

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

<b>Acknowledgements</b>	<b>6</b>
<b>1 Introduction</b>	<b>17</b>
1.1 Motivation and Problem Statement . . . . .	17
1.2 Research Objectives . . . . .	17
1.3 Thesis Structure and Contributions . . . . .	17
<b>2 Background and Literature Review</b>	<b>18</b>
2.1 Accelerator Physics Fundamentals . . . . .	18
2.1.1 Transverse Beam Dynamics as Quasi-Harmonic Oscillators . . . . .	18
2.1.2 Nonlinear Transverse Beam Dynamics and Momentum Coupling . . . . .	19
2.1.3 Principal Magnet Types in Synchrotrons . . . . .	20
2.1.4 Betatron tune and beta function . . . . .	22
2.1.5 Dipole Field Errors and Closed Orbit Distortion . . . . .	22
2.1.6 Magnet errors and resonances . . . . .	23
2.1.7 Longitudinal Beam Dynamics . . . . .	24
2.1.8 Beam Injection . . . . .	25
2.1.9 Beam Extraction . . . . .	27
2.2 The CERN Super Proton Synchrotron . . . . .	30
2.2.1 Principal Machine Design . . . . .	32
2.2.2 Principal Magnet Families . . . . .	32
2.2.3 Beam Instrumentation and Diagnostics . . . . .	34
2.3 Normal Conducting Accelerator Magnets . . . . .	35
2.3.1 Magnetic Field Generation in Normal Conducting Magnets . . . . .	36
2.3.2 Physical Mechanisms of Hysteresis . . . . .	37
2.3.3 Sources of Field Errors and Mitigation . . . . .	39
2.4 Hysteresis Models . . . . .	39
2.4.1 Operator-type Models . . . . .	40
2.4.2 ODE-type Models . . . . .	45
2.4.3 Machine Learning Approaches . . . . .	50
2.5 Eddy Current Decay . . . . .	50
2.5.1 Eddy Currents in Conducting Materials . . . . .	50
2.5.2 Eddy Currents in Synchrotron Components . . . . .	52
2.5.3 Beam Impact of Eddy Current Decay . . . . .	52
2.5.4 Mitigation Strategies . . . . .	52
<b>3 Accelerator Magnets and Magnetic Measurements</b>	<b>53</b>
3.1 Accelerator Magnets . . . . .	53
3.1.1 Multipole Field Expansion and Field Quality . . . . .	53
3.2 Measurement Techniques . . . . .	54
3.2.1 Local, Integral, and Average Measurement . . . . .	55
3.2.2 Hall Sensors . . . . .	55

3.2.3	NMR Probes . . . . .	55
3.2.4	Stretched Wire Measurement . . . . .	55
3.2.5	Fluxmeter Measurement . . . . .	55
3.2.6	Rotating Coils . . . . .	55
3.3	Laboratory Measurement Systems . . . . .	55
3.3.1	SPS Main Dipole Measurement Bench . . . . .	55
3.3.2	SPS Main Quadrupole Measurement Bench . . . . .	55
3.4	Online B-Train Measurement System . . . . .	55
3.4.1	B-Train Systems at CERN . . . . .	55
3.4.2	SPS B-Train . . . . .	55
<b>4</b>	<b>Beam Operation at CERN and the SPS</b>	<b>56</b>
4.1	Magnetic Sequences as Cycles . . . . .	56
4.2	The CERN SPS . . . . .	56
4.3	Accelerator Controls Infrastructure at CERN . . . . .	56
4.3.1	Timing . . . . .	56
4.3.2	Middleware . . . . .	56
4.3.3	Logging . . . . .	56
4.3.4	LSA . . . . .	56
4.3.5	UCAP . . . . .	56
4.3.6	FGCs . . . . .	56
4.4	Magnetic Field Control in the CERN Injectors . . . . .	56
4.4.1	Hierarchical Field Control . . . . .	56
4.4.2	B-Train for Field Regulation in the PS and PSB . . . . .	56
4.5	Magnetic Precycle and Degaussing in Operation . . . . .	56
4.5.1	Dedicated Magnetic Precycle in the SPS (MD1) . . . . .	56
4.5.2	Quasi-Degaussing of Higher Order Magnets in the SPS . . . . .	56
4.5.3	Degaussing and Mitigation in the PS . . . . .	56
4.5.4	Operational Limitations from Hysteresis . . . . .	56
4.6	Machine Development Time . . . . .	56
<b>5</b>	<b>Field Modeling with Machine Learning</b>	<b>57</b>
<b>6</b>	<b>Field Compensation Strategies in the SPS</b>	<b>58</b>
<b>7</b>	<b>Data-Driven Hysteresis Modeling</b>	<b>59</b>
7.1	Problem Formulation . . . . .	59
7.2	Model Architectures . . . . .	59
7.2.1	Baseline Models . . . . .	59
7.2.2	Transformer-Based Models . . . . .	59
7.2.3	Other Approaches . . . . .	59
7.3	Prediction Strategies . . . . .	59
7.3.1	Autoregressive vs. Direct Prediction . . . . .	59
7.3.2	Chunked Prediction Approach . . . . .	59
7.4	Measurement Design and Execution . . . . .	59
7.5	Training Methodology . . . . .	59
7.5.1	Dataset Preparation . . . . .	59
7.5.2	Pretraining Strategy . . . . .	59
7.5.3	Modeling Constraints and Considerations . . . . .	59

---

<b>8</b>	<b>Eddy Current Decay</b>	<b>60</b>
8.1	Eddy Current Decay in the SPS Main Dipoles . . . . .	60
8.2	Modeling Eddy Current Decay in the SPS . . . . .	60
8.2.1	Eddy Current Decay as ODE . . . . .	60
8.2.2	Measuring Eddy Current Decay in the SPS . . . . .	60
8.2.3	Physics Constraints and Parameter Restrictions . . . . .	60
8.3	Data-Driven Eddy Current Modeling . . . . .	60
8.4	Qualitative Evaluation . . . . .	60
8.5	Eddy Current Decay in the SPS Main Quadrupoles . . . . .	60
8.5.1	Tune Decay Measurements in the SPS . . . . .	60
8.5.2	Field Decay on Lab Magnet Measurement . . . . .	60
8.5.3	XSuite Simulation . . . . .	60
<b>9</b>	<b>Control System Implementation</b>	<b>61</b>
9.1	Control Law . . . . .	61
9.2	Feedforward Control Architecture . . . . .	61
9.2.1	Deployment to Control Room . . . . .	61
9.3	Operational Results . . . . .	61
9.3.1	Field Compensation Performance . . . . .	61
9.3.2	Energy and Time Savings . . . . .	61
<b>10</b>	<b>Experimental Results - Hysteresis Compensation</b>	<b>62</b>
10.1	Measurement of Impact of Precycle . . . . .	62
10.1.1	Chromaticity Changes . . . . .	62
10.1.2	Injection Field Shift Characterization . . . . .	62
10.1.3	Beam Tolerance Analysis . . . . .	62
10.2	Hysteresis Excitation from Supercycle Change . . . . .	62
10.3	Main Dipole Field Prediction . . . . .	62
10.3.1	Model Performance Comparison . . . . .	62
10.4	Main Dipole Field Compensation . . . . .	62
10.4.1	Fixed Target Flat Top Compensation . . . . .	62
10.4.2	Injection Field Compensation with Transformer . . . . .	62
10.4.3	Eddy Current Compensation at Injection . . . . .	62
10.4.4	Injection Field Compensation on MD Cycle . . . . .	62
10.5	Operational Validation . . . . .	62
10.6	Main Quadrupole Analysis . . . . .	62
10.6.1	Measurement Challenges . . . . .	62
10.6.2	Beam Impact Studies . . . . .	62
<b>11</b>	<b>Discussion and Future Work</b>	<b>63</b>
11.1	Current Limitations . . . . .	63
11.1.1	Measurement Accuracy . . . . .	63
11.1.2	Data Variety Limitations . . . . .	63
11.1.3	Generalization Challenges . . . . .	63
11.2	Operational Requirements Analysis . . . . .	63
11.3	Extension to Other Magnet Families . . . . .	63
11.3.1	Transfer Learning for Higher Order Magnet Families . . . . .	63
11.3.2	Other Magnet Families . . . . .	63
11.4	Future Developments . . . . .	63

**12 Conclusions**

**64**

# List of Figures

2.1	Idealized transverse magnetic field lines in the magnet aperture for a dipole, quadrupole, and sextupole. . . . .	21
a	Dipole . . . . .	21
b	Quadrupole . . . . .	21
c	Sextupole . . . . .	21
2.2	Tune diagram illustration up to 3rd order resonance with a sample working point.	23
2.3	Layout of the CERN SPS lattice, showing the six main sections (sextants) and the location of the main magnet families. . . . .	31
2.4	Simplified representation of a FODO cell in the SPS lattice. . . . .	31
2.5	Cross-sectional views of the MBA and MBB dipole magnets used in the Super Proton Synchrotron (SPS). . . . .	33
a	MBA dipole cross section. the magnet is 6,2 m long. . . . .	33
b	MBB dipole cross section. The magnet is 6,22 m long. . . . .	33
2.6	Super Proton Synchrotron (SPS) main quadrupole cross section. . . . .	34
2.7	Schematic of the dipole magnet powering topology in the Super Proton Synchrotron (SPS). . . . .	34
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. . . . .	36
2.9	Schematic of a typical $M - H$ curve for a ferromagnetic material, showing the nonlinear response and saturation behavior. . . . .	37
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. . . . .	38
2.11	Illustration of a hysteron operator $R_{\alpha\beta}$ which activates when the input signal crosses the upper threshold $\alpha$ , and deactivates when the signal decreases below $\beta$ . 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 $-1$ . . . . .	41
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 $\mu(\alpha, \beta)$ . . . . .	41
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). . . . .	43
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 $u_r$ . . . . .	43

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. . . . .	43
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. . . . .	44
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. . . . .	45
2.16	Model output with labeled parameters. . . . .	46
2.17	Model output with labeled parameters $M_s, H_c, B_r$ , optionally minor loops to illustrate drifting. . . . .	47
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. . . . .	47
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 $\tau = L/R$ . . . . .	51
2.20	Schematic cross-section of a laminated iron yoke, showing eddy current loops confined within individual lamination sheets. . . . .	52

# List of Tables

2.1	Principal parameters of the European Organization for Nuclear Research (CERN) Super Proton Synchrotron (SPS). . . . .	33
-----	--	----

# Acronyms

**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

**BPM** Beam Position Monitor

**CERN** European Organization for Nuclear Research

**CNAO** National Center for Oncological Hadrontherapy

**COSE** Constant Orbit Slow Extraction

**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

**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

**PI** Prandtl-Ishlinskii

**PS** Proton Synchrotron

**QD** Quadrupole Defocusing

**QF** Quadrupole Focusing

**RF** Radio Frequency

**RFKO** RF Knock-Out

**RL** Reinforcement Learning

**SC** Superconducting

**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,  $\rho$  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  $\delta$  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  $\xi$ , 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  $L_{\text{eff}}$  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.  $\beta$  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  $\delta B$  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  $\Delta B$ , 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  $\Delta B$  at a given field  $B$  to changes in orbit radius  $\Delta r$  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  $f_{\text{RF}}$ , 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  $f_{\text{rev}}$  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  $\phi_s$  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  $\Delta E$ , 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  $\phi_s$ , 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  $\phi_s$  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  $\eta$  is the slip factor defined as  $\eta = \alpha_c - \frac{1}{\gamma^2}$ , with  $\beta = v/c$ , and  $\gamma$  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 [47]. 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 [22, 26]. 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  $\delta$ . 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  $\delta$ , or corresponding  $\Delta r$ , 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  $\delta$  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,  $\beta$  and  $\gamma$  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 [22, 26]. 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 it 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  $\Delta x'_{\text{kicker}}$  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 spearing 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. [7, 31, 32]

### 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. [4, 50] 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 [46], 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  $\Delta r$  per three-turn period, which in the thin-lens approximation scales as

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

where  $\epsilon_{\text{res}}$  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  $\Delta r$ , 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 [14, 36], 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 Super Proton Synchrotron

The SPS [52] 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 [19] and as an electron and positron pre-injector for the Large Electron-Positron Collider (LEP) [16] 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 [39], 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 [2, 54], 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 mentioned in 1, interleaved with injecting to the LHC. In 2026, they include the Advanced Wakefield Experiment (AWAKE) [27], High Radiation to Materials (HiRadMat) [28], and various North Area experiments [5], with preparations underway for the SHiP experiment [2]. 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.



**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.

### 2.2.1 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. [52] 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  $f_{\text{rev}}$  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  $3 \times 10^{13}$  protons per batch, and up to 288 bunches per cycle for LHC injection, as well as heavy ion beams with intensities of up to  $1.5 \times 10^{11}$  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  $14 \text{ GeV}/c$ , whereas for heavy ion operation, Q26 beam is injected at  $17 \text{ GeV}/u$ . Q20 beam is injected at  $26 \text{ GeV}/c$  for protons.

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



### 2.2.2 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  $300 \text{ GeV}/c$ , the machine has been adapted to reach the  $450 \text{ GeV}/c$  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

**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	–

**(a)** MBA dipole cross section. the magnet is 6,2 m long.**(b)** MBB dipole cross section. The magnet is 6,22 m long.**Fig. 2.5:** Cross-sectional views of the MBA and MBB dipole magnets used in the SPS.


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. [44] An illustration of the cross-section of the MBA and MBB magnets is shown in Fig. 2.5. **[TODO: Magnet lamination sentence]**

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



**Fig. 2.6:** SPS main quadrupole cross section.



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

ferromagnetic iron, and suffer from the same effects as the main dipole and quadrupole magnets. **[TODO: Sextupole design parameters?]**

### 2.2.2.1 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).

### 2.2.3 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

MonitorBeam Position Monitors (BPMs), Beam Current Transformers (BCTs), and to some extent the Beam Loss MonitorBeam Loss Monitors (BLMs). [24]

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. [10, 57] 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 [12], 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. [33, 53] **[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. [58] 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. [15] 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.

## 2.3 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 2.2.2. 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



**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.

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.

### 2.3.1 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  $\chi_m$  is the magnetic susceptibility. Additionally the magnetic permeability  $\mu$  is defined as  $\mu = \mu_0 \mu_r$ . For ferromagnetic materials such as the iron yoke, however,  $\mu_r$  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.



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


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 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  $\mu_r$  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 2.3.2, and further in Section 2.4 for hysteresis modeling.

### 2.3.2 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. [9] 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.[30] 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 [6], 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 micromagnetic dynamics is the Landau-Lifshitz-Gilbert (LLG) equation [25, 37],

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

where  $\mathbf{M}$  is the local magnetization vector,  $\gamma_0$  is the gyromagnetic ratio,  $\alpha$  is the phenomenological damping constant,  $M_s$  is the saturation magnetization, and  $\mathbf{H}_{\text{eff}}$  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]**

### 2.3.3 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.

## 2.4 Hysteresis Models

As established in Section 2.3.2, 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 [9, 40, 42]. 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 2.4.3, which replace structural assumptions with flexible function approximators trained directly on measurement data.

### 2.4.1 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.

#### 2.4.1.1 The Preisach Model

The Preisach model, originally introduced to describe magnetic aftereffect in ferromagnetic materials [45] and later given a rigorous mathematical foundation by Mayergoyz [40], 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  $\alpha$  and  $\beta$  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  $\alpha$ , the hysteron switches to the  $+1$  state; when  $u(t)$  falls below  $\beta$ , it switches to the  $-1$  state; when  $u(t)$  lies between  $\beta$  and  $\alpha$ , 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,

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



**Fig. 2.11:** Illustration of a hysteron operator  $R_{\alpha\beta}$  which activates when the input signal crosses the upper threshold  $\alpha$ , and deactivates when the signal decreases below  $\beta$ . 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  $\mu(\alpha, \beta)$ .

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  $(\alpha_i, \beta_i)$ , 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 [9]. 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 [49]. 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 [48] have shown promise, and will be discussed in more detail in Section 2.4.3.

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 [40]. 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  $\beta$  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 in iron yokes are not congruent but depend on the magnetization level at which the reversal occurs [9]. 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



**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).

effective input by a quantity proportional to the current output, thereby making the minor loop shape dependent on the magnetization state [9]. 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 2.5.3. 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 2.4.3. **[TODO: more preisach extensions?, mention our recent work on neural network parametrization of the weight function]**

#### 2.4.1.2 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. [13, 41] 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  $r \geq 0$ . [43, 55] 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.

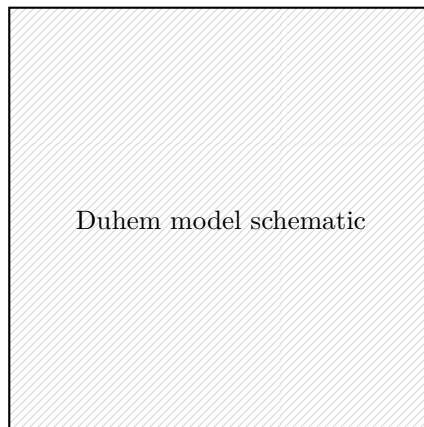
The PI model output is the weighted superposition of play operator outputs over a continuous density  $p(r) \geq 0$ ,

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

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. [1] 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. 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.



**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.

### 2.4.2 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 [20, 55],

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

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 [21], the Jiles–Atherton model [35], 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 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 [17, 59]. 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 [40]. 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 [38]. 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



**Fig. 2.16:** Model output with labeled parameters.

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.

### 2.4.2.1 The Flatley Model

The Flatley model [21] 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

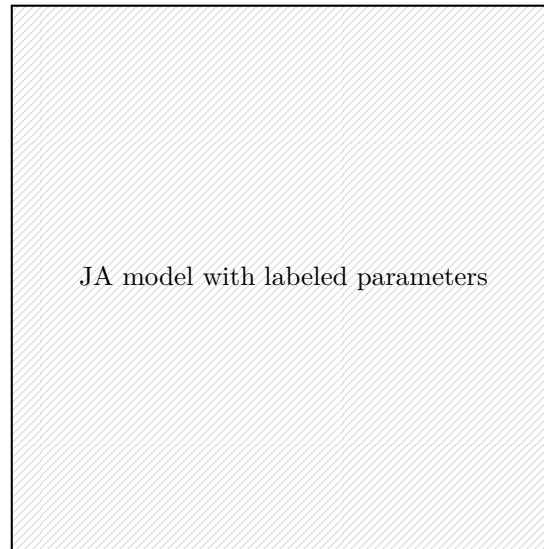
$$k = \frac{1}{H_c} \tan\left(\frac{\pi B_r}{2 B_s}\right), \quad (2.47)$$

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


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

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



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



**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

beyond the conditions under which it was fitted. Despite its limitations, the Flatley model serves as a useful baseline for model class.

### 2.4.2.2 The Jiles–Atherton Model

The Jiles–Atherton (JA) model [34, 35] 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  $M_{\text{irr}}$ , associated with domain-wall pinning, and a reversible component  $M_{\text{rev}}$ , 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  $M_{\text{an}}$  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,

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

evaluated at the effective field  $H_e = H + \alpha M$ , which accounts for interdomain coupling through the mean-field parameter  $\alpha$ , 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 [8]. 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 [17, 59]. 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 [38]. Like all Duhem models, the JA model produces congruent minor loops and is therefore unable to reproduce non-congruent behaviour [41]. 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 [17, 59]. 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 [38]. Despite these limitations, the Jiles–Atherton model remains the most widely used physics-based ODE hysteresis model in the magnetics literature.

### 2.4.2.3 Other ODE-type Models

The Fröhlich model [23] 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  $\mu_0$  is the initial permeability,  $\chi_\infty$  is the high-field susceptibility, and  $\beta$  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 [51] 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 [18, 29] 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  $\alpha$  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 [11, 56] 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  $A$ ,  $\beta$ ,  $\gamma$ , 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  $\beta$  and  $\gamma$ . 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 [59] 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 2.5. 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.

### 2.4.3 Machine Learning Approaches

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

## 2.5 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, as seen in (3.13) 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.

### 2.5.1 Eddy Currents in Conducting Materials

A time-varying magnetic flux  $\Phi$  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  $\rho$  is the electrical resistivity,  $\mu$  is the magnetic permeability, and  $\omega$  is the angular frequency of the field variation. At low frequencies or after the driving change has ceased, the currents penetrate further into the material and decay resistively.

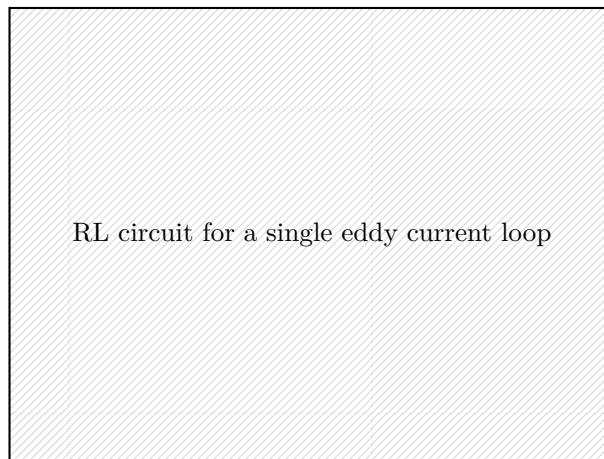
The decay of eddy currents after a step change in excitation can be described by analogy with a series Reinforcement Learning (RL) circuit, in which the inductance  $L$  represents the magnetic energy stored in the eddy current loop and the resistance  $R$  represents the ohmic dissipation. **[TODO: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

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

which has the solution

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

where  $I_0$  is the eddy current amplitude at the start of the flat excitation and  $\tau$  is the decay time constant. The secondary magnetic field contribution is proportional to  $I_{\text{eddy}}$ , so the field perturbation in the aperture decays with the same time constant  $\tau$ . **[TODO: Cite eddy current equation for RLC]**

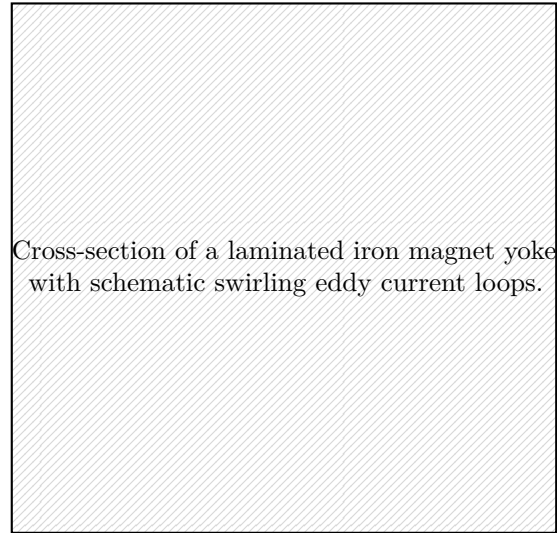


**Fig. 2.19:** Equivalent RL circuit model for a single eddy current loop in a conducting material surrounding an accelerator magnet. The inductance  $L$  represents the magnetic energy stored in the loop and  $R$  represents the ohmic dissipation. Following a ramp in the excitation current, the eddy 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,

$$B_e(t) = \sum_k A_k \exp\left(-\frac{t}{\tau_k}\right), \quad (2.60)$$

where  $A_k$  and  $\tau_k$  are the amplitude and time constant of the  $k$ -th eddy current mode. This multi-exponential character motivates the ODE-based modeling approach developed in Chapter 8.



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

## 2.5.2 Eddy Currents in Synchrotron Components

### 2.5.3 Beam Impact of Eddy Current Decay

### 2.5.4 Mitigation Strategies

[TODO: Tune compensation w.r.t. sextupoles by Jurg] [3]

# Chapter 3

## Accelerator Magnets and Magnetic Measurements

### 3.1 Accelerator Magnets

#### 3.1.1 Multipole Field Expansion and Field Quality

The magnetic field in the aperture of any accelerator magnet can be expanded as a superposition of pure multipole components. In a two-dimensional current-free region inside the magnet aperture, the transverse field  $\mathbf{B}_\perp(x, y) = (B_x, B_y)$  satisfies  $\nabla \times \mathbf{B} = 0$  and  $\nabla \cdot \mathbf{B} = 0$ , hence it is convenient to represent it by a complex multipole expansion. Introducing the complex coordinate  $z = x + iy$ , one may write for the local field

$$B_y(x, y) + iB_x(x, y) = \sum_{n=0}^{\infty} (B_n + iA_n) \left( \frac{z}{R_{\text{ref}}} \right)^n \quad (3.1)$$

where  $B_n$  and  $A_n$  are, respectively, the normal and skew harmonic coefficients, and  $R_{\text{ref}}$  is an arbitrarily chosen reference radius. With this convention,  $n = 0$  corresponds to dipole,  $n = 1$  to quadrupole,  $n = 2$  to sextupole, and so forth.

Evaluated on the magnetic axis ( $x = y = 0$ ), the  $n$ th transverse derivative of the vertical field is related to the corresponding normal harmonic by

$$\left. \frac{\partial^n B_y}{\partial x^n} \right|_{x=0, y=0} = \frac{n!}{R_{\text{ref}}^n} B_n \quad (3.2)$$

and analogously for the skew terms using  $B_x$  and  $A_n$ .

It is often useful to introduce the dimensionless normalized multipole coefficients (field quality coefficients)  $b_n$  and  $a_n$  by normalizing to a reference field  $B_{\text{ref}}$ ,

$$b_n = \frac{B_n}{B_{\text{ref}}}, \quad a_n = \frac{A_n}{B_{\text{ref}}} \quad (3.3)$$

For beam dynamics, the normalized integrated strength  $k_n$  associated with the normal  $(n + 1)$ -pole component is defined from the on-axis derivative as

$$k_n \equiv \frac{1}{B\rho} \frac{1}{n!} \left. \frac{\partial^n B_y}{\partial x^n} \right|_{x=0, y=0} \quad (3.4)$$

where  $B\rho$  is the magnetic rigidity from Eq. 2.1. Substituting Eq. (3.2) yields

$$k_n = \frac{1}{B\rho} \frac{B_n}{R_{\text{ref}}^n} \quad (3.5)$$

$$= \frac{B_{\text{ref}}}{B\rho} \frac{b_n}{R_{\text{ref}}^n} \quad (3.6)$$

For a normal quadrupole, the leading term is  $n = 1$  and the on-axis gradient is

$$g \equiv \left. \frac{\partial B_y}{\partial x} \right|_{x=0, y=0} = \frac{1!}{R_{\text{ref}}} B_1. \quad (3.7)$$

For normal dipole and quadrupole magnets, the normalized field strengths are therefore

$$k_0 = \frac{1}{B\rho} B_0 \quad (3.8)$$

$$= \frac{B_{\text{ref}}}{B\rho} b_0 \quad (3.9)$$

$$k_1 = \frac{1}{B\rho} \left. \frac{\partial B_y}{\partial x} \right|_0 \quad (3.10)$$

$$= \frac{g}{B\rho} \quad (3.11)$$

$$= \frac{1}{B\rho} \frac{B_1}{R_{\text{ref}}} \quad (3.12)$$

$$= \frac{B_{\text{ref}}}{B\rho} \frac{b_1}{R_{\text{ref}}}. \quad (3.13)$$

## 3.2 Measurement Techniques

In accelerator magnet measurement practice, an additional factor  $10^4$  is frequently included so that  $b_n$  and  $a_n$  are reported in “units” (1 unit corresponds to  $10^{-4}$  of the reference field at  $R_{\text{ref}}$ ), but the dimensionless definition in Eq. (3.1.1) is used here unless stated otherwise.

**3.2.1 Local, Integral, and Average Measurement**

**3.2.2 Hall Sensors**

**3.2.3 NMR Probes**

**3.2.4 Stretched Wire Measurement**

**3.2.5 Fluxmeter Measurement**

**3.2.6 Rotating Coils**

**3.3 Laboratory Measurement Systems**

**3.3.1 SPS Main Dipole Measurement Bench**

**3.3.2 SPS Main Quadrupole Measurement Bench**

**3.4 Online B-Train Measurement System**

**3.4.1 B-Train Systems at CERN**

**3.4.2 SPS B-Train**

# **Chapter 4**

## **Beam Operation at CERN and the SPS**

### **4.1 Magnetic Sequences as Cycles**

### **4.2 The CERN SPS**

### **4.3 Accelerator Controls Infrastructure at CERN**

#### **4.3.1 Timing**

#### **4.3.2 Middleware**

#### **4.3.3 Logging**

#### **4.3.4 LSA**

#### **4.3.5 UCAP**

#### **4.3.6 FGCs**

### **4.4 Magnetic Field Control in the CERN Injectors**

#### **4.4.1 Hierarchical Field Control**

#### **4.4.2 B-Train for Field Regulation in the PS and PSB**

### **4.5 Magnetic Precycle and Degaussing in Operation**

#### **4.5.1 Dedicated Magnetic Precycle in the SPS (MD1)**

#### **4.5.2 Quasi-Degaussing of Higher Order Magnets in the SPS**

#### **4.5.3 Degaussing and Mitigation in the PS**

#### **4.5.4 Operational Limitations from Hysteresis**

### **4.6 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**