cycling Synchrotrons	57
2.6.4 Mitigation Strategies	58
3 Accelerator Magnets and Magnetic Measurements	60
3.1 Magnetic Field in the Magnet Aperture	60
3.1.1 Multipole Expansion in Polar Coordinates	60
3.1.2 Cartesian Representation	61
3.1.3 Field Quality Conventions	62

3.1.4	Forbidden Harmonics and Magnet Symmetry	63
3.1.5	Fringe Fields and the Full-Length Measurements	64
3.2	Measurement Techniques	65
3.2.1	Local, Integral, and Average Measurements	66
3.2.2	Hall Probes	66
3.2.3	NMR Probes	68
3.2.4	Stretched Wire	69
3.2.5	Fluxmeters	71
3.2.6	Rotating Coils	75
3.3	Calibration and Accuracy Limits	77
3.3.1	Coil Area Calibration	77
3.3.2	Higher-Order Gradient Measurements	77
3.3.3	Integration Drift in Fluxmeters	77
3.3.4	Summary: Accuracy Limits and Implications	77
3.4	Magnetic Measurements at CERN	77
3.4.1	SPS Main Dipole Measurement Bench	77
3.4.2	SPS Main Quadrupole Measurement Bench	77
3.5	Online B-Train Measurement Systems	77
3.5.1	Operating Principle of a B-train	77
3.5.2	B-Train Implementations	78
3.5.3	SPS B-Train	79
4	Beam Operation at CERN and the SPS	80
4.1	Accelerator Operation in Cycles	80
4.2	Main SPS Cycles	82
4.2.1	SPS North Area Fixed Target Proton Cycle (SFTPRO)	83
4.2.2	LHC-Type Proton Cycle	84
4.2.3	SFTION	88
4.2.4	Machine Development Cycles	90
4.3	Typical SPS Supercycles	90
4.4	Accelerator Controls Infrastructure at CERN	90
4.4.1	LSA	90
4.4.2	Timing	90
4.4.3	Middleware	90
4.4.4	Logging	90
4.4.5	UCAP	90
4.4.6	FGCs	90
4.5	Magnetic Field Control in the CERN Injectors	90
4.5.1	Hierarchical Field Control	90
4.5.2	B-Train Regulation in the PS and PSB	90
4.5.3	SPS B-Train in Operation	90
4.6	Magnetic Precycle and Degaussing in Operation	90
4.6.1	Dedicated Magnetic Precycle in the SPS (MD1)	90
4.6.2	Quasi-Degaussing of Higher Order Magnets in the SPS	90
4.6.3	Operational Limitations from Hysteresis	90
4.7	Machine Development Time	90
5	Field Modeling with Machine Learning	91

6	Field Compensation Strategies in the SPS	92
7	Data-Driven Hysteresis Modeling	93
7.1	Problem Formulation	93
7.2	Model Architectures	93
7.2.1	Baseline Models	93
7.2.2	Transformer-Based Models	93
7.2.3	Other Approaches	93
7.3	Prediction Strategies	93
7.3.1	Autoregressive vs. Direct Prediction	93
7.3.2	Chunked Prediction Approach	93
7.4	Measurement Design and Execution	93
7.5	Training Methodology	93
7.5.1	Dataset Preparation	93
7.5.2	Pretraining Strategy	93
7.5.3	Modeling Constraints and Considerations	93
8	Eddy Current Decay	94
8.1	Eddy Current Decay in the SPS Main Dipoles	94
8.2	Modeling Eddy Current Decay in the SPS	94
8.2.1	Eddy Current Decay as ODE	94
8.2.2	Measuring Eddy Current Decay in the SPS	94
8.2.3	Physics Constraints and Parameter Restrictions	94
8.3	Data-Driven Eddy Current Modeling	94
8.4	Qualitative Evaluation	94
8.5	Eddy Current Decay in the SPS Main Quadrupoles	94
8.5.1	Tune Decay Measurements in the SPS	94
8.5.2	Field Decay on Lab Magnet Measurement	94
8.5.3	XSuite Simulation	94
9	Control System Implementation	95
9.1	Control Law	95
9.2	Feedforward Control Architecture	95
9.2.1	Deployment to Control Room	95
9.3	Operational Results	95
9.3.1	Field Compensation Performance	95
9.3.2	Energy and Time Savings	95
10	Experimental Results - Hysteresis Compensation	96
10.1	Measurement of Impact of Precycle	96
10.1.1	Chromaticity Changes	96
10.1.2	Injection Field Shift Characterization	96
10.1.3	Beam Tolerance Analysis	96
10.2	Hysteresis Excitation from Supercycle Change	96
10.3	Main Dipole Field Prediction	96
10.3.1	Model Performance Comparison	96
10.4	Main Dipole Field Compensation	96
10.4.1	Fixed Target Flat Top Compensation	96
10.4.2	Injection Field Compensation with Transformer	96

10.4.3	Eddy Current Compensation at Injection	96
10.4.4	Injection Field Compensation on MD Cycle	96
10.5	Operational Validation	96
10.6	Main Quadrupole Analysis	96
10.6.1	Measurement Challenges	96
10.6.2	Beam Impact Studies	96
11	Discussion and Future Work	97
11.1	Current Limitations	97
11.1.1	Measurement Accuracy	97
11.1.2	Data Variety Limitations	97
11.1.3	Generalization Challenges	97
11.2	Operational Requirements Analysis	97
11.3	Extension to Other Magnet Families	97
11.3.1	Transfer Learning for Higher Order Magnet Families	97
11.3.2	Other Magnet Families	97
11.4	Future Developments	97
12	Conclusions	98

List of Figures

2.1	Idealized transverse magnetic field lines in the magnet aperture for a dipole, quadrupole, and sextupole.	23
a	Dipole	23
b	Quadrupole	23
c	Sextupole	23
2.2	Tune diagram illustration up to 3rd order resonance with a sample working point.	25
2.3	Layout of the CERN SPS lattice, showing the six main sections (sextants) and the location of the main magnet families.	33
2.4	Simplified representation of a FODO cell in the SPS lattice.	33
2.5	Cross-sectional views of the MBA and MBB dipole magnets used in the Super Proton Synchrotron (SPS).	35
a	MBA dipole cross section. the magnet is 6,2 m long.	35
b	MBB dipole cross section. The magnet is 6,22 m long.	35
2.6	Super Proton Synchrotron (SPS) main quadrupole cross section.	35
2.7	Schematic of the dipole magnet powering topology in the Super Proton Synchrotron (SPS).	36
2.8	Schematic of a simple magnetic circuit with a C-shaped iron yoke and an air gap with height h . The coil produces a magnetomotive force NI that drives the magnetic flux \mathbf{B} through the circuit, with the field strength \mathbf{H} determined by the material properties and geometry.	38
2.9	Schematic of a typical $M - H$ curve for a ferromagnetic material, showing the nonlinear response and saturation behavior.	39
2.10	Schematic major $B-H$ hysteresis loop for a soft ferromagnetic material, indicating the remanent field B_r and coercive field H_c . The arrows indicate the direction of traversal as H is cycled.	40
2.11	Illustration of a hysteron operator $R_{\alpha\beta}$ which activates when the input signal crosses the upper threshold α , and deactivates when the signal decreases below β . The units of the illustrated input signal are arbitrary as the hysteron operator is scale-invariant, and the output is a binary state of either $+1$ or -1	43
2.12	Schematic of the Preisach plane, showing the partition into S^+ and S^- regions by the staircase interface determined by the input history. The color intensity represents the weight function $\mu(\alpha, \beta)$	43
2.13	Schematic illustration of the First Order Reversal Curve (FORC) measurement procedure and the resulting Preisach weight function. The First Order Reversal Curve (FORC) distribution in the right panel is a direct experimental estimate of the Preisach density $\mu(\alpha, \beta)$ entering Eq. (2.40).	45
a	Schematic family of First Order Reversal Curves (FORCs), obtained by driving the input to saturation and reversing to a sequence of return values u_r before returning to saturation. Each curve corresponds to a distinct return point u_r	45

b	Schematic First Order Reversal Curve (FORC) distribution $\rho(u_r, u)$ computed via Eq. (2.42), plotted on the Preisach half-plane $\alpha \geq \beta$. The distribution is proportional to the Preisach weight function $\mu(\alpha, \beta)$, with high-density regions indicating a large population of hysterons with similar switching thresholds.	45
2.14	Illustration of the play operator \mathcal{P}_r with threshold r . The output follows the input with a dead-band of width r on each side, producing a continuous piecewise-linear hysteretic response. Arrows indicate the direction of traversal under a full input cycle.	46
2.15	Schematic of the Duhem model structure; a B - H plot showing ascending/descending branches, the anhysteretic curve C_a , and the upper/lower bounds C_u, C_l . This one figure introduces the whole class.	47
2.16	Model output with labeled parameters.	48
2.17	Model output with labeled parameters M_s, H_c, B_r , optionally minor loops to illustrate drifting.	49
2.18	Schematic showing repeated asymmetric minor loops drifting away from the major loop, contrasted with a Preisach model's closed minor loops. This explains why ODE models are limited for complex cycling histories.	50
2.19	Equivalent Reinforcement Learning (RL) circuit model for a single eddy current loop in a conducting material surrounding an accelerator magnet. The inductance L represents the magnetic energy stored in the loop and R represents the ohmic dissipation. Following a ramp in the excitation current, the eddy current decays exponentially with time constant $\tau = L/R$	54
2.20	Schematic cross-section of a laminated iron yoke, showing eddy current loops confined within individual lamination sheets.	55
3.1	Schematic longitudinal profile of the main dipole field component $B_1(z)$, normalised to the body value $B_1(0)$. The physical length L_{phys} is defined by the iron geometry. The effective magnetic length L_{eff} is obtained by integrating the full profile including the fringe regions at each end, indicated by the shaded areas.	65
3.2	Photograph of a typical Hall probe assembly mounted on a PCB, illustrating the millimetre-scale active sensing element.	67
3.3	Schematic placement of a Hall probe in a dipole magnet aperture cross-section, indicating the sensitivity of the measurement to transverse position (x, y) , longitudinal position z along the magnet axis, and angular misalignment of the probe axis relative to the field direction.	67
3.4	Schematic illustration of the NMR measurement principle, showing a sample volume in a homogeneous dipole field, excitation and detection coil, and the relation between resonance frequency and magnetic flux density.	69
3.5	Schematic illustration of the stretched wire technique, showing the wire path along the magnet axis, the return conductor outside the field region, and the measurement of the induced voltage during controlled motion of the wire.	70
3.6	Schematic placement of a fluxmeter coil inside a magnet aperture, indicating the coil length, transverse position, and relation to the magnet yoke.	72
3.7	Top-down schematic view of a multi-coil fluxmeter set. Each element is placed at a distinct radius $R_{\text{coil},k}$ and samples a different linear combination of multipole components. Compensation windings, formed by connecting elements in series with opposite polarity, suppress the dominant harmonic and isolate higher-order contributions.	72

3.8	Schematic cross-section of a magnet aperture with a rotating coil assembly, showing the coil element at radius R_{coil} and the rotation angle φ	75
4.1	Sample machine cycles for the Proton Synchrotron Booster (PSB) (left), Proton Synchrotron (PS) (center), and Super Proton Synchrotron (SPS) (right), showing the momentum programme as a function of time, or equivalently the main dipole field. Colored bands below each curve indicate the beam processes: pre-injection, injection plateau, ramp, flat top, extraction, and ramp-down. All cycle durations are integer multiples of the 1,2s basic period.	81
	a PSB cycle	81
	b PS cycle	81
	c SPS cycle	81
4.2	Schematic timing diagram of the Proton Synchrotron Booster (PSB), Proton Synchrotron (PS), and Super Proton Synchrotron (SPS) supercycles on a common time axis. Arrows indicate beam transfer moments between machines. Named cycle slots correspond to the representative cycles shown in Figure 4.1; remaining slots are shown as unlabelled rectangles. ZERO cycle slots are indicated by white rectangles. The diagram illustrates that the Proton Synchrotron (PS) executes multiple cycles within one Super Proton Synchrotron (SPS) cycle period, enabling multiple injections into the Super Proton Synchrotron (SPS) before the acceleration ramp begins.	83
4.3	The SPS Fixed Target Proton (SFTPRO) machine cycle. Left: main dipole field $B(t)$ as a function of cycle time, overlaid with the beam intensity measured by the Beam Current Transformer (BCT). The two-injection pattern is visible at 200 ms and 1400 ms, followed by the acceleration ramp, the 4800 ms slow extraction spill, and the ramp-down. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	84
	a Main dipole field and normal operational beam intensity	84
	b Quadrupole, sextupole, and octupole circuits	84
4.4	Detail of the SPS Fixed Target Proton (SFTPRO) cycle during the slow extraction flat top, showing the slow ramp of the main dipole field and the corresponding adjustments of the quadrupole circuits to maintain constant optics. The slow extraction spill occurs during this phase, and therefore the field reproducibility requirements are most stringent here.	85
4.5	The Large Hadron Collider (LHC) Pilot and machine cycle (13 200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the single injection at C1015, the acceleration ramp to 450 GeV/c, and the fast extraction. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	86
	a Main dipole field and beam intensity	86
	b Quadrupole, sextupole, and octupole circuits	86
4.6	The Large Hadron Collider (LHC) 4-injection () machine cycle (20 400 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing four injections at 2400 ms intervals starting at C1015, the acceleration ramp to 450 GeV/c, and the fast extraction. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	86
	a Main dipole field and beam intensity	86

	b	Quadrupole, sextupole, and octupole circuits	86
4.7		The Advanced Wakefield Experiment (AWAKE) machine cycle (7200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the single injection at C1015, the acceleration ramp to 400 GeV/ c , and the fast extraction to the Advanced Wakefield Experiment (AWAKE) beamline. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	86
	a	Main dipole field and beam intensity	86
	b	Quadrupole, sextupole, and octupole circuits	86
4.8		The High Radiation to Materials (HiRadMat) Pilot machine cycle (8400 ms[TODO: verify]). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the single injection at C1015, the acceleration ramp to 440 GeV/ c , and the fast extraction to the High Radiation to Materials (HiRadMat) facility. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	87
	a	Main dipole field and beam intensity, Pilot cycle	87
	b	Quadrupole, sextupole, and octupole circuits, Pilot cycle	87
4.9		The High Radiation to Materials (HiRadMat) 4-injection machine cycle ([TODO: verify 20 400 ms]). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing four injections at 2400 ms intervals starting at C1015, the acceleration ramp to 440 GeV/ c , and the fast extraction to the High Radiation to Materials (HiRadMat) facility. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	88
	a	Main dipole field and beam intensity, 4-injection cycle	88
	b	Quadrupole, sextupole, and octupole circuits, 4-injection cycle	88
4.10		The Search for Hidden Particles (SHiP) machine cycle (7200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the single injection at C1015, the acceleration ramp to 400 GeV/ c , and the slow extraction spill. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	88
	a	Main dipole field and beam intensity	88
	b	Quadrupole, sextupole, and octupole circuits	88
4.11		The SPS Fixed Target Ion (SFTION) 300 GeV/ c machine cycle ([TODO: verify cycle length]). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the four injections at C735, the intermediate energy plateau, the acceleration ramp, and the slow extraction spill. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis.	89
	a	Main dipole field and beam intensity, 300 GeV/ c cycle	89
	b	Quadrupole, sextupole, and octupole circuits, 300 GeV/ c cycle	89

- 4.12 The SPS Fixed Target Ion (SFTION) 33 GeV/ c machine cycle (**[TODO: verify cycle length]**). Left: main dipole field $B(t)$ overlaid with beam intensity from the Beam Current Transformer (BCT), showing the injections at C735, the intermediate energy plateau, the acceleration ramp, and the slow extraction spill. Right: Quadrupole Focusing (QF) or Quadrupole Defocusing (QD) quadrupole, sextupole, and octupole circuit functions on the same time axis. 89
- a Main dipole field and beam intensity, 33 GeV/ c cycle 89
- b Quadrupole, sextupole, and octupole circuits, 33 GeV/ c cycle 89

List of Tables

2.1	Principal parameters of the European Organization for Nuclear Research (CERN) Super Proton Synchrotron (SPS).	34
3.1	Allowed normal harmonics b_n for the principal SPS magnet families, derived from the rotational antisymmetry condition (3.14). All other harmonics are forbidden and vanish in a geometrically perfect magnet.	63
3.2	Classification of magnetic measurement techniques by the field quantity they provide directly.	66

Acronyms

AD	Antiproton Decelerator
AGS	Alternating Gradient Synchrotron
ALPS	A Logarithmic Position System
AWAKE	Advanced Wakefield Experiment
BCT	Beam Current Transformer
BLM	Beam Loss Monitor
BNL	Brookhaven National Laboratory
BP	Basic Period
BPM	Beam Position Monitor
CERN	European Organization for Nuclear Research
CNAO	National Center for Oncological Hadrontherapy
COSE	Constant Orbit Slow Extraction
DBS	Dynamic Beam Scheduling
ELENA	Extra Low Energy Antiproton Ring
FDI	Fast Digital Integrator
FGC	Functional Generator / Controller
FODO	Focusing-Defocusing
FORC	First Order Reversal Curve
HiRadMat	High Radiation to Materials
HL-LHC	High-Luminosity Large Hadron Collider
JA	Jiles-Atherton
LEIR	Low Energy Ion Ring
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LINAC	Linear Accelerator

LLG Landau-Lifshitz-Gilbert

LS Long Shutdown

LSS Long Straight Section

MTE Multi-Turn Extraction

NASA The National Aeronautics and Space Administration

NC Normal Conducting

NMR Nuclear Magnetic Resonance

ODE Ordinary Differential Equation

PCB Printed Circuit Board

PI Prandtl-Ishlinskii

PPM Per-Pulse Modulation

PS Proton Synchrotron

PSB Proton Synchrotron Booster

QD Quadrupole Defocusing

QF Quadrupole Focusing

RF Radio Frequency

RFKO RF Knock-Out

RL Reinforcement Learning

SC Superconducting

SFTION SPS Fixed Target Ion

SFTPRO SPS Fixed Target Proton

SHiP Search for Hidden Particles

SPS Super Proton Synchrotron

WCM Wall Current Monitor

Chapter 1

Introduction

1.1 Motivation and Problem Statement

1.2 Research Objectives

1.3 Thesis Structure and Contributions

Chapter 2

Background and Literature Review

This chapter provides an overview of the fundamental accelerator physics principles relevant to synchrotron accelerators, with a focus on the impact of magnetic field errors and hysteresis on beam dynamics and stability. The chapter begins with a review of the basic concepts of transverse and longitudinal beam dynamics, including the role of dipole, quadrupole, and sextupole magnets in shaping the beam trajectory and focusing properties. The following chapters will build on this foundation to develop challenges with beam operation at European Organization for Nuclear Research (CERN) and the Super Proton Synchrotron (SPS), the challenges in attaining high precision magnetic fields, and the methods to mitigate these effects, such as using surrogate models for correction.

2.1 Accelerator Physics Fundamentals

Synchrotron accelerators accelerate charged particles to high energies through the combined action of Radio Frequency (RF) electric fields and precisely controlled magnetic fields. The magnetic field configuration in a synchrotron, also known as a lattice, serve two principal functions: to establish and maintain an equilibrium circular orbit for the particles, and to provide transverse focusing to contain the beam within acceptable apertures. Furthermore, additional more specialized magnets may be used to achieve a stable beam. This section outlines the physics principles governing beam dynamics in synchrotrons, emphasizing why small deviations in magnetic field strength, at the level of $\mathcal{O}(10^{-4})$ relative to main field, can produce unacceptable changes in beam parameters and beam quality.

2.1.1 Transverse Beam Dynamics as Quasi-Harmonic Oscillators

Modern synchrotrons employ a separated function design in which the guiding bending and focusing forces are provided by distinct magnet families. Dipolar magnets provide the main bending field to establish a closed orbit along the reference trajectory, satisfying the relation:

$$B\rho = \frac{p}{q} \quad (2.1)$$

where B is the magnetic field strength perpendicular to the beam trajectory, ρ is the bending radius of the ring, p is the particle momentum, and e is the particle charge. The relation $B\rho$ is also known as the magnetic rigidity of the beam.

For particles with small transverse displacements from the design orbit, additional focusing forces are required to achieve stable motion around the ring. The movement of particles in the transverse plane can be modeled as quasi-harmonic oscillators, which describe the evolution of small horizontal and vertical offsets from the reference trajectory, where the restoring forces are provided by magnetic field gradients. In a standard synchrotron lattice, the primary elements providing transverse beam focusing are dedicated quadrupole magnets, and in the absence of coupling between planes, the horizontal and vertical motions can be treated independently.

The transverse movement is described by Hills equation, which in the linear approximation for small displacement x from the design orbit take the form:

$$\frac{d^2x}{ds^2} + K(s)x = 0 \quad (2.2)$$

where s is the azimuthal coordinate around the ring and $K(s)$ is the focusing function, which is determined by the quadrupole field gradient and local bending geometry. An analogous equation holds for vertical displacements y .

A convenient parameterization introduces the beta function $\beta(s)$ as an amplitude envelope and the phase function $\mu(s)$ as an accumulated phase advance. One may write

$$x(s) = A_x \sqrt{\beta_x(s)} \cos(\mu_x(s) + \phi_{x,0}), \quad (2.3)$$

with an analogous expression in the vertical plane. The constant A_x is set by initial conditions. The term betatron oscillation denotes this quasi harmonic transverse oscillation about a local equilibrium orbit.

2.1.2 Nonlinear Transverse Beam Dynamics and Momentum Coupling

The equations of motion in Eq. 2.2 assumes that the transverse restoring force is proportional to displacement and that the particle momentum equals the design momentum. In reality a bunched beam contains particles with a small relative momentum deviation $\delta = \Delta p/p$. A momentum deviation changes the curvature required to follow the reference orbit through Eq. 2.1, which produces an equilibrium horizontal offset in regions with bending. This effect is described by the dispersion function $D(s)$ through

$$x(s) = x_\beta(s) + D(s) \delta, \quad (2.4)$$

which is the solution to the inhomogeneous equation of motion including momentum deviation,

$$\frac{d^2x}{ds^2} + K(s)x = \frac{\delta}{\rho(s)}, \quad (2.5)$$

where $x_\beta(s)$ denotes the betatron oscillation about the local equilibrium orbit. The dispersion describes how particles with different momenta follow different equilibrium orbits, is periodic in a closed ring, and is determined uniquely by the same focusing lattice that sets the betatron motion. In particular, dispersion is only generated in regions with non-zero bending radius $\rho(s)$.

Through dispersion, the focusing strength experienced by a particle appears in the restoring force as $k(s) \propto 1/(B\rho)$. A particle with momentum deviation δ experiences a focusing strength that differs from the design value, causing its betatron tune to shift. This momentum dependence of the tune is quantified by the chromaticity ξ , defined as

$$\xi_x = \frac{dQ_x}{d\delta}, \quad \xi_y = \frac{dQ_y}{d\delta}. \quad (2.6)$$

In synchrotrons it is also common to define the normalized chromaticity as

$$Q' = \frac{\xi}{Q} \quad (2.7)$$

Which gives the fractional tune shift per unit relative momentum deviation, easily comparable across machines with different tunes.

For a lattice consisting only of bending and focusing elements, the natural chromaticity is typically non-zero, which leads to a spread in betatron tune across the momentum distribution of the beam. This in turn can cause beam loss if the tune spread overlaps with resonance frequencies, but is correctable in the first order by sextupole magnets which produce nonlinear fields that provide momentum-dependent focusing, which will be described further below.

The implications of Q' for transverse beam stability are discussed in Section 2.1.7.2.

2.1.3 Principal Magnet Types in Synchrotrons

The transverse beam dynamics introduced above are realized in practice by specially designed accelerator magnets, each optimized for a specific function, which will be discussed more technically in detail in Chapter 3. For now, we can write the ideal transverse magnetic field in the mid-plane $y = 0$ of a synchrotron magnet as a series expansion:

$$B_y(x, 0) = B_y(0, 0) + \left. \frac{\partial B_y}{\partial x} \right|_0 x + \frac{1}{2!} \left. \frac{\partial^2 B_y}{\partial x^2} \right|_0 x^2 + \frac{1}{3!} \left. \frac{\partial^3 B_y}{\partial x^3} \right|_0 x^3 + \dots \quad (2.8)$$

Factoring out the ideal dipole field on the reference orbit, denoted by B_0 and equal to the dipole field B in Eq. 2.1, defines coefficients b_n by

$$B_y(x, 0) = B_0 \sum_{n=1}^{\infty} b_n x^{n-1}, \quad (2.9)$$

where we can immediately identify that b_1 corresponds to the dipole term, b_2 to the quadrupole term, and b_3 to the sextupole term. Equating Eq. 2.8 and Eq. 2.9 gives

$$B_0 b_n = \frac{1}{(n-1)!} \left. \frac{\partial^{n-1} B_y}{\partial x^{n-1}} \right|_0. \quad (2.10)$$

The normalized strengths used in transverse beam dynamics, denoted by k_n , are defined from the on-axis derivatives as

$$k_{n-1} = \frac{1}{B\rho} \frac{1}{(n-1)!} \left. \frac{\partial^{n-1} B_y}{\partial x^{n-1}} \right|_0 = \frac{B_0}{B\rho} b_n \quad (2.11)$$

which removes the dependence on particle momentum through $B\rho$.

2.1.3.1 Dipole magnets

A well designed accelerator dipole magnet generates an uniform transverse field in the vacuum chamber, which through the Lorentz force providing a net curvature to the reference trajectory:

$$\mathbf{F} = q \mathbf{v} \times \mathbf{B} \quad (2.12)$$

For a particle with magnetic rigidity $B\rho$, the bend angle produced by a magnet is set by the integrated field,

$$\theta = \frac{1}{B\rho} \int B dl \approx \frac{B_0 L_{\text{eff}}}{B\rho} \quad (2.13)$$

where L_{eff} the effective magnetic length defined so that the so called fringe-field contributions, that is the stray fields at the ends of the magnets, are captured by an equivalent hard-edge magnet.

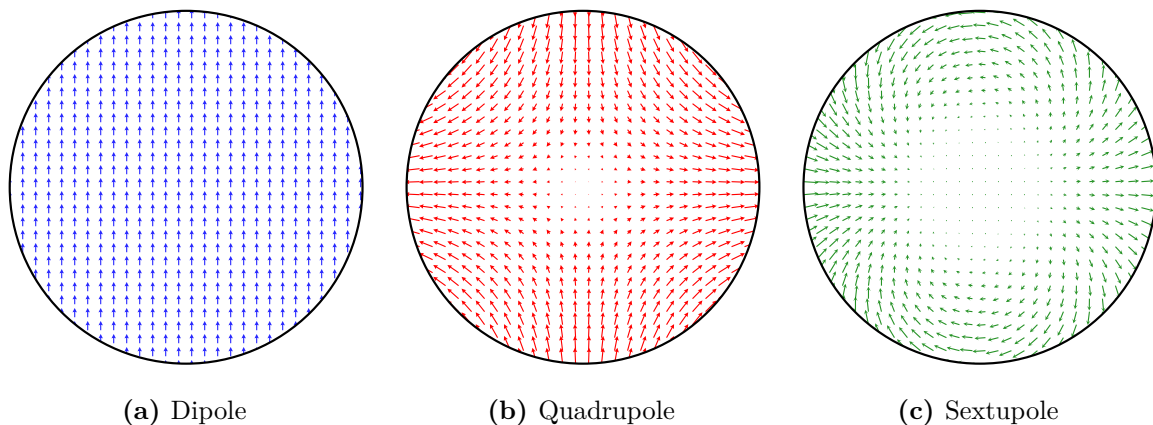


Fig. 2.1: Idealized transverse magnetic field lines in the magnet aperture for a dipole, quadrupole, and sextupole.

While exact calculation of the particle trajectory in a dipole requires integrating through the full field profile including fringe fields, one can in many cases use the B_0 as an effective representation of B to estimate the bend angle. For additional simplification with sector dipoles, the effective length can be approximated by the physical length of the magnet, in which case the bend angle is simply

$$\theta \approx \frac{L}{\rho} \quad (2.14)$$

Using Eq. 2.11, the normalized dipole strength is

$$k_0 = \frac{B_0}{B\rho} = \frac{1}{\rho} \quad (2.15)$$

Which shows that the dipole strength is directly related to the bending radius of the reference orbit.

2.1.3.2 Quadrupole magnets

A normal quadrupole produces a field linear in transverse displacement, which is recovered from Eq. 2.9 by retaining the $n = 2$ term,

$$B_y(x, 0) \approx B_0 b_2 x, \quad (2.16)$$

with the corresponding normalized strength $k_1 = (B_0/B\rho)b_2$ from Eq. 2.11.

A quadrupole focuses in one transverse plane while defocusing in the other; alternating focusing/defocusing quadrupoles yield net strong focusing around the ring.

2.1.3.3 Sextupole magnets

A normal sextupole produces a quadratic transverse field dependence, recovered from Eq. 2.9 by retaining the $n = 3$ term,

$$B_y(x, 0) \approx B_0 b_3 x^2, \quad (2.17)$$

with the corresponding normalized strength $k_2 = (B_0/B\rho)b_3$ from Eq. 2.11.

In practice, dedicated sextupoles are primarily used to correct chromaticity, with their effect relying on nonzero dispersion so that off-momentum particles sample an effective quadrupole component. Sextupoles also introduce amplitude-dependent tune shifts and can drive nonlinear resonances, so their strength and reproducibility affect dynamic aperture and beam lifetime.

2.1.4 Betatron tune and beta function

As the beam circulates around the ring, the transverse oscillations accumulate phase advance $\mu_{x,y}$ in each plane according to local focusing strength. The total phase advance accumulated around one complete revolution is defined as:

$$\mu_{x,y} = \int_0^C \frac{ds}{\beta_{x,y}(s)} \quad (2.18)$$

where the integral is taken over all magnet elements in the ring where C is the ring circumference. β is the beta function, which characterizes the local amplitude envelope of the betatron oscillations and is determined solely by the placement and strength of the magnetic elements in the lattice. Specifically, it is defined as the solution to the envelope equation

$$\frac{1}{2}\beta''(s) - \frac{1}{4}\beta'(s)^2 + K(s)\beta(s) = 1 \quad (2.19)$$

From the total phase advance, the betatron tune of the beam is defined as the number of oscillations per turn:

$$Q_x = \frac{\mu_x}{2\pi} \quad (2.20)$$

with analogous definition for the vertical plane. The dispersion function $D(s)$, which governs the equilibrium offset of off-momentum particles and was introduced in Section 2.1.2, is determined by the same lattice that sets $\beta(s)$.

2.1.5 Dipole Field Errors and Closed Orbit Distortion

A dipole magnet with a field error, either a deviation in field magnitude δB or misalignment, acts as an unintended deflecting kick. The deflection angle is $\theta = (\delta B)L/(B\rho)$, where L is the dipole length. With small field errors on individual magnets, the small deflection drives oscillations of the closed orbit around the design orbit. In the case of macroscopic field changes where each magnet experiences the same field deviation ΔB , the closed orbit changes according to the definition of the momentum compaction factor

$$\alpha_c = \frac{\Delta C/C}{\Delta p/p} = \frac{1}{C} \oint \frac{D(s)}{\rho(s)} ds \quad (2.21)$$

Which is a measure of how the path length changes with momentum, and is determined by the lattice design through the dispersion and bending radius, and not related to any beam parameters like energy or intensity. Similarly this can be related to a change in orbit radius r and magnetic field B through the beam rigidity in Eq. 2.1 as

$$\alpha_c = \frac{\Delta r/r}{\Delta B/B} \quad (2.22)$$

By rewriting the equation, we can easily relate field errors ΔB at a given field B to changes in orbit radius Δr at a reference radius R :




Fig. 2.2: Tune diagram illustration up to 3rd order resonance with a sample working point.

$$\Delta r = \alpha_c R \frac{\Delta B}{B} \quad (2.23)$$

2.1.6 Magnet errors and resonances

Accelerator magnets are designed to produce a specific multipole field component with high precision, but in practice they always contain small field errors caused by manufacturing imperfections, mechanical misalignments, and finite magnet geometry. As the beam circulates through the ring, these field errors can drive resonant behavior in the beam dynamics if the betatron tune is close to resonance conditions. The resonance condition for a normal multipole error of order p is given by the relationship

$$|m_x|Q_x + |m_y|Q_y = p \quad (2.24)$$

where m_x and m_y are integers with $|m_x| + |m_y| = p$. The relationship can be visualized in a tune diagram, where the horizontal and vertical tunes are plotted on the axes, and resonance lines corresponding to different orders of multipole errors are drawn, as seen in Fig. 2.2. At resonance conditions, the beam experiences coherent oscillations that can grow in amplitude, leading to beam loss if the resonance is strong enough. In practice, not all resonance lines are equally strong, as the strength of the resonance depends on the specific multipole error and its distribution around the ring. Nevertheless, for any synchrotron, a tune working point chosen as (Q_x, Q_y) must be carefully selected to avoid beam instability.

To identify the strength of the resonance lines, loss maps are often used, where the beam is intentionally excited at different tunes and the resulting beam loss is measured. In many cases, the integer tune is primarily set by the lattice design to achieve stable focusing and required focusing strength, while the fractional tune is adjusted to avoid strong resonances.

To avoid or reduce the impact the resonances, many mitigations can be used, such as changing the fractional tune by changing the quadrupole strengths, or by using dedicated correction magnets to compensate for the field errors that drive the resonance. Additionally the beam can be stabilized by transverse dampers, which apply a feedback kick to the beam to suppress coherent oscillations, or by using octupole magnets to introduce amplitude-dependent tune shifts that can detune the resonance for particles with larger oscillation amplitudes.

2.1.7 Longitudinal Beam Dynamics

In addition to transverse confinement defined by the lattice, a synchrotron must control longitudinal motion. Longitudinal beam dynamics in an accelerator are governed by the interaction between the beam and the accelerating RF cavities, which provide the accelerating voltage to increase the beam energy turn by turn.

The RF cavity maintains an oscillating electric field at frequency f_{RF} , which must satisfy the synchronism condition to ensure that particles receive the correct energy kick at the right time.

$$f_{\text{RF}} = h \cdot f_{\text{rev}} = h \cdot \frac{v}{2\pi R} \quad (2.25)$$

The RF frequency must be an integer multiple h of the revolution frequency f_{rev} of the particles in the ring, and h is called the harmonic number of the accelerator. Depending on the accelerator, more than one RF circuit maybe be used, for instance with a fundamental RF system at harmonic h and a second harmonic system at $2h$ to provide additional longitudinal focusing.

A particle arriving at the RF cavity at the synchronous phase ϕ_s gains a consistent energy per turn:

$$\Delta E = eV \sin \phi_s \quad (2.26)$$

where V is voltage amplitude of the RF cavity. As the particle gains energy ΔE , both the magnetic field and RF frequency must increase together to preserve synchronism. The beam rigidity in Eq. 2.1 requires that the bending dipole field must increase at a rate determined by the energy gain:

$$\frac{dB}{dt} \propto \frac{dp}{dt} \propto eV \sin \phi_s \quad (2.27)$$

Any between the actual RF frequency and the required synchronism frequency, which can arise from an RF frequency error, or equivalently from an unexpected bending field error, causes off-momentum particles to drift in RF phase. When uncorrected, this leads to degradation in beam quality due to longitudinal beam heating and emittance growth.

2.1.7.1 Synchrotron Oscillations and Synchrotron Tune

For ideal calculations, one often assumes that all particles in the beam arrive at the RF cavity with the same phase ϕ_s , however in reality, each particle in a bunched beam has slightly different energy and phase relative to the synchronous particle. Particles that arrive at the RF cavity with a phase slightly different from ϕ_s experience a different energy gain, causing them to drift in both phase and energy relative to the synchronous particle. This gives rise to synchrotron oscillations, which are small oscillations in energy and phase around the synchronous particle, analogous to betatron tune which oscillates in the transverse plane.

The frequency of synchrotron oscillations is characterized by the synchrotron tune Q_s , which is the fractional number of longitudinal oscillations per revolution:

$$Q_s = \sqrt{\frac{h|\eta|eV \cos \phi_s}{2\pi\beta^2\gamma mc^2}} \quad (2.28)$$

where η is the slip factor defined as $\eta = \alpha_c - \frac{1}{\gamma^2}$, with $\beta = v/c$, and γ the Lorentz factor. The synchrotron tune is typically much smaller than the betatron tune, often $Q_s \ll 1$ while $Q_x, Q_y > 1$. Fundamentally, this means that the longitudinal dynamics can often be treated independently from the transverse dynamics, as the timescales of the two motions are well separated. However in some cases, the transverse and longitudinal dynamics can become coupled, for instance through chromatic effects as discussed in Section 2.1.2.

2.1.7.2 Chromatic Effects and Beam Stability

The chromaticity Q' introduced in Section 2.1.2 and Eq. 2.7 has direct consequences for the transverse stability of the beam. In a bunched beam, particles at the head and tail of the bunch occupy different longitudinal positions and, through synchrotron oscillations, continuously exchange their longitudinal positions over the synchrotron period. If the chromaticity is negative, particles with higher momentum oscillate at a lower betatron tune, and the resulting phase difference between the head and tail of the bunch can couple to transverse wake fields generated by the passage of the beam through the resistive machine environment such as the vacuum chamber and accelerator components.

This coupling drives the head-tail instability, in which transverse oscillation amplitudes grow exponentially on a timescale much shorter than the synchrotron period, leading to rapid beam loss. Maintaining a small positive chromaticity is therefore standard operational practice in synchrotrons, as it reverses the phase relationship between head and tail and suppresses this instability. However, the chromaticity must remain within a controlled range, since excessively large values increase the tune spread and can push the working point onto resonance lines identified in Section 2.1.6. A hysteretic contribution to the sextupole field component b_3 , arising from the magnetic history of the lattice sextupoles, produces an uncontrolled shift of Q' that can drive the machine into an unstable regime. This effect is one of the primary operational motivations for the field prediction and correction strategies developed in subsequent chapters.

2.1.8 Beam Injection

Injection into a synchrotron can proceed by several methods depending on the source of the beam. When filling from a Linear Accelerator (LINAC), the beam is typically continuous or macro-pulsed and may be accumulated over many turns through phase space painting [cite] or charge-exchange injection [rennerCommissioningNewCERN2022]. When the source is an upstream synchrotron, the beam is already bunched and has a well-defined momentum, emittance, and longitudinal structure. In this case, a single-turn or multi-turn bunched injection scheme is employed, in which the entire bunch or a train of bunches is transferred through a transfer line and captured by the receiving ring in a single injection event [goddardIntroductionBeamTransfer2018, fraserFastExtractionSingle2018]. The receiving ring must satisfy a set of matching conditions for the injected beam to be captured without emittance growth or beam loss, which are discussed in the following subsections.

2.1.8.1 Momentum Matching

The fundamental requirement for successful injection is that the dipole field of the receiving ring matches the beam rigidity of the incoming beam. A relative momentum deviation $\delta = \Delta p/p$

between the injected beam and the design momentum of the ring causes the beam to circulate on a displaced closed orbit given by

$$x_{\text{co}}(s) = D(s) \delta, \quad (2.29)$$

where $D(s)$ is the dispersion function from Eq. 2.4.

Since the dispersion is non-zero in the bending regions of the lattice, a momentum offset produces a globally distorted closed orbit whose amplitude scales with $D(s)$ and δ . At injection, the beam is transferred from the transport line into the ring through an injection septum, a thin magnetic element that deflects the incoming beam onto the ring orbit while the circulating beam passes on the other side of the septum blade. The injected beam therefore enters the ring at a fixed transverse position defined by the septum geometry, and a closed orbit distortion shifts the design orbit relative to this fixed entry point. If the orbit offset is sufficiently large, the incoming beam no longer clears the septum blade, leading to partial interception and injection losses. Additionally, at injection flat bottom the beam rigidity is at its lowest point in the magnetic cycle, so that absolute dipole field errors constitute a comparatively large relative deviation δ , or corresponding Δr , as discussed in Section 2.1.5.

2.1.8.2 RF Capture and Longitudinal Matching

Beyond the transverse orbit, the injected bunch must be captured into a stable RF bucket of the receiving ring. The half-width of the RF bucket in momentum is

$$\delta_{\text{max}} = \sqrt{\frac{2e\hat{V}}{\pi h |\eta| E}}, \quad (2.30)$$

as derived from the longitudinal equations of motion in Section 2.1.7. The RF system must be pre-phased so that its synchronous momentum coincides with the injected momentum, ensuring δ remains within the acceptance given by Eq. 2.30. In practice, the dipole field, RF frequency, and closed orbit are not independent but are regulated jointly, so that a residual momentum mismatch propagates into both the transverse and longitudinal planes through the control loop; this coupling is discussed in detail in Chapter 4.

2.1.8.3 Tune Matching

After injection, the betatron tune of the receiving ring must be set to a working point that avoids the resonance conditions identified in Section 2.1.6. The quadrupole strengths of the receiving ring are adjusted so that the fractional tunes Q_x and Q_y are placed sufficiently far from low-order resonance lines in the tune diagram, accounting for any tune shift introduced by quadrupole field errors at flat bottom. A further consideration is the incoherent space charge tune shift, which at injection energy is given by

$$\Delta Q_{x,y} = -\frac{r_0 N}{4\pi \varepsilon_{x,y} \beta^2 \gamma^3 B_f}, \quad (2.31)$$

where r_0 is the classical proton radius, N is the number of particles per bunch, $\varepsilon_{x,y}$ is the transverse emittance, β and γ are the relativistic factors, and B_f is the bunch form factor. Since the space charge force is defocusing in both planes, it shifts the tune downward, and the working point must be chosen such that the depressed tune $Q_{x,y} + \Delta Q_{x,y}$ remains away from resonance lines.

Beyond space charge, any deviation in the quadrupole field strength from its nominal value produces a tune shift that follows from differentiating the tune definition with respect to the focusing function,

$$\Delta Q_{x,y} = \frac{1}{4\pi} \oint \beta_{x,y}(s) \Delta k(s) ds, \quad (2.32)$$

where $\beta_{x,y}(s)$ is the beta function at the location of the error and $\Delta k(s)$ is the local error in normalised quadrupole strength. Since the beta function weights the contribution of each error by the local beam size, errors located at high-beta positions have a disproportionate influence on the tune. The integer part of the tune is set by the overall lattice design and determines the periodicity of the closed orbit response to dipole errors, as well as the order of the lowest resonances that can be driven by systematic field errors around the ring. The fractional part is then adjusted by varying the quadrupole strengths, and the working point must be chosen with sufficient margin from resonance lines to accommodate both the space charge tune shift of Eq. 2.31 and any residual quadrupole field errors present at injection.

2.1.9 Beam Extraction

Analogous to the previous subsection, beam extraction transfers the circulating beam, or a fraction of it, from the synchrotron into a downstream beamline. The choice of extraction scheme is determined primarily by the downstream application: machines requiring a single, well-defined bunch transfer to another accelerator employ fast extraction, while experiments requiring a quasi-continuous particle flux over timescales of seconds employ slow extraction [[fraserFastExtractionSingle2018](#), [goddardIntroductionBeamTransfer2018](#)]. The following subsections describe the main extraction techniques, concluding with third-order resonance slow extraction that is directly relevant to the field control problem addressed in this thesis.

2.1.9.1 Fast Extraction

Fast extraction transfers the beam out of the synchrotron in one or at most a few turns, and is the standard technique for machine-to-machine transfer in an accelerator complex and to beam dumps. The main hardware elements are a fast pulsed kicker magnet, which delivers a transverse angular kick $\Delta x'_{\text{kick}}$ to the entire beam within a single revolution period, and a magnetic septum, which separates the deflected beam from the circulating beam and routes it into the extraction line. Prior to extraction, slow closed-orbit bumper magnets displace the beam transversely toward the septum blade, reducing the required kick strength. The spatial separation achieved at the septum depends on the integrated kicker field, the beta functions at the kicker and septum locations, and the betatron phase advance between them, which is ideally $\pi/2$ to maximize the deflection.

In single-turn fast extraction, all circulating bunches are transferred to the downstream machine in a single kicker pulse, preserving the bunch structure of the source machine. At CERN, the SPS extracts particles to the Large Hadron Collider (LHC) using this scheme, with a closed-orbit bump bringing the beam close to the septum blade before the kicker fires.

Alternatively, multi-turn fast extraction can be used to transfer the beam across several consecutive turns. By spreading the extraction over multiple turns, the instantaneous beam intensity at the septum blade is reduced, which can lower beam losses and residual activation of the extraction region. Additionally, splitting the beam in the transverse phase space before extraction allows the receiving machine to be filled with a more uniform transverse distribution,

which reduces space charge effects at injection into the downstream machine. The CERN Proton Synchrotron (PS) Multi-Turn Extraction (MTE) scheme achieves this by using sextupole and octupole magnets to adiabatically trap beam into stable islands in horizontal phase space near $Q_x = 4/5$, partitioning the beam into a central core and four surrounding islands. These are then extracted sequentially over five turns into the SPS, with the phase-space gap between islands coinciding with the septum blade location. [huschauerTransverseBeamSplitting2017, huschauerAdvancingCERNProton2019, barnesCERNPSMultiturn2006]

2.1.9.2 Slow Extraction

In contrast to fast extraction, slow extraction delivers a quasi-continuous flux of particles to downstream experiments over timescales of hundreds of milliseconds to several seconds, spanning many thousands to millions of turns around the ring. Unlike fast extraction, slow extraction gradually peels particles from the circulating beam by driving them into a region of transverse phase space that is rendered unstable by multipole magnets, after which they drift outward until they reach the blade of an electrostatic septum. [arrutiasotaOptimisationSlowExtraction2020, taylorSlowExtractionUpgrades2024] The electrostatic septum imparts a transverse deflection to particles that cross its wire blade, directing them into a downstream magnetic septum that routes them into the extraction beamline. The remaining circulating beam then passes on the field-free side of the blade without perturbation. Prior to extraction, the RF system is typically switched off or adiabatically reduced, debunching the beam into a coasting distribution that fills the longitudinal phase space uniformly, so that the extracted particle flux is not modulated at the RF frequency or its harmonics. The primary figure of merit for slow extraction is spill quality, which while facility dependent, is commonly defined by the uniformity of the extracted particle flux as a function of time $I(t)$. However, the spill quality can also be characterized by, or in combination with, the length of the extracted spill and the stability of the beam position and size at the extraction point. Regardless of definition, a high quality spill is critical for downstream fixed-target experiments requiring a stable and controlled beam intensity.

Several distinct mechanisms exist for driving particles into the unstable phase-space region. The simplest approach to resonant slow extraction is tune-driven, or equivalently quadrupole-driven, in which the horizontal betatron tune is intentionally ramped toward a resonance condition. The stable area in phase space progressively reduces particles at the boundary of the beam distribution become unstable and spiral outward to the septum blade. The extraction rate is controlled by the tune ramp speed, making the spill sensitive to tune ripple from power supply noise and to tune spread controlled by the sextupole field, which modulate the instantaneous size of the stable phase-space area, and therefore extraction rate. However, the tune-driven approach is straightforward to implement and requires no additional hardware beyond standard machine quadrupoles and sextupoles.

In betatron-core-driven slow extraction, used at medical synchrotrons such as National Center for Oncological Hadrontherapy (CNAO) and MedAustron [pulliaBetatronCoreDriven2016], a dedicated electromagnetic induction element accelerates the beam longitudinally, effectively sweeping the beam momentum and thus its tune through the resonance. This method provides an intrinsic suppression of high-frequency tune ripples, however the hardware implementation is more complex and less flexible than tune-driven approach, and is not suitable for high-intensity beams due to the limited accelerating voltage of the induction element.

In RF Knock-Out (RFKO) extraction, a transverse RF excitation signal, typically a frequency chirp or colored noise applied near the betatron tune sideband, directly heats the transverse emittance of the circulating beam, diffusing particles from the stable beam core into the unstable resonance region. The RFKO technique leaves the machine optics unchanged during extraction

and allows fast, flexible control of the extraction rate through the RF drive amplitude. Therefore, it is well-suited to applications requiring rapid intensity modulation, and it is employed for proton and ion slow extraction at the CERN PS East Area.

The principal disadvantage shared by all resonant slow extraction schemes is the unavoidable beam loss at the electrostatic septum blade. There, a fraction of particles are intercepted by the finite-thickness wire array rather than deflected cleanly into the extraction channel. At high-intensity facilities, losses at the septum lead to residual radioactivation of the extraction region, constraining the maximum deliverable annual intensity and imposing significant operational and maintenance costs. A particularly important variant of resonant slow extraction exploits the momentum dependence of the betatron tune to select which particles are extracted at a given moment, which is described in the following section.

2.1.9.3 Chromatic slow extraction

In chromatic slow extraction, the machine optics remain nominally constant throughout the spill and no active tune or emittance manipulation is applied during extraction. Instead, the third-integer resonance $3Q_x = p$ according to Eq. 2.24 is established by dedicated sextupole magnets prior to the spill, generating a triangular separatrix in horizontal phase space that encloses a finite stable area. The sextupoles generate a triangular separatrix in horizontal phase space enclosing a stable area that shrinks as the tune approaches the resonant value. The working tune is set close to, but offset from, the resonance, so that the chromaticity $Q' = dQ_x/d\delta$ maps each momentum slice $\delta = \Delta p/p_0$ to a distinct tune value. Only particles satisfying

$$Q_x + Q'\delta = \frac{p}{3} \quad (2.33)$$

are resonant and therefore unstable, while the remainder of the beam, at momenta not satisfying Eq. 2.33, remains confined within the separatrix.

Unstable particles advance outward along the three arms of the separatrix by a radial increment known as the spiral step Δr per three-turn period, which in the thin-lens approximation scales as

$$\Delta r \propto \frac{\epsilon_{\text{res}}}{h}, \quad (2.34)$$

where ϵ_{res} is the resonance driving term strength determined by the sextupole integrated fields and their beta-function-weighted distribution around the ring, and $h = p/3 - Q_x - Q'\delta$ is the tune distance from the resonance. A sufficiently large spiral step is required so that a particle crosses the full wire thickness of the electrostatic septum in a single three-turn pass, as a step smaller than the septum wire thickness results in direct interception and beam loss. As the dipole bending field evolves during the extraction flat top, the mean radial position of the circulating beam shifts according to Eq. 2.23, and is continuously displacing the resonant momentum slice of Eq. 2.33. This sustains a steady outward particle flux without requiring any active tune manipulation during the spill.

Chromatic slow extraction is more sensitive to field errors is compared to other slow extraction techniques. Since the extraction rate is set by the interplay between Q' , the tune offset h , and the radial position of the beam, each of these quantities must remain stable at a level with respect to the width of the resonant momentum slice, which can be as narrow as $\delta \sim 10^{-3}$ [cite]. Errors in Q' from residual sextupole field deviations alter both the width of the resonant slice and the spiral step Δr , perturbing the instantaneous extraction rate in a manner not compensable by the dipole ramp alone. Dipole field deviations shift the closed orbit at the electrostatic septum, modifying the aperture available to the extracted beam and increasing interception losses at the septum blade. Quadrupole field errors shift Q_x relative to the resonance, changing the stable

area enclosed by the separatrix and modulating the spill intensity. The extracted spill quality is therefore a direct and sensitive probe of small relative field errors, at the level of $\Delta B/B \sim 10^{-4}$ to 10^{-3} , across the sextupole, quadrupole, and dipole families of the lattice.

Compared to the extraction techniques discussed in the previous subsection, chromatic slow extraction places significantly more stringent requirements on the magnetic field quality of the lattice. Tune-driven and betatron-core-driven schemes actively manipulate the tune or momentum during the spill, providing a degree of real-time control authority that can partially absorb slow field drifts. The RFKO technique is similarly robust, as the extraction rate is governed by the RF drive amplitude rather than by the precise value of any lattice parameter. In chromatic slow extraction, by contrast, the extraction rate emerges passively from the static interplay between the chromaticity, the resonance driving term, and the dipole ramp, with no active handle to compensate for field errors on a turn-by-turn basis. This passivity makes the technique highly reproducible and operationally simple besides the need to dedicated extraction sextupoles, but transfers the burden of spill quality entirely onto the reproducibility of the magnetic fields in the lattice.

Chromatic slow extraction is employed at several high-energy fixed-target facilities including the CERN SPS [**bryantResonantExtraction2018**, **kainResonantSlowExtraction2019**], the Fermilab Delivery Ring, and the Brookhaven National Laboratory (BNL) Alternating Gradient Synchrotron (AGS) [**cite**]. At the SPS, a refinement known as Constant Orbit Slow Extraction (COSE) is operationally deployed [**cite**], in which all machine settings are co-scaled with beam rigidity following the extracted beam momentum, fixing the separatrix orientation in normalized phase space throughout the spill and eliminating the optics drift inherent to quadrupole-only ramping schemes.

2.2 The CERN Accelerator Complex

2.3 The CERN Super Proton Synchrotron

The SPS [**300GeVProgramme1972**] is the second-largest accelerator in the CERN accelerator complex, and is housed in a circular tunnel of 6.9km circumference under the Franco-Swiss border near Geneva. The SPS has served multiple roles throughout its history after its commission in 1976, including operation as a high-energy proton–antiproton collider [**deraadCERNSPSPProtonAntiproton1985**] and as an electron and positron pre-injector for the Large Electron-Positron Collider (LEP) [**cern.genevaLEPDesignReport1990**] up until 2000. Currently, the SPS operates as the injector for the LHC, as well as several fixed-target experimental areas, and accelerates protons and heavy ions, such as lead [**malwitzCERNHeavyionFacility1993**], oxygen [**cite**] and neon [**cite**], to high energies. It is currently also undergoing upgrades to serve as injector to the High-Luminosity Large Hadron Collider (HL-LHC) and to provide high-intensity beams for the Search for Hidden Particles (SHiP) experiment [**vanherwijnenSHiPSearchHidden2018**, **alekhinFacilitySearchHidden2015**], both to be commissioned in 2030 and 2032, respectively. The operational context of the current, injector chain, cycle structure of the SPS and CERN control system designs are discussed in detail in Chapter 4.

The SPS receives beam from the PS at an injection momentum of 14 GeV/ c or 26 GeV/ c for protons and 17 GeV/ u for heavy ions, and accelerates to extraction momenta of up to 450 GeV/ c for LHC injection or up to 400 GeV/ c for fixed-target physics. As with other injectors at CERN, the SPS is a multi-cycling synchrotron, which means it accelerates beam with different beam properties, such as intensity, particle type, injection and extraction energy, tune, acceleration, and other RF configuration. In practice this means serving various fixed-target experiments, as




Fig. 2.3: Layout of the CERN SPS lattice, showing the six main sections (sextants) and the location of the main magnet families.




Fig. 2.4: Simplified representation of a FODO cell in the SPS lattice.

mentioned in 1, interleaved with injecting to the LHC. In 2026, they include the Advanced Wakefield Experiment (AWAKE) [gschwendtnerAWAKEAdvancedProton2016], High Radiation to Materials (HiRadMat) [hardenHiRadMatFacilityRealms2019], and various North Area experiments [banerjeeNorthExperimentalArea2021], with preparations underway for the SHiP experiment [alekhinFacilitySearchHidden2015]. The next subsections summarize the principal machine parameters and magnet families of the SPS, with relevant details for the field prediction and correction strategies developed in this thesis.

2.3.1 Principal Machine Design

The SPS is functionally divided into six main sections, or sextants. each containing two arcs and a Long Straight Section (LSS). The SPS lattice is designed using Focusing-Defocusing (FODO) cells, with a total of 744 bending dipoles and 432 quadrupoles arranged in a separated-function lattice. Each FODO cell is made up of two quadrupoles, one focusing and one defocusing, separated by four dipole magnets, with a total cell length of approximately 60 m, shown in

Fig. 2.4. The focusing and defocusing quadrupoles use the same magnet design, but are powered with opposite polarity to achieve the required focusing in the horizontal and vertical planes. The main dipoles use two different magnet designs, MBA and MBB, which differ in length and yoke geometry but are powered in two series circuits to produce the same magnetic field. The SPS also contains 216 main sextupole magnets in focusing or defocusing configurations for chromaticity control, as well as octupole and skew quadrupole magnets for nonlinear correction and transverse coupling compensation, as well as dedicated extraction sextupoles for chromatic slow extraction. Similarly to the main quadrupoles, the sextupole magnets are powered in two series circuits for focusing and defocusing configurations, with 108 magnets in each circuit, however not with same magnet design used for both configurations. In total, 108 FODO cells make up the SPS lattice, with 16 cells per arc, and 2 per LSS, and the remaining 8 cells distributed in the straight sections for injection, extraction, and RF systems. [300GeVProgramme1972] A layout of the SPS lattice is shown in Fig. 2.3.

Acceleration in the SPS is provided by several 200 MHz RF cavities installed in LSS3, with a harmonic number of 4620. This corresponds to a fundamental frequency of approximately 43 345 Hz at 14 GeV injection, and 43 376 Hz at 450 GeV extraction. The small range of RF frequency variation is a consequence of the the beam already being ultra-relativistic at injection, so that the velocity v is close to the speed of light c and the revolution frequency f_{rev} becomes close to constant across the ramp. A complimentary set of 800 MHz RF cavities is installed in LSS3 for bunch rotation, longitudinal emittance control, as well as improving the chromatic slow extraction spill quality [cite] .

The SPS regularly accelerates proton beams with intensities of up to 3×10^{13} protons per batch, and up to 288 bunches per cycle for LHC injection, as well as heavy ion beams with intensities of up to 1.5×10^{11} lead ions per batch, and up to 144 bunches per cycle for LHC injection. The beam emittance at injection is typically around $3 \mu\text{m}$ for protons and $1,5 \mu\text{m}$ for lead ions, with a momentum spread of approximately 10^{-4} . The SPS lattice currently operates at integer tune 20, also referred to as Q20, so called LHC-type beam, as well as Q26 for slow extraction fixed target operation. For proton operation, Q26 beam is injected at 14 GeV/ c , whereas for heavy ion operation, Q26 beam is injected at 17 GeV/u. Q20 beam is injected at 26 GeV/ c for protons.

The principal parameters of the SPS are summarized in Table 2.1.

Tab. 2.1: Principal parameters of the CERN SPS.

Parameter	Value	Unit
Circumference	6911.5	m
Injection momentum (protons)	14 & 26	GeV/ c
Maximum extraction momentum	450	GeV/ c
Injection momentum (heavy ions)	17	GeV/u
Number of dipole magnets	744	–
Number of quadrupole magnets (F/D)	216/216	–
Number of sextupole magnets (F/D)	108/108	–
Total number of magnets	1317	–
Harmonic number	4620	–
RF frequency	200.2 & 800.8	MHz
Nominal betatron tune (Q_x/Q_y)	20.x / 20.x or 26.13 / 26.18	–




Fig. 2.5: Cross-sectional views of the MBA and MBB dipole magnets used in the SPS.




Fig. 2.6: SPS main quadrupole cross section.

2.3.2 Principal Magnet Families

The previous subsection presented the lattice structure of the SPS and the main magnet families that realize the required optical functions. While the SPS was originally designed to accelerate beam to 300 GeV/ c , the machine has been adapted to reach the 450 GeV/ c required for LHC injection by increasing the dipole field strength, while maintaining the same lattice structure and magnet designs. Originally designed for a maximum field of approximately 1,8 T, the main dipoles now operate at 2,0 T at 5800 A for LHC beams, close to the practical limit for the normal-conducting magnet design. [cite]

Unchanged since the SPS was designed and constructed in the 1970s, all main magnet families are normal-conducting, operate at room temperature, and are cooled by water circulation through the copper coils. The main dipoles in the SPS ring are evenly divided into the MBA and MBB types, which reach the same nominal integrated field at the same current, but differing yoke geometry. Specifically, the MBA magnets have an aperture width of 153 mm and height of 39 mm, compared to 129 mm width and 52 mm height for the MBB magnets. The lengths differ by 20 mm as the MBA magnets are 6,2 m compared to 6,22 m of the MBB magnets. The different magnet geometry is motivated by different beam sizes in each planes as they evolve through the lengths of the magnets. Consequentially, the vacuum chambers used for the magnet families are different, with rectangular vacuum chambers with rounded corners with different heights to match the magnet geometries. [pintoCarbonCoatingSPS2013] An illustration of the cross-section of the MBA and MBB magnets is shown in Fig. 2.5. [TODO: Magnet lamination sentence]

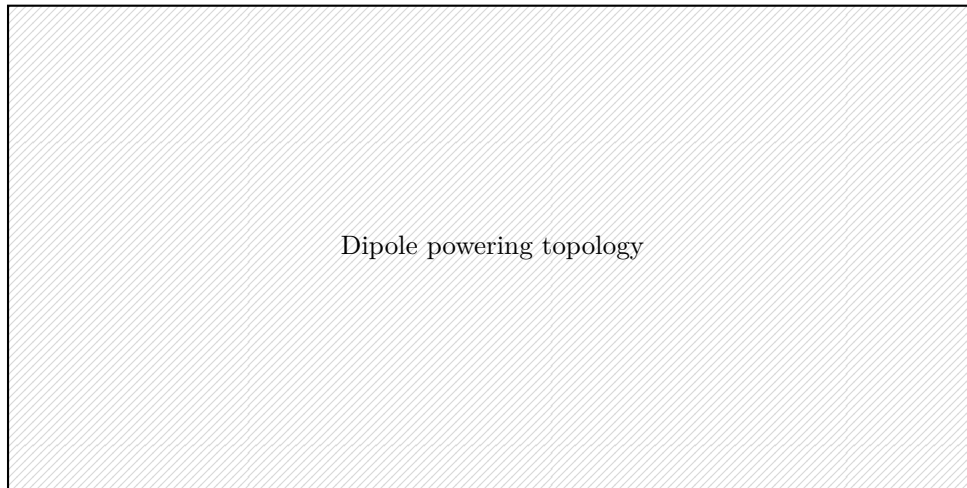


Fig. 2.7: Schematic of the dipole magnet powering topology in the SPS.

The transverse focusing in the SPS is provided by the Quadrupole Focusing (QF) and Quadrupole Defocusing (QD) magnet families, which are powered in separate circuits to provide alternating horizontal and vertical focusing along the lattice. The magnets have a common design, with a symmetric aperture of 88 mm width and height and 3,05 m long shown in Fig. 2.6. The magnets can reach a maximum field of 1,0 T at the pole tips for an equivalent field gradient of 22 T/m, used at the highest beam energy 450 GeV/u with Q26 optics. The vacuum chambers of the quadrupoles are elliptic, to reflect the beam shape of the beam passing through, and is rotated 90° to change between focusing and defocusing **TBV**.

Higher-order correction elements include sextupole magnets for chromaticity control, and octupole and skew quadrupole magnets for nonlinear correction and transverse coupling compensation. The main sextupole family are denoted LSF and LSD for focusing, and defocusing respectively, with LSE denoting the dedicated slow extraction sextupoles. The LSF, LSD and LSE all have slightly different geometric parameters, however still remain constructed with ferromagnetic iron, and suffer from the same effects as the main dipole and quadrupole magnets. **[TODO: Sextupole design parameters?]**

2.3.2.1 Magnet Powering and Control

The SPS main magnets are designed to be powered in series circuits. The powering topology of the main dipoles can be seen in Fig. 2.7, where the MBA and MBB magnets are powered in series with 12 power stations supplies referred to as SMDs, forming one global dipole circuit around the ring. Similarly for the quadrupoles... **[TODO: Magnet voltage readings]**

Each magnet family is regulated in current, below 10 ppm precision by the power converters known at CERN as Functional Generator / Controller (FGC).

2.3.3 Beam Instrumentation and Diagnostics

The SPS is equipped with a large variety of instruments for beam diagnostics and machine protection. The principal beam instrumentation relevant to this thesis includes Beam Position Monitor Beam Position Monitors (BPMs), Beam Current Transformers (BCTs), and to some extent the Beam Loss Monitor Beam Loss Monitors (BLMs). **[gasiorIntroductionBeamInstrumentation2016]**

The BPMs are used to measure the transverse beam position at various locations around the ring, and was recently upgraded to the A Logarithmic Position System (ALPS) system during the Long

Shutdown (LS)2. [wendtTechnologyFirstBeam2019, boccardiCommissioningALPSNew2021] Various types of BPMs are installed in the SPS, with a total of 108 shoe-box type in the horizontal plane next to the focusing quadrupoles, referred to as BPH, and 108 of the same in the vertical plane next to the defocusing quadrupoles, referred to as BPV. Further BPMs are installed at special locations where two-plane information is required, such as in the straight sections for injection and extraction **TBV**. ALPSs can resolve position down to $290\ \mu\text{m}$ in the horizontal plane, and $150\ \mu\text{m}$ in the vertical plane, for single bunch beams. The system is used for regulation loops such as the closed orbit feedback [brandtNewClosedorbitCorrection1990], as well as for beam-based measurements of the machine optics and magnetic fields, by measuring turn-by-turn beam position data, or average position over many turns across the beam life time.

The BCT are commonly toroidal transformers to indirectly measure the beam current by measuring the change in magnetic field caused by circulating charged particles in the beam pipe. The reading of the instrument directly translate to beam intensity, and therefore number of particles circulating in a synchrotron, and enables real-time intensity monitoring without disrupting the beam. In the SPS, the BCT measurements are realised by several different types of devices, including fast AC current transformers for single bunch measurements, and DC current transformers for average beam intensity measurements. [jakob40MHzBunchBunch2003, unserBeamCurrentTransformer1969] [cite] The beam intensity measurements are primarily used for monitoring of beam quality, losses, and diagnostics such as bunch-by-bunch profiles, as well as performance during injection, acceleration, and extraction.

BLMs are used to measure beam losses, by measuring secondary radiation from beam interactions with the beam pipe or other machine components, and are critical for machine protection and diagnostics. [wittenburgBeamLossMonitors2020] Beam losses can be caused by a variety of factors, including beam instabilities, mis-steering, or interactions with residual gas in the vacuum chamber, and can lead to damage of machine components or excessive radiation levels if not properly monitored and mitigated. In the SPS, 286 ionization chambers are installed in the main ring, and further 144 in various extraction lines. BLMs readings are primarily used for machine protection, by triggering beam dumps when losses exceed predefined thresholds, as well as for diagnostics of beam quality, efficiency and performance during injection, acceleration, and extraction. [calvogiraldoevaSPSBeamLoss2026] In the context of this thesis, BLMs are primarily relevant for beam manipulations or studies where the beam is significantly perturbed, either intentionally, or by mistake. [TODO: Beam loss sector-by-sector]

Other beam instrumentation such as Wall Current Monitor (WCM), wire scanners, are installed in the SPS for specific measurements of the longitudinal beam profile and transverse beam profile, respectively, but are not relevant for the field prediction and correction strategies developed in this thesis, and are therefore not described in detail here.

2.4 Normal Conducting Accelerator Magnets

The principal magnet types in synchrotrons introduced in Section 2.1.3 are realized in practice by electromagnets, as is the case for the SPS as described in Section 2.3.2. The two most common types of electromagnets used in synchrotrons are Normal Conducting (NC) and Superconducting (SC) magnets. NC magnets operate at room temperature and rely on resistive coils to generate the magnetic field, while SC magnets operate at cryogenic temperatures and use superconducting coils to achieve higher field strengths with lower power consumption. This section aims to briefly introduce normal conducting magnets, which are used in the CERN SPS, sources of systematic magnetic field errors such as hysteresis and eddy current decay. Chapter 3 will focus more in depth on accelerator magnets design, field tolerances and quality.




Fig. 2.8: Schematic of a simple magnetic circuit with a C-shaped iron yoke and an air gap with height h . The coil produces a magnetomotive force NI that drives the magnetic flux \mathbf{B} through the circuit, with the field strength \mathbf{H} determined by the material properties and geometry.

2.4.1 Magnetic Field Generation in Normal Conducting Magnets

A normal conducting magnet consists of a ferromagnetic yoke and a resistive coil wound around it. The yoke, typically made of laminated silicon steel, serves to concentrate and guide the magnetic flux through the magnet aperture, shaping the field profile seen by the beam. The coil carries a direct or slowly varying current, which drives the magnetic circuit and determines the field strength in the aperture. The relationship between the excitation current and the resulting magnetic field is governed by Ampère's law, which states that the line integral of the magnetic field intensity \mathbf{H} around a closed path equals the total enclosed current,

$$\oint \mathbf{H} \cdot d\mathbf{l} = NI, \quad (2.35)$$

where N is the number of turns in the coil and I is the current per turn. The magnetic flux density \mathbf{B} is related to the magnetic field intensity \mathbf{H} through the constitutive relation

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}), \quad (2.36)$$

where $\mu_0 = 4\pi \times 10^{-7}$ H/m is the permeability of free space and \mathbf{M} is the magnetization of the material, defined as the magnetic dipole moment per unit volume. In this thesis, we will use the term magnetic field to refer to the magnetic flux density \mathbf{B} unless otherwise specified, as it is the quantity that directly interacts with the beam through the Lorentz force in Eq. 2.12. Furthermore, the scalar B will be used to refer to the magnitude of the magnetic flux density \mathbf{B} for simplicity and alignment with common accelerator physics notation, while the vector \mathbf{B} will be used when the direction of the field is relevant.

Eq. 2.36 is exact and holds for all materials. In vacuum or air, $\mathbf{M} = 0$ and $\mathbf{B} = \mu_0 \mathbf{H}$. For a linear, isotropic magnetic material where the permeability is uniformly distributed, the magnetization is proportional to the applied field, $\mathbf{M} = \chi_m \mathbf{H}$, and Eq. 2.36 reduces to

$$\mathbf{B} = \mu_0 (1 + \chi_m) \mathbf{H} = \mu_0 \mu_r \mathbf{H}, \quad (2.37)$$

where $\mu_r = 1 + \chi_m$ is the relative permeability, and is a material constant independent of the applied field, and χ_m is the magnetic susceptibility. Additionally the magnetic permeability μ is defined as $\mu = \mu_0 \mu_r$. For ferromagnetic materials such as the iron yoke, however, μ_r is not a material constant but depends nonlinearly on \mathbf{H} as $\mu_r = \mu_r(\mathbf{H})$, and the general form of Eq. 2.36 must be retained.

While most magnetic theory work with macroscopic magnetization \mathbf{M} and subsequently \mathbf{B} , the underlying physical mechanism of magnetization is related to the alignment of microscopic magnetic moments in the material. The macroscopic magnetization \mathbf{M} of the iron yoke originates




Fig. 2.9: Schematic of a typical $M - H$ curve for a ferromagnetic material, showing the nonlinear response and saturation behavior.

from the collective behavior of microscopic magnetic domains, regions in which the magnetic moments of the constituent atoms are aligned in a common direction due to quantum mechanical exchange interactions. In an unmagnetized material, domains are oriented randomly such that the net macroscopic magnetization vanishes. When an external field \mathbf{H} is applied, the magnetic moments within the domains tend to align with the field, and the net macroscopic magnetization increases. At sufficiently high excitation, all domains are aligned with the applied field and the magnetization \mathbf{M} reaches its saturation value M_s and cannot increase further. This condition is known as magnetic saturation. At this condition, the effective permeability μ_r drops significantly, such that the magnetic field B in the magnet aperture no longer increases linearly with the applied current. This causes the $M - H$ curve to deviate from a linear response, as seen in Fig. 2.9, and leads to a reduction in the effective field strength for a given coil current.

When the applied field is subsequently reduced, the domain configuration does not fully retrace its prior state, leaving a remanent magnetization that depends on the previous excitation history of the magnet. It is this irreversibility that gives rise to magnetic hysteresis, which is discussed in Section 2.4.2, and further in Section 2.5 for hysteresis modeling.

2.4.2 Physical Mechanisms of Hysteresis

The irreversibility in domain configuration manifests macroscopically as magnetic hysteresis. The relationship between the applied field strength H driven by excitation current I and the resulting flux density B is not single-valued but depends on the prior excitation history of the material. Rather than responding instantaneously and reversibly to the applied field, the magnetization of the yoke lags behind changes in H , such that the same value of H corresponds to different values of B depending on whether the excitation is increasing or decreasing. [bertottiHysteresisMagnetismPhysicists1998] When H is cycled, B traces a closed loop in the $B-H$ plane, as shown schematically in Fig. 2.10. Two scalar quantities characterise the loop: the remanent field B_r , which is the flux density that persists when H is returned to zero, and the coercive field H_c , which is the reverse field required to reduce B to zero. The area enclosed by the loop is proportional to the energy dissipated per excitation cycle, reflecting the thermodynamically irreversible nature of the magnetization process.




Fig. 2.10: Schematic major B – H hysteresis loop for a soft ferromagnetic material, indicating the remanent field B_r and coercive field H_c . The arrows indicate the direction of traversal as H is cycled.

The physical origin of the lagging and irreversible response lies in the interaction of magnetic domain walls with structural defects in the material. [hubertMagneticDomainsAnalysis1998] Within each domain, the atomic magnetic moments are aligned by short-range quantum mechanical exchange interactions, while the boundaries between domains, known as Bloch walls, are transition regions across which the magnetization rotates continuously. When an external field H is applied, energetically favourable domains grow at the expense of unfavourable ones through domain wall motion. However, lattice dislocations, grain boundaries, voids, and impurity atoms create local minima in the magnetostatic energy landscape, which act as pinning sites that impede the free motion of domain walls. A domain wall resting at a pinning site requires a finite applied field to overcome the local energy barrier and advance to the next metastable configuration. Crucially, when the driving field is subsequently reduced, the wall does not retrace its path but remains trapped at a new pinning site, since the energy landscape is traversed in a different sequence on the return excursion. This path-dependence means that the magnetization state of the yoke at any instant encodes the full prior excitation history, not merely the current value of H . The resulting irreversible, discontinuous domain wall motion is known as the Barkhausen effect [barkhausenZweiMitHilfe1919], and is the microscopic mechanism responsible for the hysteretic behaviour of the macroscopic B – H relationship. Because the density and spatial distribution of pinning sites are set by the material microstructure and vary between individual magnets and with thermal and mechanical history, the hysteretic response is both nonlinear and non-unique: no functional relationship $B = f(H)$ exists that is valid across different excitation histories.

A first-principles description of hysteresis would require resolving the spatial configuration of all domain walls and their interactions with the defect landscape. The governing equation of micro-magnetic dynamics is the Landau-Lifshitz-Gilbert (LLG) equation [landauTheoryDispersionMagnetic1992, gilbertClassicsMagneticsPhenomenological2004],

$$\frac{\partial \mathbf{M}}{\partial t} = -\gamma_0 \mathbf{M} \times \mathbf{H}_{\text{eff}} + \frac{\alpha}{M_s} \mathbf{M} \times \frac{\partial \mathbf{M}}{\partial t}, \quad (2.38)$$

where \mathbf{M} is the local magnetization vector, γ_0 is the gyromagnetic ratio, α is the phenomenological damping constant, M_s is the saturation magnetization, and \mathbf{H}_{eff} is an effective field that

includes the applied field, the demagnetising field, magnetocrystalline anisotropy, and exchange interactions. Eq. (2.38) must be solved on a spatial grid with a cell size comparable to the exchange length, which is of order nanometres for iron. A metre-scale accelerator magnet yoke would therefore require a spatial discretisation with an astronomically large number of cells, rendering direct numerical integration of Eq. (2.38) computationally intractable. Consequently, no first-principles model capable of predicting the macroscopic B – H relationship of an accelerator magnet yoke for an arbitrary excitation history currently exists. This fundamental limitation motivates the phenomenological and data-driven modeling approaches reviewed in the following section.

[TODO: Rate-dependence and rate-independent hysteresis]

2.4.3 Sources of Field Errors and Mitigation

As briefly discussed in Section 2.1, particle accelerators require highly precise magnetic fields to maintain beam stability and quality. In practice, the realized field differs from the ideal design due to a combination of systematic effects shared by a magnet family or by all magnets under similar operating conditions and random effects that vary from unit to unit. This subsection gives a qualitative overview of the dominant error sources in accelerator magnets, with emphasis on normal-conducting iron-dominated magnets, and summarizes typical mitigation strategies in design, production, and operation.

2.5 Hysteresis Models

As established in Section 2.4.2, the macroscopic B – H relationship of a ferromagnetic yoke is nonlinear, history-dependent, and not reducible to a single-valued function. Predictive modeling of this relationship therefore requires models that are capable of representing path-dependent behavior under arbitrary excitation histories. The models reviewed in this section operate at the macroscopic level, treating the yoke as a lumped element characterized by its bulk B – H response, rather than resolving the underlying domain microstructure. Classical phenomenological approaches fall into two broad classes: operator-type models, which represent the hysteretic map as a weighted superposition of elementary bistable operators, and differential models, which describe the evolution of magnetization through a system of ordinary differential equations driven by the applied field. More recently, data-driven machine learning models have been applied to hysteresis modeling, learning the input to output mapping directly from measurement data without an explicit physical parametrization, which we will also review in this section.

Although hysteresis models are conventionally formulated in terms of the material quantities H and M , or equivalently H and B , the same mathematical structure applies when the excitation current I is used as the input and the flux density B in the magnet aperture as the output, since I and H are related through Ampère’s law in Eq. 2.35 and the geometry of the magnetic circuit. In this thesis, models are formulated directly in terms of I and B , as these are the quantities accessible through measurement.

Despite the diversity of available phenomenological models, no single model provides a universal solution that simultaneously achieves high accuracy, broad generalizability and compact parameterization across arbitrary magnets and excitation histories. Each model class has structural assumptions about the form of the hysteretic response, whether it be a superposition of bistable operators, a low-dimensional differential equation, or a parametric curve. These assumptions are typically well-matched only to the class of excitation histories and material behaviours for which the model was designed [mayergoyzMathematicalModelsHysteresis1991, bertottiHysteresisMagnetismPmoreeReviewHysteresisModels2023]. Parameter identification is performed on a finite and

necessarily incomplete set of measured cycles, so that a model calibrated on major loop data may systematically misrepresent the response under minor loop excitations or multi-reversal histories not represented in the identification dataset. Furthermore, the compactness that makes phenomenological models computationally attractive directly limits their representational capacity, and a model with few free parameters cannot, in general, capture the full complexity of the history-dependent magnetization state of a real iron-yoke magnet. These inherent limitations motivate the data-driven approaches reviewed in 2.5.3, which replace structural assumptions with flexible function approximators trained directly on measurement data.

2.5.1 Operator-type Models

Operator-type models represent the hysteretic input–output map as a functional acting on the full excitation history, rather than on the instantaneous value of the input alone. The history dependence is encoded through a collection of elementary bistable operators, each of which switches irreversibly between two output states at threshold values of the input, and the aggregate output is formed as a weighted superposition of the states of all operators. Because the state of each operator depends on the sequence of past inputs rather than only on the current value, the model is inherently history-aware without requiring an explicit memory variable. The two most widely used operator-type models in the magnetics literature are the Preisach model and the Prandtl–Ishlinskii model. The Preisach model assigns a continuous weight distribution over a two-dimensional space of operator thresholds, and is the most general and mathematically well-characterized member of this class. The Prandtl–Ishlinskii model is a special case in which the weight distribution is constrained such that the model admits a closed-form analytical inversion, which is advantageous in control applications where the inverse map is required.

2.5.1.1 The Preisach Model

The Preisach model, originally introduced to describe magnetic aftereffect in ferromagnetic materials [[preisachUeberMagnetischeNachwirkung1935](#)] and later given a rigorous mathematical foundation by Mayergoyz [[mayergoyzMathematicalModelsHysteresis1991](#)], is the most widely used classical model for magnetic hysteresis. It decomposes the hysteretic response of a macroscopic material into a superposition of elementary bistable units called hysteron. Each hysteron represents a microscopic switching event such as the irreversible motion of a domain wall past a pinning site.

Each hysteron is characterized by two threshold values α and β with $\alpha \geq \beta$, which define the input levels at which the unit switches between its two output states. When the input $u(t)$ rises above α , the hysteron switches to the +1 state; when $u(t)$ falls below β , it switches to the –1 state; when $u(t)$ lies between β and α , the hysteron retains its current state. Formally, the hysteron output is defined as

$$\hat{R}_{\alpha\beta}[u](t) = \begin{cases} +1 & \text{if } u(t) > \alpha, \\ -1 & \text{if } u(t) < \beta, \\ \hat{R}_{\alpha\beta}[u](t^-) & \text{otherwise,} \end{cases} \quad (2.39)$$

where t^- denotes the previous time instant.

The Preisach plane is the half-plane $\{(\alpha, \beta) : \alpha \geq \beta\}$, and each point in this plane corresponds to a distinct hysteron. The full Preisach model assigns a weight $\mu(\alpha, \beta) \geq 0$ to each point in this plane, known as the Preisach density or weight function, and defines the model output as the weighted integral of all hysteron states,




Fig. 2.11: Illustration of a hysteron operator $R_{\alpha\beta}$ which activates when the input signal crosses the upper threshold α , and deactivates when the signal decreases below β . The units of the illustrated input signal are arbitrary as the hysteron operator is scale-invariant, and the output is a binary state of either $+1$ or -1 .




Fig. 2.12: Schematic of the Preisach plane, showing the partition into S^+ and S^- regions by the staircase interface determined by the input history. The color intensity represents the weight function $\mu(\alpha, \beta)$.

$$B(t) = \iint_{\alpha \geq \beta} \mu(\alpha, \beta) \hat{R}_{\alpha\beta}[u](t) d\alpha d\beta. \quad (2.40)$$

The weight function $\mu(\alpha, \beta)$ encodes the distribution of switching thresholds across the ensemble of hysterons and is the only free function of the model. Physically, a region of high density in the Preisach plane corresponds to a large population of hysterons with similar switching thresholds, which in turn reflects a concentration of pinning sites with similar energy barriers in the material microstructure. The Preisach plane is therefore a material representation such that different magnets with different microstructures will in general have different weight functions. For example, soft magnetic materials with low coercivity are expected to have a weight function concentrated near the diagonal $\alpha \approx \beta$, with mostly a linear magnetic response.

At any instant t , the Preisach plane is partitioned into two regions, S^+ and S^- , in which the hysterons are in the $+1$ and -1 states respectively. The boundary between these two regions is

a staircase interface whose vertices are determined by the sequence of past input extrema, as shown schematically in Fig. 2.12. Each upward step in the staircase corresponds to a past local maximum of the input, and each downward step to a past local minimum. The model output at time t is determined entirely by this staircase geometry. Concretely, the output is proportional to the excess weight in S^+ over S^- , since hysterons in S^+ contribute positively and those in S^- negatively to the integral in Eq. (2.40). As the input evolves, the staircase interface moves and the balance between S^+ and S^- shifts, producing the hysteretic B - I response.

While in its original formulation, the Preisach model is continuous and defined by an integral over the Preisach plane, in practice it is often implemented in a discretized form for numerical evaluation and parameter identification. Specifically, the Preisach plane is discretized onto a finite mesh of N points (α_i, β_i) , and the continuous integral in Eq. (2.40) is replaced by a weighted sum

$$B(t) \approx \sum_{i=1}^N \mu_i \hat{R}_{\alpha_i \beta_i}[u](t), \quad (2.41)$$

where $\mu_i = \mu(\alpha_i, \beta_i)$ are the mesh weights to be identified from data.

The classical experimental technique for identifying the Preisach weight function is the measurement of First Order Reversal Curves (FORCs). In a FORC measurement, the sample is first driven to positive saturation to establish a well-defined initial state, after which the input is reduced to a return value u_r and then increased back to saturation while the output $B(u_r, u)$ is recorded. Repeating this procedure for a sequence of return values u_r yields a family of curves that systematically probe the irreversible switching events in the material. The FORC distribution is defined as the mixed second partial derivative of the measured surface,

$$\rho(u_r, u) = -\frac{\partial^2 B(u_r, u)}{\partial u_r \partial u}, \quad (2.42)$$

Alternatively, $\mu(\alpha, \beta)$ may be approximated analytically as a product of two probability density functions whose parameters are fitted to the major hysteresis loop [bertottiHysteresisMagnetismPhysicists1991]. Both approaches become impractical as mesh resolution increases. The FORC method requires a number of reversal measurements that scales with the number of mesh points, and the analytical approximation loses accuracy for soft magnetic materials with complex minor loop behavior [szaboImplementationIdentificationPreisach2016]. For a mesh of N points, direct least-squares identification is an $\mathcal{O}(N^2)$ problem that rapidly becomes computationally intractable. More recent approaches that use of gradient-based optimization over differentiable implementations of the model [rousseauDifferentiablePreisachModeling2022] have shown promise, and will be discussed in more detail in Section 2.5.3.

The Preisach model satisfies two properties that Mayergoyz showed to be both necessary and sufficient to characterize the class of hysteretic systems representable by a Preisach operator [mayergoyzMathematicalModelsHysteresis1991]. The first is the wiping-out property, also called the return-point memory property. This means that whenever the input reaches a new local maximum, all previous staircase vertices with a β value smaller than that maximum are erased, and similarly for local minima. Consequently, only dominant past extrema of the input are retained in the model memory, while minor fluctuations that are subsequently exceeded are forgotten. The second is the congruency property, which states that all minor loops traced between the same pair of input reversal values u_1 and u_2 are congruent, meaning they have the same shape and vertical extent regardless of the magnetization level at which they are initiated. The wiping-out property is satisfied to good approximation by soft ferromagnetic materials, but the congruency property is often violated: experimentally measured minor loops

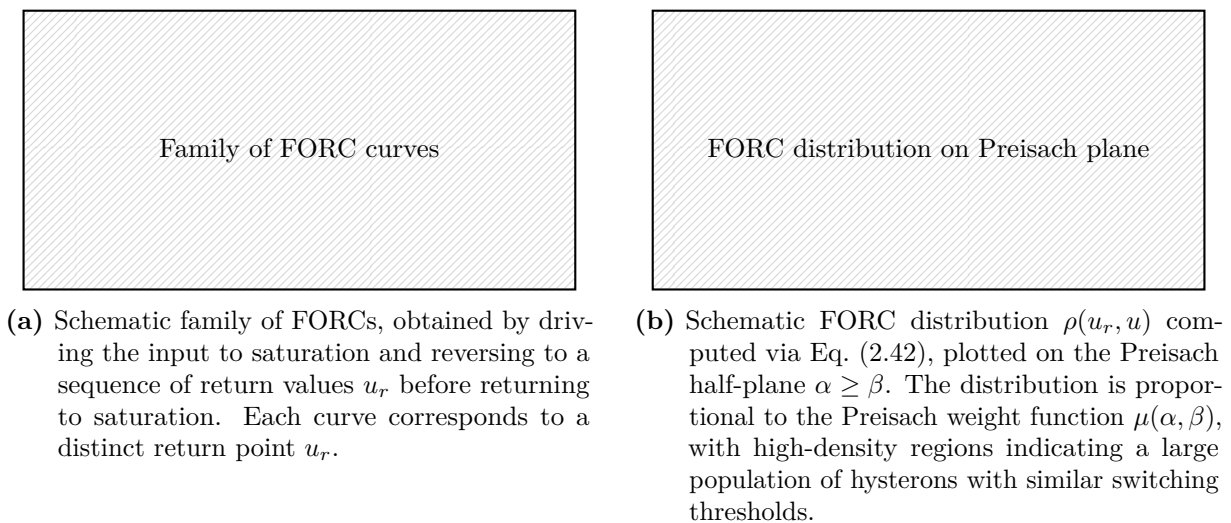


Fig. 2.13: Schematic illustration of the FORC measurement procedure and the resulting Preisach weight function. The FORC distribution in the right panel is a direct experimental estimate of the Preisach density $\mu(\alpha, \beta)$ entering Eq. (2.40).

in iron yokes are not congruent but depend on the magnetization level at which the reversal occurs [bertottiHysteresisMagnetismPhysicists1998]. This is a fundamental limitation of the classical Preisach model that cannot be resolved by refining the mesh or improving the identification procedure.

Several extensions of the classical model have been proposed to address the congruency violation. The moving Preisach model introduces a mean-field feedback term that shifts the effective input by a quantity proportional to the current output, thereby making the minor loop shape dependent on the magnetization state [bertottiHysteresisMagnetismPhysicists1998]. The nonlinear Preisach model replaces the fixed weight function with one that depends on the current output value, introducing state-dependent weighting at the cost of a more complex identification procedure. Despite these extensions, the classical and extended Preisach models share a common practical limitation, where the weight function is a static object identified from a fixed set of excitation cycles. Consequently, it cannot adapt to changes in the magnet's magnetic history or operating conditions without re-identification. Furthermore, the model does not account for rate-dependent hysteresis, nor rate-dependent in the macroscopic magnetic field such as eddy current contributions at higher ramp rates, which are discussed separately in Section 2.6.3. More recently, a differentiable implementation of the Preisach model with numerical parametrizations of the weight function has enabled gradient-based identification from arbitrary measurement sequences, which is discussed further in Section 2.5.3. **[TODO: more preisach extensions?, mention our recent work on neural network parametrization of the weight function]**

2.5.1.2 The Prandtl–Ishlinskii Model

The Prandtl–Ishlinskii model is a special case of the Preisach model in which the two-dimensional threshold distribution is replaced by a one-dimensional density over a single threshold parameter. [mayergoyzGeneralizedPreisachModel1988, brokatePhaseTransitionsHysteresis1996] Rather than employing bistable hysteron operators $R_{\alpha\beta}$, the PI model is built from play operators \mathcal{P}_r , each parametrized by a scalar threshold $r \geq 0$. [visintinDifferentialModelsHysteresis1994,

moreeReviewPlayPreisach2023] A play operator with threshold r maps an input $u(t)$ to an output $\mathcal{P}_r[u](t)$ according to

$$\mathcal{P}_r[u](t) = \max(u(t) - r, \min(u(t) + r, \mathcal{P}_r[u](t^-))), \quad (2.43)$$

where t^- denotes the previous time instant. The operator output follows the input only when the absolute difference between input and output exceeds r , and remains frozen otherwise. This produces a continuous, piecewise-linear output rather than the binary switching of the hysteron. A schematic illustration of the play operator is shown in Fig. 2.14.



Fig. 2.14: Illustration of the play operator \mathcal{P}_r with threshold r . The output follows the input with a dead-band of width r on each side, producing a continuous piecewise-linear hysteretic response. Arrows indicate the direction of traversal under a full input cycle.

The PI model output is the weighted superposition of play operator outputs over a continuous density $p(r) \geq 0$,

$$\Gamma[u](t) = \int_0^\infty p(r) \mathcal{P}_r[u](t) dr, \quad (2.44)$$

which is the one-dimensional analog of the Preisach integral in Eq. 2.40. The density $p(r)$ is the only free function of the model and is analogous to the Preisach weight function $\mu(\alpha, \beta)$, but lives on the positive real line rather than on a two-dimensional half-plane. Parameter identification therefore reduces to fitting a one-dimensional function from the measured major loop and a small set of minor loops, without requiring the full FORC measurement grid needed for the Preisach model.

The main practical advantage of the PI model is that its inverse can be computed analytically in closed form for the classical rate-independent case with a symmetric density. **[aljanaidehAnalyticalGeneralizedPrandtl2011]** This makes it attractive for feedforward compensation of hysteresis in smart actuators such as piezoelectric and magnetostrictive devices, where the inverse map is applied directly as a pre-filter. **[cite]** In the context of the present thesis, however, this advantage does not apply, since the compensation strategy is based on a data-driven surrogate model whose predictions are used in a feedforward control loop, and no closed-form inversion of a physics-based model is required.

Additionally, as a strict subclass of the Preisach model, the PI model inherits the congruency property, which states that all minor loops traced between the same pair of reversal values are congruent regardless of the magnetization state at which they are initiated. As discussed in the previous section, this property is experimentally violated in iron-yoke magnets, and represents a fundamental limitation that cannot be overcome by improving the density identification.

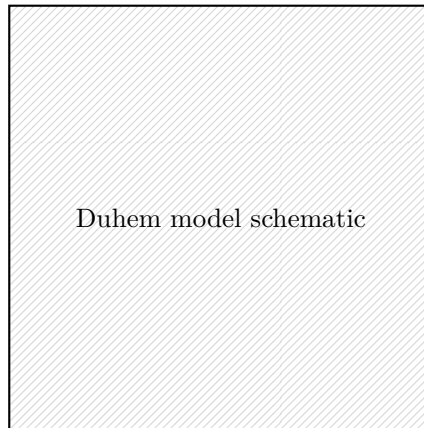


Fig. 2.15: Schematic of the Duhem model structure; a B – H plot showing ascending/descending branches, the anhysteretic curve C_a , and the upper/lower bounds C_u , C_l . This one figure introduces the whole class.

Furthermore, the classical PI model requires the hysteresis loop to be symmetric with respect to the input axis, so that ascending and descending branches are mirror images of each other. Generalized extensions that introduce asymmetric envelope functions can relax the symmetry constraint, but at the cost of losing the analytical invertibility that constitutes the primary motivation for choosing the PI model over the Preisach model. For these reasons, the PI model offers no advantage over the Preisach framework for the application considered here, and is not pursued further in this thesis.

2.5.2 ODE-type Models

ODE-type hysteresis models describe the magnetic state through a first-order ordinary differential equation in which the rate of change of magnetisation with respect to the applied field depends on the current state and the direction of the field sweep. The general form is given by the Duhem model [duhemDauerndenAnderungenUnd1900, visintinDifferentialModelsHysteresis1994],

$$\frac{dM}{dH} = \begin{cases} f^+(M, H) & \dot{H} > 0, \\ f^-(M, H) & \dot{H} < 0, \end{cases} \quad (2.45)$$

where f^+ and f^- are model-specific functions governing the ascending and descending branches respectively. An illustration of the model structure is shown in Fig. 2.15.

Hysteresis arises from the asymmetry $f^+ \neq f^-$, and the model possesses local memory represented only by the current state (M, H) . The future trajectory is then fully determined by state (M, H) and the sign of \dot{H} , with no dependence on the prior input history beyond what is encoded in M . This stands in contrast to the operator-type models, whose output is determined by a distributed memory structure over the full input history. The models discussed in the following subsections such as the Flatley model [flatleyMagneticHysteresisModel1995], the Jiles–Atherton model [jilesTheoryFerromagneticHysteresis1984], and several others are all special cases of (2.45), differing in the choice of f^\pm .

The main advantages of the ODE-type models are computational simplicity, a compact parameterisation, and, in some variants, a direct correspondence between parameters and material properties. On the other hand, several limitations apply for the general class of models. First, numerical integration of (2.45) over long or complex excitation sequences accumulates




Fig. 2.16: Model output with labeled parameters.

truncation and rounding errors, causing the simulated trajectory to drift from the true hysteresis loop. Additionally, no natural reset mechanism exists unless one is explicitly introduced [chwastekModellingDynamicHysteresis2009, zirkaViscositybasedMagnetodynamicModel2006]. Second, ODE models of the Duhem form produce congruent minor loops by construction, since the trajectory on any reversal branch is determined solely by the current state. They therefore cannot reproduce the non-congruent minor loop behaviour observed in many soft magnetic materials [mayergoyzMathematicalModelsHysteresis1991]. Third, even with few model parameters, parameter identification is a non-convex optimisation problem, and gradient-based methods are not guaranteed to find a global minimum [ledererParameterIdentificationApplication1999]. Finally, the formulation in (2.45) is by definition rate-independent, so rate-dependent hysteresis or dynamic effects arising from eddy currents, cannot be modeled with this form of models. The following subsections discuss specific models in order of increasing physical complexity and based on relevance to the topic of this thesis.

2.5.2.1 The Flatley Model

The Flatley model [flatleyMagneticHysteresisModel1995] is one of the simplest concrete realisations of the Duhem class, developed at The National Aeronautics and Space Administration (NASA) for simulating the hysteretic response of magnetic materials in aerospace applications. Its design was motivated by the geometric observation that the ascending and descending branches of a measured hysteresis loop resemble shifted arctangent curves. The major loop branches are therefore defined directly as

$$B_{\text{lim}}(H) = B_s \frac{2}{\pi} \arctan(k(H \pm H_c)), \quad (2.46)$$

where B_s is the saturation field, H_c is the coercive field, and k is a shape parameter controlling the slope of the transition region, and is defined as

$$k = \frac{1}{H_c} \tan\left(\frac{\pi B_r}{2 B_s}\right), \quad (2.47)$$

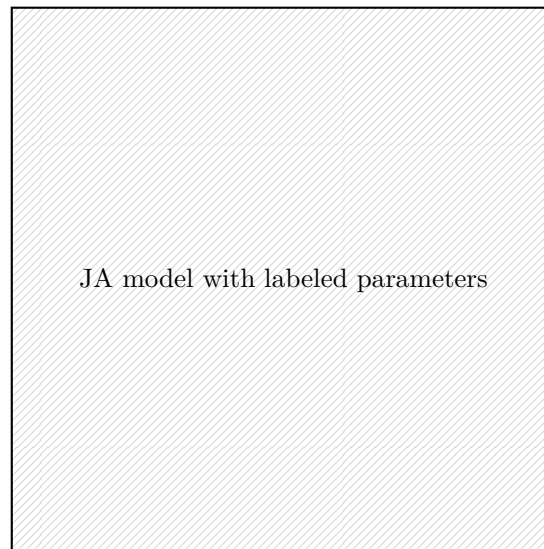


Fig. 2.17: Model output with labeled parameters M_s , H_c , B_r , optionally minor loops to illustrate drifting.

so the model is fully determined by three measurable quantities (B_s , B_r , H_c), all readable directly from the major loop. The so called interior trajectories, followed when a field reversal occurs before saturation is reached, are governed by the ODE

$$\frac{dB}{dH} = \frac{2}{\pi} k B_s \cos^2\left(\frac{\pi B}{2 B_s}\right) \left(\frac{H - H_{\text{lim}}(B)}{2H_c}\right)^2, \quad (2.48)$$

where $H_{\text{lim}}(B)$ is the inverse of (2.46). The model therefore drives any interior trajectory toward the nearest major loop branch, reproducing the qualitative behaviour of return paths without any additional parameters. An illustration of the model output can be seen in Fig. 2.16.

The model is popular among researchers primarily for its simplicity and the direct identifiability of its three parameters from standard major loop measurements, without numerical optimisation. The limitations follow from the same source, as the arctangent form is a geometric ansatz with no grounding in domain physics. Consequently, the model cannot be expected to generalise beyond the conditions under which it was fitted. Despite its limitations, the Flatley model serves as a useful baseline for model class.

2.5.2.2 The Jiles–Atherton Model

The Jiles–Atherton (JA) model [jilesTheoryFerromagneticHysteresis1984, jilesNumericalDeterminatio] derives the magnetisation response from a physical picture of domain-wall motion in a medium with pinning sites. The total magnetisation M is decomposed into an irreversible component M_{irr} , associated with domain-wall pinning, and a reversible component M_{rev} , associated with domain-wall bowing without depinning,

$$M = M_{\text{irr}} + M_{\text{rev}} = M_{\text{irr}} + c(M_{\text{an}} - M_{\text{irr}}) \quad (2.49)$$

where $c \in [0, 1]$ is the reversibility coefficient and M_{an} is the anhysteretic magnetisation, representing the equilibrium state in the absence of pinning. Consequently it is also a representation of the idealised non-hysteretic response of the material, containing saturation and nonlinearity but no history dependence. The anhysteretic curve is modelled by a Langevin function,




Fig. 2.18: Schematic showing repeated asymmetric minor loops drifting away from the major loop, contrasted with a Preisach model's closed minor loops. This explains why ODE models are limited for complex cycling histories

$$M_{\text{an}}(H_e) = M_s \left[\coth\left(\frac{H_e}{a}\right) - \frac{a}{H_e} \right], \quad (2.50)$$

evaluated at the effective field $H_e = H + \alpha M$, which accounts for interdomain coupling through the mean-field parameter α , which accounts for interdomain coupling. The irreversible magnetisation evolves according to

$$\frac{dM_{\text{irr}}}{dH} = \frac{M_{\text{an}} - M_{\text{irr}}}{k\delta - \alpha(M_{\text{an}} - M_{\text{irr}})}, \quad (2.51)$$

where k is the pinning coefficient controlling coercivity and $\delta = \pm 1$ is the sign of \dot{H} , which enforces the Duhem direction dependence of (2.45). Combining (2.49) and (2.51) yields the full ODE for M , which is integrated numerically given the internal state $(M_{\text{irr}}, M_{\text{rev}})$ at the start of the excitation sequence $H(t)$. However, this internal state is not directly observable, as a single measurement of (M, H) does not uniquely determine the decomposition in (2.49). Therefore, initialising the model at an arbitrary point on a minor branch requires additional assumptions about the prior magnetic history [**bergqvistMagneticVectorHysteresis1997**]. In practice the model is therefore typically initialised from a known reference state such as saturation or the demagnetised state, which limits its applicability when the excitation history is not fully controlled.

The JA model reproduces the qualitative features of major loops and FORCs, and its parameters can in principle be related to independently measurable material properties. However, several practical difficulties arise. The denominator of (2.51) can change sign near reversal points, producing non-physical negative differential susceptibility and numerical instability in the integration [**chwastekModellingDynamicHysteresis2009, zirkaViscositybasedMagnetodynamicModel2006**]. Parameter identification requires a non-convex optimisation over five coupled parameters, and different local minima can yield loops that are visually similar on the major loop but diverge substantially on minor branches [**ledererParameterIdentificationApplication1999**]. Like all Duhem models, the JA model produces congruent minor loops and is therefore unable to reproduce non-congruent behaviour [**mayergoyzGeneralizedPreisachModel1988**]. Numerical drift under repeated cycling, accumulates over long excitation sequences and limits the practical use of the model as an accurate surrogate model without periodic re-initialisation. Inherent minor loop drift as illustrated in Fig. 2.18 is a fundamental limitation of the ODE model structure, and cannot be resolved by improving numerical integration or parameter identification.

The five-parameter structure and physical derivation of the JA model the most generic model for magnetic hysteresis within the ODE class. However, the fundamental limitations of drift, congruency, and initialisation ambiguity remain, and are not resolved by the additional physical grounding. Despite these limitations, the JA model remains the most widely used physics-based ODE hysteresis model in the magnetics literature.

The Jiles–Atherton model reproduces the qualitative features of major loops and FORCs, and its parameters can in principle be related to independently measurable material properties. However, the denominator of (2.51) can change sign near reversal points, producing non-physical negative differential susceptibility and numerical instability in the integration [chwastekModellingDynamicHysteresis2009, zirkaViscositybasedMagnetodynamicModel2006]. Parameter identification requires a non-convex optimisation over five coupled parameters, and different local minima can yield loops that are visually similar on the major loop but diverge substantially on minor branches [ledererParameterIdentificationApplication1999]. Despite these limitations, the Jiles–Atherton model remains the most widely used physics-based ODE hysteresis model in the magnetics literature.

2.5.2.3 Other ODE-type Models

The Fröhlich model [frohlichMagneticCircuitContactless2008] is one of the earliest analytical descriptions of differential permeability in ferromagnetic materials. It expresses the differential permeability as a rational function of the magnetisation,

$$\frac{dM}{dH} = \frac{(\mu_0 - \chi_\infty)}{(1 + \beta M)^2}, \quad (2.52)$$

where μ_0 is the initial permeability, χ_∞ is the high-field susceptibility, and β is a shape parameter. In its original form the model is single-valued and describes only the anhysteretic curve, and hysteresis can be introduced by applying the rational permeability separately to the ascending and descending branches according to (2.45). The model is of primarily historical interest, but its rational functional form reappears in several later ODE formulations.

The Tellinen model [tellinenSimpleScalarModel1998] takes a different approach by avoiding a fixed functional form entirely. Instead, the differential permeability at any interior state (B, H) is obtained by interpolating between the upper and lower major loop branches $B^+(H)$ and $B^-(H)$,

$$\frac{dB}{dH} = \frac{dB^\pm}{dH} \cdot \frac{B^\pm(H) - B}{B^+(H) - B^-(H)}, \quad (2.53)$$

where the sign is chosen according to \dot{H} . Since the limiting curves enter as arbitrary inputs rather than as fitted parametric forms, the model can in principle accommodate any measured major loop without restricting the loop shape. The Tellinen model is simple to implement, but retains the congruency constraint of all Duhem models and its accuracy on minor branches depends entirely on the quality of the measured limiting curves.

The Coleman–Hodgdon model [colemanConstitutiveRelationRateindependent1986, hodgdonApplication1986] was developed specifically for soft magnetic materials subject to the unipolar or quasi-unipolar excitation typical of power engineering applications. The ODE takes the form

$$\frac{dB}{dH} = \alpha(f(H) - B) + f'(H), \quad (2.54)$$

where $f(H)$ is a user-supplied saturation function and α is a dissipation coefficient controlling the width of the hysteresis loop. The model is rate-independent and congruent, and has been

incorporated into several finite-element codes due to its simple functional form and straightforward parameter identification from the major loop.

The Bouc–Wen model [**boucForcedVibrationMechanical1967**, **wenMethodRandomVibration1976**] originated in structural mechanics as a description of hysteretic restoring forces in nonlinear oscillators. The governing equation introduces an auxiliary hysteresis variable z whose evolution is coupled to the input displacement u ,

$$\dot{z} = \dot{u} \left(A - (\beta \operatorname{sgn}(\dot{u}) z + \gamma z) |z|^{n-1} \right), \quad (2.55)$$

where A , β , γ , and n are shape parameters. Its flexible functional form allows it to reproduce a wide range of loop shapes by adjusting the exponent n and the relative weights of β and γ . In the magnetics context, the correspondence $u \leftrightarrow H$ and $z \leftrightarrow M$ maps the mechanical formulation directly onto a hysteresis ODE and has been used to model

The models discussed above are all rate-independent, in the sense that the output path depends only on the input trajectory and not on the rate at which it is traversed. Rate-dependent extensions can be constructed by adding a term proportional to \dot{H} or \dot{B} to the right-hand side of the Duhem ODE in (2.45), allowing the loop width to increase with the sweep rate. The viscosity-based model of Zirka and Moroz [**zirkaViscositybasedMagnetodynamicModel2006**] is a widely cited example, augmenting the JA framework with a dynamic excess loss term derived from statistical loss theory. It should be noted that rate dependence at the macroscopic ODE level is conceptually distinct from the microscopic spin-dynamics described by the LLG equation in (2.38), which operates at a scale below that of the domain-level models considered here. In the context of iron-yoke accelerator magnets, dynamic field errors are primarily dominated by eddy current effects treated separately in Section 2.6. Additionally to reduce complexity of the method built in this thesis, we focus on the rate-independent component of the hysteretic response, and the ODE models discussed above are all of the rate-independent form.

2.5.3 Machine Learning Approaches

Preisach-NN, NARX, Neural ODEs, PINN, all approaches with LSTMs etc.

2.6 Eddy Current Decay

When particle beams are accelerated in synchrotrons, the electromagnets are ramped by rapidly increasing the excitation to achieve the desired magnetic field for beam steering and focusing. The rapid change in current produces a time-varying magnetic flux that induces circulating currents in all conducting materials threaded by that flux. These induced currents, known as eddy currents, generate a secondary magnetic field that opposes the applied flux change according to Faraday’s law of induction. Once the current ramp ends and the excitation is held constant, the eddy currents dissipate resistively and the secondary field decays toward zero. The result is a time-decaying perturbation of the magnetic field in the magnet aperture, which persists for timescales ranging from milliseconds to several seconds depending on the geometry and conductivity of the surrounding materials.

This transient field perturbation appears after and during a fast magnet ramp, such as at injection after a previous magnet ramp-down, during acceleration, and right after the ramp-up is complete. During this period, decaying multipole field components, which frequently includes strong dipole and quadrupole components, can measurable impact the beam. The dipole component displaces the closed orbit as introduced in Section 2.1.5, while a decaying quadrupole component shifts the betatron tune and chromaticity which can affect beam stability. Additionally,

coupling between the dipole component, or in other terms beam energy, implicitly causes tune drifts as the effective focusing strength is normalized by beam rigidity. A qualitatively distinct class of dynamic field perturbations worth mentioning, is caused by the beam itself. As a charged bunch of particles travels through the resistive vacuum chamber, it excites electromagnetic wakefields that act back on trailing particles, driving transverse and longitudinal instabilities. [cite] Eddy current decay and wakefields are therefore complementary effects, but acting in significantly different timeframes. For this thesis we consider only the dynamic effects caused by the power supply history that can be feed-forward controlled.

2.6.1 Eddy Currents in Conducting Materials

A time-varying magnetic flux Φ threading a conducting loop induces an electromotive force according to a special case of Faraday's law, which drives a circulating current in the conductor.

$$\mathcal{E} = -\frac{d\Phi}{dt} \quad (2.56)$$

Here the factor N familiar from wound coils is absent since eddy current loops in a bulk conductor form single closed paths. According to Lenz's law [cite], this induced current flows in a direction that opposes the change in flux, producing a secondary magnetic field that partially cancels the applied field change. In a bulk conductor of finite resistivity, the induced currents are distributed continuously throughout the material volume and are referred to as eddy currents.

The spatial distribution of eddy currents is not uniform throughout the conductor. At high rates of change, currents are confined to a surface layer characterised by the skin depth,

$$\delta = \sqrt{\frac{2\rho}{\omega\mu}}, \quad (2.57)$$

where ρ is the electrical resistivity, μ is the magnetic permeability, and ω is the angular frequency of the field variation. A smaller skin depth confines eddy currents to a thin surface layer, reducing the effective loop area and increasing the effective resistance of the eddy current path. This penetration scale is the physical basis for lamination as a mitigation strategy, which is discussed in Section 2.6.2. At low frequencies or after the driving change has ceased, the currents penetrate further into the material and decay resistively. [TODO: Reformulate]

The decay of eddy currents after a step change in excitation can be described by analogy with a series Reinforcement Learning (RL) circuit, in which the inductance L represents the magnetic energy stored in the eddy current loop and the resistance R represents the ohmic dissipation. [TODO: Cite Buzio CAS] The governing equation for the current I_{eddy} in such a loop is

$$L \frac{dI_{\text{eddy}}}{dt} + R I_{\text{eddy}} = 0, \quad (2.58)$$

which has the solution

$$I_{\text{eddy}}(t) = I_0 \exp\left(-\frac{t}{\tau}\right), \quad \tau = \frac{L}{R}, \quad (2.59)$$

[TODO: Cite eddy current equation for RLC] where I_0 is the eddy current amplitude at the start of the flat excitation and τ is the decay time constant. The secondary magnetic field contribution in the aperture is proportional to the eddy current,

$$\Delta B_e(t) = \alpha I_{\text{eddy}}(t), \quad (2.60)$$

where α is a geometry-dependent proportionality constant that depends on the conductor geometry and its position relative to the aperture. [cite] The secondary magnetic, or in other words eddy current driven, field contribution is proportional to I_{eddy} , so the field perturbation in the aperture decays with the same time constant τ . The amplitude I_0 is set by the particular solution of the inhomogeneous Ordinary Differential Equation (ODE) driven by the excitation ramp. For a constant ramp rate $\frac{dI}{dt}$, the steady-state driven solution gives $I_0 \propto \tau \frac{dI}{dt}$, so a faster preceding ramp produces a larger initial eddy current amplitude. However, this decay picture is valid only when the excitation is held constant. During an arbitrary time-varying excitation, or even small excitation ramps follow a larger one, the full inhomogeneous ODE must be solved and the field perturbation does not take a simple exponential form. Finally, the time constants τ_k are not in general fixed material properties when discussing eddy current decay in the iron yoke, but depend on the local magnetic permeability μ of the iron, which varies with the excitation level through the nonlinear B - H curve. During low excitation I , or equivalently H , the time constant is approximately constant. However as the iron approaches saturation, μ decreases and the time constant τ becomes $\tau = \tau(\mu)$ [praegMethodsComputingTime1967].

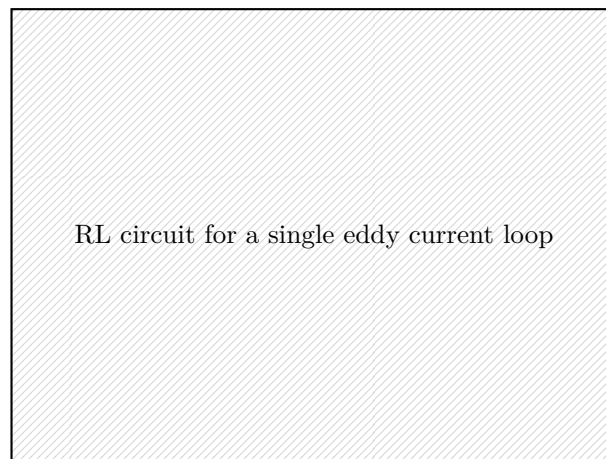


Fig. 2.19: Equivalent RL circuit model for a single eddy current loop in a conducting material surrounding an accelerator magnet. The inductance L represents the magnetic energy stored in the loop and R represents the ohmic dissipation. Following a ramp in the excitation current, the eddy current decays exponentially with time constant $\tau = L/R$.


In practice, a real conductor contains many such loops of differing geometry, each with its own time constant. The total eddy current field contribution B_e is therefore a superposition of exponentially decaying contributions,

$$B_e(t) = \sum_k A_k \exp\left(-\frac{t}{\tau_k}\right), \quad (2.61)$$

where A_k and τ_k are the amplitude and time constant of the k -th eddy current mode. This multi-exponential character will be further developed in Chapter 8.

2.6.2 Eddy Currents in Synchrotron Components

The preceding section introduced eddy currents as a general phenomenon in conducting materials, however in a synchrotron, the time-varying magnetic flux threads not only the iron yoke of the magnet but every conducting structure within and around the magnet assembly.



Cross-section of a laminated iron magnet yoke with schematic swirling eddy current loops.

Fig. 2.20: Schematic cross-section of a laminated iron yoke, showing eddy current loops confined within individual lamination sheets.

Eddy currents therefore arise in every conducting material threaded by the flux, not only in the iron yoke itself. The relevant conductors in a normal-conducting synchrotron magnet assembly include the laminated iron yoke, the beam vacuum chamber, and other mechanical support structures. Each contributes an eddy current mode with a characteristic amplitude and time constant determined by its geometry, conductivity, and proximity to the beam aperture. [moritzEddyCurrentsAccelerator2011]

2.6.2.1 Iron Yoke

The iron yoke carries the bulk of the magnetic flux and thus experiences the largest flux variation during a current ramp. Its high relative permeability $\mu_r \gg 1$ results in a small skin depth δ from (2.57), which would confine eddy currents to thin surface layers were the yoke a solid conductor. In practice, iron yokes in accelerator magnets are constructed from thin laminated sheets, electrically insulated from one another [moritzEddyCurrentsAccelerator2011]. Each sheet forms an independent eddy current loop with a cross-sectional area proportional to the sheet thickness squared, which increases the effective resistance R of each loop and reduces both the eddy current amplitude and the time constant $\tau = L/R$. This increases the total ohmic resistance R of each loop and reduces both the amplitude and the time constant $\tau = L/R$, resulting in eddy current decay in the iron on timescales of order milliseconds or shorter [taupadelmaurusLinearWeightOptimization2025]. As a consequence, although the iron yoke is the most strongly driven conductor, its eddy current contribution decays rapidly and becomes negligible on longer time scales.

2.6.2.2 Vacuum Chamber

The vacuum chamber runs through the magnet aperture and encloses a large fraction of the flux path, while being physically closer to the aperture center than the iron yoke. Unlike the iron yoke, the vacuum chamber is not laminated and is typically fabricated from stainless steel, aluminium, or copper, materials with low relative permeability $\mu_r \approx 1$. [TODO: Why these materials, need forward reference from Accelerator Physics section + SPS vacuum chamber specifics] The low permeability results in a large skin depth, so the full

wall thickness participates in eddy current conduction and forms a large effective loop with low resistance. Consequently, the resulting time constants can range from tens of milliseconds to several seconds depending on the chamber material, wall thickness, and cross-sectional geometry [**podobedovEddyCurrentShielding2009**]. These slow components are therefore the dominant contribution to eddy current decay on timescales relevant for beam dynamics, and are the primary focus of this thesis.

The transverse cross-sectional geometry of the vacuum chamber determines which multipole components the induced eddy currents produce in the aperture. For a circular cross-section, by symmetry, a dipole excitation induces eddy currents that generate only a dipole field component inside the chamber [**podobedovEddyCurrentShielding2009**], as will also be explained in more detail in Chapter 3. When the cross-section deviates from circular symmetry, as is the case for elliptical, rectangular, or flat chambers that are common in synchrotrons to accommodate the horizontal beam size, the eddy current distribution loses this symmetry. The asymmetric current distribution then generates higher-order multipole components in addition to the fundamental dipole, with the specific multipoles that arise determined by the symmetry class of the cross-section [**leeMultipoleExpansionField1990**]. For example, an elliptical or rectangular chamber with one axis of symmetry along the dipole field direction generates a sextupole component from a pure dipole excitation [**podobedovEddyCurrentShielding2009**, **tsoupasEddyCurrentEffect1999**]. Finally, each multipole mode carries its own time constant, so the overall eddy current contribution from a realistic chamber cross-section is a multi-exponential decay with mode-dependent multipole content.

While individual vacuum chambers can carry different eddy current modes, in practice, the full ring of a synchrotron is composed of many adjacent vacuum chamber sections joined together by bellows and flanges. If these joints are electrically conducting, the eddy current path is no longer confined to a single chamber section but can extend across longer sections, or even around the full ring, forming a large low-resistance loop with a correspondingly long time constant. In practice, RF shields or insulating flanges are sometimes used to break such longitudinal current paths and limit the effective loop length.

The vacuum chamber material and geometry must simultaneously satisfy competing requirements from field transparency, beam coupling impedance, and grounding continuity. A thicker or more conductive chamber wall increases both the eddy current time constant and the induced field amplitude, which is undesirable from a field quality perspective, but a thinner or more resistive wall increases the beam coupling impedance of the chamber [**podobedovEddyCurrentShielding2009**, **kaltenbacherMachineElementContribution2016**]. Therefore, the choice of vacuum chamber material and wall thickness involves a fundamental design tradeoff. Vacuum chambers in modern fast-cycling synchrotrons are therefore commonly fabricated from stainless steel, which offers a compromise between sufficient resistivity to limit eddy current amplitudes and adequate mechanical rigidity for the required aperture geometry. Where lower impedance is required, chambers are sometimes made from aluminium or copper with a thin resistive coating on the inner surface, or from ceramic with a thin metallic layer, so that the eddy current path is geometrically constrained while the beam sees a low-impedance wall. The grounding scheme of the chamber further complicates this tradeoff, as a continuous conducting path along the beam pipe allows eddy currents to propagate over long distances, whereas segmenting the chamber with insulating flanges reduces this effect at the cost of increased impedance at each junction, as was the case in the SPS [**kaltenbacherCharacterizationShieldingCERNSPS2017**, **karpovivanExperimentalConfirmationImpedance23**].

2.6.2.3 Other Conducting Structures

While the vacuum chamber and magnet yoke are the dominant sources of eddy currents in synchrotron magnets, other mechanical structures in and around the magnet aperture contribute additional induced current paths that can perturb the field seen by the beam.

Even in the magnets, the laminated yokes are typically clamped by solid steel end plates at each magnet end. Unlike the laminated body, these end plates are bulk conducting structures in which eddy currents can circulate freely, and their time constant can be significantly longer than that of the laminations. [**moritzEddyCurrentsAccelerator2011**, **taniEddyCurrentEffect2004**]

Magnet support girders and steel alignment fixtures that mechanically connect adjacent magnets can form closed conducting loops in the vicinity of the yoke. The contribution of these paths to the total eddy current response is difficult to characterise reproducibly, as it depends on the contact resistance at bolted joints, which varies with assembly and may change over time. [**moritzEddyCurrentsAccelerator2011**]

Bus bars and current leads that connect the magnet coils to the power converter run along the length of the magnet and carry the same time-varying current as the main excitation circuit. Although their distance from the beam axis limits their direct influence on the field in the aperture, inductive coupling between the leads and the surrounding conducting structures can drive additional circulating currents during rapid field ramps. [**moritzEddyCurrentsAccelerator2011**]

Finally, in-vacuum instrumentation such as BPMs electrodes introduces local conducting inserts into the vacuum chamber wall at discrete locations around the ring. These inserts locally perturb the otherwise uniform eddy current distribution in the chamber wall, and can introduce asymmetric field multipole contributions that are not captured by a simple one-dimensional model of the chamber. [**clevastefanoButtontypeBeamPosition2026**]

2.6.3 Beam Impact of Eddy Current Decay in Fast Cycling Synchrotrons

2.6.3.1 Beam Impact of Eddy Current Decay

Eddy current decay in the conducting structures of a synchrotron lattice produces time-dependent errors in all principal field components, which in turn perturb the main beam parameters. Since the orbit, momentum, tune, and chromaticity of the beam are coupled through the lattice optics, as established in Section 2.1, a perturbation to any one field component propagates into the others. The resulting beam parameter drift can therefore significantly impact injection efficiency, beam stability during acceleration, and extraction quality, or even cause full beam loss if the field perturbation is large enough to drive the beam outside the dynamic aperture or cause instabilities.

Eddy current induced field decay in the dipole component causes the actual bending field to fall below the programmed value after a change in excitation, which shifts the closed orbit of the beam as described in Section 2.1.5. At injection flat bottom, this manifests as an orbit mismatch between the incoming beam and the receiving ring, since the dipole field has not yet reached to its stationary value, and the beam arrives on a displaced orbit relative to the design trajectory.

The decay in quadrupolar field components, whether arising from the main quadrupole magnets or from asymmetric eddy current distributions in surrounding conducting structures, causes the effective focusing strength of the lattice to drift over time, producing a corresponding drift in the betatron tune [**lachaizeEstimateEddyCurrent2007**]. Additionally, since the beam circulates on a displaced closed orbit due to the dipole field decay, it experiences a different quadrupole gradient through the dispersion function in (2.4), introducing an additional tune shift through chromaticity, as described in Section 2.1.7.2. A simultaneous drift in the sextupolar field component, which can be induced by eddy currents in the vacuum chambers with specific shapes,

can further perturb the chromaticity independently of the quadrupole decay. The interplay between these effects is further complicated when an RF feedback loop is active, as the loop responds to the momentum deviation caused by the dipole field decay by adjusting the RF frequency, which in turn shifts the closed orbit through the momentum compaction factor and couples the orbit, tune, and chromaticity drifts in a non-trivial way.

In fast-cycling synchrotrons, beam is typically injected immediately after the preceding ramp-down to minimise machine downtime, so the conducting structures of the lattice have not fully reached their unperturbed remanent state. The residual eddy currents from the preceding cycle therefore set the initial conditions for the stationary decay at injection flat bottom, making the beam parameters at injection sensitive to the history of the machine cycle [lachaizeEstimateEddyCurrent2007, lu:ipac2025-wean2]. This inter-cycle dependence is particularly pronounced in the quadrupole and sextupole circuits, where the decaying field components directly perturb the tune and chromaticity on a timescale comparable to the injection and early acceleration period.

During a fast field ramp, eddy currents in the conducting structures of the lattice cause the actual field to lag behind the programmed excitation, with the magnitude of the lag determined by the amplitude and time constants of the eddy current modes [lachaizeEstimateEddyCurrent2007, chenStudyEddyCurrent2019]. Fast eddy current modes that are negligible during the stationary flat-bottom decay become significant during rapid ramps, and if the resulting field mismatch between the programmed and actual field is not compensated, it can be sufficient to prevent beam acceleration entirely [arimateaDipoleFieldTune2007]. This intra-cycle effect is not confined to the main magnets, as any conducting structure in the lattice, including vacuum chambers, support elements, and instrumentation, contributes to the total lagging field [moritzEddyCurrentsAccelerator2011]. Furthermore, chromaticity errors induced by eddy currents during the ramp can drive the head-tail instability, as has been observed during commissioning of rapid cycling synchrotrons where extra sextupole correction was not applied. [arimateaDipoleFieldTune2007]

Upon completing acceleration and entering the extraction flat top, a stationary eddy current decay reappears with initial conditions set by the preceding ramp, analogous to the inter-cycle effect at injection. An uncompensated field drift at extraction flat top perturbs the beam orbit and optics at the moment of extraction, and both fast and slow extraction schemes require the beam parameters to match the extraction channel geometry to within tight tolerances, and can therefore become degraded if not properly compensated for.

Since in the CERN SPS the main beam parameters are coupled through the orbit, momentum, tune, and chromaticity, maintaining reproducible optics requires that the ratios between dipole, quadrupole, and sextupole fields remain constant throughout the cycle. This has been achieved by applying manually determined current offsets to the quadrupole and sextupole circuits during the ramp, compensating the effect of eddy current decay and preserving cycle-to-cycle reproducibility of the beam optics [arimateaDipoleFieldTune2007].

2.6.4 Mitigation Strategies

Mitigation of eddy current effects can be addressed at the hardware, operational, and software levels, and in practice multiple approaches are combined depending on the magnitude of the effect relative to the field tolerance required for beam operation.

At the hardware level, lamination of the iron yoke is the primary mitigation, as discussed in Section 2.6.1, and the same principle can be extended to the end plates that clamp the laminated yoke, by replacing bulk steel clamps with laminated or non-conducting composite structures to suppress their slow-decaying field contribution [moritzEddyCurrentsAccelerator2011]. The vacuum

chamber wall can be slotted or segmented to break the continuous eddy current path along the beam pipe, or insulating flanges can be used to sectorise the conducting path into shorter grounded sections, as implemented in the SPS with one ground connection per half FODO cell of approximately 32 m [kaltenbacherCharacterizationShieldingCERNSPS2017]. As discussed in Section 2.6.2, both approaches reduce the effective loop area and induced field amplitude, but each discontinuity in the chamber wall introduces a contribution to the beam coupling impedance that must be accounted for in the machine impedance model [kaltenbacherCharacterizationShieldingCERNSPS2017, karpovivanExperimentalConfirmationImpedance23].

At the operational level, increasing the dwell time between successive machine cycles allows eddy currents from the preceding ramp to decay further before injection, directly reducing the inter-cycle initial condition sensitivity described in the previous subsections. This comes at the cost of reduced machine duty time, and is therefore most applicable in machines where the cycle repetition rate is not a primary operational constraint. Choosing a slower field ramp rate reduces the amplitude of intra-cycle lagging effects, since the induced eddy current amplitude scales with the rate of change of excitation, but this likewise conflicts with the demand for high cycle efficiency in fast-cycling machines.

Where the eddy current decay is sufficiently reproducible from cycle to cycle, its effect on beam parameters can be compensated by applying pre-computed feed-forward current offsets to the affected magnet circuits as a function of time within the cycle [arimateaDipoleFieldTune2007]. This approach has been applied in the CERN SPS, where manually determined offsets to the quadrupole and sextupole circuits during the ramp preserve cycle-to-cycle reproducibility of the beam optics [arimateaDipoleFieldTune2007]. Its validity is however conditional on a controlled magnetic history, since the correction waveforms are derived under specific pre-cycle conditions and lose accuracy when the supercycle structure or cycle sequence changes.

Where greater flexibility or precision is required, a physics-based or data-driven model of the eddy current decay can be used to predict the time evolution of the field error and generate a dynamic compensation waveform that adapts to varying machine histories [lachaizeEstimateEddyCurrent2007, chenStudyEddyCurrent2019]. The development and deployment of such a model for the CERN SPS main dipoles and quadrupoles is the subject of Chapter 8. Finally, if the eddy current contribution to the field error is small relative to the field tolerance required for beam operation, no active mitigation is applied and the effect is absorbed into the overall field uncertainty budget, which is typically the case for higher-order conducting structures such as instrumentation inserts and bus bars [moritzEddyCurrentsAccelerator2011].

Chapter 3

Accelerator Magnets and Magnetic Measurements

The previous chapter introduced the transverse magnetic field in an accelerator magnet as a Taylor expansion of B_y along the horizontal midplane in (2.9). This expression was sufficient to derive the equations of beam motion but restricted to the on-axis gradient structure of upright magnets. This chapter develops the full two-dimensional description of the field inside the magnet aperture, starting from Maxwell's equations in the current-free interior region, which is needed to describe the magnetic measurement methods later on. The resulting multipole expansion, expressed in polar coordinates, is a natural description of magnet design and measurements, compared to the cartesian formulation preferred by beam physicists, This form additionally makes explicit both the normal field components that enter beam dynamics and the skew components that arise from magnet imperfections.

The chapter then surveys the principal techniques used to measure these the main magnetic fields and multipole components in practice, and briefly evaluating them against the field accuracy requirements established in Chapter 2. A particular emphasis is placed on integral measurement methods, which directly yield the field quantity relevant for beam dynamics, and the systematic limitations that prevent any single technique from achieving drift-free, absolute, multipole-resolved measurements at the required level during pulsed synchrotron operation.

The chapter concludes by describing the specific laboratory measurement systems and the online magnetic monitoring system used in CERN synchrotron, including the SPS, which provide the data underpinning the field models developed in Chapters 6 and 7. The measurement accuracy limitations identified here are the primary constraint on the quality of that training data, and consequently on the compensation performance discussed in Chapter 10.

3.1 Magnetic Field in the Magnet Aperture

In the current-free interior of a magnet aperture, the transverse field is fully described by the solution to Laplace's equation, which naturally decomposes into a set of multipole basis functions. The following subsections derive this expansion and establish its connection to the beam dynamics quantities introduced in Chapter 2.

3.1.1 Multipole Expansion in Polar Coordinates

The magnetic field in the aperture interior is governed by Maxwell's equations. In the absence of free currents, Ampère's law in (2.35) in differential form reduces to

$$\nabla \times \mathbf{H} = 0 \tag{3.1}$$

and in the non-magnetic vacuum of the aperture, where $\mathbf{B} = \mu_0 \mathbf{H}$ from the constitutive relation in (2.36), this gives

$$\nabla \times \mathbf{B} = 0 \quad (3.2)$$

Which permits expressing the \mathbf{B} field through the scalar potential ϕ , $\mathbf{B} = -\mu_0 \nabla \phi$. Substituting into Gauss' law for magnetism, $\nabla \cdot \mathbf{B} = 0$, then yields Laplace's equation

$$\nabla^2 \phi = 0 \quad (3.3)$$

In polar coordinates (r, θ) , separable regular solutions of Laplace's equation are given by

$$\phi_n^{(c)}(r, \theta) = r^n \cos(n\theta), \quad \phi_n^{(s)}(r, \theta) = r^n \sin(n\theta) \quad (3.4)$$

for integer $n \geq 0$, since regularity at the origin eliminates all negative powers of r . The general solution for the transverse field components can be written compactly by introducing the complex position variable $\zeta = r e^{i\theta}$ and an arbitrarily chosen reference radius R_{ref} as

$$B_y(r, \theta) + iB_r(r, \theta) = \sum_{n=0}^{\infty} (B_n + iA_n) \left(\frac{\zeta}{R_{\text{ref}}} \right)^n, \quad (3.5)$$

where B_n and A_n are the normal and skew harmonic coefficients of order n . The normal coefficients B_n describe field components produced by an upright magnet, while the skew coefficients A_n describe components rotated by $90^\circ/n$ in the transverse plane, which arise normally from magnet tilts, assembly asymmetries, or manufacturing imperfections. With this convention, $n = 0$ corresponds to a dipole field, $n = 1$ to a quadrupole, and $n = 2$ to a sextupole.

The reference radius R_{ref} is chosen to lie within the good-field region of the aperture, where the field harmonics can be evaluated for coefficient identification. **[TODO: Write about why magnet designers prefer polar coordinates]** Through this choice, the expansion in (3.5) converges throughout the entire source-free interior,

3.1.2 Cartesian Representation

The polar expansion in (3.5) is the complete and exact description of the transverse field in the aperture. While Chapter 2 introduced only a special case of this expansion in cartesian coordinate, B_y , along the horizontal midplane $y = 0$ this subsection provides the complete multipole expansion.

Any transverse position can be written equivalently in polar coordinates (r, θ) or Cartesian coordinates (x, y) , related by $x = r \cos \theta$ and $y = r \sin \theta$. Introducing the Cartesian complex coordinate $z = x + iy$, the polar variable $\zeta = r e^{i\theta}$ and z are identical, since $r e^{i\theta} = r \cos \theta + ir \sin \theta = x + iy$. The expansion (3.5) therefore reads equivalently as

$$B_y(x, y) + iB_x(x, y) = \sum_{n=0}^{\infty} (B_n + iA_n) \left(\frac{z}{R_{\text{ref}}} \right)^n \quad (3.6)$$

where B_r in the polar form has become B_x in the Cartesian form through the standard rotation between radial and Cartesian field components.

Setting $y = 0$, so that $z = x$ is real, and restricting to an upright magnet with $A_n = 0$, the imaginary part of (3.6) vanishes and the vertical field reduces to

$$B_y(x, 0) = \sum_{n=0}^{\infty} B_n \left(\frac{x}{R_{\text{ref}}} \right)^n \quad (3.7)$$

which is precisely the Taylor series in (2.9) from Chapter 2. The normal harmonics B_n are therefore related to the on-axis derivatives of B_y by

$$B_n = \frac{R_{\text{ref}}^n}{n!} \left. \frac{\partial^n B_y}{\partial x^n} \right|_0 \quad (3.8)$$

and the normalized strengths k_n entering the equations of motion follow directly from (2.11) and (3.8) as

$$k_n = \frac{1}{B\rho} \frac{B_n}{R_{\text{ref}}^n} = \frac{B_{\text{ref}}}{B\rho} \frac{b_n}{R_{\text{ref}}^n} \quad (3.9)$$

which removes the dependence on particle momentum through $B\rho$.

Finally, the midplane projection in (3.7) captures only the normal components of an upright magnet. The full expansion in (3.6) carries two additional structures that the Chapter 2 treatment omits. First, the horizontal field component B_x is not independent but is determined by the same coefficients through the imaginary part of (3.6), and is non-zero for any particle displaced vertically from the midplane. Second, the skew coefficients A_n are zero only for a perfectly upright magnet with full reflective symmetry about the horizontal midplane. Manufacturing tolerances, mechanical misalignment, and residual magnetization from asymmetric excitation histories all generate non-zero skew components in practice. A non-zero A_n rotates the n -th multipole pattern by $90^\circ/n$ in the transverse plane relative to the corresponding normal component, which couple the two planes of beam motion. The implications of forbidden and skew harmonics for magnet symmetry and field quality are discussed further in Section 3.1.4.

3.1.3 Field Quality Conventions

The coefficients B_n and A_n entering (3.6) carry physical dimensions of magnetic flux density, and their magnitude depends on the choice of reference radius R_{ref} . For magnet characterisation and comparison across different magnet families, it is convenient to work with dimensionless field quality coefficients, defined by normalising to the principal field component B_{ref} ,

$$b_n = \frac{B_n}{B_{\text{ref}}}, \quad a_n = \frac{A_n}{B_{\text{ref}}}, \quad (3.10)$$

where B_{ref} is the main harmonic of the magnet type under consideration: B_0 for a dipole, B_1 for a quadrupole, and so on, following the same convention as in Chapter 2. The coefficients b_n and a_n thus express each multipole component as a fraction of the design field evaluated at R_{ref} .

By convention in accelerator magnet metrology, b_n and a_n are reported in units where 1 unit = 10^{-4} , so that a coefficient of 1 unit corresponds to a field error of 10^{-4} of the reference field at R_{ref} . The normalized strengths k_n used by accelerator physicists are related to b_n through (3.9) as

$$k_n = \frac{B_{\text{ref}}}{B\rho} \frac{b_n}{R_{\text{ref}}^n}, \quad (3.11)$$

which makes explicit that accelerator physicists require the absolute value of B_{ref} and the beam rigidity $B\rho$ to translate the dimensionless b_n into a focusing strength, whereas magnet builders work entirely with the relative coefficients b_n and a_n at a fixed R_{ref} . In practice, beam physicists evaluate the equations of motion using the longitudinal integral $\int B_n dz$, a quantity that the normalized strength k_n in (3.9) represents by implicitly assuming the field is concentrated over an effective hard-edge length. This distinction between relative coefficients reported by measurement systems and absolute integrated strengths required by beam dynamics drives the

accuracy requirements on the measurement techniques discussed the next section about Magnetic Measurements.

3.1.4 Forbidden Harmonics and Magnet Symmetry

The multipole expansion in (3.5) admits, in principle, all harmonic orders $n \geq 0$. In practice, the geometry of a magnet imposes discrete symmetries on the scalar potential ϕ in the aperture, which directly constrain which coefficients B_n and A_n in (3.5) can be non-zero.

The general solution to Laplace's equation in polar coordinates is a superposition of the basis functions from (3.3), each corresponding to a distinct multipole order n . A $2p$ -pole magnet has iron pole geometry that is invariant under rotation by π/p , combined with a reversal of the field direction, since the poles alternate in sign around the aperture. Consequently, the scalar potential must satisfy

$$\phi\left(r, \theta + \frac{\pi}{p}\right) = -\phi(r, \theta). \quad (3.12)$$

Substituting the basis functions of (3.3) into (3.12) gives

$$r^n \cos\left(n\theta + \frac{n\pi}{p}\right) = -r^n \cos(n\theta), \quad (3.13)$$

which is satisfied if and only if $\cos(n\pi/p) = -1$, that is, when $n\pi/p$ is an odd multiple of π , which requires

$$\frac{n}{p} = 2k - 1, \quad k = 1, 2, 3, \dots, \quad (3.14)$$

or equivalently $n = p(2k - 1)$. Any harmonic whose order does not satisfy (3.14) produces a basis function that cannot be consistent with the field antisymmetry imposed by the magnet geometry, and its coefficient must therefore vanish identically in a geometrically perfect magnet. These are the forbidden harmonics of the magnet.

For a dipole, $p = 1$, and the condition (3.14) requires $n = 1, 3, 5, 7, \dots$, so all even harmonics are forbidden and only odd normal harmonics are allowed. For a quadrupole, $p = 2$, giving $n = 2, 6, 10, 14, \dots$, so the sextupole b_3 , the octupole b_4 , and all other non- $(4k - 2)$ harmonics are forbidden. For a sextupole, $p = 3$, giving $n = 3, 9, 15, \dots$. These allowed sets and the corresponding lowest-order error harmonics are summarized in Table 3.1.

Tab. 3.1: Allowed normal harmonics b_n for the principal SPS magnet families, derived from the rotational antisymmetry condition (3.14). All other harmonics are forbidden and vanish in a geometrically perfect magnet.

Magnet type	Allowed normal harmonics	Lowest allowed error harmonic
Dipole	$n = 1, 3, 5, 7, \dots$	b_3 (sextupole)
Quadrupole	$n = 2, 6, 10, 14, \dots$	b_6 (duodecapole)
Sextupole	$n = 3, 9, 15, \dots$	b_9 (18-pole)

The operationally relevant consequence of this structure is that the lowest-order field error in a dipole is the sextupole component b_3 . Since b_3 is allowed by symmetry, it is generated whenever the iron yoke departs from a perfectly uniform magnetization state, whether from manufacturing tolerances, mechanical asymmetry, or, most pertinently for this thesis, from the history-dependent domain configuration of the iron. The consequences of b_3 variation on the beam in the SPS dipoles are discussed in Chapter 10.

For the SPS main quadrupoles, the symmetry forbids b_1 , b_3 , b_4 , b_5 , and all non- $(4k - 2)$ components. The lowest allowed error is b_6 , a duodecapole, whose beam dynamics effect scales as $(r/R_{\text{ref}})^5$ relative to the main quadrupole gradient, making it negligible for beams well within the aperture.

The harmonic structure identified here also determines which field components are accessible to a given measurement technique. This becomes relevant when evaluating the accuracy limits of integral measurements along the magnet axis, as discussed in Section 3.2.

3.1.5 Fringe Fields and the Full-Length Measurements

The multipole expansion of Section 3.1.2 is strictly valid in the transverse cross-section of a magnet with infinite length. In a real magnet of finite physical length L_{phys} , the field does not terminate abruptly at the iron boundaries but decays continuously over a fringe-field region whose longitudinal extent is set by the aperture size and pole geometry. Each multipole coefficient B_n therefore becomes a function of the longitudinal coordinate z , with a profile $B_n(z)$ that differs from the constant body value in both magnitude and multipole content.

From the perspective of beam dynamics, a particle traverses the full length of the magnet including the fringe regions, and the physically relevant quantity is the longitudinal integral of each multipole component,

$$\int B_n \, d\ell = \int_{-\infty}^{+\infty} B_n(z) \, dz. \quad (3.15)$$

The integrated normalized strength entering the equations of motion follows from Eq. (3.11) as

$$\int k_n \, d\ell = \frac{\int B_n \, d\ell}{(B\rho) R_{\text{ref}}^n}, \quad (3.16)$$

which extends the hard-edge definition of k_n to the full field profile. The integral quantities are often referred to as $B \, d\ell$, or simply $B \, dL$, and $k_n \, d\ell$ or $k_n L$ in the accelerator physics literature and spoken convention. The effective magnetic length is then defined as

$$L_{\text{eff}} = \frac{\int B_n \, d\ell}{B_n(0)}, \quad (3.17)$$

so that $\int k_n \, d\ell = k_n \cdot L_{\text{eff}}$. In the hard-edge approximation, L_{eff} is replaced by L_{phys} , which is exact only when the fringe contribution vanishes. For long magnets such as the SPS dipoles from Section 2.3.2, the fringe region is short relative to the magnet body and the approximation $L_{\text{eff}} \approx L_{\text{phys}}$ is typically acceptable for the main dipole component B_1 . The approximate average field,

$$\langle B_n \rangle = \frac{\int B_n \, d\ell}{L_{\text{phys}}} = B_n(0) \frac{L_{\text{eff}}}{L_{\text{phys}}}, \quad (3.18)$$

provides a single-number characterisation equivalent to $\int B \, d\ell$ once L_{phys} is known. The schematic longitudinal profile of the dipole field is illustrated in Fig. 3.1.

In addition to capturing nonuniform variations of the longitudinal field profile, the fringe field also modifies the higher-order multipole content of the integrated field. In the magnet body, the allowed harmonics are constrained by the rotational antisymmetry of the pole geometry, as established in Section 3.1.4. Near the magnet ends this symmetry is broken by the three-dimensional field structure [**russenschuckFieldComputationAccelerator2010**], and additional multipole components appear in the fringe region even in a geometrically perfect magnet. The integrated sextupole coefficient $\int B_3 \, d\ell$ of a dipole is therefore non-zero even in the absence of any iron imperfection or magnetization error. For higher-order multipoles, the fringe contribution can

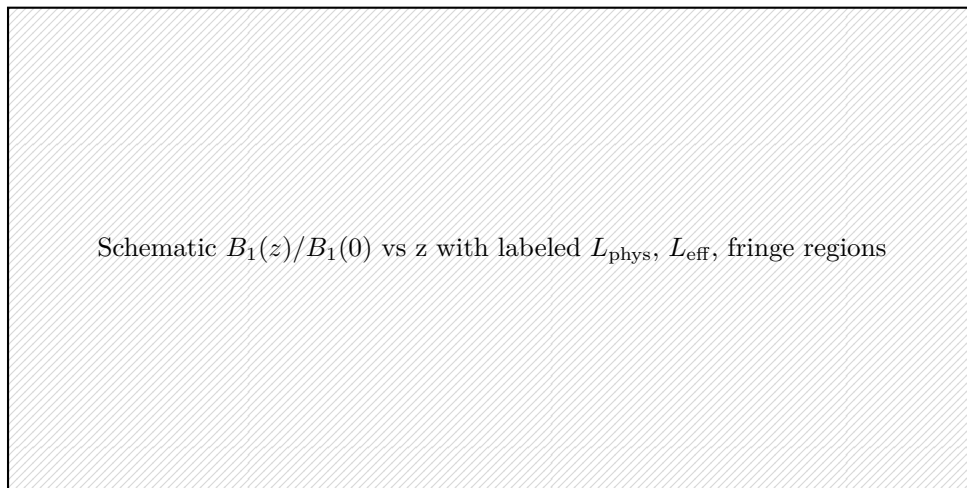


Fig. 3.1: Schematic longitudinal profile of the main dipole field component $B_1(z)$, normalised to the body value $B_1(0)$. The physical length L_{phys} is defined by the iron geometry. The effective magnetic length L_{eff} is obtained by integrating the full profile including the fringe regions at each end, indicated by the shaded areas.

constitute a significant fraction of the total integrated value, since the body contribution to b_3 is small in a well-designed dipole.

The geometry of the dipole further determines the character of the fringe field through the distinction between sector and Cartesian dipoles. A sector dipole has pole faces perpendicular to the reference orbit, so the beam enters and exits at normal incidence and the fringe field profile is symmetric about the pole face. A Cartesian, or rectangular, dipole has straight pole faces parallel to the magnet axis, so the fringe field profile and its higher-order multipole content differ from the sector case [lindbergFringeFieldMaps2023]. The SPS MBA and MBB dipoles are of the Cartesian type, with differing aperture cross-sections as described in Section 2.3.2, which consequently can produce slightly different fringe-field extents and consequently different values of L_{eff} at the same excitation current, but is not covered in this thesis.

Remember to cite [taupadelmaurusLinearWeightOptimization2025] field profiles somewhere.

3.2 Measurement Techniques

Magnetic measurements of accelerator magnets serve a range of purposes, including the characterisation of field quality and multipole content during magnet production and subsequent validation, and the determination of magnetic axis position and alignment. In contrast, accelerator physicists are interested in the measurement of integrated field strengths for beam optics calculations, and the monitoring of field reproducibility during accelerator operation. Each application places different requirements on the measurement technique. Production characterisation requires multipole resolution across the full aperture, alignment requires knowledge of field centre position, and is often performed at a single axial position, with fixed and reproducible excitation. On the other hand, beam dynamics requires the integrated field along the particle trajectory with non-static excitation, and operational monitoring similarly requires pulsed excitation and calibration stability outside of a controlled laboratory environment, in addition. No single instrument satisfies all of these requirements simultaneously, and in practice a combination of complementary techniques is used. A further distinction exists between absolute accuracy, which quantifies how

closely a measurement reflects the true field, and cycle-to-cycle reproducibility, which quantifies whether the same excitation history yields the same measured signal independent of drift or offset errors. These two notions of measurement quality place different and not always compatible demands on the instrumentation. The following subsections review the principal measurement techniques, organised by the physical quantity they provide and their compatibility with pulsed operation.

3.2.1 Local, Integral, and Average Measurements

Because the beam-dynamics quantity of interest is $\int B_n d\ell$ from Eq. (3.15), magnetic measurement techniques are classified by which field quantity they directly provide. A local measurement yields the field $B_n(x, y, z)$ at a single point or cross-section, and it cannot determine L_{eff} or the fringe-field multipole contributions without an additional axial scan of the full magnet length. An integral measurement yields $\int B_n d\ell$ directly from a single measurement along the z axis, and an average measurement yields $\langle B_n \rangle$ over the coil cross-section and length, related to $\int B_n d\ell$ through Eq. (3.18). The principal instruments belonging to each class are listed in Table 3.2.

Tab. 3.2: Classification of magnetic measurement techniques by the field quantity they provide directly.

Class	Quantity provided	Principal instruments
Local	$B_n(x, y, z)$	Hall probe, NMR probe
Integral	$\int B_n d\ell$ over coil length	Rotating coil, fluxmeter
Integral	$\int B_n d\ell$ over full length	Stretched wire, long fluxmeter

Integral measurements are the primary methods for beam dynamics characterisation, since they directly yield $k_n L_{\text{eff}}$. Local measurements are simpler, more flexible, and better suited to pulsed operation, and complement integral methods for qualitative field checks, transverse profile mapping, and axial scans to reconstruct the longitudinal field distribution. The accuracy and practical limitations of each technique are discussed in Section 3.3. The following subsections describe each instrument in turn, primarily with focus on their application for magnetic measurements suitable for accelerator beam dynamics.

3.2.2 Hall Probes

A Hall probe exploits the Hall effect, where a current-carrying semiconductor placed in a transverse magnetic field B develops a voltage V_H perpendicular to both the current and field directions according to

$$V_H = R_H \frac{IB}{d}, \quad (3.19)$$

where R_H is the Hall coefficient of the semiconductor material, I is the bias current fed into the semiconductor, and d is the thickness of the sensing element. The output voltage is proportional to the local flux density B at the probe location, making the Hall probe a local point measurement device. A typical probe assembly consists of the semiconductor chip mounted on a small printed circuit board, with total dimensions on the order of a few millimetres, which allow straightforward installation into magnet apertures, either on coil windings and iron yokes directly, or on mechanical supports for fixed-position monitoring. Additionally, a typical hall probe can be sampled at rates of several kilohertz, or even megahertz for specialised fast probes, which make them very interesting for pulsed operation and field monitoring during accelerator ramps.

The combination of small physical size, mechanical flexibility, low cost, and the ability to sample the field continuously at rates well above any synchrotron ramp rate makes Hall probes among the most practical field monitoring tools available. A probe installed at a fixed position can track the absolute field through an entire machine cycle, capturing cycle-to-cycle variations including those arising from hysteresis and eddy current decay, which make them a candidate for the magnetic measurements required for this thesis.

An illustration of a Hall probe assembly and its placement in a magnet aperture is shown in Figs. 3.2 and 3.3.

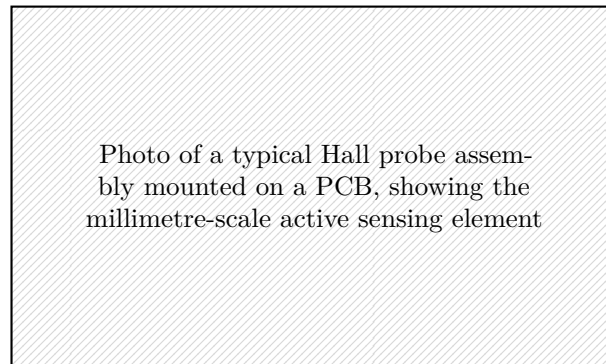


Fig. 3.2: Photograph of a typical Hall probe assembly mounted on a PCB, illustrating the millimetre-scale active sensing element.

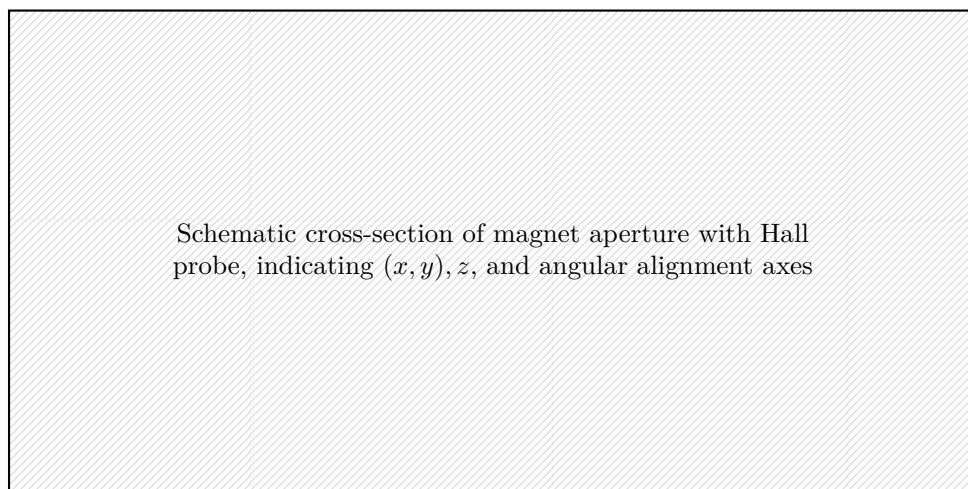


Fig. 3.3: Schematic placement of a Hall probe in a dipole magnet aperture cross-section, indicating the sensitivity of the measurement to transverse position (x, y) , longitudinal position z along the magnet axis, and angular misalignment of the probe axis relative to the field direction.

Despite these practical advantages, achieving absolute field accuracy below the 1×10^{-4} T level with a Hall probe is intrinsically challenging, due to a combination of calibration, thermal, and mechanical limitations that each contribute errors comparable to the target accuracy and precision. Firstly, Hall probes require external calibration against an independently established reference field, typically provided by Nuclear Magnetic Resonances (NMRs) or a primary field standard [sanfilippoHallProbesPhysics2011]. The calibration must be performed individually for each probe, since the Hall coefficient, offset, and temperature sensitivity differ between units

even within the same production batch. Additionally, the Hall coefficient R_H depends on both the applied field magnitude, introducing a nonlinearity of typically $\leq \pm 0.1\%$ up to 1.5 T for well-characterised devices [**A**sensor:**HE244X**], which consequently require calibration at multiple field points across the measurement range to achieve the high accuracy. Temperature variations on the semiconductor can also produce a sensitivity drift of approximately -150 ppm K for GaAs-based sensors [**cite**] such as those discussed in 3.4. In controlled laboratory conditions, temperature can be kept approximately constant, however in an accelerator tunnel, temperature gradients of several kelvin can exist along the magnet string, and the local yoke temperature itself varies with the machine cycle due to resistive heating of the coils, which would make the effective probe temperature difficult to track. A tunnel environment further introduces ionising radiation, which is known to affect the characteristics of semiconductor devices over time, causing gradual shifts in the calibration, or are outright destructive to the sensor. Separate from sensitivity drift, the zero-field offset of the probe output is generally not stable over time, and is particularly problematic near zero field, which is relevant for cycled magnets that leave a small remanent field at the end of the cycle. Beyond intrinsic sensor errors, the measurement is also sensitive to mechanical alignment and positioning, with the severity depending on the local field gradient at the probe location [**sanfilippoHallProbesPhysics2011**].

Because the probe provides only a local field value at a single point, an additional limitation arises when it is used as an absolute field reference. This relies on the assumption that the local field at the probe position is representative of the average field in the magnet body, which holds in the uniform central region of a long dipole. However this assumption but breaks down near fringe regions, for higher-order multipole components, and whenever the field distribution across the aperture is non-uniform. The absolute accuracy achievable in this configuration is therefore bounded both by the probe limitations described above and by the validity of this representativeness assumption.

3.2.3 NMR Probes

NMR probes provide an absolute reference for the magnetic flux density by exploiting the Larmor precession of nuclear spins in a magnetic field. [**cite**] In a homogeneous field with magnitude B , the nuclear precession frequency f_0 is proportional to the field strength according to

$$f_0 = \frac{\gamma}{2\pi} B, \quad (3.20)$$

where γ is the gyromagnetic ratio of the chosen nucleus, typically the proton for magnet metrology. Modern frequency standards determine f_0 with very high precision, such that NMR probes can reach absolute field accuracies of order 1×10^{-7} T in well controlled laboratory conditions. A schematic illustration of the NMR measurement principle is shown in Fig. 3.4.

The main advantage of NMR probes compared to Hall probes is that the measured quantity is a frequency, rather than an electrical voltage proportional to the Hall coefficient. As a consequence, NMR measurements are largely immune to temperature drifts, long term gain changes, and ageing of the electronics, and do not require frequent recalibration once the gyromagnetic ratio and frequency reference are fixed. In terms of attainable absolute accuracy under static conditions, NMR therefore provides an upper bound that is typically one to two orders of magnitude better than what can realistically be achieved with Hall probes on the same accelerator magnet measurement bench.

The advantage of the NMR however come at significant tradeoffs. To obtain a narrow and well defined resonance line, the magnetic field must be highly homogeneous across the active sample volume and essentially constant over the full measurement time, which typically ranges




Fig. 3.4: Schematic illustration of the NMR measurement principle, showing a sample volume in a homogeneous dipole field, excitation and detection coil, and the relation between resonance frequency and magnetic flux density.

from milliseconds to seconds depending on the signal-to-noise ratio and the required precision. NMR are therefore generally incompatible with pulsed operation in an accelerator ramp, since the temporal variations of the field during the ramp are too rapid to allow for a stable resonance signal. An additional consequence of the homogeneity requirement is that NMR probes are not suitable for measuring higher-order multipole components, since the field gradients across the sample volume broaden the resonance and reduce the signal quality. Classical NMR teslameters therefore operate in quasi static conditions in dedicated reference magnets, and are not a universal high accuracy field measurement solution for accelerator magnets, but rather serve as calibration references or highly accurate field monitors at fixed fields.

Recent commercial developments partially address these limitations. Some modern NMR teslameters integrate three dimensional Hall sensors and advanced locking schemes, using the Hall readings to steer the excitation frequency close to the expected resonance and to track slow field variations. [cite] Cayler and Metrolab tech Such hybrid systems extend the range of fields and operating conditions where an NMRs lock can be maintained, and are being adopted in accelerator magnet metrology at the time of writing. However, the measurement bench employed in this thesis uses an earlier generation NMR system without these capabilities, and the main focus of this work is on fast cycling synchrotron operation where the field dynamics remain incompatible with continuous NMR tracking. For these reasons, Hall probes are used as the primary local field sensor on the pulsed measurement bench, while NMR defines the static reference accuracy that cannot be reached during dynamic operation.

3.2.4 Stretched Wire

The stretched wire technique is one of the simplest and most fundamental methods for measuring magnetic fields in accelerator magnets. A taut conducting wire is placed along the magnet aperture and moved in a static field, and the induced voltage is integrated to obtain the magnetic field integral along the prescribed trajectory. This makes the technique a simple and straightforward reference for integrated field measurements in magnet metrology and establishes a direct link between the mechanical geometry of the magnet and the accelerator-physics quantities used in optics and beam dynamics.

In its most basic configuration, a single conducting wire is stretched along the nominal magnet axis and forms a closed loop with a return conductor outside the field region. When the wire is displaced or translated in a static magnetic field, Faraday's law in Eq. (2.56) relates the induced voltage to the time derivative of the magnetic flux through the loop. By integrating the voltage signal over the motion, one directly obtains the line integral of the relevant field component along the wire path, such as the vertical field integral

$$\int B_y dl \quad (3.21)$$

for a horizontal-bend dipole or the gradient integral for a quadrupole. A longitudinal scan of the wire at fixed transverse position provides the integrated main field and the effective magnetic length of the magnet, while a transverse scan across the aperture yields the transverse distribution of the integrated field. Using the effective length definition in Eq. (3.17), this integral gives

$$\int B_1(z) dz = B_{1,0} L_{\text{eff}}, \quad (3.22)$$

where $B_{1,0}$ is the nominal body field and L_{eff} is the effective magnetic length from the previous subsection.

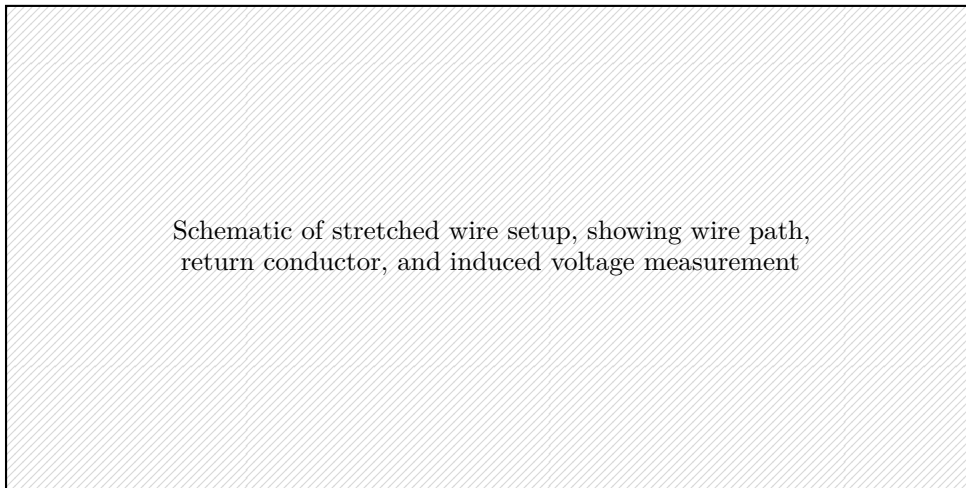


Fig. 3.5: Schematic illustration of the stretched wire technique, showing the wire path along the magnet axis, the return conductor outside the field region, and the measurement of the induced voltage during controlled motion of the wire.

The stretched wire method is also directly related to the multipole formalism introduced in Section 3.1.2. In the complex representation of the transverse field, the field components are expanded as

$$B_y(x, y, z) + iB_x(x, y, z) = \sum_{n=1}^{\infty} (B_n(z) + iA_n(z)) (x + iy)^{n-1}, \quad (3.23)$$

where B_n and A_n denote the normal and skew multipole coefficients as a function of the longitudinal coordinate. A stretched wire longitudinal scan at fixed transverse position directly measures the integrated multipoles

$$\int B_n(z) dz, \quad \int A_n(z) dz, \quad (3.24)$$

which are proportional to the integrated strengths $k_n L_{\text{eff}}$ that enter the description of linear and nonlinear optics. Transverse scans of the wire across the aperture at selected longitudinal positions provide the dependence of the integrated field on transverse position and allow the determination of the magnetic axis and roll from symmetry conditions on the measured signal.

In practice, stretched wire measurements are used to determine the integrated field strength $k_n L$, effective magnet length L_{eff} , and magnetic axis of accelerator magnets in static conditions. The technique is also applied to fiducialization [cite], by relating the magnetic axis and roll to external mechanical reference features, and to the determination of low-order integrated multipole errors from transverse scans of the integrated signal. The achievable accuracy is set by mechanical uncertainties in the wire position, straightness, and sag, which limit the resolution for higher-order multipoles and require well-controlled supports and metrology to reach relative errors at the level of at most a few 10^{-5} on the main component under optimal conditions.

The method intrinsically relies on controlled mechanical motion of the wire and integration of the induced voltage during that motion. As a result, it is restricted to static and is not used to follow fast ramps, or characterize field decays like eddy currents. For systematic harmonic analysis beyond the lowest orders, and for time-resolved measurements during dynamic cycles, rotating coils and coil-based fluxmeter systems presented in the next sections are preferred. In this thesis, the stretched wire technique mainly provides a reference for integrated field strength and magnet axes in static conditions, while the dynamic characterisation of the SPS magnet system relies on fluxmeter measurements described in the following subsection.

3.2.5 Fluxmeters

Fluxmeters are integral magnetic measurement devices that infer the longitudinal field integral from the voltage induced in a sensing coil according to Faraday's law. For a coil with N turns and instantaneous flux $\Phi(t)$ through each turn, the induced voltage is

$$V(t) = -N \frac{d\Phi(t)}{dt} \quad (3.25)$$

so that an ideal electronic integrator with output signal $U(t)$ realises

$$U(t) = - \int_{t_0}^t V(t') dt' = N[\Phi(t) - \Phi(t_0)] \quad (3.26)$$

The instrument therefore provides the change of magnetic flux linked to the coil between an initial reference time t_0 and the measurement time t . In accelerator magnet metrology, the coil is arranged to span the length of the magnet aperture so that the flux is expressed as a longitudinal integral over the field components of interest [russenschuckFieldComputationAccelerator2010]. A schematic placement of a typical fluxmeter coil in a magnet aperture is shown in Fig. 3.6.

In practice, a fluxmeter system for accelerator magnet characterisation consists not of a single coil but of an array of elements placed at different transverse positions within the aperture, as illustrated in Fig. 3.7. A minimal system for dipole field measurement uses a single element whose calibrated area yields the average field, while multipole characterisation requires at least two elements connected in a compensated arrangement that suppresses the dominant harmonic and isolates the target order. More complete systems can employ five or more at distinct radii, and high-resolution implementations reach twenty or more elements, enabling simultaneous extraction of several harmonic orders from a single excitation cycle [russenschuckFieldComputationAccelerator2010]. The coils are traditionally hand-




Fig. 3.6: Schematic placement of a fluxmeter coil inside a magnet aperture, indicating the coil length, transverse position, and relation to the magnet yoke.

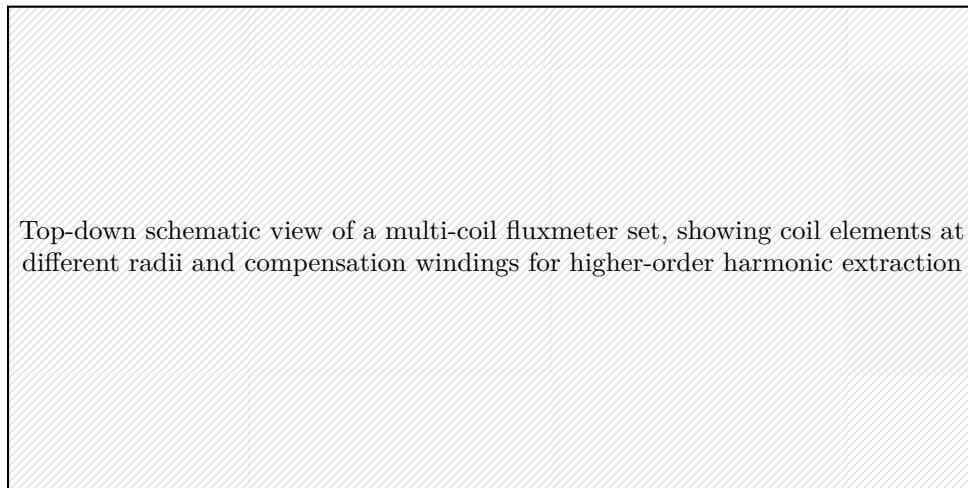


Fig. 3.7: Top-down schematic view of a multi-coil fluxmeter set. Each element is placed at a distinct radius $R_{\text{coil},k}$ and samples a different linear combination of multipole components. Compensation windings, formed by connecting elements in series with opposite polarity, suppress the dominant harmonic and isolate higher-order contributions.

wound on mechanical supports, but modern implementations increasingly use Printed Circuit Board (PCB) fluxmeters to achieve higher precision and reproducibility in the coil geometry.

The flux linked to element k , placed at radius $R_{\text{coil},k}$ with N_k turns, is related to the multipole expansion introduced in Section 3.1.2. Using

$$B_y(x, y, z) + iB_x(x, y, z) = \sum_{n=1}^{\infty} [B_n(z) + iA_n(z)] \left(\frac{x + iy}{R_{\text{ref}}} \right)^{n-1}, \quad (3.27)$$

and neglecting small variations of B_n across the coil thickness, the contribution of the normal multipole of order n to the flux through element k scales as [russenschuckFieldComputationAccelerator2010]

$$\Phi_{n,k} \approx N_k R_{\text{coil},k}^n \int_0^{L_{\text{coil}}} \frac{B_n(z)}{R_{\text{ref}}^{n-1}} dz \quad (3.28)$$

Each element therefore measures a different linear combination of integrated multipole strengths, weighted by its geometry. By constructing and inverting the matrix of geometric weights across all elements, the individual integrated multipoles $\int B_n dz$ can be recovered simultaneously up to the order supported by the number of independent elements [russenschuckFieldComputationAccelerator2010]. For a pure dipole with a nearly uniform field over the coil cross-section, the expression simplifies to the total effective area $A_{\text{tot},k} = N_k A_{\text{eff},k}$, and the change of average dipole field follows directly as

$$\Delta B_{\text{avg}} = \frac{\Delta U_k}{A_{\text{tot},k}} \quad (3.29)$$

where ΔU_k denotes the change of the integrator output. The absolute field is then recovered by accumulating changes from a reference time t_0 at which the field is independently known,

$$B_{\text{avg}}(t) = \sum_{t_0}^t \frac{\Delta U_k(t')}{A_{\text{tot},k}} + B_{\text{avg}}(t_0) \quad (3.30)$$

Because the induced voltage $V(t)$ is measurable in real time during magnet ramps, fluxmeters are naturally suited to pulsed magnetic measurements compared to the previous techniques. The voltage from the coils is digitised and integrated numerically at sampling rates from the kilohertz regime up to several hundred megahertz, depending on the excitation bandwidth and noise floor. Dedicated Fast Digital Integrators (FDIs) are employed when high sampling rates and long integration windows are required simultaneously, but the underlying principle remains the time integration of Eq. (3.25). This dynamic capability allows fluxmeters to follow complete synchrotron cycles and capture static and dynamic field effects that depend on the full excitation history, which is the primary operational mode for the measurements described in subsequent chapters.

The accuracy of a fluxmeter measurement is ultimately limited by how well the effective area $A_{\text{tot},k}$ of each element is known. In a single-element system, a relative calibration error $\delta A_{\text{tot}}/A_{\text{tot}}$ propagates directly as a relative error of the same magnitude in the inferred field. In a multi-element array, each element must be calibrated independently, since manufacturing tolerances in winding geometry, conductor placement, and substrate thickness produce element-to-element variations in both the effective area and the effective radius $R_{\text{coil},k}$. From Eq. (3.28), a relative error $\delta R_{\text{coil},k}/R_{\text{coil},k}$ in the effective radius of element k introduces a relative error of order $n \delta R_{\text{coil},k}/R_{\text{coil},k}$ in the inferred integrated multipole of order n [russenschuckFieldComputationAccelerator2010]. To reach a precision of order 10^{-4} for a sextupole component with $n = 3$, the effective radius must therefore be known to a few

parts in 10^{-5} , which illustrates how geometric tolerances tighten progressively with multipole order. For PCB-based coil sets, the inter-element reproducibility of the calibrated areas further determines the floor on multipole accuracy, since the matrix inversion that extracts individual harmonics amplifies any systematic discrepancy between elements. **[TODO: why for PCB / high number of channels only]**

Misalignment of the coil set with respect to the magnetic axis introduces an additional source of systematic error by coupling different multipole orders. For a coil offset by δy in the vertical direction inside a dipole that carries a residual normal quadrupole component B_2 , the measured vertical field becomes

$$B_y(0, \delta y, z) \approx B_1(z) + B_2(z) \frac{\delta y}{R_{\text{ref}}} + \dots \quad (3.31)$$

so that the inferred dipole field acquires a systematic bias proportional to δy . If $B_2/B_1 \sim 10^{-2}$, a relative misalignment $\delta y/R_{\text{ref}} \sim 10^{-2}$ already introduces a relative dipole error of order 10^{-4} **[russenschuckFieldComputationAccelerator2010]**. Analogous couplings arise at all multipole orders, so that the required transverse alignment tolerance scales inversely with the ratio of the dominant to the target harmonic. Additionally, if the coil is shorter than the magnet, displaced longitudinally, or tilted with respect to the magnetic axis, the measured flux samples only a specific axial window of $B_n(z)$ and applies a modified weighting to the fringe-field regions, and in practice these effects are minimised by mechanical design and absorbed into the calibrated scale factors.

As Eq. (3.26) shows, the integrator output $U(t)$ is defined only up to an additive constant set by the reference flux $\Phi(t_0)$, so a fluxmeter provides only relative measurements of field changes. Every measurement therefore requires a well-defined integration marker, a point in the cycle at which the integrated field is assigned a known value. In static calibration this reference is established by driving the magnet through a reproducible precycle and anchoring the integrator output to an auxiliary high-accuracy sensor such as an NMR probe in a dedicated reference magnet. In pulsed operation like a synchrotron cycle, the marker is imposed at a reproducible point in the machine cycle, after which the fluxmeter tracks relative changes until the next marker. In addition to the marker requirement, practical measurements are affected by small DC offsets and leakage currents at the integrator input. Denoting the effective offset by V_{off} , the measured signal in Eq. (3.26) becomes

$$U_{\text{meas}}(t) = N[\Phi(t) - \Phi(t_0)] - V_{\text{off}}(t - t_0) \quad (3.32)$$

where the second term represents a linear drift that is indistinguishable from a genuine slow field variation. Over a machine cycle of duration Δt , the relative error introduced by the offset scales as

$$\frac{\Delta\Phi_{\text{drift}}}{\Delta\Phi_{\text{sig}}} \approx \frac{V_{\text{off}} \Delta t}{\int_{t_0}^{t_0+\Delta t} |V(t')| dt'} \quad (3.33)$$

which sets a practical lower bound on accuracy even if the coil geometry is perfectly known. In real integrators V_{off} varies slowly with variables like temperature, and external conditions, so the drift is rarely linear and cannot be removed exactly by subtraction over long integration windows. The large loop areas and high turn counts of multi-element fluxmeter sets also make them sensitive to external magnetic disturbances and mechanical vibrations, which couple additional parasitic contributions into $V(t)$ that are indistinguishable from the magnetic signal. These effects are mitigated by rigid mechanical supports, stable cable routing, and twisted-pair wiring, but can rarely be eliminated completely outside of tightly controlled environments.

In compensated coil arrangements that suppress the dominant harmonic, drift, noise, and external pickup become progressively more limiting relative to the residual output as the target

multipole order increases. Rotating-coil systems, discussed in the following subsection, employ the same underlying flux integration principle but use the azimuthal motion of the coil to extract static multipole content with high sensitivity and reduced susceptibility to long-term integration drift.

3.2.6 Rotating Coils

A rotating coil magnetometer measures the magnetic field in the aperture by exploiting the direct correspondence between the azimuthal Fourier harmonics of the induced voltage signal and the multipole coefficients of the polar expansion introduced in Section 3.1.1. When an induction coil rotates at angular velocity $\dot{\varphi}$ inside the aperture, the instantaneous voltage induced in a single element k placed at radius $R_{\text{coil},k}$ with N_k turns is

$$V_k(t) = N_k \sum_{n=1}^{\infty} n \dot{\varphi}(t) S_{n,k} [B_n \sin(n\varphi(t)) + A_n \cos(n\varphi(t))] \quad (3.34)$$

where $S_{n,k} \propto R_{\text{coil},k}^n$ is the geometric sensitivity of element k to the n -th harmonic and $\varphi(t)$ is the instantaneous angular coil position. Integrating V_k over one complete revolution and applying a discrete Fourier transform yields all normal and skew harmonic coefficients B_n and A_n simultaneously from a single rotation. The azimuthal periodicity that defines the multipole basis in Eq. (3.5) is thus directly resolved by the angular sweep of the coil, which makes rotating coils the a direct experimental measurement of the polar description [arpaiaRotatingCoilMagnetometerScanning2019, russenschuckFieldComputationAccelerator2019].




Fig. 3.8: Schematic cross-section of a magnet aperture with a rotating coil assembly, showing the coil element at radius R_{coil} and the rotation angle φ .

Comparing to the fluxmeter presented in the previous subsection, the rotating coil generalises the fluxmeter array from a fixed radial arrangement to a continuous azimuthal sweep. A fluxmeter samples a set of discrete radial positions simultaneously, whereas the rotating coil samples all azimuthal directions sequentially and recovers the full harmonic spectrum through the Fourier decomposition. Compensated coil arrangements that suppress the dominant harmonic, analogous to the compensated fluxmeter described in Section 3.2.5, are used when high sensitivity to weaker harmonics is required and are designed specifically for each magnet aperture and field order.

The principal advantage of rotating coils over fluxmeters is their resistance to integrator drift. Since the coil returns to its starting position at the end of each revolution, any linear DC offset in the integrator electronics produces a contribution that is periodic in φ and cancels from the extracted Fourier coefficients. The effective drift therefore accumulates only over one rotation period, typically of order one to a few seconds, rather than over a full machine cycle as in the fluxmeter case.

This makes rotating coils the new gold standard technique for production characterisation and acceptance testing of accelerator magnets [arpaiaRotatingCoilMagnetometerScanning2019, russenschuckFieldComputationAccelerator2010].

While the rotating coil is superior to the fluxmeter for static measurements, it has several limitations that make it unsuitable for the pulsed, time-resolved measurements required in this thesis. First, the Fourier decomposition requires the instantaneous angular position to be known precisely throughout the rotation, which is classically provided by a high-resolution optical encoder rigidly coupled to the coil shaft. Speed variations during rotation introduce harmonic mixing in the Fourier analysis, and reducing this below the 10^{-4} level demands encoder resolutions below approximately 0,1 mrad [rogackiEncoderLessAcquisitionSystem2024]. Encoder-less acquisition schemes that reparametrise flux samples by the integrated voltage rather than by angle have been developed to address this, but require careful calibration of the sensitivity coefficients [rogackiEncoderLessAcquisitionSystem2024].

Beyond mechanical limitations, the coil length fixes the axial integration window over which the flux is accumulated. Extending the coil beyond the physical iron ends to capture the fringe-field regions discussed in Section 3.1.5 is standard practice and recovers the full integrated multipole $\int B_n dz$, provided the coil reaches a region where $B_n(z)$ is negligibly small [russenschuckFieldComputationAccelerator2010]. For long magnets, however, constructing a mechanically rigid coil that spans the full effective length while maintaining the geometric tolerances required for accurate harmonic extraction becomes a significant engineering challenge [russenschuckFieldComputationAccelerator2010]. Coil sag, torsional compliance, and the difficulty of maintaining a well-defined rotation axis over several metres all degrade the harmonic accuracy, and in practice the measurement bench must be purpose-built to the specific magnet length and aperture.

The main limitation with the respect to this thesis, is that the Fourier decomposition requires the field to be static or slowly varying on the timescale of one full rotation. During a real synchrotron ramp the field can change significantly within a single rotation period, so the extracted harmonic coefficients acquire systematic errors that grow with the rate of field change [russenschuckFieldComputationAccelerator2010]. A non-rotating single-shot flux integration remains possible during a ramp, but this degrades to a fluxmeter-type measurement and forfeits the harmonic resolution and drift-free advantages of the rotating mode.

Rotating coils complement fluxmeter systems by offering improved harmonic resolution and reduced drift for static excitation. However, they are not suitable for the pulsed, time-resolved measurements over complete synchrotron cycles required in this thesis, and are therefore not expanded further here.

3.3 Calibration and Accuracy Limits

3.3.1 Coil Area Calibration

3.3.2 Higher-Order Gradient Measurements

3.3.3 Integration Drift in Fluxmeters

3.3.4 Summary: Accuracy Limits and Implications

3.4 Magnetic Measurements at CERN

3.4.1 SPS Main Dipole Measurement Bench

3.4.2 SPS Main Quadrupole Measurement Bench

3.5 Online B-Train Measurement Systems

In a synchrotron, the RF frequency must track the beam revolution frequency in real time according to the synchronism condition in Eq. (2.25). As the beam accelerates, the momentum and therefore the required RF frequency evolve continuously over the cycle, and maintaining this synchronism requires knowledge of the instantaneous average bending field $B(t)$. For machines in which the particle velocity changes substantially during the cycle, the RF frequency sweep is large enough that deriving the required frequency from a programmed current function alone is insufficient at the required precision. In those cases, a direct, continuous measurement of the bending field is operationally necessary.

The concept of measuring this field in real time from an induction coil wound on a reference magnet was pioneered at CERN for PS, where the combined-function lattice makes a direct field measurement essential [regentreifCERNProtonSynchrotron1962, caspersAlternativeClassicalRealtime1997]. The measurement system that performs this function is commonly referred to as the B-train, which originally referred the continuous real-time signal chain that distributed the average bending field estimate to the RF system, the power converters, and the beam instrumentation [caspersAlternativeClassicalRealtime1997, buzioDevelopmentUpgradedMagnetic2010]. The principle has since been adopted at accelerator facilities worldwide, including medical synchrotrons where cycle-to-cycle field reproducibility is a direct requirement for dose precision [feldmeierFirstMagneticField2012, franziniFinalDesignFeatures2010, margreiterMagneticFieldCo

3.5.1 Operating Principle of a B-train

For a standard B-Train, a sensing coil is wound on a reference magnet powered in series with the main ring dipole circuit, so the reference magnet receives the same excitation as the ring dipoles. The universally used field measurement instrument for a B-Train is the fluxmeter coil, which provides a continuous measurement of the field integral as described in Section 3.2.5. For most machines using a B-Train, the average dipole field from Eq. (3.18) suffices to describe the beam dynamics and to derive the required RF frequency. However for machines with significant fringe fields or non-trivial design such as sector magnets, the full integral field $B d\ell$ from Eq. (3.15) is used.

In practical designs, an FDI digitises the coil voltage at high sampling rates at several hundred kilohertz, or even megahertz, and performs the numerical integration in real time. This makes the instantaneous field estimate $B_{\text{avg}}(t)$ available to the RF control system and power converters at rates sufficient to follow the field ramp at ramp rates exceeding 1 T/s while integrating out irregular small amplitude noise in the voltage measurements [buzioDevelopmentUpgradedMagnetic2010,

caspersAlternativeClassicalRealtime1997]. As synchrotrons operate cycle-by-cycle, the B-Train provides a cycle-by-cycle measurement of the field, and anchors the integration with a field marker that fires when the field passes through a calibrated reference value. This bounds the accumulation of integrator drift to the interval between successive resets. At CERN, all deployed B-trains currently use NMR-based markers for this purpose [**albrightFirstOperationalResults2023**, **buzioDevelopmentUpgradedMagnetic2010**], and marks before injection to avoid large field jumps when beam is present and the integrator is reset to zero. The accuracy and drift limitations that govern the B-train measurement remains those of the fluxmeter integration chain discussed in Section 3.3.3.

An additional limitation of the B-train in addition to the fluxmeter-bound ones, is the requirement for near-constant uptime. In practice, the effective voltage offset V_{off} of the integration chain seen in Eq. (3.32) is not stable over time but drifts on timescales ranging from minutes to hours depending on temperature, electronics load, and cable and external conditions [**grechErrorCharacterizationCalibration2021**]. While offline calibration suffices for laboratory measurements, constant, or near-constant uptime of the B-train system means that the offset correction must be determined and maintained in real time in order to avoid accumulated drift that exceeds beam dynamics tolerances. The standard procedure is to create an electronic short-circuit at the integrator input, disconnecting the sensing coil and measuring the residual voltage bias directly, and this offset is then subtracted from subsequent integrations [**albrightFirstOperationalResults2023**].

3.5.2 B-Train Implementations

At CERN, B-train systems are deployed on all synchrotrons in the injector chain, distributing the real-time field estimate to the RF system, the power converters, and the beam instrumentation over optical White Rabbit [**serranoWhiteRabbitProject2009**] timing links [**albrightFirstOperationalResults2023**, **buzioDevelopmentUpgradedMagnetic2010**]. The B-train output serves as the primary input to the RF phase and radial position loops, which adjust the RF frequency to maintain synchronism with the circulating beam. In machines where field-based rather than current-based regulation is implemented, the B-train closes a feedback loop on the main magnet power converter on the dipole field exclusively, which suppresses cycle-to-cycle variability in the bending field such as those caused by hysteresis. Additionally, the B-train coil is installed inside the vacuum chamber of the reference magnet so the measured integral includes the contribution of eddy currents in the vacuum chamber wall, as discussed in Section 2.6.1. In this arrangement, the measured field integral represents the field seen by the beam rather than the field in the bare magnet aperture.

The PS is the original B-train machine and the most operationally known implementation [**caspersAlternativeClassicalRealtime1997**, **regenstreifCERNProtonSynchrotron1962**]. The reference magnet is a combined-function unit identical to the main ring units. Specific to the PS is that it cannot operate on a simulated field function alone, relating to the complexity of the combined-function design, and therefore uses field regulation. The Proton Synchrotron Booster (PSB), Low Energy Ion Ring (LEIR), and Antiproton Decelerator (AD) all use B-trains for both RF control and power converter field regulation of the dipole circuit [**albrightFirstOperationalResults2023**, **grechErrorCharacterizationCalibration2021**]. **[cite]** PSB, LEIR design reports The Extra Low Energy Antiproton Ring (ELENA) ring, which decelerates antiprotons to 100 keV, operates at very low bending fields that approach the noise floor of standard induction coils. Therefore, in ELENA, a curved fluxmeter is employed, in which the fluxmeter follows the arc of the dipole aperture, in order to increase the effec-

tive coil area and the signal-to-noise ratio at low field [jorgensenExtraLowENergy2014, grechDevelopmentOptimisationBtrain2020].

The SPS B-train serves a different role compared to the smaller synchrotrons, because the injection energy is sufficiently high that the revolution frequency is nearly constant over the cycle, the B-train output is not required for RF control, and static or simulated fields can be used to set the RF frequency. It is neither used for field-regulation, as the SPS main magnets are regulated in current. The SPS B-train serves as a core part of the data acquisition for this thesis, and is more thoroughly presented in its own Section 3.5.3.

Beyond CERN, B-train systems are deployed at medical synchrotrons where cycle-to-cycle reproducibility of the bending field is a direct requirement for dose precision. At CNAO, the main synchrotron uses a B-train on the bending dipole circuit for both RF frequency control and field reproducibility monitoring [feldmeierFirstMagneticField2012, franziniFinalDesignFeatures2010]. At MedAustron, a B-train was part of the original design. However, drift accumulation over the long cycles and lack of flexibility of where to apply the NMR field marker led to the system being superseded by a Hall-probe-based real-time field measurement, which avoids integrator drift at the cost of requiring an accurate sensor model and regular calibration [margreiterMagneticFieldControl2025, margreiterFieldControlDynamically2025]. The MedAustron experience illustrates a general boundary of the B-train approach, as machines with long cycles, demanding absolute accuracy requirements, or infrequent marker resets, integrator drift can become the dominant limitation and alternative real-time measurement strategies may be preferable.

3.5.3 SPS B-Train

[TODO: Deferred until spec sheets for SPS B-train can be found]

Chapter 4

Beam Operation at CERN and the SPS

[TODO: where to put idle current (155 A) and field (0.625 T) for the main dipoles, or concept of idle current?]

The CERN accelerator complex operates a chain of synchrotrons that each prepare and deliver beam to a variety of experiments and downstream machines. At the center of this chain for the work presented in this thesis is the SPS, which functions as a multi-cycling accelerator serving several different beam destinations within a single repeating sequence of machine cycles. Each cycle accelerates beam to a different energy or for a different purpose, and these cycles are interleaved in a fixed supercycle that repeats continuously during operation.

The accelerator control system that governs this operation is a layered architecture spanning settings management, timing, communication middleware, data logging, postmortem, power converters and hardware controllers. This architecture determines how cycle settings are defined, structured, and propagated from high-level beam physics quantities down to magnet hardware, how operational data is recorded and made available for offline analysis, and how runtime corrections can be computed and applied within the control chain. These aspects of the control system are directly relevant to the work of this thesis, as they set the practical constraints on data acquisition, model design, and the deployment of the compensation strategy developed in the following chapters.

Section 4.1 introduces the cycle and supercycle as the fundamental units of SPS operation. The main SPS cycle families and their role in the machine programme are described in Sections 4.2 and 4.3. The accelerator controls infrastructure is presented in Section 4.4, covering settings management, timing, middleware, logging, and hardware controllers. The field control hierarchy in the injector chain and the operational procedures used to maintain magnetic, and therefore beam, reproducibility are covered in Sections 4.5 and 4.6. The machine development framework within which experimental validation in Chapter 10 is carried out is introduced in Section 4.7.

4.1 Accelerator Operation in Cycles

At CERN, each accelerator in the injector chain operates according to a programmed sequence of settings that is executed repeatedly. One complete execution of this sequence is referred to as a machine cycle. A cycle defines the settings of all accelerator equipment for one period of operation, including the magnets, radio frequency systems, beam instrumentation, and timing devices. It is therefore not a purely magnetic concept but a complete operational programme for the machine, even though the term magnetic cycle often appears in the literature and is used colloquially to refer to the same concept. Within a cycle, the sequence is divided into logical phases called beam processes, which correspond to distinct operational stages such as pre-injection, flat bottom / injection, acceleration ramp, flat top / extraction, and ramp-down. Different settings can in principle be assigned independently to each beam process within the settings management system, described in Section 4.4.1, while the hardware layer is agnostic to this and operates without knowledge of the process boundaries.

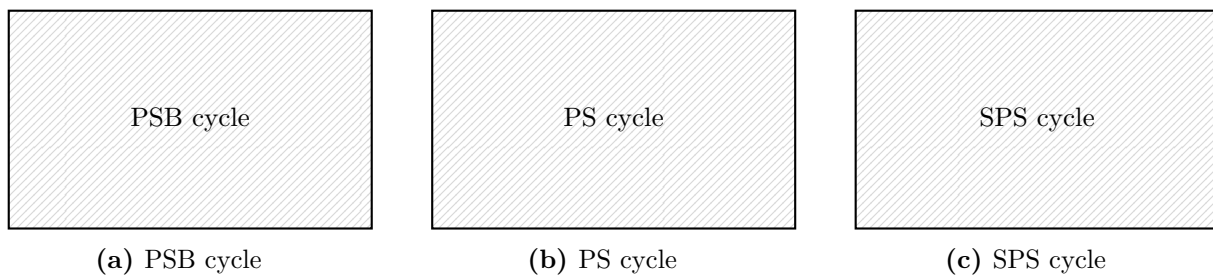


Fig. 4.1: Sample machine cycles for the PSB (left), PS (center), and SPS (right), showing the momentum programme as a function of time, or equivalently the main dipole field. Colored bands below each curve indicate the beam processes: pre-injection, injection plateau, ramp, flat top, extraction, and ramp-down. All cycle durations are integer multiples of the 1,2s basic period.

The CERN injector chain, introduced in Section 2.2, consists of several synchrotrons that each operate according to their own cycle structure. The duration of a machine cycle is always an integer multiple of the so called Basic Period (BP) of 1,2s. The number of basic periods per cycle depends on the machine and the operational requirements and beam parameters, and can vary from cycle to cycle within the same machine. In the proton injector chain, LINAC4 and the PSB operate synchronously with cycles of one basic period. PS cycles span one or two basic periods depending on the beam destination, energy and other beam parameters, with the two-period cycles typically used for the 26 GeV/ c LHC-type beam. SPS cycles consist of three or more basic periods, with operational cycles reaching up to approximately 50 basic periods for heavy ion operation. The ion injector chain follows the same basic period convention, with LINAC3 and LEIR operating with two to three basic periods per cycle, and the PS and SPS using the same cycle structures as for protons but with different settings and timing. Figure 4.1 illustrates these relationships for representative cycles of each injector in the proton chain, with the beam processes annotated for each.

In addition to real machine cycles, a dedicated placeholder cycle called the ZERO cycle occupies basic period slots in the supercycle where no beam is required, or not desired for operational reasons. This could also be caused by desire to keep the supercycle at a fixed length, or to reduce stress on the machine by inserting a break in the beam delivery. In a ZERO cycle, all magnets are held at their idle current rather than ramped, and most equipment is placed in a standby or ready state instead of cycled. The ZERO cycle is always one basic period long, regardless of the machine in which it appears, and serves as the standard filler slot across the PSB, PS, and SPS supercycles.

In addition to the standard supercycle mechanism with cycles, the control system supports the concept of normal and spare cycles. A normal cycle can be paired with a designated spare cycle of the same length, and the two can be exchanged dynamically with some advance notice, without requiring a full supercycle change. This mechanism allows beam clients to request a change at shorter notice than the supercycle turnaround time would otherwise permit. As an example, a facility such as the AD [bairdDesignStudyAntiproton1997], which receives beam every 2 minutes, may use a nearby spare slot to activate the spare cycle in the PS supercycle rather than waiting for the next supercycle iteration. The normal and spare mechanism normally used in both the SPS and the PS, and the SPS case will be elaborated in Section 4.3.

The ordered sequence of cycles executed by a machine is called the Supercycle [bauManagingRealtimeBeha lewisCentralBeamCycle1999, lewisEvolutionCERNSPS2003]. Each position in the supercycle is referred to as a cycle slot, and one supercycle repeats continuously during operation.

The composition of the supercycle is determined by the beam requests of all downstream users across the accelerator complex, and is coordinated at the level of the full injector chain. The structure and length of the supercycle is primarily set by the SPS, since it operates the longest individual cycles, and typically ranges between 30 s and 60 s, with occasional extensions up to approximately 75 s depending on the operational programme. Changes to a supercycle composition are always initiated manually by a human operator, and changing which supercycle is played similarly requires operator action. In practice, the supercycle composition is therefore static during operation, and for the SPS, different supercycles exist for different operational programmes, such as fixed-target physics, LHC injection, or machine development, and is described in more detail in Section 4.3. Additionally since 2025, a system known as Dynamic Beam Scheduling (DBS) [[irannejadConstraintbasedOptimizationParticle2025](#)] automates the filling and optimization of PS and PSB supercycle slots to improve machine usage when new supercycles are created, while the SPS supercycle composition or templates remains under manual operator control. Notably, the LHC does not operate on a cycle structure in the same sense, as it is a storage ring that holds beam continuously, and injection time and phase is already timed relative to the SPS supercycle, so the LHC orchestration is not considered in the supercycle context besides the accelerated beam in the SPS

Each cycle slot in the supercycle is associated at runtime with a user identifier, referred to simply as a user. The user is the runtime label under which devices execute their settings and publish their data for that cycle slot. In the settings management system, a cycle is the object that holds the complete configuration of all devices for one slot, and a user is mapped to a cycle before execution. This distinction between the user as a runtime concept and the cycle as a settings concept is used throughout this thesis when referring to logged operational data, and is discussed further in Section 4.4.1.

The ability to hold independent settings per user simultaneously is referred to as Per-Pulse Modulation (PPM), or multiplexing, and devices or parameters that support this are described as PPM-capable or enabled, or simply as multiplexed. A PPM-capable device can apply a different configuration for each user as the supercycle advances, which is the hardware mechanism that makes multi-cycling operation possible. Devices or parameters that are not PPM-capable hold a single global setting regardless of which user is currently playing, and must therefore be set to values compatible with all cycles in the supercycle simultaneously. The distinction between PPM and non-PPM parameters is relevant for the compensation strategy developed in this thesis, as it determines which settings can be varied independently per cycle and which cannot, and is discussed further in Section 4.4.1.

Figure 4.2 illustrates the timing relationship between the PSB, PS, and SPS supercycles, showing how transfers between machines are aligned and how the PS executes multiple cycles within one SPS cycle period.

4.2 Main SPS Cycles

As discussed in Chapter 2, the SPS serves a variety of beam destinations and experiments, each with distinct requirements on beam energy, intensity, emittance, and extraction mode. These requirements translate directly into different momentum curves, optics configurations, and timing structures, and the resulting cycle families are the basic building blocks of the SPS supercycles. The cycles are presented in this section by proton and ion cycles, as they differ in injector chain, magnetic rigidity, optics, and injection timing, in addition to being run in different operational periods.

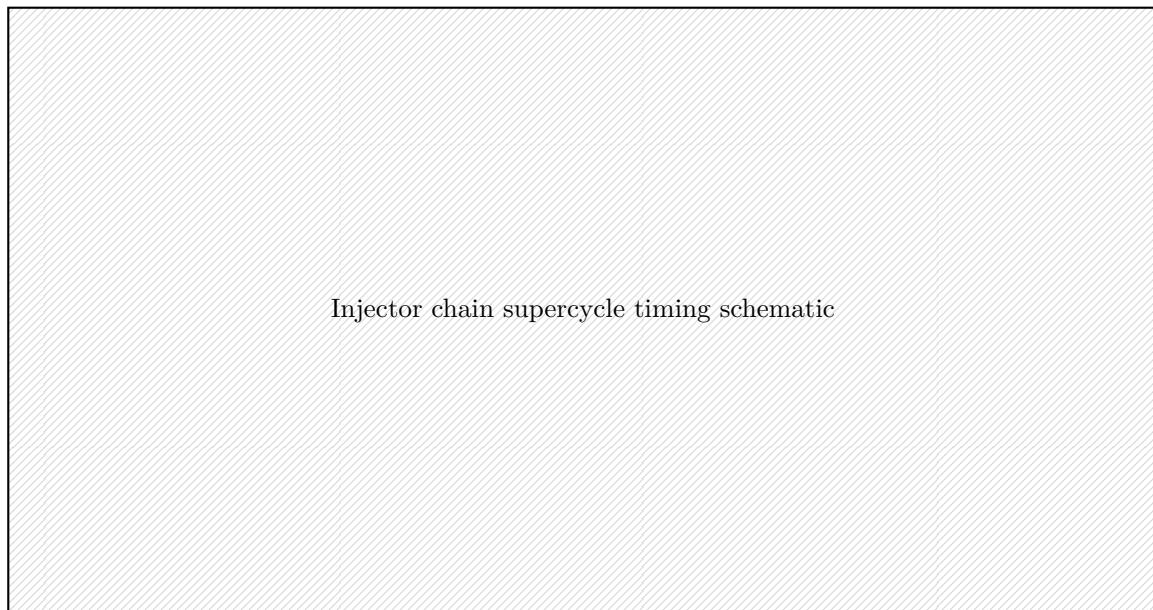


Fig. 4.2: Schematic timing diagram of the PSB, PS, and SPS supercycles on a common time axis. Arrows indicate beam transfer moments between machines. Named cycle slots correspond to the representative cycles shown in Figure 4.1; remaining slots are shown as unlabelled rectangles. ZERO cycle slots are indicated by white rectangles. The diagram illustrates that the PS executes multiple cycles within one SPS cycle period, enabling multiple injections into the SPS before the acceleration ramp begins.

For each cycle, the beam requirements, the resulting magnetic cycle shape, the principal magnet circuit parameters, and the extraction mode are described. So called machine development cycles, which can in principle take any cycle shape, are not necessarily tied to a specific particle species or beam parameters, but a few common examples will be presented nonetheless, and their usage explained in Section 4.7. A dedicated magnetic precycle MD1, which precedes fixed-target proton cycles in the standard supercycle, is introduced in the context of the SPS Fixed Target Proton (SFTPRO) cycle and described in full in Section 4.6.1. The composition of representative SPS supercycles from these cycle families is presented later in Section 4.3.

4.2.1 SPS North Area Fixed Target Proton Cycle (SFTPRO)

The North Area Fixed Target beam, or commonly referred to as SFTPRO or SFT, cycle is the primary fixed-target proton cycle of the SPS, delivering beam to the North Area experiments. It operates with integer tune 26 (Q26) optics and receives protons from the PS via MTE at an injection momentum of $14 \text{ GeV}/c$. The fractional tune working points are 26.58 in the vertical plane, and 26.62 in the horizontal plane. The cycle length fixed at 10 800 ms. In standard operation, two injections are performed, with the first at 200 ms and the second at 1400 ms after cycle start, reaching a nominal peak intensity of approximately 1.3×10^{13} protons per cycle. In each injection, 5 macro-bunches are injected from the MTE, at up to 5×10^{12} protons per bunch, with a total of **[TODO: XX bunches]** per injection. The scheme of injecting an MTE beam with Q26 optics at $14 \text{ GeV}/c$ is normally referred to as an SFT-type beam. The acceleration ramp begins at 1500 ms, and the flat top is reached at 4460 ms. Beam is extracted at $400 \text{ GeV}/c$ via chromatic slow extraction, described in Section 2.1.9.3, over a spill duration of 4800 ms ending at 9260 ms, after which the magnets ramp down. Beam leaves the SPS through the TT20 transfer line to the




Fig. 4.3: The SFTPRO machine cycle. Left: main dipole field $B(t)$ as a function of cycle time, overlaid with the beam intensity measured by the BCT. The two-injection pattern is visible at 200 ms and 1400 ms, followed by the acceleration ramp, the 4800 ms slow extraction spill, and the ramp-down. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.

North Area. The SFT cycle has had up to 9600 ms slow extraction [cite], with only the flat top elongated. Under normal and optimal machine operation, the extracted spill is stable and with approximately constant extracted intensity over time, over the full spill duration, and with minimal losses at the electrostatic septa. With poor main field reproducibility, the extracted spill rate can change, as shown in [velottiInvestigationRemanentField2017], which can trigger operator retuning of the beam to satisfy the experimental area targets. Furthermore, a spill with a beam mismatched with the electrostatic septa can risk damaging the septum, and high-intensity beam is often dumped early. The quality of the extracted spill is computed as [TODO: Figure out exact metric] and is used as the main performance indicator for optimization in the cycle.

The main dipole field at injection is approximately 0,625 T, which is the lowest point in the SPS with beam. The magnetic cycle reaches a peak field of approximately 1,8 T at the flat top. The QF and QD quadrupole circuits follow the same momentum programme through the ramp to maintain the optics, and The sextupole circuits are ramped in proportion to momentum to maintain the chromaticity, although it can be seen in Figure 4.3 corrections are applied on top of the nominal ramp, which are used to optimize beam transmission. The octupole circuits provide Landau damping [cite] throughout the cycle.

For RF control loops, the synchro-loop and phase-loop are enabled during injection of both bunches, with the radial loop engaging after the second injection, and following the through the entirety of the ramp. At the start of extraction, all RF loops are disabled, and the beam is debunched for the slow extraction.

The use of chromatic slow extraction places stringent requirements on the reproducibility of the dipole, quadrupole, and sextupole fields during the flat top, as established in Section 2.1.9.3. During this phase, the dipole field is ramped slowly to sweep the resonant momentum slice through the circulating beam distribution, and the quadrupole circuits are adjusted proportionally to maintain constant optics throughout the spill. The quantitative field tolerances and their measured beam impact are part of this thesis and presented in Chapter 10. The SFTPRO cycle is in practice always preceded in the standard magnetic sequence by the MD1 precycle, which establishes a reproducible magnetic state before injection and is described in Section 4.6.1.




Fig. 4.4: Detail of the SFTPRO cycle during the slow extraction flat top, showing the slow ramp of the main dipole field and the corresponding adjustments of the quadrupole circuits to maintain constant optics. The slow extraction spill occurs during this phase, and therefore the field reproducibility requirements are most stringent here.

4.2.2 LHC-Type Proton Cycle

The SPS delivers proton beam to the LHC using a family of cycles that share the same ramp and extraction beam processes but differ in the length of the injection flat bottom, which accommodates different numbers of injections. All LHC-type proton cycles operate with Q20 optics under standard conditions in Run 3, although the optics can be adjusted to Q26 for special operational requirements, as was the case during proton-oxygen collisions in 2025 [cite]. Injection is received from the PS at 26 GeV/c at C1015, corresponding to a main dipole field of approximately 0,2 T. Each normal injection delivers 72 bunches, with successive injections separated by 2400 ms, and up to five injections are possible per cycle. The beam is accelerated to 450 GeV/c and extracted to the LHC via fast extraction at a fixed point within the extraction beam process, regardless of cycle length.

Three principal cycle variants are used in standard LHC proton operation. The INDIV cycle carries a single bunch at very low intensity [TODO: specify intensity] and is primarily used for machine commissioning and optics measurements. The LHC Pilot cycle carries a single injection at low intensity and serves as the standard commissioning beam. The LHC 4-injection cycle uses the () beam scheme [cite] [karlsen-baebkirkemilComparisonBCMSStandard2025] and carries four injections of [TODO: XX intensity] at full physics intensity in Run 3. In terms of main cycle programs, including magnetic cycles, the INDIV and Pilot cycles have identical main field shapes, where the main differences between them are in the intensity-dependent corrections applied for collective effects. The Pilot and INDIV cycles have a length of 13 200 ms (11 basic periods), while the 4-injection cycle has a length of 20 400 ms (17 basic periods). The magnetic cycles of LHC INDIV / Pilot and LHC 4 injection cycles can be seen in figures 4.5 and 4.6. The scheme of injecting a bunched beam with Q20 optics at 26 GeV/c is normally referred to as an LHC-type beam.

4.2.2.1 AWAKE

The AWAKE cycle delivers a proton driver beam to the AWAKE experimental facility [gschwendtnerAWAKE]. It operates with Q20 optics and receives a single-bunch beam from the PS at 26 GeV/c, with injection at C1015, same as LHC [TODO: Add injection bunch parameters.] The beam is accelerated to 400 GeV/c and delivered to the AWAKE beamline via fast extraction. The cycle length is 7200 ms (6 basic periods), and only one AWAKE cycle variant exists in standard




Fig. 4.5: The LHC Pilot and machine cycle (13 200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the single injection at C1015, the acceleration ramp to 450 GeV/ c , and the fast extraction. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.

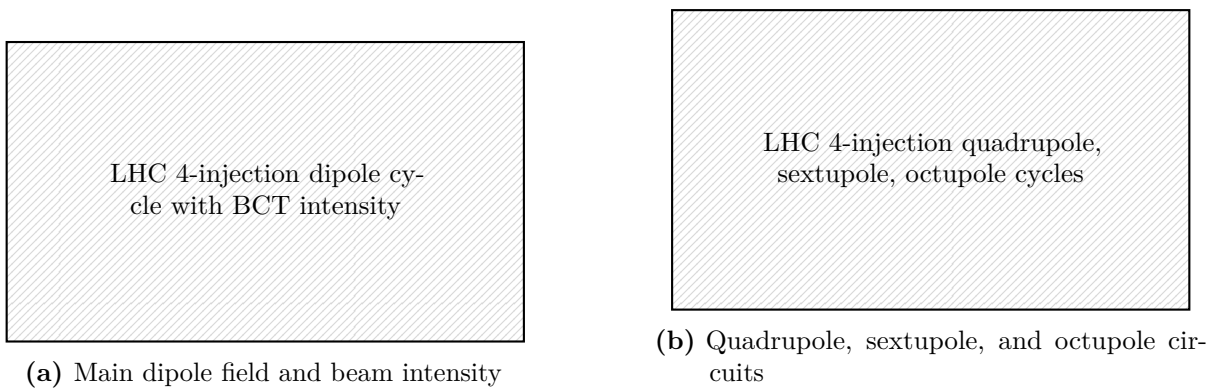


Fig. 4.6: The LHC 4-injection () machine cycle (20 400 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing four injections at 2400 ms intervals starting at C1015, the acceleration ramp to 450 GeV/ c , and the fast extraction. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.




Fig. 4.7: The AWAKE machine cycle (7200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the single injection at C1015, the acceleration ramp to 400 GeV/ c , and the fast extraction to the AWAKE beamline. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.

operation. The ramp profile is similar to the LHC-type proton cycles described in Section 4.2.2, with extraction at 400 GeV/ c rather than 450 GeV/ c . The higher-order magnet circuits follow the momentum programme analogously to other Q20 proton cycles, and all magnetic cycles can be seen in Figure 4.7


4.2.2.2 HIRADMAT

The HiRadMat cycle delivers high-intensity proton pulses to the High Radiation to Materials facility [**hardenHiRadMatFacilityRealms2019**]. The fixed target has previously received heavy lead ions from the SPS, but in Run 3 only received proton beam. The HiRadMat proton beam operates with an LHC-type beam at Q20 optics and receives beam from the PS at 26 GeV/ c with injection at C1015. The beam is extracted at 440 GeV/ c via fast extraction, and HiRadMat delivers the highest intensity per pulse of any SPS cycle, with up to 288 bunches per pulse. **[TODO: verify this because it's the same as a nominal LHC cycle]**


Two cycle variants exist of the cycle, analogous in structure to the LHC-type proton cycles described in Section 4.2.2. The Pilot cycle has a length of 8400 ms (7 basic periods) **[TODO: verify]**, and carries a single injection. The 4-injection cycle has a length of **[TODO: verify 20 400 ms (17 basic periods)]**. Similarly to the LHC the ramp and extraction beam processes are identical between variants, with the main cycle difference being the length of the injection flat bottom, which accommodates the different number of injections. The two cycles can be seen in Figures 4.8 and 4.9.

4.2.2.3 SHiP (Future)

The SHiP cycle is a fixed-target proton cycle planned to deliver beam to the SHiP experiment [**alekhinFacilitySearchHidden2015, vanherwijnenSHiPSearchHidden2018**] in the North Area via the TT20 transfer line, the same transfer line used by the SFTPRO cycle. It operates with Q26 optics and receives a two MTE injections at 14 GeV in a similar fashion to the SFT cycle from Section 4.3, at C200 and C1400. The beam is accelerated to 400 GeV/ c and extracted via chromatic slow extraction, over a spill duration of 1 s. The cycle length is 7200 ms (6 basic periods), and the target intensity is approximately 4×10^{13} protons per cycle, significantly higher than the nominal SFTPRO intensity on a shorter cycle. The ramp and optics configuration follow the same general structure as the SFTPRO cycle, with adjustments to the

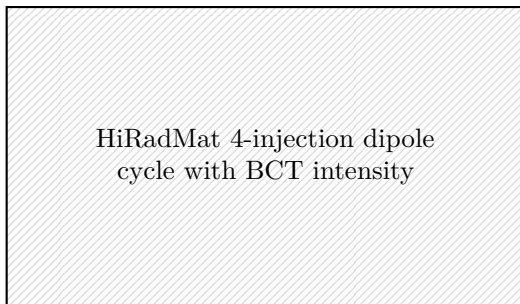


(a) Main dipole field and beam intensity, Pilot cycle

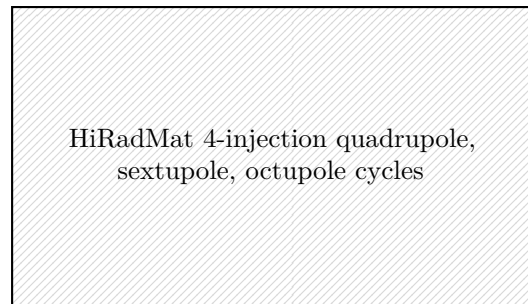


(b) Quadrupole, sextupole, and octupole circuits, Pilot cycle

Fig. 4.8: The HiRadMat Pilot machine cycle (8400 ms[**TODO: verify**]). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the single injection at C1015, the acceleration ramp to 440 GeV/ c , and the fast extraction to the HiRadMat facility. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.



(a) Main dipole field and beam intensity, 4-injection cycle



(b) Quadrupole, sextupole, and octupole circuits, 4-injection cycle

Fig. 4.9: The HiRadMat 4-injection machine cycle ([**TODO: verify 20 400 ms**]). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing four injections at 2400 ms intervals starting at C1015, the acceleration ramp to 440 GeV/ c , and the fast extraction to the HiRadMat facility. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.

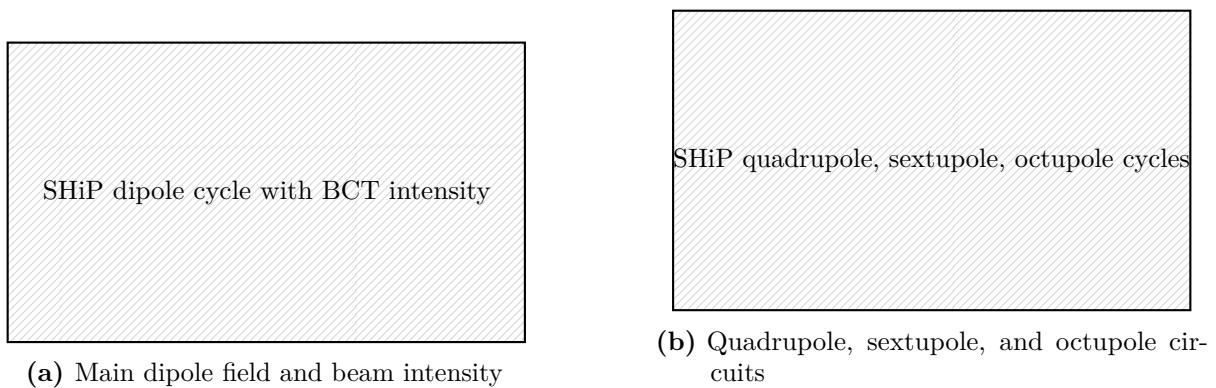


Fig. 4.10: The SHiP machine cycle (7200 ms). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the single injection at C1015, the acceleration ramp to 400 GeV/ c , and the slow extraction spill. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.


magnet circuit functions to improve beam transmission at the higher intensity. The cycle was fully tested in 2025 [cite] but will not enter routine operation until Run 4 around 2030, when the SHiP experiment comes online. It is included here because it represents a future operational target to which the field prediction and compensation methods developed in this thesis are intended to apply. Since the SHiP cycle uses chromatic slow extraction at the same energy with Q26 optics, the same requirements on dipole, quadrupole, and sextupole field reproducibility apply as for the SFTPRO cycle, as described in Section 4.2.1.

4.2.3 SFTION


The SPS Fixed Target Ion (SFTION) cycles deliver heavy-ion beam to the North Area fixed-target experiments via the TT20 transfer line, the same transfer line used by the SFTPRO cycle. While the SFTION cycle family and the LHCION cycles in the next subsection are not the main focus of this thesis, they are included here for completeness and because some of the studied cycles and future extensions involve ion running. In Run 3, two variants exist, accelerating heavy-ion beam to 300 GeV/ c and 33 GeV/ c proton-equivalent momentum, respectively. Both cycles operate with Q26 optics and receive ions from the PS at 13 GeV/u with injection at C735, and use chromatic slow extraction described in Section 2.1.9.3. [TODO: something about intensity (depending on particle species)] The 300 GeV/ c cycle takes four injections and was used for lead and oxygen ion running, with a cycle length of [TODO: verify cycle length in ms and basic periods] . The 33 GeV/ c cycle was used for lead ion running only, with a cycle length of [TODO: verify cycle length in ms and basic periods] .

The main dipole field at injection is approximately [TODO: verify injection field in T] for both cycles, which is below the idle current field of the main dipoles at [TODO: value] . Both cycles include an intermediate energy plateau before the final ramp, at which [TODO: describe RF manipulation at intermediate plateau] . Due to the use of the same chromatic slow extraction as the SFTPRO cycle, the same requirements on dipole, quadrupole, and sextupole field reproducibility therefore apply, though the experimental requirements on spill quality are typically less stringent than for protons.

Finally as an operational limitation, ion and proton fixed-target operation are interlocked such that SFTION and SFTPRO cycles cannot coexist in the same scheduled magnetic sequence




(a) Main dipole field and beam intensity, 300 GeV/ c cycle




(b) Quadrupole, sextupole, and octupole circuits, 300 GeV/ c cycle

Fig. 4.11: The SFTION 300 GeV/ c machine cycle ([**TODO: verify cycle length**]). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the four injections at C735, the intermediate energy plateau, the acceleration ramp, and the slow extraction spill. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.



(a) Main dipole field and beam intensity, 33 GeV/ c cycle



(b) Quadrupole, sextupole, and octupole circuits, 33 GeV/ c cycle

Fig. 4.12: The SFTION 33 GeV/ c machine cycle ([**TODO: verify cycle length**]). Left: main dipole field $B(t)$ overlaid with beam intensity from the BCT, showing the injections at C735, the intermediate energy plateau, the acceleration ramp, and the slow extraction spill. Right: QF or QD quadrupole, sextupole, and octupole circuit functions on the same time axis.

in the SPS. In practice, this is not a limitation, as the proton and ion run normally does not overlap.

4.2.3.1 LHC Ion Injection Cycles

4.2.4 Machine Development Cycles

4.3 Typical SPS Supercycles

4.4 Accelerator Controls Infrastructure at CERN

4.4.1 LSA

4.4.2 Timing

4.4.3 Middleware

4.4.4 Logging

4.4.5 UCAP

4.4.6 FGCs

4.5 Magnetic Field Control in the CERN Injectors

4.5.1 Hierarchical Field Control

4.5.2 B-Train Regulation in the PS and PSB

4.5.3 SPS B-Train in Operation

4.6 Magnetic Precycle and Degaussing in Operation

4.6.1 Dedicated Magnetic Precycle in the SPS (MD1)

4.6.2 Quasi-Degaussing of Higher Order Magnets in the SPS

4.6.3 Operational Limitations from Hysteresis

4.7 Machine Development Time

Chapter 5

Field Modeling with Machine Learning

Chapter 6

Field Compensation Strategies in the SPS

Chapter 7

Data-Driven Hysteresis Modeling

7.1 Problem Formulation

7.2 Model Architectures

7.2.1 Baseline Models

7.2.2 Transformer-Based Models

7.2.3 Other Approaches

7.3 Prediction Strategies

7.3.1 Autoregressive vs. Direct Prediction

7.3.2 Chunked Prediction Approach

7.4 Measurement Design and Execution

7.5 Training Methodology

7.5.1 Dataset Preparation

7.5.2 Pretraining Strategy

7.5.3 Modeling Constraints and Considerations

Chapter 8

Eddy Current Decay

8.1 Eddy Current Decay in the SPS Main Dipoles

8.2 Modeling Eddy Current Decay in the SPS

8.2.1 Eddy Current Decay as ODE

8.2.2 Measuring Eddy Current Decay in the SPS

8.2.3 Physics Constraints and Parameter Restrictions

8.3 Data-Driven Eddy Current Modeling

8.4 Qualitative Evaluation

8.5 Eddy Current Decay in the SPS Main Quadrupoles

8.5.1 Tune Decay Measurements in the SPS

8.5.2 Field Decay on Lab Magnet Measurement

8.5.3 XSuite Simulation

Chapter 9

Control System Implementation

9.1 Control Law

9.2 Feedforward Control Architecture

9.2.1 Deployment to Control Room

9.3 Operational Results

9.3.1 Field Compensation Performance

9.3.2 Energy and Time Savings

Chapter 10

Experimental Results - Hysteresis Compensation

10.1 Measurement of Impact of Precycle

10.1.1 Chromaticity Changes

10.1.2 Injection Field Shift Characterization

10.1.3 Beam Tolerance Analysis

10.2 Hysteresis Excitation from Supercycle Change

10.3 Main Dipole Field Prediction

10.3.1 Model Performance Comparison

10.4 Main Dipole Field Compensation

10.4.1 Fixed Target Flat Top Compensation

10.4.2 Injection Field Compensation with Transformer

10.4.3 Eddy Current Compensation at Injection

10.4.4 Injection Field Compensation on MD Cycle

10.5 Operational Validation

10.6 Main Quadrupole Analysis

10.6.1 Measurement Challenges

10.6.2 Beam Impact Studies

Chapter 11

Discussion and Future Work

11.1 Current Limitations

11.1.1 Measurement Accuracy

11.1.2 Data Variety Limitations

11.1.3 Generalization Challenges

11.2 Operational Requirements Analysis

11.3 Extension to Other Magnet Families

11.3.1 Transfer Learning for Higher Order Magnet Families

11.3.2 Other Magnet Families

11.4 Future Developments

Chapter 12

Conclusions