



SAPIENZA
UNIVERSITÀ DI ROMA

PHYSICS DEPARTMENT

PH.D. COURSE IN ACCELERATOR PHYSICS

CYCLE XXXVIII

Longitudinal limitations for high-intensity beams in the upgraded Proton Synchrotron Booster

Author:

Mariangela MARCHI

ID number: 154762

Internal Supervisor:

Prof. Mauro MIGLIORATI

External Supervisor:

Dr. Simon ALBRIGHT

*Thesis submitted in fulfillment of the requirements for the
degree of Doctor of Philosophy in Accelerator Physics*

ACADEMIC YEAR 2024/2025

Longitudinal limitations for high-intensity beams in the upgraded Proton Synchrotron Booster

Ph.D. thesis, Sapienza University of Rome

Final Ph.D. defence: 7 May 2026

© 2026 Mariangela Marchi. All rights reserved.

This thesis has been typeset by L^AT_EX.

Author's e-mail: mariangela.marchi@uniroma1.it
mariangela.marchi@cern.ch

To Andrea, Caterina, Paolo and Trilly.

Abstract

During CERN's Long Shutdown 2 (LS2), the Proton Synchrotron Booster (PSB) underwent a major upgrade program as part of the LHC Injectors Upgrade (LIU) project. Key modifications include the integration of Linac4 as the new injector, increased injection and extraction energies, and the installation of new wideband Finemet-loaded RF cavities. Along with the new cavities, the digital low-level RF (LLRF) was upgraded with multi-harmonic control loops aimed at compensating beam loading. While these upgrades enhanced operational flexibility and strengthened the PSB's potential for high-intensity beam production, they also substantially altered the machine impedance and, consequently, the longitudinal stability landscape. This called for a renewed experimental assessment of the PSB longitudinal performance limits.

This thesis addresses the need to identify the dominant longitudinal limitations and their influence on the ultimate high-intensity performance of the upgraded PSB. A systematic investigation of longitudinal instability thresholds is carried out for different RF configurations through dedicated beam-based measurements. Additionally, the work combines experimental studies with advanced modelling of the main impedance contributions and RF feedback systems, with particular emphasis on the role of the Finemet cavities and beam-loading compensation mechanisms have been performed. The measured results are used to validate and refine the models and to provide quantitative benchmarks for longitudinal simulations, ultimately improving the predictive capability of beam dynamics tools and the understanding of operational stability margins.

Finally, the thesis presents the outcomes of high-intensity Machine Development (MD) campaigns, performed to demonstrate the machine's potential and to establish a solid experimental basis supporting future PSB upgrades. Overall, this work provides a consistent and experimentally anchored description of the post-LS2 longitudinal limitations of the PSB, contributing to the definition of operational strategies and modelling approaches for high-intensity performance.

Acknowledgements

A challenging and rewarding chapter of my life has come to an end.

When I first began my studies at university, I could not have imagined that pursuing a Ph.D. would one day become part of my path. Yet, through unexpected circumstances, fortunate coincidences, and the encouragement of people I met along the way, I have arrived here, at the end of a journey marked by intense work, reflection, personal growth, and deep gratitude.

This thesis would not have been possible without the support and kindness of many remarkable people whom I had the privilege to meet throughout these years.

Firstly, I would like to express my deepest gratitude to my supervisors, Dr. Simon Albright and Prof. Mauro Migliorati, for giving me the opportunity to pursue my Ph.D. under their guidance, and for supporting both me and my work throughout this time. It has been a privilege to benefit from your expertise, encouragement, and trust.

I sincerely thank Dr. Carlo Zannini and Dr. Fumihiko Tamura for accepting to review this manuscript, and for their insightful comments and valuable suggestions.

My heartfelt thanks also go to all the people with whom I had the opportunity to collaborate across different areas of accelerator physics, in particular to my colleagues in the ABP and RF groups at CERN. In this regard, I would like to thank Dr. Foteini Asvesta, Dr. Gian Piero Di Giovanni, Dr. Tirsi Prebibaj, Giulia Gnemmi, and the PSB Operations Group, whose support and collaboration were essential for the development of this work.

I am especially grateful to Dr. Diego Barrientos and Dr. Wolfgang Hofle for helping me to understand the LLRF systems of synchrotrons in great depth. It has been both extremely valuable and genuinely enjoyable to discuss LLRF systems, feedback loops, and beam control with you.

Of course, the members of the SY-RF-BR section have played an important role in both my professional and personal life during these years. I would like to acknowledge all the people I had the chance and honour to meet in this context, where I had the fortune to develop my work surrounded by extraordinary and brilliant colleagues.

My sincere thanks go to Dr. Heiko Damerau, for his endless expertise, his passion for RF, and his practical help; to Dr. Alexandre Lasheen, for his remarkable brilliance, kindness, and encouragement; and to Dr. Ivan Karpov, for his thoughtful help, curiosity, and interest in my work.

Thanks to my Italian colleagues in the section, Dr. Michela Neroni and Dr. Leandro Intelisano, for their friendship, support, and invaluable presence throughout these years. Your encouragement helped me to face the challenges of this journey and contributed greatly to my professional and personal growth.

I am grateful to everyone in the section for making it such an inspiring, welcoming, and enjoyable environment in which to work.

Un ringraziamento speciale va a Francesco, Manuel e Davide, tre grandi amici, un tecnico, un ingegnere e un fisico, incontrati in modo fortuito ma diventati essenziali nella mia vita. Grazie per aver creduto in me, per avermi aiutata, per avermi fatto sorridere e per essermi stati accanto.

Un enorme e sentito ringraziamento va alla mia famiglia, per aver sempre creduto in me, per avermi sostenuta e sopportata, per avermi dato la possibilità di studiare e di arrivare fino a questo traguardo, e per aver saputo colmare ogni distanza con la propria presenza.

Contents

1	Introduction	1
1.1	The CERN accelerator complex	2
1.2	The upgraded Proton Synchrotron Booster	4
1.3	Outline of the thesis	6
1.4	Author's contribution	7
2	Longitudinal beam dynamics in synchrotrons	9
2.1	Synchronism and longitudinal coordinates	10
2.1.1	Reference particle and synchronism	10
2.1.2	Longitudinal phase space variables	10
2.2	Longitudinal equations of motion	11
2.2.1	RF energy gain	11
2.2.2	Slip factor and phase drift	11
2.2.3	Continuous-time form and small-amplitude synchrotron frequency	11
2.2.4	Synchrotron oscillations and coherent bunch modes	12
2.3	Hamiltonian and RF bucket parameters	13
2.3.1	Single-particle Hamiltonian	13
2.3.2	Bucket height and bucket area	13
2.3.3	Longitudinal emittance	14
2.3.4	Filling factor	15
2.4	Intensity effects: wakefields, impedance, beam loading, and space charge	15
2.4.1	Wakefields and induced voltage	15
2.4.2	Potential-well distortion	16
2.4.3	Beam loading in RF cavities	17
2.4.4	Longitudinal space charge below transition	18
2.5	Double-harmonic RF and bunch shaping	19
2.5.1	Voltage waveform and relative phase	19
2.5.2	Bunch lengthening mode (BLM) and bunch shortening mode (BSM)	20

2.6	Summary	20
3	PSB RF system	23
3.1	Wideband Finemet cavities	24
3.1.1	Operating principle of wideband magnetic-alloy cavities	24
3.1.2	Layout and cellular cavity concept	26
3.1.3	Frequency response and voltage availability	30
3.2	LL RF fundamentals and PSB multi-harmonic digital feedback	31
3.2.1	Fundamentals of digital LLRF systems	32
3.2.2	PSB multi-harmonic digital feedback implementation	34
3.3	Characterisation of the PSB multi-harmonic cavity loops . .	37
3.3.1	Measurement setup and procedure	37
3.3.2	Measured closed-loop response	38
3.3.3	Comparison of the measured transfer functions . . .	38
3.3.4	Amplitude characteristics	40
3.3.5	Phase and group delay behaviour	40
3.4	Summary	44
4	Longitudinal Instability Studies	47
4.1	Stability Predictions Pre-Run 3	48
4.2	Instability Thresholds in Single-Harmonic Operation	49
4.2.1	Constant Voltage Configuration	51
4.2.2	Constant Bucket Area Configuration	55
4.3	Mode Response to LLRF Configuration	58
4.3.1	Qualitative dependence on loop configuration	60
4.3.2	Time-frequency analysis of induced oscillations . . .	62
4.4	Characterisation in Double RF Operation	62
4.5	Summary	65
5	Modelling of the PSB LLRF cavity feedback loops	69
5.1	Modelling strategy and development	70
5.1.1	Model overview and time-base separation	73
5.1.2	Input signal construction	73
5.1.3	Multi-harmonic local oscillator (MHLO)	75
5.1.4	CIC filtering and harmonic-dependent decimation . .	76
5.1.5	I/Q regulation, RF drive reconstruction, and multi-harmonic coupling	79
5.1.6	Diagnostics: I/Q constellations and controller-domain observables	82
5.1.7	Frequency-domain formulation	83

5.2	Model validation and benchmarking	85
5.3	Summary	91
6	Quadrupole synchrotron frequency shifts as a probe of the reactive longitudinal impedance	93
6.1	Principle of the method	94
6.1.1	Quadrupole oscillations and large-amplitude effects .	95
6.1.2	Synchrotron frequency for particles with large oscillation amplitude	95
6.1.3	Incoherent synchrotron frequency shift	96
6.2	PSB longitudinal reactive impedance	99
6.3	Quadrupole synchrotron frequency shifts at PSB flat bottom	101
6.3.1	Measurement and simulation setup	101
6.4	Summary	106
7	High-intensity studies in the PSB and possible future applications	109
7.1	High-intensity operation and strategy to push the intensity .	111
7.2	DC current demand of the RF power supply as an observable	112
7.2.1	Engineering interpretation of I_{DC}	113
7.2.2	Frequency dependence and multi-harmonic operation	113
7.2.3	Sampling at peak acceleration demand	115
7.3	Impact of the magnetic cycle ramp-rate on the DC current demand	118
7.4	High-intensity studies at 1.4 GeV: demonstrated single-ring reach	119
7.5	Summary	121
8	Conclusions and outlook	125
8.1	Summary of results	125
8.2	Outlook	127
A	Beam-based RF voltage calibrations during PSB commissioning	129
A.1	Principle of tomography-based voltage calibration	129
A.2	Convergence study and acquisition requirements	131
A.3	PSB Finemet RF voltage calibration	132
A.4	Summary	132
	Bibliography	134

List of Figures

1.1	The CERN accelerator complex [5].	3
2.1	Schematic of a synchrotron orbit and the curvilinear reference coordinate s used to define longitudinal motion. The RF cavity provides discrete energy kicks once per turn, while longitudinal dynamics is described relative to the synchronous trajectory.	14
2.2	Conceptual illustration of a source particle exciting a wake-field and a witness particle sampling the induced field at a delay τ . Wakefields provide the time-domain description of the longitudinal beam environment interaction.	16
2.3	Illustrative decomposition of RF voltage, induced voltage, and total voltage experienced by the beam.	17
2.4	Illustration of bunch shaping using a double-harmonic system. The stationary line density is shown for single-harmonic RF (SRF) and for double-harmonic RF (DRF) with the second harmonic operated either in counter-phase at the bunch position (BLM) or in-phase (BSM).	21
3.1	One Finemet RF station in a PSB straight section.	27
3.2	Basic accelerating cell layout.	28
3.3	Differential output (push-pull) amplifier output circuit configuration for one PSB wideband RF cell, reproduced from Ref. [23].	29
3.4	Measured longitudinal coupling impedance of one accelerating cell, including the attenuation provided by the local fast RF feedback of the Finemet power amplifier. The impedance is predominantly resistive over most of the operational frequency range and exhibits a pronounced peak around 14 MHz.	30

3.5	Block diagram of the test-stand used to characterise the PSB multi-harmonic feedback system. A VNA injects a swept excitation through Port 1 into the RF drive path of the beam and cavity controller, while Port 2 measures the transmitted response. The cavity is replaced by a jumper cable connected through power splitters, so that the closed-loop transfer function S_{21} of the digital cavity controller can be measured without beam.	37
3.6	Comparison of the measured closed-loop transfer functions for h_1 and h_8 at flat-bottom and flat-top. The dashed lines and shaded region indicate the ± 2 kHz synchrotron sideband interval considered for the comparison.	39
3.7	Group delay extracted from the measured phase responses for the four representative cases. The dashed lines and shaded region indicate the ± 2 kHz synchrotron sideband interval.	41
3.8	Detailed comparison of the h_1 servoloop response at flat-bottom and flat-top, with the ± 2 kHz synchrotron sideband interval highlighted.	42
3.9	Detailed comparison of the h_8 servoloop response at flat-bottom and flat-top, with the ± 2 kHz synchrotron sideband interval highlighted.	43
4.1	Constant voltage programs (10, 15, and 20 kV) used for the study of instability thresholds in single RF operation.	50
4.2	Constant bucket area programs corresponding to fixed filling factors of 20, 50, and 70 %.	51
4.3	Full Width Half Maximum bunch length τ_L computed at flat top energy for the Constant Voltage setup with feedback disabled. Clear thresholds in intensity define fast blow-up for all the three considered voltage amplitudes, with the intensity being proportional to the square root of the voltage peak amplitude.	53
4.4	FWHM bunch length τ_L computed at flat top energy for the Constant Voltage setup with feedback enabled. In this case, the instability mechanism is not characterised by sharp thresholds but by a gradual and uncontrolled emittance growth.	54
4.5	FWHM bunch length τ_L computed at flat top energy for the CBA setup with feedback disabled.	55

4.6	FWHM bunch length τ_L measured at flat top energy for the CBA configuration with the cavity feedback loops enabled. In contrast to open-loop operation, no sharp instability threshold is observed; instead, the bunch length increases progressively with intensity, indicating a gradual longitudinal emittance dilution that is most pronounced for the lowest filling factor ($FF = 20\%$).	56
4.7	Mountain ranges plotted for each constant bucket area study-case, showing the impact of the filling factor on the start of loops-induced perturbations.	57
4.8	Waterfall plots of 400 longitudinal bunch profiles for the CBA case with $FF = 20\%$ and all cavity servoloops enabled (CL), shown for increasing bunch intensity.	59
4.9	Tomographic reconstructions of the longitudinal phase space in the CBA configuration with $FF = 20\%$ and all cavity servoloops enabled, for two close intensity settings.	59
4.10	Waterfall plots of 400 longitudinal line-density profiles (CBA, $FF = 20\%$, $N_{p,1}$) for different multi-harmonic LLRF loop configurations.	61
4.11	Spectrograms of the FWHM bunch-length time series extracted from tomoscope profiles (CBA, $FF = 20\%$, $N_{p,1}$), shown for representative multi-harmonic loop configurations. The persistent spectral lines correspond to coherent longitudinal components whose amplitude and onset depend on which revolution-harmonic servoloops are enabled.	63
4.12	Measured bunch-oscillation amplitude versus intensity at flat top energy (2 GeV). A strong coherent dipole oscillation is observed in BSM OL, whereas it is suppressed in BSM CL by the cavity feedback loops.	65
4.13	Bunch length versus intensity at flat top energy (2 GeV). Comparing BLM and BSM in closed-loop highlights the higher sensitivity of BSM to oscillations, whereas BLM remains more robust.	66

5.1	Signal-processing chain implemented in the numerical PSB multi-harmonic LLRF model. The clock-domain cavity gap voltage is down-converted to baseband by I/Q mixing with MHLO references generated by an NCO locked to hf_{rev} . The resulting raw baseband components are filtered and decimated by two CIC stages (CIC1–CIC2), yielding controller-domain I/Q signals. A discrete-time I/Q controller generates correction signals that are expanded back to the clock domain, remodulated and summed over harmonics to obtain the RF drive contribution.	72
5.2	Clock-domain bunch current $I_b(t)$ over one turn at injection for the parabolic profile used in the model. The waveform is repeated every T_{rev} over N_{turn} turns.	74
5.3	One-turn decomposition of the model input at injection: programmed RF V_{prog} , beam-induced voltage V_{ind} computed from the Finemet impedance, and the resulting open-loop gap voltage $V_{\text{gap,open}} = V_{\text{prog}} + V_{\text{ind}}$, which enters the MHFB demodulation chain.	75
5.4	Wrapped phase evolution of the NCO phase accumulator for selected harmonics ($h = 1, 2, 8$) at injection and extraction. In both cases, the phase slope increases with harmonic number, consistent with $f_h = hf_{\text{rev}}$. The steeper evolution at extraction reflects the higher revolution frequency later in the cycle.	77
5.5	Example of I/Q signals through the demodulation chain for $h = 1$ at injection: mixer output (clock domain), after CIC1, and after CIC2 (controller domain). The CIC stages progressively suppress out-of-band components and decimate to the controller sampling rate.	78
5.6	Controller-domain sampling rate $f_{\text{ctrl},h} = f_{\text{clk}}/(R_{\text{CIC1},h}R_{\text{CIC2},h})$ as a function of harmonic number at injection and extraction. The variation with h reflects the harmonic-dependent choice of down-conversion parameters.	81
5.7	Time-coloured I/Q constellations for $h = 8$ at injection: after mixing (left), after CIC1 (centre), and after CIC2 (right). The colourbar represents time (or sample index) and highlights transient evolution and turn-periodic structures.	82
5.8	Time-coloured I/Q constellations for $h = 8$ at extraction: after mixing (left), after CIC1 (centre), and after CIC2 (right).	82

5.9	Controller-domain signals at CIC2 rate for $h = 1$ at injection: measured I/Q (and implied amplitude/phase), errors (e_I, e_Q), and controller outputs (u_I, u_Q).	83
5.10	Controller-domain signals at CIC2 rate for $h = 2$ at injection: measured I/Q (and implied amplitude/phase), errors (e_I, e_Q), and controller outputs (u_I, u_Q).	84
5.11	Combination of the first and second CIC filters present in the system demodulator, with the application of the gain compensation, shown for representative injection and extraction conditions.	86
5.12	Example oscilloscope acquisition for a given beam and voltage setting: gap-sum voltage V_{gap} (red), RF drive V_{drive} (blue), and phase pick-up signal (green), shown over a $\sim 100 \mu\text{s}$ time window.	86
5.13	Comparison between measured and simulated harmonic amplitudes for the first five revolution harmonics, for $N = 140 \times 10^{10}$ [p/b]. Left column: $\tau_l = 350$ ns. Right column: $\tau_l = 655$ ns. Blue: measurement. Red: model.	88
5.14	Time-domain comparison of the summed multi-harmonic RF drive at the DAC sampling rate. The measured signal (blue) acquired from the oscilloscope is overlaid with the model prediction (red). The model reproduces the overall amplitude build-up, envelope evolution, and modulation content of the RF drive, with small discrepancies attributable to measurement offsets and analogue chain effects.	89
5.15	Closed-loop magnitude response around the first revolution harmonic at injection. The simulated response (blue) is compared with the VNA measurement (red).	90
6.1	Schematic representation of quadrupole (bunch length) oscillations in longitudinal phase space ($\tau, \Delta E$). The mismatched portion of the distribution, relative to the part contributing to quadrupole motion, is shown in violet, while the matched area is shown in black.	96
6.2	Estimation of the PSB longitudinal space charge impedance in the cycle-time C275-C805 for a 2 GeV cycle.	100

-
- 6.3 Comparison of measured quadrupole oscillations in closed-loop and open-loop conditions and for different intensities. The defocusing and intensity-dependent effect due to the strong longitudinal space charge effect is clearly visible when the Finemet cavity reactive impedance is compensated by the enabled action of the LLRF loops. 103
- 6.4 BLoND simulations of the quadrupole synchrotron frequency f_{s2} at flat bottom as a function of bunch intensity, comparing a pure space-charge model with a combined space-charge + Finemet impedance model for two injected bunch lengths. . 104
- 6.5 Comparison between measured and simulated quadrupole synchrotron frequency shifts in the open-loop condition. The comparison is performed for different bunch lengths. . . 105
- 7.1 Single-ring ISOLDE beam intensity in 2018, before the LIU upgrades, and in 2023, before and after performance optimisations. The red trace shows pre-LIU performance, while the green and blue traces show post-LIU performance, with the blue one corresponding to the highest intensity achieved in 2023. *Courtesy of Tirsi Prebibaj*. 110
- 7.2 Example of frequency dependence of the I_{DC} observable when exciting individual harmonics of the Finemet RF system. The shaded regions indicate the frequency spans covered by successive harmonics ($h \times [f_{inj}, f_{ext}]$), highlighting that the RF power demand (and therefore I_{DC}) depends on harmonic frequency through the frequency-dependent cavity/station impedance. 114
- 7.3 DC current requested by the PSB RF Finemet power supply during the 2 GeV cycle. The signal is acquired via an oscilloscope in the CERN control room and is used as a reference for the peak RF power demand as the beam intensity is increased. The yellow marker indicates the current value at peak acceleration demand used for the intensity scan. 115

- 7.4 DC current I_{DC} requested by the PSB Finemet RF power stage at the point of peak acceleration as a function of beam intensity for the 2 GeV cycle. The top panel shows a single-ring scan, while the bottom panel overlays measurements from two rings to demonstrate reproducibility. The black dashed line indicates the nominal interlock level and the red dashed line the increased threshold temporarily authorised for the MD; data acquisition stops once the interlock is reached. Star markers show extrapolated current values at the target intensity, used to estimate the over request relative to the interlock and to evaluate the feasibility and justification of an RF power stage upgrade. In this campaign, the peak RF voltage availability did not represent a bottleneck. 117
- 7.5 Comparison of the magnetic field program $B(t)$ (top) and its time derivative $\dot{B}(t)$ (bottom) for the nominal magnetic cycle and an early-extraction configuration (extraction advanced by 100 ms). 118
- 7.6 DC current I_{DC} requested by the PSB Finemet RF power stage at the point of peak acceleration as a function of beam intensity for the 2 GeV cycle, comparing the nominal magnetic cycle (triangles) with a modified configuration (circles). The modified cycle reduces the current demand starting for intermediate intensities. Star markers indicate the extrapolated I_{DC} values at the target intensity of 1.6×10^{13} p/b for the two configurations, highlighting the expected difference in required headroom. 119
- A.1 Principle of longitudinal phase space tomography used for voltage calibration. Measured bunch profiles are combined with a tracking model (parameterised by machine settings, including the RF voltage) to reconstruct the phase space. The RF voltage is computed by minimising the discrepancy between measured profiles and the projections obtained from the reconstructed distribution. 130
- A.2 Convergence study of the tomography-based voltage calibration. The obtained voltage error stabilises once the acquisition includes approximately one synchrotron period, which was adopted as a criterion for the calibration measurements. 131

- A.3 Tomography-based RF voltage calibration results for the four PSB rings. For each ring, the voltage error is shown for sectors 5LI, 7LI and 13LI and for the summed voltage (SUM), comparing the 2023 and 2024 measurement campaigns. . . . 133

List of Tables

1.1	Run 3 PSB Beam Specifications.	5
3.1	Main quantities extracted from the representative closed-loop transfer functions.	40
3.2	Phase and delay related quantities extracted inside the ± 2 kHz synchrotron sideband interval.	41
4.1	Small-amplitude synchrotron frequency f_{s0} , synchrotron-frequency spread Δf_s , bunch intensity N_p , and full bunch length τ_{full} for the four configurations at flat top energy (2 GeV).	64
5.1	Digital down-conversion parameters at injection for each harmonic channel.	79
5.2	Digital down-conversion parameters at extraction for each harmonic channel.	80
7.1	Summary of high-intensity performance levels for the 1.4 GeV and 2 GeV cycles. The four-ring equivalent is obtained by multiplying the single-ring highest value by four.	120

Chapter 1

Introduction

The study of the fundamental constituents of matter is essential to deepen the human understanding of nature and the laws governing physical phenomena. Today, this framework is formalised in the Standard Model of particle physics [1], which has delivered precise and experimentally validated predictions across a wide range of measurements. Over the last seventy years, a large fraction of the progress in nuclear and particle physics has been enabled by particle accelerators. The operation and development of these complex machines has, in turn, been tightly connected with advances in the understanding and practical exploitation of electricity and electromagnetism, from their theoretical foundations to the implementation of technologies capable of generating and controlling intense electromagnetic fields.

Particle accelerators were initially designed to provide well regulated charged particle beams with defined particle type, energy, intensity, and time structure for nuclear and particle physics experiments [2]. Over time, accelerator technology has matured into a broader application platform, supporting uses ranging from industrial processing to medical treatments and particle therapy.

The diversity of objectives and operating regimes has led to a wide range of accelerator types [3, 4]:

- Electrostatic accelerators (e.g. Van de Graaff), valued for stable beams at low-to-medium energies.
- Linacs (linear accelerators), used as standalone machines (e.g. medical linacs) or as injectors for circular machines.
- Cyclotrons, compact circular accelerators frequently employed for isotope production and proton therapy.

- Synchrotrons and storage rings, enabling flexible acceleration and beam manipulation.
- Colliders and high-power target facilities, optimised respectively for maximum collision energy or for intensity and reliability to generate secondary beams.

Across this spectrum, performance is set not only by the maximum achievable energy, but also by intensity, beam quality, and stability. Reaching demanding operating regimes therefore requires an accurate and systematic understanding of beam dynamics and of the radio-frequency (RF) and feedback systems that control and stabilise the beam [3, 4]. Accelerator physics addresses these challenges by studying the production, transport, and manipulation of charged particle beams, and by linking beam behaviour to electromagnetic design, RF engineering, beam optics, and control systems. In practice, magnetic fields guide and focus the beam to maintain confinement along a desired trajectory, while energy is provided by electric fields, most commonly through RF electromagnetic fields sustained in resonant structures (RF cavities) [3, 4]. In circular machines, the reuse of the same RF systems turn after turn allows very high energies to be reached, while simultaneously increasing the relevance of intensity-dependent collective effects.

As modern facilities push toward higher beam intensities to meet performance goals and physics throughput, accelerators operate close to physical and technological limits. Among the most critical limitations are collective effects arising from beam-induced electromagnetic fields (wake-fields), which can drive coherent instabilities as well as incoherent emittance growth. These effects can degrade beam quality through emittance growth, bunch lengthening, or coherent oscillations and, in severe cases, lead to partial or total beam loss. Such limitations are particularly critical in injector chains, where beams must be delivered within tight specifications for downstream acceleration.

1.1 The CERN accelerator complex

CERN, the European Organisation for Nuclear Research, operates a chain of accelerators that progressively increases the beam energy and distributes beams to multiple experimental programmes. At the end of the chain, the Large Hadron Collider (LHC) is a 27 km circular collider where two counter-rotating beams are brought into collision in four main experiments:




Figure 1.1: The CERN accelerator complex [5].

ATLAS, CMS, ALICE, and LHCb. Figure 1.1 shows the current CERN accelerator complex.

For proton operation, the injector chain starts from H^- ions produced by the source and brought to a kinetic energy of 160 MeV in the linear injector Linac4. The ions are stripped during injection into the Proton Synchrotron Booster (PSB), which consists of four superimposed synchrotrons, each with a circumference of about 157 m. The PSB prepares the beam and raises its energy to 2 GeV before transfer to the Proton Synchrotron (PS), which brings it to 26 GeV. The Super Proton Synchrotron (SPS) then boosts the proton beam to 450 GeV and serves both the LHC and several fixed-target facilities. Finally, the LHC ramps the injected protons to the collision energy; during Run 3 it operates at 6.8 TeV per beam (13.6 TeV centre-of-mass energy) [6].

For heavy-ion operation, the first stages are different: ions are produced in a dedicated source and injected into the linear accelerator Linac3, which provides ~ 4.2 MeV/u at its output [7]. The beam is then accumulated and cooled in the Low Energy Ion Ring (LEIR) and extracted at 72 MeV/u [8,

9]. After transfer to the PS, the ions reach 5.9 GeV/u and are then injected into the SPS, which takes them to 177 GeV/u, corresponding to the LHC heavy-ion injection energy [7, 9].

A characteristic of the CERN accelerator complex is the need to serve diverse users and produce multiple beam types, implying a wide operational parameter space and frequent cycle changes. Any limitation in an upstream accelerator propagates to the subsequent ones: if one stage is not able to provide stable beam with the required intensity, emittance, and time structure, downstream machines cannot meet their performance targets.

For collider operation, the physics output is governed by the luminosity L , which links the observed event rate R to the production cross section σ via $R = L \sigma$. For two colliding beams, the luminosity scales with the bunch populations N_1 and N_2 ; for identical bunches $N_1 = N_2 = N_b$ one has $L \propto N_b^2$ [10]. Increasing the bunch intensity is therefore crucial to raise the number of recorded events and improve the sensitivity to rare processes.

This is one of the central drivers of the High-Luminosity LHC (HL-LHC) project, whose objective is to increase the integrated luminosity by roughly one order of magnitude beyond the LHC design programme [11]. In this context, the LHC Injectors Upgrade (LIU) programme [12, 13] was defined to enable the injector chain to deliver LHC beams with approximately doubled intensity and brightness with respect to pre-upgrade operation, which necessarily pushed the machines closer to intensity-dependent collective limitations and motivated dedicated upgrades and consolidation across the complex.

1.2 The upgraded Proton Synchrotron Booster

The PSB is the first synchrotron in the LHC injector chain and delivers a wide variety of proton beams to the CERN complex, serving both collider and fixed-target users. During Long Shutdown 2 (LS2, 2019–2021), the PSB underwent an extensive upgrade within the framework of the LIU programme. The LIU objectives were successfully reached in the PSB, with high transmission, operational flexibility, and reliability.

A core element of the upgrade was the replacement of Linac2 by Linac4 [14], raising the injection energy from 50 MeV protons to 160 MeV H^- ions and introducing charge-exchange injection into the PSB. The higher injection energy mitigates space charge effects at low energy, while the H^- injection scheme enables efficient accumulation of high-intensity beams. Beyond the injection scheme, the accelerator was extensively upgraded: the

injection region was rebuilt, new power converters and wideband Finemet RF systems were installed, and the transfer lines to the PS were adapted to the upgraded beam parameters.

In this thesis, the time coordinate along a PSB cycle is expressed using the PSB Cycle-Time convention: cycle instants are denoted by the letter C followed by the time in milliseconds (e.g. C275 for 275 ms).

After LS2, the performance of the upgraded PSB has been systematically monitored using CERN's Accelerator Fault Tracking (AFT) system [15]. Despite the increased complexity associated with the new injection scheme and extensive hardware interventions, the PSB achieved an excellent level of reliability, comparable to or exceeding pre-LS2 operation.

From an operational perspective, the upgraded injection system provides significant margin for high-intensity production through dedicated optimisations. In particular, transverse phase-space painting in the horizontal plane is performed using four dedicated Kicker SloW (KSW) magnets, generating a closed-orbit bump of up to 35 mm and allowing tailored injection conditions for the different beam types. A clear illustration of the improvement is given by high-intensity fixed-target beams: prior to the LIU upgrades, reaching the required intensities could involve substantial early-cycle losses, whereas after the upgrades the same intensity could be achieved with transmission above 99%. These results underline that the upgraded PSB not only fulfils present operational demands, but also provides headroom for future intensity increases.

Table 1.1: Run 3 PSB Beam Specifications.

	AD	TOF	MTE	ISOLDE	LHCINDIV	LHC 25ns	BCMS
L4 E_{spread} [keV rms]	280	280	450	280	100	450	
Extr. Intensity [$\times 10^{10}$ p/p]	2100	350–850	50–2400	50–3800	0.5–26	325	165
Kin. energy at extr. [GeV]	2.0			1.4	2.0		
Q_H/Q_V inj.	4.39/4.46 (ring 3) 4.22/4.30 (others)	4.24/4.41	4.11/4.34	4.22/4.36	4.17/4.23	4.32/4.44	4.28/4.34
Q_H/Q_V extr.	4.17/4.23	4.27/4.30	4.17/4.23				
Harmonics at inj.	$h = 1, 2$			$h = 1, 2$ and 3	$h = 1$	$h = 1, 2$ and 3	
Harmonics at extr.	$h = 2$ (ring 3) $h = 1$ (others)	$h = 1$	$h = 2$	$h = 1$			
ϵ_H/ϵ_V [mm mrad]	5–6/3–4	8/7	10/4	7–8/6–7	1–2/1–2	1.8	1.43
Bunch length at extr. [ns]	130	215	110	250	70	210	160
ϵ_l [eVs]	2.0	2.4	1.3	2.0	0.3	3.0	1.5
dp/p	1.3	1.1	1.3	1.1	0.5	1.0	1.0

Table 1.1 summarises Run 3 operating points for the PSB users, highlighting the wide range of beam parameters that must be produced operationally by the accelerator. The extracted intensity spans almost four orders of magnitude, from very low-intensity beams for LHCINDIV ($0.5\text{--}26 \times 10^{10}$ protons per pulse [p/p]) up to high-intensity cycles for ISOLDE and AD (up to 3.8×10^{13} [p/p] and 2.1×10^{13} [p/p], respectively). With the ex-

ception of ISOLDE, which is extracted at 1.4 GeV, most beams are delivered at 2.0 GeV, reflecting the post-LIU energy upgrade.

The table also illustrates the diversity of optics and RF configurations. The working point is adjusted between users, with different injection tunes and a comparatively uniform extraction tune for most beams. Multi-harmonic RF operation is routinely employed at injection (typically $h = 1, 2$, and for some beams $h = 3$), whereas extraction is generally performed on a single harmonic. The transverse emittances and bunch show the trade-offs between brightness and intensity: LHC beams require small emittances (order 1–2 mm mrad) and comparatively short bunches, while high-intensity fixed-target beams tolerate larger transverse emittances and longer bunches. Finally, the listed longitudinal emittance and momentum spread quantify the required longitudinal beam quality at extraction.

1.3 Outline of the thesis

This thesis is organised as follows. After this introductory chapter, which motivates the high-intensity context of the upgraded PSB and its role in the CERN injector complex, the subsequent chapters address progressively more specific aspects of longitudinal beam dynamics, RF systems, and intensity limitations in the post-LS2 PSB.

Chapter 2 introduces the longitudinal dynamics framework used throughout the thesis, including synchronism, longitudinal phase space variables, the equations of motion, RF bucket concepts, and the main intensity effects (wakefields/impedance, beam loading, and longitudinal space charge below transition). Double-harmonic RF operation and its use for bunch shaping is also discussed.

Chapter 3 presents the upgraded PSB RF system, with emphasis on the wideband Finemet cavities and the multi-harmonic digital LLRF architecture. The chapter includes the experimental characterisation of the cavity feedback loops.

Chapter 4 reports dedicated longitudinal instability studies in the PSB after LS2, focusing on experimentally determined instability thresholds and on the dependence of coherent motion on RF configuration (single-harmonic and double-harmonic) and on the multi-harmonic cavity-loop settings.

Chapter 5 develops and validates a numerical model of the PSB multi-harmonic cavity feedback loops, reproducing the main digital signal-processing chain and benchmarking the model against measured intermediate observables and closed-loop transfer functions. The resulting model

is a tool for quantitative studies of beam-cavity-feedback interaction.

Chapter 6 uses quadrupole synchrotron-frequency shifts as an experimental probe of the reactive longitudinal impedance at PSB flat bottom, benchmarking the relative roles of longitudinal space charge and the Finemet cavity impedance and clarifying the effect of open-loop versus closed-loop cavity operation.

Chapter 7 summarises the outcomes of high-intensity campaigns in the PSB and discusses operational strategies and observables (including RF power-demand indicators) to push intensity, together with perspectives for future performance improvements.

1.4 Author's contribution

All work presented in this thesis was carried out by the author. This includes the conception, setup design and execution of the measurement campaigns, the analysis of experimental data, the full development and validation of the numerical and analytical models, and the interpretation of the results in the context of PSB longitudinal performance after LS2. When results build on operational activities or collaborative machine study periods, the author's contribution consists of the complete analysis workflow, the RF expertise and the modelling/interpretation that connect the measurements to the conclusions drawn in this manuscript.

Chapter 2

Longitudinal beam dynamics in synchrotrons

Longitudinal beam dynamics describes the evolution of particle energy and arrival time with respect to the RF waveform. In synchrotrons, the RF system provides both the net energy gain required by the magnetic cycle and the longitudinal focusing that confines particles into stable regions of phase space (RF buckets). The resulting synchrotron motion determines key operational beam parameters such as bunch length, momentum spread, and longitudinal emittance, and it governs the response to controlled RF manipulations (capture, bunch rotation, splitting, and multi-harmonic bunch shaping).

At high intensity, the beam no longer behaves as a passive probe of the RF field. The circulating bunch excites electromagnetic fields in the surrounding structures; these wakefields act back on the beam through an induced voltage, which can distort the effective RF potential well, shift the synchronous phase, modify the synchrotron-frequency distribution, and ultimately drive longitudinal instabilities. These mechanisms are central for the PSB, where most operational cycles accelerate a single proton bunch and performance is often constrained by intensity-dependent effects.

This chapter introduces (i) the synchronism condition and the longitudinal equations of motion, (ii) the Hamiltonian framework and RF bucket parameters (height, area, emittance, filling factor), and (iii) intensity effects through longitudinal impedance, beam loading, and longitudinal space charge. The chapter concludes with a compact description of double-harmonic RF operation and its two principal configurations: bunch lengthening mode (BLM) and bunch shortening mode (BSM), which will be used later to interpret stability measurements.

2.1 Synchronism and longitudinal coordinates

2.1.1 Reference particle and synchronism

In a synchrotron, dipole magnets define the reference orbit through the magnetic rigidity

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

where p_s is the synchronous momentum and q the particle charge. The revolution period of the synchronous particle is

$$T_{\text{rev}} = \frac{C}{\beta_s c}, \quad \omega_{\text{rev}} = \frac{2\pi}{T_{\text{rev}}}, \quad (2.2)$$

with machine circumference C , relativistic factor β_s , and the speed of light c .

Stable acceleration requires that the beam samples the RF waveform at a reproducible phase each turn. This is enforced by the synchronism condition

$$\omega_{\text{rf}} = h \omega_{\text{rev}}, \quad (2.3)$$

with harmonic number $h \in \mathbb{N}$. This condition implies that RF frequency programs must follow the magnetic cycle whenever β_s changes appreciably.

2.1.2 Longitudinal phase space variables

Two equivalent coordinate pairs are commonly used:

- RF phase ϕ and energy deviation ΔE ,
- arrival time offset τ and energy deviation ΔE .

They are related by $\phi = \omega_{\text{rf}}\tau$ (up to an arbitrary reference offset). In this thesis we use $(\phi, \Delta E)$ for analytical developments and $(\tau, \Delta E)$ when discussing measured bunch profiles and tracking-style interpretation.

We define deviations with respect to the synchronous particle:

$$\Delta E \equiv E - E_s, \quad \Delta\phi \equiv \phi - \phi_s, \quad \tau \equiv \frac{\Delta\phi}{\omega_{\text{rf}}}. \quad (2.4)$$

2.2 Longitudinal equations of motion

2.2.1 RF energy gain

The energy gain per turn from the RF system is

$$\Delta E_{\text{rf}} = q V_{\text{rf}} \sin \phi, \quad (2.5)$$

with total RF voltage amplitude V_{rf} and RF phase ϕ at cavity passage. The synchronous particle satisfies

$$\Delta E_s = q V_{\text{rf}} \sin \phi_s, \quad (2.6)$$

where ΔE_s is set by the magnetic ramp (or $\Delta E_s = 0$ on a flat bottom/top plateau).

2.2.2 Slip factor and phase drift

The fractional revolution-frequency change due to momentum deviation is described by the slip factor η :

$$\frac{\Delta \omega_{\text{rev}}}{\omega_{\text{rev}}} = -\eta \frac{\Delta p}{p_s}, \quad \eta = \alpha_c - \frac{1}{\gamma_s^2}, \quad (2.7)$$

with momentum compaction α_c and relativistic factor γ_s . The transition energy is defined by $\eta = 0$ (i.e. $\gamma_s = \gamma_t$).

2.2.3 Continuous-time form and small-amplitude synchrotron frequency

In a smooth approximation, the longitudinal equations of motion can be written as

$$\dot{\Delta \phi} = h \omega_{\text{rev}} \eta \delta, \quad (2.8)$$

$$\dot{\Delta E} = \frac{q V_{\text{rf}}}{T_{\text{rev}}} \left[\sin(\phi_s + \Delta \phi) - \sin \phi_s \right], \quad (2.9)$$

where $\delta \simeq \Delta p/p_s \simeq \Delta E/(\beta_s^2 E_s)$.

Linearising for $|\Delta \phi| \ll 1$ yields harmonic motion with synchrotron angular frequency

$$\omega_{s0} = \omega_{\text{rev}} \sqrt{\frac{h |\eta| q V_{\text{rf}} \cos \phi_s}{2\pi \beta_s^2 E_s}}, \quad f_{s0} = \frac{\omega_{s0}}{2\pi}. \quad (2.10)$$

Longitudinal stability requires $-\eta \cos \phi_s > 0$.

2.2.4 Synchrotron oscillations and coherent bunch modes

The RF focusing causes particles to perform synchrotron oscillations around the stable fixed point $(\Delta\phi, \Delta E) = (0, 0)$. In the small-amplitude limit, the motion is harmonic with angular frequency ω_{s0} (Eq. (2.10)), corresponding to a synchrotron tune $Q_s = f_{s0}/f_{\text{rev}}$.

Dipole and quadrupole oscillations

Experimental diagnostics observe the longitudinal bunch distribution rather than single-particle trajectories, so it is useful to distinguish coherent modes through low-order moments of the line density $\lambda(\tau, t)$. Two key observables are the bunch centroid (dipole moment)

$$\langle \tau \rangle(t) = \frac{\int \tau \lambda(\tau, t) d\tau}{\int \lambda(\tau, t) d\tau}, \quad (2.11)$$

and the bunch length (quadrupole moment)

$$\tau_L^2(t) = \frac{\int (\tau - \langle \tau \rangle)^2 \lambda(\tau, t) d\tau}{\int \lambda(\tau, t) d\tau}. \quad (2.12)$$

A dipole oscillation corresponds to a coherent displacement of the bunch within the RF bucket: the centroid $\langle \tau \rangle(t)$ oscillates around the stable fixed point, while the internal bunch shape is only weakly affected. To first order, the distribution in phase space translates (and rotates) as a whole, so that the bunch length $\tau_L(t)$ remains approximately constant. Dipole motion is therefore the natural signature of a phase/arrival-time mismatch.

A quadrupole oscillation corresponds to a coherent breathing of the bunch: the centroid stays close to the fixed point, but the bunch alternately enlarges and compresses in longitudinal phase space, producing an oscillation of the bunch length $\tau_L(t)$. Quadrupole motion is typically excited by a mismatch of the bunch shape with respect to the RF bucket (e.g. voltage mismatch).

Mode number and multiples of the synchrotron frequency

Small coherent oscillations can be classified by an integer mode index m , with $m = 1$ for dipole and $m = 2$ for quadrupole. In the linear limit, the

characteristic frequency of mode m is close to an integer multiple of the single-particle synchrotron frequency:

$$f_{s,m} \simeq m f_{s0}, \quad (m = 1, 2, \dots), \quad (2.13)$$

so that quadrupole oscillations are expected near $2f_{s0}$. In realistic RF buckets, nonlinearity makes the single-particle synchrotron frequency amplitude-dependent, which generates an incoherent synchrotron-frequency spread Δf_s . This spread provides Landau damping [16] through phase mixing and is therefore important for longitudinal stability.

2.3 Hamiltonian and RF bucket parameters

2.3.1 Single-particle Hamiltonian

The longitudinal dynamics can be expressed through a Hamiltonian of the form

$$H(\Delta\phi, \Delta E) = \frac{h\omega_{\text{rev}}\eta}{2\beta_s^2 E_s} (\Delta E)^2 + \frac{qV_{\text{rf}}}{2\pi} \left[\cos(\phi_s + \Delta\phi) - \cos \phi_s + \Delta\phi \sin \phi_s \right], \quad (2.14)$$

where constant offsets are chosen such that $H(\Delta\phi = 0, \Delta E = 0) = 0$. Closed contours of H correspond to bounded synchrotron motion; the separatrix defines the RF bucket boundary.

2.3.2 Bucket height and bucket area

For a single-harmonic RF system without intensity effects, the longitudinal dynamics can be described by the Hamiltonian in Eq. (2.14). The separatrix defines the boundary of captured motion, from which the bucket height and area can be derived.

For a stationary bucket ($\phi_s = 0$), the maximum energy deviation (bucket height) is given by

$$\Delta E_{\text{sep}} = \sqrt{\frac{2\beta_s^2 E_0 q V_{\text{rf}}}{\pi h |\eta|}}, \quad (2.15)$$

showing the characteristic scaling

$$\Delta E_{\text{sep}} \propto \sqrt{V_{\text{rf}}}. \quad (2.16)$$




Figure 2.1: Schematic of a synchrotron orbit and the curvilinear reference coordinate s used to define longitudinal motion. The RF cavity provides discrete energy kicks once per turn, while longitudinal dynamics is described relative to the synchronous trajectory.

The corresponding bucket area is

$$A_{bk} = 16 \sqrt{\frac{\beta^2 E_0 q V_{rf}}{2\pi h |\eta|}}, \quad (2.17)$$

which also increases with RF voltage and depends on the slip factor η .

For accelerating buckets ($\phi_s \neq 0$) or in the presence of multi-harmonic RF systems, the separatrix becomes asymmetric and no closed analytical expression exists. In these cases, both the bucket height ΔE_{sep} and area A_b are computed numerically from the separatrix associated with Eq. (2.14).

2.3.3 Longitudinal emittance

The longitudinal emittance quantifies the area occupied by the beam in longitudinal phase space. In the $(\phi, \Delta E)$ coordinates, it is defined as

$$\varepsilon_l \equiv \oint \Delta E d\phi, \quad (2.18)$$

where the integral is taken over the phase space contour enclosing the bunch.

Using the time-energy coordinates $(\tau, \Delta E)$, the emittance can be expressed as

$$\varepsilon_l = \oint \Delta E d(\omega_{\text{rf}}\tau) = \omega_{\text{rf}} \oint \Delta E d\tau, \quad (2.19)$$

highlighting that ε_l is proportional to the area in the $(\tau, \Delta E)$ phase space.

In practice, the longitudinal emittance is often approximated as

$$\varepsilon_l \sim \Delta E_{\text{rms}} \cdot \tau_{\text{rms}}, \quad (2.20)$$

linking it to measurable quantities such as the bunch length τ and energy spread ΔE .

For conservative longitudinal motion, Liouville's theorem ensures that the emittance is preserved. However, non-adiabatic RF manipulations, injection mismatch, RF noise, and collective effects (e.g. space charge and beam coupling impedance) can lead to emittance growth through filamentation and diffusion processes.

2.3.4 Filling factor

To quantify how full a bucket is, we define the filling factor as the ratio of the bunch emittance to the bucket area:

$$\text{FF} \equiv \frac{\varepsilon_l}{A_{bk}}. \quad (2.21)$$

For fixed ε_l , increasing V_{rf} enlarges A_{bk} and therefore reduces FF. Operationally, lowering FF generally improves robustness to perturbations (more margin to the separatrix), but it can also reduce Landau damping if the potential becomes more linear and the synchrotron-frequency spread decreases. This trade-off is central to the stability observations discussed later.

2.4 Intensity effects: wakefields, impedance, beam loading, and space charge

2.4.1 Wakefields and induced voltage

As a bunch passes through accelerator components, it excites electromagnetic fields that persist behind the source particles. These wakefields act




Figure 2.2: Conceptual illustration of a source particle exciting a wakefield and a witness particle sampling the induced field at a delay τ . Wakefields provide the time-domain description of the longitudinal beam environment interaction.

on trailing particles and can be described by a longitudinal wake function $W(\tau)$.

For a bunch with N_p particles and normalised line density $\lambda(\tau)$, the induced voltage is given by the convolution

$$V_{\text{ind}}(\tau) = -qN_p \int_{-\infty}^{\infty} \lambda(\tau') W(\tau - \tau') d\tau'. \quad (2.22)$$

In frequency-domain, defining the bunch spectrum $I(\omega)$ and impedance $Z_{\parallel}(\omega)$ as Fourier transforms of λ and W , we obtain

$$V_{\text{ind}}(\omega) = -Z_{\parallel}(\omega) I(\omega). \quad (2.23)$$

Resistive parts of $Z_{\parallel}$ lead to net energy loss per turn (synchronous phase shift), while reactive parts distort the potential well and modify the synchrotron-frequency distribution.

2.4.2 Potential-well distortion

In the presence of induced voltage, the effective voltage seen by the beam becomes

$$V_{\text{tot}}(\phi) = V_{\text{rf}}(\phi) + V_{\text{ind}}(\phi), \quad (2.24)$$

and the total potential well is modified accordingly. A typical decomposition of the voltage into the one programmed in the RF cavity V_{rf} , induced V_{ind} , and total V_{tot} contributions is illustrated in Fig. 2.3.




Figure 2.3: Illustrative decomposition of RF voltage, induced voltage, and total voltage experienced by the beam.

2.4.3 Beam loading in RF cavities

Beam loading refers to the voltage induced by the beam in RF cavities. The induced cavity voltage changes the amplitude and phase of the cavity field; this must be compensated, for instance, by RF feedback loops. Even when compensation is active, the beam effectively interacts with a modified impedance determined by the cavity transfer function and the loop gain. In particular:

- in open-loop conditions, transient beam loading can imprint coherent motion and distort bunch profiles;
- in closed-loop conditions, the effective cavity impedance is reduced at specific frequencies, but loop-induced excitations can occur if the loop dynamics couples to synchrotron motion.

This distinction provides the conceptual basis for the different instability signatures observed experimentally when cavity feedback loops are enabled or disabled.

2.4.4 Longitudinal space charge below transition

At low and intermediate energies, the dominant collective interaction within a single bunch is often the longitudinal space charge force. It originates from the self-electric field generated by the longitudinal modulation of the line charge density, together with the associated magnetic field of the moving beam. In the ultra-relativistic limit the electric and magnetic contributions nearly cancel, whereas, in the non ultra-relativistic case, the cancellation is incomplete and the net longitudinal self-field scales approximately as $1/(\beta^2\gamma^2)$. Consequently, longitudinal space charge is typically a low energy effect.

Reactive impedance representation and geometry factor

In the impedance formalism, longitudinal space charge can be modelled as a reactive impedance. For frequencies well below the transverse cut-off of the vacuum chamber, it is commonly approximated by a constant $\text{Im}\{Z_{\parallel}\}/n$ (with $n = f/f_{\text{rev}}$), i.e. an inductive like interaction in the longitudinal equations of motion.

A convenient relation expresses the induced voltage from a constant reactive impedance as proportional to the derivative of the line density:

$$\Delta V(\tau) = -\frac{i}{\omega_s} \frac{Z_{\parallel}}{n} \frac{d\lambda}{d\tau}, \quad (2.25)$$

where $\lambda(\tau)$ is the normalised longitudinal line density, τ is the time coordinate within the bunch, and ω_s is the synchrotron angular frequency used for normalisation in this expression.

The space charge impedance per harmonic can be written in the form

$$\frac{Z_{\text{sc}}}{n} = -i \frac{Z_0}{\beta \gamma^2} \bar{g}, \quad (2.26)$$

where $Z_0 \simeq 377 \Omega$ is the impedance of free space and $\bar{g}$ is a geometry factor that accounts for the transverse beam distribution and the finite conducting aperture of the vacuum chamber.

Physically, $\bar{g}$ encodes how strongly the longitudinal self-field is screened by image currents and depends on the ratio between a characteristic beam size and the beam-pipe radius. The longitudinal space charge effect is strongest for particles close to the bunch center (small transverse radius), and weaker for particles at larger transverse amplitudes.

Combining Eqs. (2.25)-(2.26), we obtain an induced space charge voltage whose magnitude scales with $Z_0 \bar{g}/(\beta\gamma^2)$ and with the local slope of the line density:

$$\Delta V_{sc}(\tau) \propto -\frac{Z_0}{\beta\gamma^2} \bar{g} \frac{d\lambda}{d\tau}. \quad (2.27)$$

Impact on the RF potential well and synchrotron motion

The induced space charge voltage modifies the effective focusing around the stable fixed point, leading to potential-well distortion and an intensity-dependent change of the synchrotron-frequency distribution.

A particularly important aspect for the PSB is that operation is below transition ($\eta < 0$). In this regime, the sign of the synchrotron-frequency shift produced by a given reactive impedance differs from the above transition case. In the small-amplitude approximation, the incoherent synchrotron-frequency shift can be written in the generic form

$$\frac{\Delta\omega_s}{\omega_{s0}} \propto \frac{qN_p}{V_{rf}\omega_{rf} \cos \phi_s} Z_1, \quad (2.28)$$

where Z_1 is the effective reactive impedance.

Below transition ($\eta < 0$), phase stability requires $-\eta \cos \phi_s > 0$, which implies $\cos \phi_s > 0$. In this regime, a capacitive reactive impedance contributes with the opposite sign to the RF curvature at the synchronous phase, effectively weakening the longitudinal focusing. As a consequence, the restoring force is reduced, leading to a lower small-amplitude synchrotron frequency and, typically, to a smaller incoherent synchrotron tune spread.

In the PSB, where operation remains below transition throughout the cycle, these considerations motivate the use of multi-harmonic RF at low energy: by reshaping the RF potential and reducing the peak line density, the relative importance of space charge defocusing is mitigated and the stability margin can be improved.

2.5 Double-harmonic RF and bunch shaping

2.5.1 Voltage waveform and relative phase

With two RF systems (fundamental and higher harmonic), the total RF voltage can be written as

$$V_{rf}(\phi) = V_1 \sin(\phi + \phi_{s1}) + V_2 \sin(n\phi + n\phi_{s1} + \Phi_2), \quad (2.29)$$

where $n = h_2/h_1$ is the harmonic ratio (in the PSB typically $n = 2$), V_2/V_1 is the voltage ratio, and Φ_2 is the relative phase between the harmonics. The additional harmonic reshapes the RF potential and modifies the non-linearity, changing the amplitude-dependent synchrotron frequency and the incoherent synchrotron-frequency spread [17].

2.5.2 Bunch lengthening mode (BLM) and bunch shortening mode (BSM)

Two configurations are of particular relevance:

- BLM (counter-phase at the bunch position): the potential well is flattened around the stable phase, yielding longer bunches and reduced peak line density. This generally increases Landau damping by enhancing the synchrotron-frequency spread and reduces intensity-driven effects.
- BSM (in-phase at the bunch position): the central focusing is increased, producing shorter bunches and a higher peak line density. Compared to single-harmonic RF, the added harmonic still introduces additional nonlinearity and therefore generally increases the incoherent synchrotron-frequency spread. However, because the bunch is shorter (and the peak current larger), it can sample reactive impedance and beam loading more strongly.

The resulting steady-state bunch shaping for SRF and DRF configurations at fixed emittance is shown in Fig. 2.4.

2.6 Summary

This chapter provided the longitudinal beam dynamics fundamentals used throughout this thesis. After introducing the synchronism condition and the longitudinal equations of motion, the small-amplitude synchrotron frequency and the interpretation of coherent dipole and quadrupole oscillations were established to connect the theory to the experimental observables used later.

A Hamiltonian description was then employed to define the RF bucket and the main operational parameters: bucket height, bucket area, longitudinal emittance, and filling factor. These quantities are used in the following chapters, as design RF parameters. The chapter subsequently introduced intensity effects through the wakefield/impedance formalism,




Figure 2.4: Illustration of bunch shaping using a double-harmonic system. The stationary line density is shown for single-harmonic RF (SRF) and for double-harmonic RF (DRF) with the second harmonic operated either in counter-phase at the bunch position (BLM) or in-phase (BSM).

leading to induced voltage, potential-well distortion, and beam loading in RF cavities. Longitudinal space charge in the below transition regime was also presented, where its predominantly reactive contribution weakens longitudinal focusing and can affect synchrotron frequency spread and stability margins.

Finally, double-harmonic RF operation was presented as a controlled method to reshape the longitudinal potential and tune the synchrotron-frequency distribution. The distinction between bunch lengthening mode (BLM) and bunch shortening mode (BSM) provides the conceptual basis for the stability measurements and comparisons between RF configurations discussed in the remainder of this thesis.

Chapter 3

PSB RF system

The PSB RF upgrade accomplished during LS2, within the framework of the LIU project, had a major impact on beam production schemes, operational flexibility, and longitudinal beam stability. The former tunable ferrite systems were replaced by wideband magnetic-alloy (MA) cavities based on Finemet material, together with a new digital system capable of multi-harmonic cavity control. The wideband behaviour of the Finemet cavities, extending over approximately 1–18 MHz, enables simultaneous operation on multiple revolution harmonics and therefore provides a high degree of flexibility for longitudinal beam manipulations. This is particularly advantageous for bunch shaping, controlled longitudinal blow-up, and optimisation of beam quality for the different PSB operational cycles. At the same time, the broadband response of the cavity system implies a substantial longitudinal coupling impedance, so that beam loading becomes a key design and operational constraint. For intense beams, the beam-induced voltage can significantly perturb the accelerating field, making direct local feedback and carefully designed cavity control feedback loops essential.

During the last few years (2022-2026), the upgraded PSB RF system could be systematically studied and characterised, together with the longitudinal beam dynamics resulting from its operation. In this chapter, the upgraded PSB RF system is described in detail. The first part focuses on the High-Level RF (HLRF) system, including the principles of MA cavities and the specific implementation adopted in the PSB. The second part introduces the main concepts of digital Low-Level RF (LLRF) control and presents the multi-harmonic cavity controller developed for the PSB. Along with that a detailed characterisation of the transfer function of the LLRF cavity feedbacks is also presented.

3.1 Wideband Finemet cavities

This section introduces the HLRF system installed in the PSB. Before describing the specific hardware implementation, it is useful to recall the operating principle of wideband magnetic-alloy cavities and the reasons why Finemet-based structures are particularly well suited to the requirements of the upgraded machine.

3.1.1 Operating principle of wideband magnetic-alloy cavities

Magnetic-alloy loaded cavities were developed as an alternative to conventional ferrite-loaded accelerating systems, whose operation is generally based on a narrowband resonant response and therefore requires active tuning, typically through bias-current control of the ferrite permeability. By contrast, MA cavities are designed to operate over a broad frequency interval without the need of tuning the cavity. Their impedance is intentionally spread over a wide band, so that a single cavity can sustain a large frequency sweep and, when required, support the simultaneous excitation of several harmonics. This feature is particularly attractive in synchrotrons where the RF programme must provide not only acceleration, but also flexible longitudinal manipulations such as bunch shaping, flattening, or controlled blow-up [18].

The physical origin of this behaviour lies in the complex and frequency-dependent permeability of MA materials,

$$\mu(\omega) = \mu'(\omega) - j\mu''(\omega), \quad (3.1)$$

where μ' determines the inductive component of the response, while μ'' describes the magnetic losses. In MA materials, μ'' remains appreciable over a broad frequency range. As a consequence, the cavity quality factor remains low, typically of the order of unity, and the resulting impedance bandwidth is intrinsically broad [19].

From a hardware point of view, MA cavities are typically realised as coaxial structures in which toroidal magnetic cores provide the reactive and dissipative part of the impedance, while the accelerating voltage is developed across a ceramic gap. The broad frequency response is therefore not obtained by dynamically tuning a sharp resonance, but rather by exploiting the lossy magnetic behaviour of the core material itself. The price to pay for this tuning-free operation is that the cavity behaves as a

wideband impedance source and can therefore sustain significant beam-induced voltage over a large frequency range.

Several classes of magnetic alloys have been developed for accelerator applications, including Fe-based amorphous alloys, Co-based amorphous alloys, and nanocrystalline Fe-based alloys. Among them, nanocrystalline alloys of the Finemet family have become particularly important because they combine favourable RF performance with good industrial reproducibility and acceptable cost [18, 19]. Finemet cores are manufactured from thin ribbons wound into toroidal rings and subsequently annealed in a controlled way to produce the nanocrystalline structure responsible for their magnetic properties [18, 20]. This processing route is crucial, since the final RF performance depends not only on the alloy composition but also on the resulting domain structure and its evolution under alternating magnetic field.

For cavity design, MA materials are often characterised in terms of their parallel complex permeability, which is more directly connected to the achievable shunt impedance of the cavity. A convenient figure of merit is the impedance product, commonly written as $\mu'_p Q f$, which allows meaningful comparison between materials and production batches while remaining only weakly dependent on the core dimensions [19]. In this context, the magnetic properties of the core can be directly related to the voltage capability and the corresponding power dissipation, thereby providing practical guidance for cavity dimensioning and material selection.

The performance of a wideband MA cavity is not determined by the magnetic material alone. Core processing, mechanical integration, and thermal design also play a role. Magnetic annealing, for example, strongly influences the magnetic-domain microstructure and can significantly reduce RF losses associated with domain-wall motion. Operational experience from other facilities, such as J-PARC, has shown that suitable annealing procedures can substantially reduce core losses and increase the available voltage margin [18]. Since the power dissipated in the cores can be considerable, thermal management becomes a central design aspect. Depending on the required voltage gradient and average power, this may involve direct water cooling of the core stacks or indirect cooling through conductive structures such as copper plates or discs [18].

At system level, the final RF architecture must be conceived together with the cavity technology itself. Depending on the required power and beam-loading conditions, different facilities have adopted different driving solutions. In the case of CERN, the PSB upgrade led to a modular cavity concept based on multiple accelerating cells driven by radiation-tolerant solid-state amplifiers, chosen to maximise availability and simplify main-

tenance in the accelerator environment [18].

3.1.2 Layout and cellular cavity concept

Each PSB ring is equipped with three Finemet RF cavities located in the straight sections 5LI, 7LI and 13LI. Since the four PSB rings are vertically stacked, each of these straight sections houses four superimposed cavity systems, one per ring, together with the corresponding solid-state RF power amplifiers, cooling-water structures, cabling, and ancillary electronics.

A view of one Finemet RF station installed in the PSB tunnel is shown in Fig. 3.1.

Part of the DC power supplies, the digital LLRF electronics, and the interlock system are located in the corresponding surface equipment room. The wideband properties of the Finemet material are fully exploited through the cellular architecture adopted for the cavity design. This design was the result of dedicated optimisation studies [21, 22], and is based on the principle that each cell contributes a fraction of the total RF voltage.

The elementary accelerating cell consists of a vacuum chamber with a ceramic gap at its center and one magnetic-alloy core disk on either side. The geometry is constrained by the available transverse space in the PSB and by the dimensions of the MA cores compatible with the stacked-ring configuration. Each cavity is composed of twelve such cells, corresponding to 36 cells per ring. Taking into account the core material properties and the adopted cooling scheme, a single cell can provide up to about 700 V from a few hundred kilohertz to above 5 MHz. One cavity therefore delivers up to approximately 8 kV, and the three cavities together provide up to 24 kV per ring in the low-MHz region. Above 5 MHz the voltage linearly decrease to 4 kV at 18 MHz [22].

The layout of one elementary cell is shown in Fig. 3.2. A solid-state RF power amplifier is mounted on one side of the cavity, while the opposite side has been kept free for possible future upgrades. The amplifier drives the two sides of the accelerating gap in push-pull configuration.

From the LLRF output to the accelerating gap, the drive signal passes through the following main elements:

1. RF compensation network (+330 pF): a passive equalisation stage used to compensate the frequency-dependent attenuation of long cables and to linearise the overall frequency response. Since it is fully passive, it also introduces a reduction of the available signal power.



Figure 3.1: One Finemet RF station in a PSB straight section.

2. 500 mW RF amplifier: provides a gain of approximately 12 dB over the frequency range 1–18 MHz.
3. 4 W wideband amplifier: provides a gain of approximately 34 dB over 5 kHz–100 MHz.
4. 16-way bi-directional splitter/combiner: equally distributes the drive signal, with the first twelve outputs feeding the twelve per-cell power amplifiers of the cavity.
5. Per-cell push–pull Finemet power amplifiers: each cell is driven by a dedicated RF power stage including




Figure 3.2: Basic accelerating cell layout.

- an RF driver,
 - two 200 W input circuits,
 - sixteen 200 W output circuits,
 - control and monitoring interfaces,
 - a local feedback network acting continuously on the gap-return signal and providing a partial attenuation of the induced voltage of about 10 dB.
6. RF summing amplifier: combines the twelve cell-return signals into the cavity gap-return signal used by the control system.

The output stage of the PSB wideband amplifier is shown in Fig. 3.3. The outputs of eight 200 W modules are combined on each side and stepped up by a 1:3 RF voltage transformer to a nominal $50\ \Omega$ impedance. The two branches are then operated in anti-phase so as to produce a differential push-pull output on the cavity gap. Since each 200 W module already includes an internal 1:4 matching transformer, the overall voltage transformation from the MOSFET drain to the push-pull output is 24. To improve reliability, the DC supply voltage was reduced from 50 V to 40 V, which limits the maximum gap voltage of this configuration to about $750\ V_{\text{peak}}$.




Figure 3.3: Differential output (push-pull) amplifier output circuit configuration for one PSB wideband RF cell, reproduced from Ref. [23].

In the PSB cellular architecture, this corresponds to approximately $8 \text{ kV}_{\text{peak}}$ per cavity when twelve cells are operated in series.

The output transformation chain also determines how the cavity-gap current is reflected back to the individual MOSFET drains. As discussed in [23], the drain current can be expressed as

$$I_{\text{Drain}} = I_{\text{Gap}} \frac{n}{N}, \quad (3.2)$$

where N is the number of combined 200 W modules and n is the overall voltage transformation ratio from the amplifier output to each MOSFET drain. For the differential push-pull configuration, $N = 8$ and $n = 12$, giving

$$\frac{I_{\text{Drain}}}{I_{\text{Gap}}} = \frac{12}{8} = 1.5. \quad (3.3)$$

These relations are important for evaluating the current capability of the amplifier and for comparing the suitability of the two output configurations under the beam-loading conditions in the PSB.

3.1.3 Frequency response and voltage availability

At low frequency, the cavity response is mainly determined by the magnetic-alloy material properties, whereas at higher frequency the gap capacitance, cabling capacitance, and amplifier loading become increasingly important [21]. As a result, the available cavity voltage is highest in the low-MHz region and decreases towards the upper end of the operating band.

This voltage envelope is derived from cell impedance characterisation performed through commissioning measurements and compared with simulations based on equivalent-circuit fits, typically described by a ladder $R-L$ model with shunt capacitance C . Figure 3.4 shows the most recent measurement of the longitudinal coupling impedance of a single accelerating cell, while the action of the local, fast beam-loading compensation system is applied. Assuming all cells to be nominally identical, the equivalent impedance of one full ring can be approximated by multiplying the single-cell impedance by 36.




Figure 3.4: Measured longitudinal coupling impedance of one accelerating cell, including the attenuation provided by the local fast RF feedback of the Finemet power amplifier. The impedance is predominantly resistive over most of the operational frequency range and exhibits a pronounced peak around 14 MHz.

The impedance is predominantly resistive over most of the operational frequency range. It remains relatively flat at low frequency and then rises

sharply towards the upper end of the band, reaching a resonant peak around 14 MHz. An important feature of the PSB Finemet system is that this impedance is a property of the cavity hardware and therefore does not depend on the beam energy: at flat-bottom, during acceleration, and at flat-top, the beam always interacts with the same cavity impedance as a function of frequency.

In single-bunch operation, acceleration is performed at the fundamental RF frequency, which follows the revolution frequency and therefore almost doubles over the cycle. This frequency span is highlighted in violet in Fig. 3.4. In addition to acceleration, the RF system is routinely used for bunch shaping, typically by means of a second-harmonic voltage, with harmonic ratio $h_r = 2$ with respect to the fundamental. The corresponding frequency span is indicated in pink. To improve beam stability a controlled longitudinal emittance blow-up is performed by allocating voltage at a higher harmonic, often chosen to be the $h = 10$ (in cyan). The figure therefore illustrates not only the broadband character of the cavity system, but also the extent to which the beam samples different regions of the cavity impedance during the machine cycle.

3.2 LL RF fundamentals and PSB multi-harmonic digital feedback

The LLRF system forms the control layer of the RF plant. Its purpose is to ensure that the accelerating fields seen by the beam follow the prescribed programme in amplitude, phase, and timing despite external disturbances such as beam loading, component drifts, or changes in operating frequency throughout the cycle. In the PSB, these tasks are made more demanding by the broadband nature of the cavity system and by the requirement to control several harmonics simultaneously.

This section first recalls the main concepts underlying digital LLRF systems in particle accelerators, with particular emphasis on complex signal representation, frequency synthesis, digital down-conversion, filtering, and feedback. The discussion then focuses on the specific implementation adopted in the PSB for multi-harmonic cavity control.

3.2.1 Fundamentals of digital LLRF systems

Control objectives

At a system level, the role of the LLRF system is to regulate the RF voltage experienced by the beam. In practice, this means closing one or more feedback loops around measured RF quantities in order to control:

- the amplitude and phase of the cavity voltage,
- synchronisation with the beam and machine timing,
- disturbance rejection in the presence of beam loading, ripple, or plant variation,
- and, when required, cavity tuning for resonant systems.

In the case of broadband cavities or multi-frequency operation, control must often be performed simultaneously at several harmonic frequencies, which naturally motivates a complex-baseband, multi-channel implementation.

Complex-envelope and I/Q representation

A narrowband RF signal centred at angular frequency ω_c can be represented as

$$v(t) = \Re\{\tilde{v}(t)e^{j\omega_c t}\}, \quad (3.4)$$

where $\tilde{v}(t) = I(t) + jQ(t)$ is the complex envelope. In this notation,

$$|\tilde{v}(t)| = A(t), \quad \arg \tilde{v}(t) = \phi(t), \quad (3.5)$$

so regulating amplitude and phase is equivalent to regulating the in-phase and quadrature components I and Q .

This representation is central to digital LLRF design because it linearises many operations of interest, such as gain control, phase rotation, vector summation, and PI regulation.

Digital down-conversion and remodulation

The standard digital processing chain starts from a sampled RF or intermediate frequency (IF) signal. A numerically generated local oscillator

provides sine and cosine references at the selected frequency. By multiplying the input by these references and applying low-pass filtering (LPF), the in-phase and quadrature components are obtained:

$$I(t) \sim \text{LPF}\{v(t) \cos(\omega_{\text{LO}}t)\}, \quad (3.6)$$

$$Q(t) \sim \text{LPF}\{v(t) \sin(\omega_{\text{LO}}t)\}. \quad (3.7)$$

This operation is known as digital down-conversion. The inverse operation, digital up-conversion or remodulation, reconstructs the RF drive from the corrected complex baseband signal.

In accelerator applications, phase coherence between down-conversion and remodulation is essential. Any inconsistency directly appears as a systematic phase or amplitude error at cavity level.

Sampling, clocks, and deterministic timing

The performance of a digital LLRF system is strongly influenced by the sampling frequency, the total processing latency, and the clock distribution. In closed-loop operation, the achievable control bandwidth is limited by the total delay, including analog-to-digital and digital-to-analog conversion (ADC, DAC), digital filtering, transport, and RF-chain response.

A fixed and globally distributed clock is generally preferred for robust synchronisation across several processing nodes. Frequency programmes are then translated into numerically controlled oscillator tuning words while maintaining deterministic phase relationships through synchronous updates.

NCO/DDS frequency synthesis, filtering and decimation

Sinusoidal references are commonly generated digitally by means of a numerically controlled oscillator (NCO) and direct digital synthesiser (DDS). For a phase accumulator of width N and a clock frequency f_{clk} ,

$$f_{\text{out}} = \frac{\text{FTW}}{2^N} f_{\text{clk}}, \quad (3.8)$$

from which

$$\text{FTW} = f_{\text{out}} \frac{2^N}{f_{\text{clk}}}. \quad (3.9)$$

In synchrotrons, f_{out} is generally tied to the revolution frequency or one of its harmonics. The resolution and synchronisation of frequency updates have a direct impact on phase continuity and reproducibility.

After demodulation, low-pass filtering is required to isolate the baseband component and reject images. In FPGA-based implementations, cascaded-integrator-comb (CIC) filters are often used because they allow efficient decimation with low multiplier usage.

Digital feedback law

In complex baseband, the most common control law is of the proportional-integral (PI) type, applied independently to the in-phase and quadrature errors $e_{I,Q}$:

$$u_{I,Q}(s) = \left(K_p + \frac{K_i}{s} \right) e_{I,Q}(s). \quad (3.10)$$

In practice, this is complemented by phase rotations, gain normalisation, compensation of transport delays. Robust operation requires sufficient phase margin under worst-case delay conditions and careful handling of saturation and quantisation.

Beam loading and the need for multi-harmonic control

In a broadband cavity, the beam-induced voltage extends over a wide range of frequencies. Correcting only a single harmonic is therefore insufficient in many operational situations. A multi-harmonic LLRF architecture addresses this problem by splitting the signal into several narrowband branches, each centered on a selected harmonic of the revolution frequency, and by applying a dedicated correction to each branch before recombining them into the RF drive.

3.2.2 PSB multi-harmonic digital feedback implementation

The PSB constitutes a particularly demanding application for digital cavity control. The Finemet cavities are intrinsically broadband, the revolution frequency changes continuously during the cycle, and the beam-induced voltage can be large over several revolution harmonics. For this reason, the upgraded LLRF does not rely on a single narrow-band loop, but on a fully digital multi-harmonic cavity controller able to regulate, in parallel, up to 16 harmonics of the revolution frequency for each individual cavity system [24–26].

At a global level, the cavity controller is implemented in FPGA and follows a fixed-frequency-clock paradigm. The cavity gap signal is digitised at

approximately 122.7 MSPS, and all harmonic processing is performed synchronously with this common master clock. This design choice replaces the older sweeping-clock scheme and brings several practical advantages: the ADCs and DACs always operate at their optimal conversion rate, the analogue anti-aliasing and reconstruction filters can remain fixed, harmonic distortion is pushed out of band, and the full ADC sample stream can be exploited without sacrificing granularity at low revolution frequency. The price to pay is the need for reliable synchronous distribution of the revolution frequency and for absolute synchronism between all NCOs in the system. In the CERN platform, this is achieved by multicasting the 122.7 MHz clock together with a TAG line carrying the encoded revolution-frequency information and, when required, a double-TAG synchronisation event for phase-accumulator reset.

The overall control architecture is based on 16 parallel single-harmonic I/Q loops, all operating on the same ADC input stream. Each loop contains a demodulator, a loop filter, and a modulator, while the harmonic references are generated centrally by a multi-harmonic local oscillator (MHLO). Rather than synthesising each harmonic independently, the firmware generates the fundamental complex sinusoid with a classical numerically controlled oscillator, consisting of a phase accumulator and a COordinate Rotation Digital Computer (CORDIC) engine, and then derives the higher harmonics by cascaded complex multiplications. In this way, harmonic coherence is guaranteed by construction and FPGA resource usage remains modest. Only the first harmonic FTW is generated and distributed, while the higher harmonics are built by frequency addition in the complex plane, for example $h = 2$ from $h = 1 \times h = 1$, $h = 3$ from $h = 2 \times h = 1$, and so on, with pipelining chosen so that all harmonic LOs emerge synchronously [27].

For each harmonic h , the cavity gap signal is digitally down-converted to baseband by mixing it with the corresponding complex LO generated by the MHLO. The resulting in-phase and quadrature components are then processed by a two-stage digital filter chain. The first stage is a CIC demodulation filter operating directly at the fixed 122.7 MHz clock, with a decimation factor chosen so that its output rate tracks the revolution frequency. This filter is responsible for the main rejection of neighbouring revolution harmonics. Its gain is automatically corrected so that the demodulator response remains essentially independent of f_{rev} . A second post-filter then further reduces the data rate to a range suitable for the digital signal processing (DSP) control layer, typically between 120 and 240 kSPS. In the fixed-frequency implementation this post-filter changes its decimation exponent automatically as the revolution frequency evolves,

thereby maintaining a roughly constant processed data rate across the cycle. This organisation, shown in both the 2017 fixed-clock presentation and the servo-loop implementation note, is one of the key features that allows narrow-band harmonic control with a fixed master clock.

After demodulation, each harmonic channel enters a Cartesian feedback loop. The error signal is formed between the complex measured cavity voltage and a harmonic-dependent set-point, and is processed by a dual I/Q PI controller implemented fully in FPGA DSP resources. The set-point itself can be defined either from registers or dynamically through function channels, and is in practice constrained by several supervisory mechanisms. First, a cavity-dependent balancing gain redistributes the requested voltage according to the number of active Finemet cells in each cavity, so that cavities with reduced availability are automatically requested to deliver less voltage. Second, a frequency-dependent maximum-level limiter prevents the controller from requesting voltages outside the cavity operating envelope. Third, an overflow-driven limiter acts when ADC or DAC saturation is detected: the integrators are frozen, the requested set-point is slowly reduced, and the controller resumes normal operation only once the loop error becomes negative again.

Two independent phase rotations are applied in each harmonic branch. The first is the cavity rotator, placed after demodulation, whose role is to reconstruct the cavity voltage in a common physical reference frame, i.e. as seen by the beam. This compensates the electrical length differences of the gap-return cables and the azimuthal position of the cavity around the ring. In the PSB, where the three cavity groups are installed in sectors 5LI, 7LI and 13LI, the azimuth correction is essential to ensure that the vector sum of the RF voltages from different straight sections remains aligned for all harmonics. The second rotation is the loop rotator, which corrects the residual phase of the closed loop so as to maintain adequate loop stability margins [26].

Before remodulation, an additional feed-forward correction is applied in gain and phase. This compensation is derived from commissioning measurements of the complete chain from RF drive to gap-voltage monitor and is stored in lookup tables. For each harmonic and frequency point, the controller interpolates gain and phase corrections so as to compensate the frequency-dependent non-linearities of the plant. In this way, the closed-loop gain and phase margins can be kept more uniform over the full operating range [26]. The vector modulator then remodulates the corrected baseband signal back to RF using the same harmonic LO family. All 16 RF contributions are summed digitally and converted by the DAC into the final RF drive signal sent to the HLRF chain.

3.3 Characterisation of the PSB multi-harmonic cavity loops

After the description of the upgraded PSB RF system hardware implementation, the experimental establishment of how the cavity control loops behave in closed-loop is reported. Dedicated measurements were carried out on a test-stand in order to characterise the transfer function of the cavity controller and to provide a quantitative reference for subsequent modelling work.

3.3.1 Measurement setup and procedure

The characterisation was performed on a dedicated test-stand reproducing the cavity-control chain while omitting the physical cavity itself. In this setup, the cavity was replaced by a cable, which provides a suitable first-order approximation for the present purpose, since the Finemet cavity response is broadband and relatively flat over the frequency interval of interest.




Figure 3.5: Block diagram of the test-stand used to characterise the PSB multi-harmonic feedback system. A VNA injects a swept excitation through Port 1 into the RF drive path of the beam and cavity controller, while Port 2 measures the transmitted response. The cavity is replaced by a jumper cable connected through power splitters, so that the closed-loop transfer function S_{21} of the digital cavity controller can be measured without beam.

The test-stand and setup used for the characterisation of the PSB multi-harmonic feedback system is shown schematically in Fig. 3.5 [28]. A vector network analyser (VNA) is used both as excitation source and as measurement instrument. The swept signal generated at Port 1 is injected into the

RF drive path of the beam and cavity controller through a power splitter, thereby emulating a small external perturbation. The response of the closed-loop system is sampled at the return path through a second splitter and routed back to Port 2 of the VNA, which measures the corresponding forward transmission coefficient S_{21} .

The system was operated at fixed frequencies corresponding to selected revolution harmonics. Around each target harmonic, the VNA swept a frequency interval of ± 50 kHz, thus covering the closed-loop notch and the surrounding passband region. This procedure was repeated for different harmonics in order to evaluate how the loop response evolves across the operational frequency range.

3.3.2 Measured closed-loop response

In order to assess how the PSB multi-harmonic cavity feedback behaves under different operating conditions, four representative cases were compared: the servoloops acting on harmonics h_1 and h_8 , each evaluated at flat-bottom and flat-top. This choice makes it possible to discuss two complementary aspects. First, one can compare two different servoloops at the same fixed revolution frequency. Second, one can evaluate how the response of a given servoloop changes between flat-bottom and flat-top.

The loops were designed such that, in principle, the closed-loop attenuation and bandwidth should be broadly similar in all cases [28]. The target values were a notch depth of about -36 dB and a total -3 dB bandwidth of about 16 kHz, corresponding to approximately ± 8 kHz around the notch center. The measurements presented here show that this ideal uniformity is not fully achieved in practice.

3.3.3 Comparison of the measured transfer functions

Figure 3.6 compares the measured transfer functions for the four selected cases. The upper panel shows the magnitude response, while the lower panel shows the phase response. The horizontal axis is expressed as frequency offset from the notch center. The dashed lines and shaded region indicate the ± 2 kHz interval around the notch, which is taken here as a representative synchrotron sideband span to be compensated by the harmonic feedback.

All four cases exhibit the expected narrowband rejection centered on the controlled harmonic, confirming the qualitative functionality of the harmonic servoloops. However, the measured responses are not identi-




Figure 3.6: Comparison of the measured closed-loop transfer functions for h_1 and h_8 at flat-bottom and flat-top. The dashed lines and shaded region indicate the ± 2 kHz synchrotron sideband interval considered for the comparison.

cal. The notch depth, the effective width of the rejection region, and the phase evolution around the notch vary from case to case. This shows that, although the servoloops share the same design principle, the realised closed-loop response depends both on the harmonic number and on the operating frequency.

A first comparison can be made between h_1 and h_8 at the same operating point. This comparison shows that two different servoloops at the same fixed revolution frequency do not produce exactly the same notch. The differences are visible both in the depth of the rejection and in the detailed shape of the amplitude and phase responses. A second comparison can be made for a given servoloop between flat-bottom and flat-top. In this case, the measurements show that the response is not invariant with frequency: the same harmonic branch exhibits a different closed-loop behaviour when operated at the two points of the cycle.

3.3.4 Amplitude characteristics

The main quantitative descriptors of the rejection are the notch depth and the -3 dB bandwidth. The notch depth measures the maximum suppression achieved by the loop, while the -3 dB bandwidth characterises the frequency span over which the rejection remains significant. These quantities are summarised in Table 3.1.

Table 3.1: Main quantities extracted from the representative closed-loop transfer functions.

Case	f_0 [MHz]	Depth [dB]	BW [kHz]	Mean att. in ± 2 kHz [dB]
h_1 FB	1.000	70.65	8.40	-52.31
h_8 FB	8.000	66.08	5.89	-55.74
h_1 FT	1.810	13.91	60.65	-51.22
h_8 FT	14.480	53.74	11.20	-49.87

The measurements clearly show that the realised attenuation and bandwidth differ from the nominal design targets.

From the beam dynamics point of view, the most relevant quantity is the attenuation inside the synchrotron sideband interval. This region corresponds to the frequency span in which the loop is expected to compensate perturbations around the controlled harmonic.

3.3.5 Phase and group delay behaviour

The phase responses provide complementary information on the local dispersive behaviour of the loop. In the present case, the phase must be interpreted with care because the transmission becomes very small at the notch center. A global phase unwrapping would therefore introduce artificial 360° branch jumps. For this reason, the phase in Fig. 3.6 is displayed as a wrapped phase referenced to the off-notch region.

To quantify the phase behaviour more robustly, the group delay was also evaluated from the unwrapped phase according to

$$\tau_g = -\frac{d\phi}{d\omega}, \quad (3.11)$$

which is the standard RF and DSP definition of group delay [29]. Figure 3.7 shows the resulting group delay comparison.




Figure 3.7: Group delay extracted from the measured phase responses for the four representative cases. The dashed lines and shaded region indicate the ± 2 kHz synchrotron sideband interval.

The group delay remains close to zero away from the notch in all cases, showing that the closed-loop response is only strongly dispersive in a narrow region around the rejected harmonic. At the exact notch center, the group delay estimate exhibits a large spike because the transmission approaches zero and the phase derivative becomes numerically singular. For this reason, the value at the exact notch center is not interpreted quantitatively. Instead, the relevant quantities are extracted inside the ± 2 kHz interval around the notch.

These phase and delay related quantities are reported in Table 3.2.

Table 3.2: Phase and delay related quantities extracted inside the ± 2 kHz synchrotron sideband interval.

Case	$\Delta\phi_{\pm 2 \text{ kHz}}$ [deg]	$\langle \tau_g \rangle_{\pm 2 \text{ kHz}}$ [μs]	$ \tau_g _{\text{max}, \pm 2 \text{ kHz}}$ [μs]
h1 FB	183.3	94.650	2355.899
h8 FB	181.1	-25.128	1484.578
h1 FT	181.3	-33.935	1911.398
h8 FT	169.1	-60.582	1008.566

For completeness, Figs. 3.8 and 3.9 show the detailed responses for h_1 and h_8 , respectively, at flat bottom and flat top, again with the synchrotron sideband region highlighted.

These detailed views make it more explicit that the same servoloop




Figure 3.8: Detailed comparison of the h_1 servoloop response at flat-bottom and flat-top, with the ± 2 kHz synchrotron sideband interval highlighted.




Figure 3.9: Detailed comparison of the h_8 servoloop response at flat-bottom and flat-top, with the ± 2 kHz synchrotron sideband interval highlighted.

does not preserve exactly the same notch shape between flat bottom and flat top, and that the two harmonic branches also differ from one another at a given operating point.

3.4 Summary

This chapter presented the upgraded RF system of the CERN PSB, with emphasis on both its hardware architecture and its closed-loop control behaviour.

The first part of the chapter described the principles of wideband magnetic-alloy cavities and their implementation in the PSB. The Finemet technology was introduced as a low- Q , tuning-free alternative to ferrite-loaded cavities, in which the broad frequency response is obtained from the intrinsic complex permeability of the magnetic-alloy material. The specific PSB implementation was then presented: each ring is equipped with three cellular cavities, each made of twelve accelerating cells driven by dedicated solid-state push-pull power amplifiers. The modular cavity concept was shown to provide both operational flexibility and maintainability, while the measured impedance and voltage envelope highlighted the broadband nature of the system and its consequences for beam-RF interaction.

The second part of the chapter introduced the digital LLRF concepts required to understand the upgraded PSB control system. After recalling the principles of I/Q representation, digital down-conversion, NCO/DDS frequency synthesis, CIC filtering, and PI feedback, the chapter described the PSB multi-harmonic cavity controller. The implemented architecture is based on 16 parallel harmonic loops operating synchronously on a fixed clock, with a multi-harmonic local oscillator, harmonic-by-harmonic demodulation and remodulation, Cartesian feedback, cavity and loop phase rotations, and gain/phase feed-forward corrections. This architecture makes it possible to regulate the wideband cavity response over the full PSB frequency programme while compensating beam-induced perturbations on several harmonics simultaneously.

Finally, the chapter reported the experimental characterisation of the cavity feedback loops on a dedicated test-stand. Closed-loop transfer-function measurements were performed for representative cases corresponding to harmonics h_1 and h_8 at flat bottom and flat top. The results confirmed the expected narrow-band rejection around the controlled harmonic, but also showed that the realised loop response is not identical for all harmonics and operating points. In particular, the measured attenuation, bandwidth, and phase behaviour differ from case to case and do not fully

reproduce the nominal design goal of uniform notch performance. The phase and group delay analysis further showed that the loop is strongly dispersive only in a narrow interval around the notch, corresponding to the synchrotron sideband region of interest.

In conclusion, this chapter established the technological and experimental framework needed for the rest of the thesis. It showed that the upgraded PSB RF system combines a broadband, flexible HLRF plant with a sophisticated multi-harmonic digital control layer, and that the practical closed-loop response depends on both harmonic number and operating frequency. These results provide the basis for the numerical modelling developed in the following chapter and for the subsequent analysis of the impact of the RF system on longitudinal beam dynamics.

Chapter 4

Longitudinal Instability Studies

In the PSB, the vast majority of operational cycles accelerate a single bunch. As the proton bunch circulates, it interacts with the surrounding environment, which is defined by the electric and magnetic fields related to the various components specified in the synchrotron design. Complex linear and non-linear beam dynamics arise from these fields, coupled to the initial conditions of the particles in the bunch.

The fields generated by the beam interacting with the vacuum chamber vicinities, known as wakefields, can perturb the external designed electromagnetic fields, used to guide and accelerate the beam. As the intensity increases, the perturbation driven by these fields becomes stronger and, above a certain threshold, it leads to an instability and, often, a subsequent beam loss. Single-bunch instabilities arise from collective effects acting on the internal structure of a single bunch, as opposed to coupled-bunch instabilities that correlate the motion of many bunches through the accelerator. Over the past decades, extensive theoretical and experimental work has been devoted to understanding the underlying mechanisms and to developing effective mitigation and control strategies [30–42].

In the PSB, the upgrade of the RF system has substantially altered the longitudinal stability scenario, making a systematic characterisation of longitudinal instability patterns essential.

Such studies provide important information for validating the longitudinal impedance model and for possible optimisations of the RF program with the aim of reaching higher beam intensities [43].

In this chapter, longitudinal instability measurements are presented for various RF configurations. A first distinction is made between operation with a single-harmonic RF system (SRF) and with a double-harmonic RF system (DRF). For SRF operation, two different RF voltage setups have been

implemented during the cycle; their impact on instability thresholds and observable patterns is compared. In addition, the role of the LLRF cavity feedback loops in shaping the instability mechanisms is investigated. For DRF operation, the focus is placed on the comparison between the bunch lengthening and bunch shortening modes. It is shown and explained why the bunch lengthening mode is beneficial for longitudinal stability and for extending the achievable intensity range, whereas the bunch shortening mode is detrimental to the longitudinal beam quality.

4.1 Stability Predictions Pre-Run 3

Simulation studies performed for the PSB before LS2 [44] assume the full current PSB impedance model, with the exception of the Finemet cavity contribution, whose longitudinal impedance was re-measured and updated in the context of this thesis. For high-intensity beams, specifically with $N_p = 1.6 \cdot 10^{13}$ protons per bunch and at a kinetic energy of 2 GeV, results showed it was not possible to achieve smooth acceleration due to the occurrence of instabilities during the acceleration cycle. In addition to an overall bunch length growth, coherent dipole and quadrupole oscillations were observed. By examining the bunch distribution in longitudinal phase space at the onset of the first instability, a 21 MHz modulation was identified as perturbing the bunch [44, 45]. Based on the Finemet longitudinal impedance, it was concluded that the high-frequency induced voltage generated by the cavity could trigger a microwave instability. In fact the bunch became shorter during the acceleration cycle, inducing an interaction of the bunch spectrum with the peak of the impedance model at 19 MHz. The inclusion of negative resistors in the simulation, to implement notches at 16 multiples of the revolution frequency, had the effect to delay the insurgence of this instability, but not the complete suppression of it. This behaviour was attributed to the considered limited bandwidth of the notches (about 16 kHz), which was not sufficient to cover the larger bandwidth present at the peak of the Finemet longitudinal impedance. As a result of these studies, it was predicted that the reach of high-intensity beams would be possible under SRF operation, considering the action of the wake-cancellation feedback applied to 16 revolution-harmonics.

During LS2, dedicated operational schemes were developed for each type of beam produced by the PSB, exploiting the flexibility offered by the broadband response of the Finemet RF system. During the first beam commissioning in 2021, these schemes were put into operation [46]. However, for medium to high-intensity beams, the observed longitudinal stability

did not follow the predictions, and the production schemes had to be revised [47]. It was found that pure SRF operation was detrimental to beam stability, in contrast to expectations, whereas the inclusion of DRF, in bunch lengthening mode, restored stability even though this configuration had been predicted to be unstable.

The experimental studies presented in this chapter confirm that, under SRF conditions, the application of the LLRF beam loading compensation system alone is not sufficient to maintain longitudinal stability, already starting from intermediate bunch intensities.

4.2 Instability Thresholds in Single-Harmonic Operation

In this section, the characterisation of single-bunch instabilities under SRF operation is presented. Different voltage configurations were applied by changing only the fundamental accelerating harmonic, whose frequency increases by almost a factor of two from injection to extraction, following the revolution frequency sweep ($f_{\text{RF}} = h \cdot f_{\text{rev},0}$ with $h = 1$).

Two voltage program categories were investigated:

- (a) Constant voltage (CV) programs: three fixed RF voltage amplitudes of 10, 15, and 20 kV were used starting from cycle-time C450 to extraction. These cases allowed to isolate the dependence of the instability thresholds on the bucket height and beam loading strength.
- (b) Constant bucket area (CBA) programs: three voltage programs were designed to maintain a constant bucket area along the cycle, from C550 to flat top, corresponding to fixed filling factors of 20, 50, and 70 %. This approach enabled the study of stability at comparable longitudinal acceptance.

Each of the above categories was examined under two LLRF cavity feedback loops configurations:

- Open-Loop (OL): the action of the loops was disabled, leaving the beam–cavity interaction uncompensated;
- Closed-Loop (CL): the feedback system was active, leaving the compensation of the induced voltage from the beam current and Finemet longitudinal coupling impedance.

At the beginning of the PSB cycle, the beam is routinely captured and accelerated in DRF operation by combining the fundamental and second harmonic systems with a relative phase chosen such that the two voltages are in counter-phase at the bunch position (bunch lengthening mode, BLM). In this arrangement the effective RF potential is flattened around the stable phase, producing a longer bunch and therefore a reduced peak line density and peak current. This reduction is particularly important in the first part of the cycle, where transverse space charge forces are strongest. Keeping this configuration in the early part of the cycle thus provides a robust and reproducible longitudinal match before the instability studies are initiated.

For both the voltage program families investigated, the transition from DRF to SRF operation was therefore performed only after a sufficiently long BLM period. The SRF measurement window was opened at cycle times C450 and C550 for the CV and CBA studies, respectively. The corresponding RF voltage programs, together with the kinetic energy evolution for the 2 GeV cycle, are shown in Figs. 4.1 and 4.2.




Figure 4.1: Constant voltage programs (10, 15, and 20 kV) used for the study of instability thresholds in single RF operation.

Longitudinal bunch profiles were acquired through a wall current monitor (WCM) in the SRF region at regular time intervals. The intensity was varied by changing the number of turns injected from Linac4. Due to the injection mechanisms, this assured an approximate constant longitudinal




Figure 4.2: Constant bucket area programs corresponding to fixed filling factors of 20, 50, and 70 %.

emittance.

4.2.1 Constant Voltage Configuration

To provide a clear overview of the influence of the bunch intensity N_p and RF voltage amplitude V_{pk} on the onset of instabilities, the bunch lengths τ_L measured at flat top energy are compared. When the LLRF cavity feedback loops are disabled, a rapid emittance blow-up is observed for all three voltage amplitudes at well defined intensity thresholds, as shown in Fig. 4.3.

In the stable regime, the bunch length is nearly independent of intensity and is systematically shorter for higher RF voltages. This behaviour is consistent with the stronger longitudinal focusing associated with a larger bucket height: for an approximately constant longitudinal emittance, increasing V_{pk} raises the small-amplitude synchrotron frequency ($f_{s0} \propto \sqrt{V_{pk}}$ at fixed energy) and reduces the bunch length. The weak dependence on N_p below threshold further suggests that, in this intensity range, reactive beam-loading effects and potential-well distortion (PWD) remain secondary with respect to the externally imposed RF focusing, so that the equilibrium bunch shape is primarily set by the voltage program.

Beyond a critical intensity, the bunch length abruptly increases and

beam loss follows within a few synchrotron periods. The measured instability thresholds are $N_p \approx 2.1 \cdot 10^{12}$, $2.6 \cdot 10^{12}$, and $3.0 \cdot 10^{12}$ p/b for RF voltages of 10, 15, and 20 kV, respectively, indicating an empirical scaling close to $N_{p,\text{th}} \propto \sqrt{V_{\text{pk}}}$.

In open-loop operation, where no LLRF beam-loading compensation is applied, the bunch interacts directly with the cavity longitudinal coupling impedance. As N_p increases, the induced voltage progressively distorts the RF potential well and modifies the synchrotron-frequency distribution, reducing the effective Landau damping. Once the overlap between the bunch spectrum and the effective impedance becomes sufficiently strong, a rapid growth of the bunch length is triggered, leading to fast blow-up and beam loss. Although increasing V_{pk} shortens the bunch (and increases the peak line density), it also strengthens the longitudinal focusing and increases the rms momentum spread of the matched distribution at approximately constant longitudinal emittance. In Keil-Schnell-Boussard estimates of the microwave instability threshold, the intensity threshold scales favourably with the intrinsic energy spread, $N_{p,\text{th}} \propto \sigma_\delta^2 / |Z_\parallel / n|$ (up to form factors accounting for bunch shape, resonator bandwidth and potential-well distortion) [48]. Therefore, the stabilising effect of larger f_{s0} and larger σ_δ at higher V_{pk} can outweigh the destabilising effect of the increased peak current of a shorter bunch, consistent with the observed $N_{p,\text{th}} \propto \sqrt{V_{\text{pk}}}$ trend in open-loop.

The same measurements were repeated at flat top energy with all 16 cavity servoloops enabled at their nominal operational gains and settings, and the same analysis procedure was applied. As shown in Fig. 4.4, the abrupt blow-up observed in open-loop is no longer present, confirming that the fast instability is driven by the uncompensated Finemet cavity longitudinal coupling impedance. The digital cavity feedback cancels most of the beam-induced voltage at the fundamental component (and associated revolution-harmonic lines), thereby suppressing the rapid instability and enabling operation at substantially higher intensity than in the OL case. Compared to open-loop, the CL configuration extends the reachable range from $N_p \approx 3.5 \times 10^{12}$ to $N_p \approx 8 \times 10^{12}$ p/b (within the explored scan), but the limitation is replaced by a gradual and approximately linear increase of the bunch length with N_p , with τ_t increasing from about 140 ns at low intensity to above 210 ns at the highest intensities in Fig. 4.4, i.e. a controlled yet undesirable longitudinal emittance dilution rather than a sharp threshold.

From these observations it follows that, while the cavity loops suppress the fast open-loop blow-up, they also modify the longitudinal dynamics




Figure 4.3: Full Width Half Maximum bunch length τ_L computed at flat top energy for the Constant Voltage setup with feedback disabled. Clear thresholds in intensity define fast blow-up for all the three considered voltage amplitudes, with the intensity being proportional to the square root of the voltage peak amplitude.

in a way that produces a smooth, monotonic bunch length increase with intensity. Therefore, the feedback extends the stability domain in terms of loss-free acceleration, but at the cost of a gradual longitudinal emittance dilution. Moreover, the curves corresponding to different RF voltage levels nearly overlap, indicating that V_{pk} no longer influences the instability threshold or bunch length evolution. This confirms that the dominant limitation has shifted: the cavity impedance is largely compensated in the frequency bands where the feedback is effective, and the remaining beam quality degradation is instead ruled by residual closed-loop effects (finite bandwidth, phase lag and noise), which become more pronounced as the beam-induced signal increases with N_p .

CV case message

The main outcome of this first comparison is that, for SRF operation, a higher V_{pk} improves robustness in open-loop at moderate intensity, whereas operation at the highest intensities requires the activation of the LLRF cavity feedback system, at the cost of a gradual bunch length (and




Figure 4.4: FWHM bunch length τ_L computed at flat top energy for the Constant Voltage setup with feedback enabled. In this case, the instability mechanism is not characterised by sharp thresholds but by a gradual and uncontrolled emittance growth.

emittance) growth that must be managed operationally.

In open-loop condition, the RF voltage amplitude directly sets the longitudinal focusing through the bucket height and the small-amplitude synchrotron frequency. For an approximately constant longitudinal emittance, this determines the matched bunch length and momentum spread, and therefore the balance between the stabilising RF focusing and the destabilising beam-induced voltage. In the explored range, the measured instability threshold follows the first-order scaling

$$N_{p,\text{th}} \propto \sqrt{V_{\text{pk}}}. \quad (4.1)$$

When the cavity feedback system is enabled, the dominant components of the beam-induced voltage are actively compensated by a narrowband response centered on the revolution-harmonic lines. In those frequency bands, the effective impedance experienced by the beam is reduced. Consequently, the cavity behaves as a low impedance system for the spectral components that carry most of the single-bunch power at flat top, and the voltage seen by the beam remains close to the programmed setpoint, with a much weaker dependence on beam current than in open-loop.

4.2.2 Constant Bucket Area Configuration

In this section, the instability measurements are discussed for the CBA setup. Figure 4.5 shows the corresponding behaviour at flat top energy with the cavity feedback loops not enabled.

The three SRF voltage programs were designed to keep the RF bucket area A_{bk} approximately constant during the acceleration ramp. Since no non-adiabatic changes were introduced in the RF waveform, the longitudinal emittance ϵ_l can be assumed to remain conserved over the considered time window. It is therefore convenient to characterise the CBA settings by a constant filling factor $FF = \epsilon_l / A_{bk}$, defined as the ratio between the longitudinal emittance and the bucket area. Three values of filling factor are considered: $FF = 20\%$, $FF = 50\%$ and $FF = 70\%$.




Figure 4.5: FWHM bunch length τ_L computed at flat top energy for the CBA setup with feedback disabled.

Figure 4.5 shows that, similarly to the CV open-loop case, the transition to instability is characterised by a well defined intensity threshold, above which the bunch length increases rapidly and beam loss follows. A clear dependence on filling factor is observed: the lower filling factor ($FF = 20\%$) remains stable up to higher intensity than the $FF = 50\%$ case, indicating that operating with a larger bucket margin improves the stability threshold

in open-loop. For even higher filling factor (e.g. $FF = 70\%$), the beam is lost at comparatively low intensity and the corresponding points are therefore not included in Fig. 4.5. The achievement of a smaller filling factor requires a higher RF voltage amplitude for the same longitudinal emittance. Consequently, in open-loop operation the $FF = 20\%$ case corresponds to the highest voltage program and naturally exhibits a slightly higher instability threshold, in line with the trend already observed in the CV open-loop study where larger V_{pk} shifted the threshold to higher N_p .




Figure 4.6: FWHM bunch length τ_L measured at flat top energy for the CBA configuration with the cavity feedback loops enabled. In contrast to open-loop operation, no sharp instability threshold is observed; instead, the bunch length increases progressively with intensity, indicating a gradual longitudinal emittance dilution that is most pronounced for the lowest filling factor ($FF = 20\%$).

In closed-loop the filling factor has a clear impact on the bunch length behaviour. The $FF = 20\%$ configuration exhibits the strongest intensity dependence, with a pronounced and progressive bunch length increase as N_p rises, indicating significant longitudinal emittance dilution. In contrast, the higher filling factor cases ($FF = 50\%$ and $FF = 70\%$) remain comparatively insensitive to intensity, suggesting that the larger occupied phase space area (and the associated synchrotron frequency spread) tends to cure loop-induced perturbations. The low filling factor nevertheless provides the largest margin in longitudinal acceptance and therefore allows

the intensity to be pushed to higher values, albeit with degraded beam quality.




Figure 4.7: Mountain ranges plotted for each constant bucket area study-case, showing the impact of the filling factor on the start of loops-induced perturbations.

Figure 4.7 depicts the mountain range plots during the first 100 ms of acquisition (C550-C650). The intensities at which the profile evolutions are shown are the ones at which perturbations start (except the 70 % case), being $1.45 \cdot 10^{12}$ p/b, $1.91 \cdot 10^{12}$ p/b and $2.85 \cdot 10^{12}$ p/b from left to right. For the 20 % case, the bunch exhibits sensitivity to high-frequency components. However, a stabilising effect is observed at larger bucket areas, most likely due to the enlarged synchrotron frequency spread.

CBA case message

The CBA measurements highlight that the action of the filling factor depends critically on whether the cavity loops are active. In open-loop, for higher filling factors, the instability develops earlier, as expected from reduced bucket margin and lower synchrotron frequency. In closed-loop the picture reverses: the 20 % case develops loop-driven perturbations at lower

energies and lower intensity than the others. The higher filling factor cases (50 % and 70 % cases) are more stable in terms of loop-induced oscillations and gradual emittance growth. However, at very high intensity, beam loss appears earlier in OL, while CL allows the beam to reach higher N_p ($\approx 3 \cdot 10^{12}$ p/b) albeit with degraded quality. Overall, the cavity impedance dominates. In CL, even though the 20 % case shows early modulation, the loops still suppress the fast instability.

4.3 Mode Response to LLRF Configuration

The closed-loop CBA measurements of Section 4.2.2 showed that the FF = 20% working point is the most sensitive to loop-induced perturbations. This configuration was therefore selected for a dedicated study on the excitations sustained or shaped by the multi-harmonic LLRF beam-loading compensation system.

In the previous analysis, the feedback system was treated as a binary setting (all loops open versus all loops closed). Here, the configuration space is explored in a more differential way by selectively opening individual revolution-harmonic servoloops while keeping the remaining loops closed. This approach provides a direct experimental handle to identify which harmonic feedback paths are most strongly correlated with the emergence of coherent longitudinal motion.

As a first step, an intensity scan was performed with all cavity servoloops enabled in order to identify the onset of coherent longitudinal motion and to quantify how the corresponding energy threshold depends on the bunch intensity. Figure 4.8 shows representative waterfall plots at low and high N_p : as the intensity is increased, bunch profile modulations become visible at earlier times in the cycle, indicating that the excitation threshold in energy decreases with N_p .

Two representative intensity levels are discussed in more detail, $N_{p,1} \approx 1.75 \times 10^{12}$ p/b and $N_{p,2} \approx 2.0 \times 10^{12}$ p/b, chosen because they lie close to the region where loop-induced perturbations become clearly visible while the beam remains within the longitudinal acceptance.

To inspect the internal structure of the coherent motion, tomographic reconstructions of the longitudinal phase space were performed at these two operating points (Fig. 4.9). In both cases, the reconstructed density shows a four-lobed pattern, which is characteristic of an azimuthal octupole mode ($m = 4$). The structure is more pronounced at $N_{p,2}$, suggesting a stronger coupling between radial and azimuthal degrees of freedom as the intensity is increased. From the reconstruction, the small-amplitude




Figure 4.8: Waterfall plots of 400 longitudinal bunch profiles for the CBA case with $FF = 20\%$ and all cavity servoloops enabled (CL), shown for increasing bunch intensity.

synchrotron frequency at this time in the cycle is $f_s \approx 829$ Hz, which is used in the following to interpret the spectral lines observed in the time-frequency analysis.




Figure 4.9: Tomographic reconstructions of the longitudinal phase space in the CBA configuration with $FF = 20\%$ and all cavity servoloops enabled, for two close intensity settings.

The presence of a four-lobed structure in Fig. 4.9a–4.9b indicates that the coherent motion has a strong octupolar ($m = 4$) content. This moti-

vates complementing the usual scalar bunch-length diagnostics (FWHM or RMS) with observables that are explicitly sensitive to higher-order shape distortions. Indeed, an $m = 4$ oscillation can leave the overall bunch length nearly unchanged while producing a pronounced internal redistribution of line density and phase space density. To resolve this internal structure, the dynamics must be analysed in terms of longitudinal modes: the octupolar pattern projects onto fourth-order components of the phase space distribution, and is therefore more naturally quantified through higher-order moments (such as the fourth central moment or excess kurtosis) and, when possible, through a modal decomposition of the reconstructed distribution.

In the following, the impact of selectively opening individual servoloops is quantified using waterfall plots together with frequency-domain and time-frequency diagnostics.

4.3.1 Qualitative dependence on loop configuration

Figure 4.10 compares the longitudinal profile evolution at $N_{p,1}$ for several multi-harmonic beam-loading compensation configurations. The comparison highlights that the qualitative behaviour is mainly governed by which revolution-harmonic components of the induced voltage are actively compensated and, on the contrary, which ones are left to act back on the bunch.

When the feedback action is removed specifically at $h = 4$ and $h = 5$, while keeping the remaining harmonics closed (Fig. 4.10f), the structured modulations observed in the reference closed-loop case are strongly reduced. At this intensity the profile evolution becomes close to the fully open-loop behaviour, indicating that the coherent perturbations seen in closed loop are tightly correlated with the control action around the fourth and fifth revolution harmonics: these loops are therefore the most influential in seeding or sustaining the observed coherent motion.

A complementary trend emerges when the low-order harmonic compensation is progressively removed. Configurations that leave the low-frequency part of the beam spectrum effectively uncompensated (i.e. with several of the first harmonics opened) favour the onset of fast, growing oscillations towards flat top. In particular, a clear growth in the modulation amplitude at flat top is observed only when the first eight loops are disabled, whereas the instability disappears when they are enabled. This supports the interpretation that the fast regime is associated with insufficient compensation of the induced voltage in the low-frequency/revolution-harmonic region, where the single-bunch spectrum carries a significant




Figure 4.10: Waterfall plots of 400 longitudinal line-density profiles (CBA, $FF = 20\%$, $N_{p,1}$) for different multi-harmonic LLRF loop configurations.

fraction of its power. Enabling the beam-loading compensation in this range suppresses the blow-up, but this also implies operating with the $h = 4$ and $h = 5$ loops active, which are precisely the ones most strongly coupled to the slower, loop-driven profile perturbations.

Finally, repeating the $h = 4$ and $h = 5$ open-loop configuration at the

higher intensity $N_{p,2}$ confirmed that a smooth profile evolution can still be maintained, reinforcing the conclusion that, for this working point, the dominant closed-loop perturbation channel is linked to the control paths at the fourth and fifth revolution harmonics, whereas the suppression of the fast instability requires compensation of the lowest harmonics.

4.3.2 Time-frequency analysis of induced oscillations

Figure 4.11 summarises the time-frequency content of the FWHM bunch length observable for the main loop configurations. In the all-closed reference case (Fig. 4.11a), the spectrum is dominated by a narrow, stationary line, indicating a coherent oscillation with nearly constant frequency over the acquisition window. Opening selected loops modifies this picture in two ways: it can (i) reduce the amplitude of the coherent component by removing the loop-driven excitation path, or (ii) introduce additional low-frequency content and a broader spectral footprint, consistent with increased envelope modulation and emittance growth. In particular, the configuration with $h = 4$ and $h = 5$ open (Fig. 4.11e) exhibits a markedly reduced coherent signature with respect to the fully closed case, in line with the qualitative waterfall observations that the profile perturbations strongly diminish when these two harmonic feedback paths are disabled. Overall, the spectrograms provide a compact way to distinguish persistent coherent motion from broadband or slowly varying components, and they support the interpretation that a subset of the low-order harmonic loops plays a disproportionate role in seeding or sustaining the observed longitudinal modulation at this working point.

4.4 Characterisation in Double RF Operation

The addition of a higher-harmonic RF voltage modifies the longitudinal potential and introduces further non-linearities, increasing the incoherent synchrotron frequency spread.

In routine operation, the PSB employs a second-harmonic DRF scheme ($h = 2$) in BLM. This configuration is required early in the cycle to mitigate strong transverse space charge effects at low energy, and it remains beneficial at higher energies for longitudinal stability.




Figure 4.11: Spectrograms of the FWHM bunch-length time series extracted from to-moscope profiles (CBA, $FF = 20\%$, $N_{p,1}$), shown for representative multi-harmonic loop configurations. The persistent spectral lines correspond to coherent longitudinal components whose amplitude and onset depend on which revolution-harmonic servoloops are enabled.

Synchrotron-frequency spread and stability measurements

To quantify the stabilising mechanism, the synchrotron-frequency distribution was measured at flat top energy (2 GeV) as a function of the full

bunch length τ_{full} for different RF configurations. Measurements were performed with a fixed second-to-fundamental voltage ratio $r_V = 0.53$, chosen to avoid beam loss during the acceleration ramp. The shortest bunches observed before the onset of instability define the threshold bunch length for each configuration.

The measurements were carried out at $N_p = 1.1 \times 10^{12}$ p/b for BSM and $N_p = 3.9 \times 10^{12}$ p/b for BLM. The corresponding parameters and inferred frequency spreads are summarised in Table 4.1. Overall, BLM exhibits a substantially larger synchrotron frequency spread than BSM, both with cavity loops open and closed. This enhanced spread is expected to provide stronger Landau damping, consistent with the more robust longitudinal stability observed at high intensity.

Table 4.1: Small-amplitude synchrotron frequency f_{s0} , synchrotron-frequency spread Δf_s , bunch intensity N_p , and full bunch length τ_{full} for the four configurations at flat top energy (2 GeV).

Config.	f_{s0} [Hz]	Δf_s [Hz]	$N_p [\times 10^{12}]$	τ_{full} [ns]
BSM CL	462	59	1.1	160
BSM OL	462	59	1.1	161
BLM CL	110	222	3.9	198
BLM OL	110	240	3.9	238

Complementary stability measurements were performed in a dedicated 2 GeV cycle comparing BLM and BSM. Longitudinal profiles were acquired 225 ms after injection; afterwards, the voltage ratio was held constant at $r_V = 0.53$ to allow an intensity scan with a fixed bucket shape. The cavity loops were alternately enabled and disabled.

Figure 4.12 summarises the centroid motion at flat top for three cases: BSM with cavity loops closed (CL) and open (OL), and BLM with loops closed. The corresponding evolution of bunch length oscillations is shown in Figs. 4.13. The different intensity ranges indicate the maximum intensity reached before coherent oscillations developed; shading denotes the standard deviation around the mean.

Among the tested configurations, the operational case BLM CL is the most stable at high intensity. In this case, the maximum intensity is ultimately limited by transverse effects, driven by the impedance of the unterminated extraction kicker [49]. In contrast, BSM CL exhibits pronounced quadrupole bunch length oscillations already at low intensity while the centroid remains stable, consistent with loop-driven parametric excitation near $2f_{s0}$ combined with reduced Landau damping in BSM. When the cav-

ity loops are opened (BSM OL), the quadrupole oscillations are mitigated; however, the intensity reach becomes more limited and the observations are compatible with uncompensated beam loading and potential-well distortion, visible as an increase of the dipole oscillation amplitude with intensity and a residual breathing of the bunch length.

Overall, the behaviour observed in single-RF operation is recovered in DRF conditions: configurations producing shorter bunches and reduced frequency spread are more vulnerable to coherent motion driven by the cavity loops. In particular, BSM is even more sensitive than SRF to cavity-loop-induced excitation, whereas BLM benefits from stronger Landau damping and provides the most robust longitudinal stability over the measured intensity range.




Figure 4.12: Measured bunch-oscillation amplitude versus intensity at flat top energy (2 GeV). A strong coherent dipole oscillation is observed in BSM OL, whereas it is suppressed in BSM CL by the cavity feedback loops.

4.5 Summary

This chapter presented a systematic experimental characterisation of longitudinal single-bunch instabilities in the PSB after the RF system upgrade.




Figure 4.13: Bunch length versus intensity at flat top energy (2 GeV). Comparing BLM and BSM in closed-loop highlights the higher sensitivity of BSM to oscillations, whereas BLM remains more robust.

First, the Run 3 measurements were contrasted with the pre-LS2/LS2 stability predictions. While simulations indicated that single-harmonic operation with multi-harmonic wake-cancellation notches should enable high-intensity beams, the first commissioning measurements showed the opposite trend: pure SRF operation proved detrimental to stability, whereas DRF in bunch-lengthening mode restored robust behaviour.

For SRF operation, instability thresholds were measured for two RF program families: constant voltage (CV) and constant bucket area (CBA), both in open-loop and closed-loop LLRF conditions. In open-loop, both CV and CBA cases exhibited sharp intensity thresholds followed by rapid emittance blow-up and beam loss. The CV data showed an empirical scaling close to $N_{p,th} \propto \sqrt{V_{pk}}$, consistent with the stabilising effect of increased longitudinal focusing and intrinsic momentum spread at higher RF voltage. In closed-loop, the abrupt blow-up was suppressed, confirming that the fast open-loop instability is driven by the uncompensated Finemet cavity beam loading. However, a gradual, approximately monotonic bunch length increase with N_p is observed, indicating a progressive longitudinal emittance dilution governed by closed-loop effects.

In the CBA study, the role of the filling factor was shown to depend criti-

cally on the loop state. In open-loop, a smaller filling factor (corresponding to a higher RF voltage for fixed longitudinal emittance) increased the instability threshold, in line with the CV trend. In closed-loop, the sensitivity inverted: the $FF = 20\%$ configuration displayed the strongest loop-driven modulations and the most pronounced emittance growth, whereas higher filling factors were comparatively more tolerant.

Motivated by this sensitivity, a dedicated investigation was carried out at $FF = 20\%$ in closed-loop by selectively opening individual revolution-harmonic servoloops. Waterfall comparisons showed that the perturbations are strongly correlated with the control action around the fourth and fifth revolution harmonics: opening $h = 4$ and $h = 5$ while keeping the other loops closed strongly reduced the structured profile modulations. Tomographic reconstructions at two close intensities revealed a clear four-lobed phase-space pattern, indicating a dominant octupolar ($m = 4$) component, and therefore motivating the use of higher-order observables in addition to bunch length diagnostics. A time-frequency analysis based on spectrograms of the FWHM time series provided a compact representation of persistent coherent components and their dependence on the loop configuration.

Finally, in DRF operation the stability performance was compared between bunch lengthening and bunch shortening modes. Measurements of synchrotron frequency spread and intensity scans confirmed that bunch lengthening provides a larger incoherent tune spread and therefore stronger Landau damping, providing a more robust longitudinal stability margin at high intensity. On the other hand, bunch shortening reduces the frequency spread and increases the sensitivity to loop-driven coherent motion, leading to an earlier degradation of longitudinal beam quality.

Chapter 5

Modelling of the PSB LLRF cavity feedback loops

In Chapter 3, fundamental notions of LLRF engineering in synchrotrons and a comprehensive description of the PSB multi-harmonic LLRF beam-loading compensation system are presented (Section 3.2).

The beam-induced voltage associated with the broadband, large longitudinal coupling impedance of the Finemet cavities is a strong limiting factor for PSB high-intensity operation, driving coherent longitudinal motion and triggering single-bunch instabilities. Together with the direct RF feedback acting continuously over the operational band, the MHFB system operates at integer multiples of the revolution frequency, suppressing beam-induced components and successfully mitigating fast blow-ups.

The MHFB cavity feedback loops also have an unexpected and undesirable impact on the longitudinal beam dynamics, degrading beam quality. As a consequence, it is not sufficient to operate in SRF: it becomes compulsory to operate in DRF in BLM throughout the entire PSB cycle to preserve smooth operation and high-quality beam parameters.

The investigation of these effects can be enabled by the use of a numerical model that captures not only the cavity impedance in open loop, but also the detailed behaviour of the LLRF feedback chain in closed loop, including the digital implementation features that shape the loop response [50]. This approach is useful to achieve an optimised operation of the servoloops, tailored to the specifications of a cycle.

The PSB Finemet cavity already represents a peculiar application of RF engineering in synchrotrons, and the same can be stated for the multi-harmonic LLRF system.

The objective of this chapter is to present the numerical model developed in Python, considering the following points:

- capture the multi-harmonic nature of the RF voltage synthesis and regulation, including the fixed high-rate clock domain used for digitisation and mixing and the lower rate used for digital signal processing;
- implement the multi-harmonic demodulation and modulation structure used in the MHFB, including a numerically controlled oscillator (NCO) based on a phase accumulator with dynamic tuning-word updates derived from the revolution-frequency programme;
- consider the narrow-band selectivity and decimation behaviour of the cascaded integrator-comb (CIC) filter chain, including its dynamic configuration as a function of revolution frequency;
- provide time-domain and frequency-domain models of the closed-loop PSB RF system suitable for stability and performance studies;
- include the dominant non-idealities that shape the closed-loop response (delays, bandwidth limitations, saturation, quantisation where relevant);
- offer a framework for systematic comparison against measurements (open-loop response, step responses, beam-loading transients).

The numerical model is intended to be predictive in the frequency range relevant for longitudinal dynamics and LLRF regulation, while being computationally efficient to support parameter scans.

For this reason, the model is organised as a set of per-harmonic processing chains (one for each controlled revolution harmonic), coupled only through common timing parameters and the final summation of the harmonic drive components. This representation matches the physical structure of the MHFB and facilitates direct comparison with the transfer-function measurements performed on the dedicated test stand, as well as with intermediate signals acquired along the chain.

5.1 Modelling strategy and development

The PSB MHFB is a multi-rate digital control system: the cavity voltage is sampled in a fixed high-rate clock domain, while filtering, decimation and regulation are performed at harmonic-dependent lower rates. The numerical model mirrors this architecture by implementing the signal chain stage

by stage, preserving a one-to-one correspondence with the architecture blocks and providing relevant intermediate observables.

Figure 5.1 illustrates a simplified block diagram describing the flow followed for the development of the model. In the diagram the parts referring to a single signal and to multiple ones (parts in common for all the 16 loops) are presented differently and three categories of signals are also highlighted. In pink the reference frequency values taken as inputs in the model. The RF signals are presented in navy blue, while the baseband ones in teal. The main input is the cavity voltage signal, which can be generated in an artificial way, could be an effective gap voltage measured and acquired by one of the scopes in the control room or a signal generated in BLonD [51]. It contains the programmed voltage, with content allocated at different harmonics of the RF frequency, and the beam-induced voltage. The signal comes already defined in discrete-time domain and it is processed so that it is actually sampled at the clock frequency and available for the down-conversion. Up to now the modelling belong to the input generation stage. To extract the complex envelope of each harmonic, the signal is down-converted to baseband by an in-phase/quadrature (I/Q) demodulation. This part of the model belong to a common class called MHLO. For each harmonic h , the gap signal is multiplied by a pair of coherent reference waveforms $\cos(\phi_h[n])$ and $\sin(\phi_h[n])$, producing the raw baseband components $I_{\text{raw},h}[n]$ and $Q_{\text{raw},h}[n]$. As visible in the figure, at this point 16 individual signals have to be processed in parallel. The reference phases $\phi_h[n]$ are generated by a MHLO implemented as a numerically controlled oscillator (NCO) clocked at f_{clk} , whose tuning words are derived from the (time-varying) revolution frequency so that the demodulation remains phase-coherent throughout the cycle. In this way, the model reproduces the same digital local oscillator principle used in the real LLRF system, ensuring that each harmonic is mixed with a reference locked to hf_{rev} .

The raw I/Q streams still contain high-frequency mixing products and require anti-alias filtering prior to decimation. This is achieved using a two-stage CIC low-pass filtering and decimation chain. The first stage (CIC1) performs an initial low-pass response and reduces the sample rate by a harmonic-dependent factor, providing efficient rejection of the high-frequency content while preserving the slowly varying complex envelope. The second stage (CIC2) completes the decimation to the final baseband sampling rate used by the control and diagnostic algorithms. Because CIC filters are efficient but exhibit passband droop and transient behaviour, the model explicitly accounts for stage-dependent sampling rates, decimator factors, and the initial transient region that must be discarded before using

the signals for control purposes. The final outputs ($I_{\text{filt},h}$, $Q_{\text{filt},h}$) represent the filtered and decimated complex baseband quantities that can be compared across operating points (injection, ramp, extraction) and used as inputs to the subsequent feedback controller and performance diagnostics.




Figure 5.1: Signal-processing chain implemented in the numerical PSB multi-harmonic LLRF model. The clock-domain cavity gap voltage is down-converted to baseband by I/Q mixing with MHLO references generated by an NCO locked to hf_{rev} . The resulting raw baseband components are filtered and decimated by two CIC stages (CIC1–CIC2), yielding controller-domain I/Q signals. A discrete-time I/Q controller generates correction signals that are expanded back to the clock domain, remodulated and summed over harmonics to obtain the RF drive contribution.

Two complementary implementations are used:

- a time-domain implementation, which is natural for generating a realistic input voltage (programmed RF plus beam-induced voltage), for modelling transients, and for reproducing the internal I/Q signals along the digital chain;
- a frequency-domain implementation, where the same chain is expressed as a cascade of transfer functions, which is convenient to

describe harmonic selectivity, dynamic CIC behaviour, and closed-loop frequency responses.

This section focuses on the time-domain signal-flow description. The frequency-domain formulation is also introduced afterwards using the same block decomposition, ensuring consistency between the two views.

5.1.1 Model overview and time-base separation

A single simulation grid is defined at the LLRF clock frequency f_{clk} , yielding a clock-domain time step $\Delta t = 1/f_{\text{clk}}$. The simulated duration is set to a user-defined number of turns N_{turn} , so that

$$T_{\text{sim}} = N_{\text{turn}} T_{\text{rev}}, \quad T_{\text{rev}} = \frac{1}{f_{\text{rev}}}. \quad (5.1)$$

Two representative cycle points are considered: injection and extraction, corresponding to the minimum and maximum revolution frequency in the PSB 2 GeV cycle.

For each controlled harmonic $h = 1, \dots, 16$ the model builds a dedicated processing chain that operates on: (i) a clock-domain representation (mixing, remodulation, and harmonic summation) and (ii) a controller-domain representation after the cascaded decimation stages (CIC2 output), where discrete-time I/Q regulation is applied. The 16 chains are coupled only through the common input $V_{\text{gap}}(t)$ and through the final summation of the reconstructed harmonic drive components.

5.1.2 Input signal construction

The first step in using the numerical model is the selection of the PSB cycle point (or time interval) to be simulated. This choice determines the corresponding revolution frequency values to be applied, $f_{\text{rev}}(t)$, and the relevant portion of the programmed RF voltage waveform (i.e. the voltage reference provided to the RF system during the selected cycle segment).

The signal to be processed is generated in time domain as a sampled waveform $V_{\text{gap}}[n]$, where the index n denotes the discrete-time samples. The sampling frequency is fixed and set by the LLRF clock, $f_{\text{clk}} \approx 122.7 \text{ MHz}$, so that $V_{\text{gap}}[n] \equiv V_{\text{gap}}(n/f_{\text{clk}})$.

To reproduce realistic features of the PSB Finemet cavity voltage, the input signal includes a beam-induced contribution computed from the

longitudinal impedance and the beam current. In particular, the induced voltage is obtained in the frequency domain as

$$V_{\text{ind}}(f) = -Z_{\parallel}(f) I_b(f), \quad (5.2)$$

where $Z_{\parallel}(f)$ is the complex longitudinal impedance of the Finemet cavity system (including the effect of the direct RF feedback in the adopted impedance model) and $I_b(f)$ is the spectrum of the beam current associated with a single bunch of given longitudinal distribution, bunch length τ_L , and intensity N_p . The total gap voltage used as input to the subsequent signal-processing stages is then given by

$$V_{\text{gap}}[n] = V_{\text{prog}}[n] + V_{\text{ind}}[n], \quad (5.3)$$

with $V_{\text{prog}}(t)$ the desired, programmed RF voltage.

As an initial reference scenario, the SRF configuration with a parabolic bunch is assumed, with full bunch length τ_{full} , repeated once per turn. The bunch current $I_b(t)$ is normalised so that the charge per turn equals $Q_b = N_p e$. A one-turn view of the clock-domain bunch current is shown in Fig. 5.2.




Figure 5.2: Clock-domain bunch current $I_b(t)$ over one turn at injection for the parabolic profile used in the model. The waveform is repeated every T_{rev} over N_{turn} turns.

The beam-induced voltage is computed using the standard linear impedance relation

$$V_{\text{ind}}(\omega) = -Z_{\parallel}(\omega) I_b(\omega), \quad (5.4)$$

where $Z_{\parallel}$ is the complex longitudinal coupling impedance of the Finemet cavities. An inverse-FFT is applied to retrieve the time-domain beam-induced voltage. The complex impedance is considered for the calculation, so that both the resistive and reactive part are taken into account.

For the reference test case used to visualise the processing chain, the programmed voltage is taken as a single tone at the fundamental ($h = 1$) with peak amplitude $V_{\text{prog,pk}} = 10$ kV and synchronous phase ϕ_s (below transition). This choice provides a clean RF reference while isolating the broadband beam-induced contribution.

Figure 5.3 shows, over one turn, the three signals entering the digital chain: V_{prog} , V_{ind} , and their sum $V_{\text{gap,open}}$. This view is used as the reference for interpreting all subsequent intermediate signals.




Figure 5.3: One-turn decomposition of the model input at injection: programmed RF V_{prog} , beam-induced voltage V_{ind} computed from the Finemet impedance, and the resulting open-loop gap voltage $V_{\text{gap,open}} = V_{\text{prog}} + V_{\text{ind}}$, which enters the MHFB demodulation chain.

5.1.3 Multi-harmonic local oscillator (MHLO)

In the PSB Finemet system the cavity gap voltage is not a single narrow-band RF tone, but the superposition of several revolution-frequency harmonics. The multi-harmonic LLRF therefore requires a set of mutually phase-coherent local oscillators at the harmonic frequencies $hf_{\text{rev}}(t)$, with

$h = 1, \dots, 16$. A central MHLO provides these references and guarantees that the same phase model is used consistently for both digital down-conversion (demodulation) and up-conversion (modulation) in all harmonic channels.

The MHFB requires phase-coherent digital local oscillators at $h f_{\text{rev}}$ ($h = 1, \dots, 16$). In hardware, the references are produced by a numerically controlled oscillator (NCO) based on a $N_{\text{bit}} = 32$ phase accumulator driven by a frequency tuning word, FTW. In the model, each harmonic is generated with

$$f_h = h f_{\text{rev}}, \quad (5.5)$$

and the accumulator evolves as

$$\Phi[n+1] = (\Phi[n] + \text{FTW}[n]) \bmod 2^{N_{\text{bit}}}, \quad \phi_h[n] = 2\pi \frac{\Phi[n]}{2^{N_{\text{bit}}}}, \quad (5.6)$$

where $\text{FTW}[n] \propto f_h / f_{\text{clk}}$. The quadrature reference signals are then obtained as $\cos \phi_h[n]$ and $\sin \phi_h[n]$.

Figure 5.4 compares the wrapped phase evolution of the NCO phase accumulator for selected harmonics ($h = 1, 2, 8$) at injection and extraction. The different phase slopes reflect the harmonic scaling of the oscillator frequency, while the comparison between the two time points highlights the change in revolution frequency across the cycle.

For each harmonic h , the clock-domain gap voltage is mixed down to baseband using the MHLO quadrature references. Using the sign convention adopted consistently throughout the model, the mixer outputs are

$$I_h(t) = -2 V_{\text{gap}}(t) \sin \phi_h(t), \quad Q_h(t) = 2 V_{\text{gap}}(t) \cos \phi_h(t), \quad (5.7)$$

where $\phi_h(t)$ is the NCO phase for harmonic h . The mixer outputs contain the desired baseband component and high-frequency mixing products, which must be filtered before decimation.

5.1.4 CIC filtering and harmonic-dependent decimation

Anti-alias filtering and decimation are modelled with two cascaded single-stage CIC filters (CIC1 and CIC2). For each harmonic, the overall decimation factor R is

$$R_{\text{tot},h} = R_{\text{CIC1},h} R_{\text{CIC2},h}, \quad f_{\text{ctrl},h} = \frac{f_{\text{clk}}}{R_{\text{tot},h}}. \quad (5.8)$$




Figure 5.4: Wrapped phase evolution of the NCO phase accumulator for selected harmonics ($h = 1, 2, 8$) at injection and extraction. In both cases, the phase slope increases with harmonic number, consistent with $f_h = h f_{\text{rev}}$. The steeper evolution at extraction reflects the higher revolution frequency later in the cycle.

In the firmware configuration, CIC1 uses a differential delay $D_{\text{CIC1}} = 32$ and a decimation factor tracking f_{rev} ,

$$R_{\text{CIC1}} = \frac{f_{\text{clk}}}{8f_{\text{rev}}}, \quad (5.9)$$

while CIC2 uses $D_{\text{CIC2}} = 1$ and a power-of-two decimation determined by the platform configuration rules. Since R_{CIC1} must be an integer, rounding introduces small configuration-dependent deviations, which are reproduced in the model and are relevant when discussing harmonic rejection close to configuration transitions. A representative time-domain example of the I/Q signals through the demodulation chain (mixer output, after CIC1, after CIC2) is shown in Fig. 5.5.




Figure 5.5: Example of I/Q signals through the demodulation chain for $h = 1$ at injection: mixer output (clock domain), after CIC1, and after CIC2 (controller domain). The CIC stages progressively suppress out-of-band components and decimate to the controller sampling rate.

The resulting controller-domain sampling rate depends on the harmonic number through the decimation factors selected in the digital down-conversion chain. Figure 5.6 compares the controller-domain sampling rate $f_{\text{ctrl},h}$ as a function of harmonic number at injection and extraction. For completeness, the corresponding sets of digital down-conversion parameters $\{R_1, D_1, R_2, D_2\}$ used for each harmonic are reported in Tables 5.1 and 5.2.

Table 5.1: Digital down-conversion parameters at injection for each harmonic channel.

h	f_h [Hz]	R_1	D_1	R_2	D_2	R_{tot}	$f_{\text{ctrl},h}$ [Hz]
1	9.94×10^5	15	32	64	1	960	1.28×10^5
2	1.99×10^6	7	32	128	1	896	1.37×10^5
3	2.98×10^6	5	32	128	1	640	1.92×10^5
4	3.98×10^6	3	32	256	1	768	1.60×10^5
5	4.97×10^6	3	32	256	1	768	1.60×10^5
6	5.96×10^6	2	32	256	1	512	2.40×10^5
7	6.96×10^6	2	32	256	1	512	2.40×10^5
8	7.95×10^6	1	32	512	1	512	2.40×10^5
9	8.95×10^6	1	32	512	1	512	2.40×10^5
10	9.94×10^6	1	32	512	1	512	2.40×10^5
11	1.09×10^7	1	32	512	1	512	2.40×10^5
12	1.19×10^7	1	32	512	1	512	2.40×10^5
13	1.29×10^7	1	32	512	1	512	2.40×10^5
14	1.39×10^7	1	32	512	1	512	2.40×10^5
15	1.49×10^7	1	32	512	1	512	2.40×10^5
16	1.59×10^7	1	32	1024	1	1024	1.20×10^5

5.1.5 I/Q regulation, RF drive reconstruction, and multi-harmonic coupling

After the second CIC stage, each harmonic component is regulated in the I/Q plane by a discrete-time proportional-integral controller operating at the controller-domain sampling rate. For each harmonic channel h , the controller acts on the in-phase and quadrature errors

$$e_{I,h}[k] = I_{h,\text{sp}}[k] - I_h[k], \quad e_{Q,h}[k] = Q_{h,\text{sp}}[k] - Q_h[k], \quad (5.10)$$

where k denotes the discrete-time index at the output of the second CIC stage. The regulation is performed independently in the I and Q branches with proportional and integral gains K_P and K_I . In the numerical model, the controller outputs are updated according to

$$u_{I,h}[k] = u_{I,h}[k-1] + K_P(e_{I,h}[k] - e_{I,h}[k-1]) + K_I T_S e_{I,h}[k], \quad (5.11)$$

$$u_{Q,h}[k] = u_{Q,h}[k-1] + K_P(e_{Q,h}[k] - e_{Q,h}[k-1]) + K_I T_S e_{Q,h}[k], \quad (5.12)$$

where T_S is the controller-domain sampling period after decimation. This recursive form makes explicit the proportional action on the error variation

Table 5.2: Digital down-conversion parameters at extraction for each harmonic channel.

h	f_h [Hz]	R_1	D_1	R_2	D_2	R_{tot}	$f_{\text{ctrl},h}$ [Hz]
1	1.81×10^6	8	32	64	1	512	2.40×10^5
2	3.62×10^6	4	32	128	1	512	2.40×10^5
3	5.43×10^6	2	32	256	1	512	2.40×10^5
4	7.24×10^6	2	32	256	1	512	2.40×10^5
5	9.05×10^6	1	32	512	1	512	2.40×10^5
6	1.09×10^7	1	32	512	1	512	2.40×10^5
7	1.27×10^7	1	32	512	1	512	2.40×10^5
8	1.45×10^7	1	32	512	1	512	2.40×10^5
9	1.63×10^7	1	32	1024	1	1024	1.20×10^5
10	1.81×10^7	1	32	1024	1	1024	1.20×10^5
11	1.99×10^7	1	32	1024	1	1024	1.20×10^5
12	2.17×10^7	1	32	1024	1	1024	1.20×10^5
13	2.35×10^7	1	32	1024	1	1024	1.20×10^5
14	2.53×10^7	1	32	1024	1	1024	1.20×10^5
15	2.72×10^7	1	32	1024	1	1024	1.20×10^5
16	2.90×10^7	1	32	1024	1	1024	1.20×10^5

and the integral accumulation required to remove steady-state offsets. In addition, a fixed delay is included in the model to represent the cumulative latency associated with ADC/DAC conversion, FPGA pipeline stages, and cable propagation.


The controller outputs $(u_{I,h}, u_{Q,h})$ are then expanded back to the clock domain by a zero-order hold (ZOH) with factor $R_{\text{tot},h}$ and remodulated using the same harmonic phase reference generated by the MHLO. For each harmonic, the reconstructed drive signal is written as

$$V_{\text{drive},h}(t) = \frac{1}{2} \left(u_{Q,h}(t) \cos \phi_h(t) - u_{I,h}(t) \sin \phi_h(t) \right), \quad (5.13)$$


and the total drive waveform applied to the RF chain is obtained by summing over all controlled harmonics,

$$V_{\text{drive}}(t) = \sum_{h=1}^{16} V_{\text{drive},h}(t). \quad (5.14)$$

The model therefore reconstructs, in the time domain and at the DAC sampling rate, the multi-harmonic RF drive waveform delivered to the cavity power amplifiers.



(a) Injection



(b) Extraction

Figure 5.6: Controller-domain sampling rate $f_{\text{ctrl},h} = f_{\text{clk}}/(R_{\text{CIC1},h}R_{\text{CIC2},h})$ as a function of harmonic number at injection and extraction. The variation with h reflects the harmonic-dependent choice of down-conversion parameters.

Since each harmonic demodulator sees the total gap voltage, the 16

channels are not independent but are coupled through

$$V_{\text{gap}}(t) = V_{\text{prog}}(t) + V_{\text{ind}}(t) + V_{\text{drive}}(t). \quad (5.15)$$

5.1.6 Diagnostics: I/Q constellations and controller-domain observables

To verify internal consistency and to provide interpretable intermediate observables, the model outputs: (i) I/Q stage plots (mixer, CIC1, CIC2), (ii) time-coloured I/Q constellations at the same points, and (iii) controller-domain signals (errors and control outputs) for selected harmonics.

Figures 5.7 and 5.8 show a representative set of I/Q constellations for $h = 8$ at injection and extraction frequencies, while Figs. 5.9 and 5.10 shows the controller-domain observables for $h = 1$ and $h = 2$, at injection.




Figure 5.7: Time-coloured I/Q constellations for $h = 8$ at injection: after mixing (left), after CIC1 (centre), and after CIC2 (right). The colourbar represents time (or sample index) and highlights transient evolution and turn-periodic structures.




Figure 5.8: Time-coloured I/Q constellations for $h = 8$ at extraction: after mixing (left), after CIC1 (centre), and after CIC2 (right).




Figure 5.9: Controller-domain signals at CIC2 rate for $h = 1$ at injection: measured I/Q (and implied amplitude/phase), errors (e_I, e_Q), and controller outputs (u_I, u_Q).

5.1.7 Frequency-domain formulation

The time-domain implementation reproduces the multi-rate digital signal chain and provides access to the internal observables required for interpretation of the loop dynamics. In parallel, a frequency-domain representation is useful to analyse the harmonic selectivity of the demodulator, the passband droop introduced by the CIC filters, and the evolution of loop gain and phase with cycle point. For this reason, the same processing chain is also represented analytically as a cascade of transfer functions corresponding to the main functional blocks of the hardware.

Following the formulation adopted in the numerical model, the transfer function of each CIC stage is written as

$$H_{\text{CIC}}(f) = \left[\frac{\sin(\pi D f / f_{\text{out}})}{\sin(\pi f / (R f_{\text{out}}))} \right]^N e^{-j\phi(f)}, \quad (5.16)$$

where D is the differential delay, R the decimation factor, N the number of stages, f_{out} the output sampling frequency after decimation, and $\phi(f)$ the associated phase contribution.




Figure 5.10: Controller-domain signals at CIC2 rate for $h = 2$ at injection: measured I/Q (and implied amplitude/phase), errors (e_I, e_Q), and controller outputs (u_I, u_Q).

In the present implementation, both CIC filters are single-stage filters ($N = 1$). The first stage uses $D_{\text{CIC1}} = 32$ and a decimation factor chosen to track the revolution frequency according to

$$R_{\text{CIC1},h} \simeq \frac{f_{\text{clk}}}{8f_{\text{rev}}}, \quad (5.17)$$

with integer rounding applied in the digital implementation. The second stage uses $D_{\text{CIC2}} = 1$, while its decimation factor is selected so as to bring the controller-domain sampling rate close to the DSP rate,

$$R_{\text{CIC2},h} = 2^{\left\lfloor \log_2 \left(\frac{f_{\text{rev}} M}{2^{22}} \right) \right\rfloor + \left\lceil \log_2 8 \right\rceil}, \quad (5.18)$$

where $M = 2^{32}/f_{\text{clk}}$ is the NCO scaling factor. To make the CIC magnitude independent of frequency, each stage is multiplied by the corresponding gain-compensation factor

$$G_{\text{comp}} = R D. \quad (5.19)$$

The overall demodulator response for harmonic h is therefore written as

$$H_{\text{demod},h}(f) = G_{\text{comp},1,h} H_{\text{CIC1},h}(f) G_{\text{comp},2,h} H_{\text{CIC2},h}(f). \quad (5.20)$$

After demodulation, the sampled baseband signal is processed by the digital PI controller. Its discrete-time transfer function is

$$H_{\text{PI}}(f) = \frac{(K_P + K_I T_S) - K_P e^{-j2\pi f T_S}}{1 - e^{-j2\pi f T_S}}, \quad (5.21)$$

where K_P and K_I are the proportional and integral gains, respectively, and T_S is the controller-domain sampling period. The digital loop also includes a fixed delay term accounting for ADC/DAC conversion, FPGA pipeline latency, and cable propagation:

$$H_{\text{delay}}(f) = \left(\frac{1}{e^{j2\pi f}} \right)^\tau, \quad (5.22)$$

where $\tau = T_{\text{delay}} f_{\text{clk}}$ is the delay expressed in clock samples.

The open-loop response of the h -th harmonic channel is then assembled as

$$H_{\text{OL},h}(f) = H_{\text{demod},h}(f) H_{\text{PI}}(f) H_{\text{delay}}(f) G_S, \quad (5.23)$$

where G_S is a free parameter accounting for the overall loop gain. The corresponding closed-loop transfer function is

$$H_{\text{CL},h}(f) = \frac{H_{\text{OL},h}(f)}{1 + H_{\text{OL},h}(f)}. \quad (5.24)$$

This formulation provides a compact description of the harmonic feedback loop in terms of the same parameters used in the time-domain implementation.

Figure 5.11 illustrates the combined response of the first and second CIC filters, including gain compensation, for representative cases at injection and extraction. The plot highlights the evolution of the demodulator selectivity with cycle point and harmonic number, and provides the frequency-domain counterpart of the controller-domain behaviour discussed in the previous subsection.

5.2 Model validation and benchmarking

The validation of the numerical model is carried out in successive stages, starting from the internal digital signal processing signals and then moving towards the reproduction of measured closed-loop responses. As a first step, the time-domain implementation is benchmarked against intermediate observables of the LLRF chain. This approach makes it possible to verify that the model reproduces correctly the sequence of operations.




Figure 5.11: Combination of the first and second CIC filters present in the system demodulator, with the application of the gain compensation, shown for representative injection and extraction conditions.




Figure 5.12: Example oscilloscope acquisition for a given beam and voltage setting: gap-sum voltage V_{gap} (red), RF drive V_{drive} (blue), and phase pick-up signal (green), shown over a $\sim 100 \mu\text{s}$ time window.

From a dedicated oscilloscope application available in the CERN control room, the key analogue LLRF observables were recorded for a range of voltage set-points and beam conditions, providing a direct dataset for time-domain benchmarking of the numerical model. Figure 5.12 shows one of the acquisitions: the measured gap-sum voltage V_{gap} (red), the cor-

responding RF drive V_{drive} (blue), and the phase pick-up signal (green) over a $\sim 100 \mu\text{s}$ window. The traces illustrate the transient evolution following the excitation, together with the steady oscillatory component, and serve as a reference for comparing the simulated waveform dynamics (amplitude, phase evolution, and transient settling) under the same operating conditions.

The amplitudes of the first five harmonics for this configuration of bunch length and intensity are reported in Fig. 5.13. The measured amplitudes (blue) are compared with the corresponding model predictions (red) as a function of time over the acquisition window. In order to account for the triggering conditions of the oscilloscope acquisition, the measured traces are offset by $+25 \mu\text{s}$ with respect to the simulated data. For the dominant harmonics ($h = 1$ and $h = 2$), which carry the largest fraction of the beam-induced voltage, the agreement between measurement and simulation is satisfactory. The model reproduces both the transient growth and the peak amplitude, as well as the subsequent decay toward the steady-state regime. For harmonics $h = 3$ to $h = 5$, the absolute amplitudes are significantly smaller, yet the qualitative behaviour remains consistent between model and measurement. Small discrepancies become more visible relative to the reduced signal level and can be attributed to measurement noise, residual unmodelled delays, and simplifications in the cavity and loop representation.

Harmonics $h = 6$ to $h = 8$ are intentionally omitted from the comparison. For this particular beam configuration, the signal-to-noise ratio in the measured V_{gap} data is too low at those harmonics, resulting in noisy reconstructed amplitudes and correspondingly noisy model outputs, for which no meaningful quantitative comparison can be established.

The benchmarking of the RF drive is shown in the time domain in Fig. 5.14. The plot compares the measured multi-harmonic RF drive signal, acquired from the oscilloscope, with the corresponding output of the numerical model. Both signals are represented at the DAC sampling rate to ensure a consistent comparison. The model reproduces the overall amplitude build-up, the envelope evolution, and the modulation content of the RF drive with good fidelity. In particular, $V_{\text{drive,model}}$ can be observed to generally match the profile of V_{drive} :

- $V_{\text{drive,model}}$ begins from a zero value,
- rises at approximately the correct time,
- and increases with an approximately linear slope towards V_{setpoint} .




Figure 5.13: Comparison between measured and simulated harmonic amplitudes for the first five revolution harmonics, for $N = 140 \times 10^{10}$ [p/b]. Left column: $\tau_l = 350$ ns. Right column: $\tau_l = 655$ ns. Blue: measurement. Red: model.

The progressive growth of the envelope and the superimposed multi-harmonic oscillatory structure are accurately captured, demonstrating that the demodulation, digital regulation, and remodulation stages are correctly




Figure 5.14: Time-domain comparison of the summed multi-harmonic RF drive at the DAC sampling rate. The measured signal (blue) acquired from the oscilloscope is overlaid with the model prediction (red). The model reproduces the overall amplitude build-up, envelope evolution, and modulation content of the RF drive, with small discrepancies attributable to measurement offsets and analogue chain effects.

implemented in the numerical chain.

However, differences can also be identified. In particular, a small DC offset appears to be present across the measured signal, which is not reproduced by the model. This offset may originate from analogue chain asymmetries, acquisition bias, or calibration effects in the oscilloscope or DAC path. Additionally, minor discrepancies in the envelope thickness and high-frequency ripple content are visible, which can be attributed to measurement noise, residual phase misalignment, and simplified modelling of the analogue front-end and amplifier dynamics.

Overall, the comparison confirms that the time-domain implementation provides a realistic reconstruction of the summed multi-harmonic RF drive delivered to the cavity, while highlighting second-order effects that are external to the idealised digital model.

The frequency-domain benchmarking of the model is presented in Fig. 5.15, which shows the closed-loop magnitude response around the first revolution harmonic at injection. The simulated response is compared with the corresponding VNA measurement performed on the test-stand of




Figure 5.15: Closed-loop magnitude response around the first revolution harmonic at injection. The simulated response (blue) is compared with the VNA measurement (red).

the LLRF system.

The simulation reproduces both the depth and the bandwidth of the suppression with good accuracy. In particular, the position of the notch is correctly centered at the nominal revolution frequency, confirming the correct implementation of the demodulation frequency tracking and harmonic alignment in the model.

The agreement in notch depth demonstrates that the loop gain and phase characteristics, including the CIC filtering stages, PI controller, and delay modelling, are realistically represented in the frequency-domain formulation. Small deviations between measurement and simulation are visible in the skirts of the response and in the asymmetry of the side lobes. These differences can be attributed to unmodelled analogue chain effects, residual calibration uncertainties, and frequency-dependent behaviour of the amplifier and cavity system that are approximated in the simplified transfer-function representation.

This comparison validates the analytical transfer-function implementation of the multi-harmonic feedback loop and confirms its suitability for stability analysis and loop-parameter optimisation studies.

5.3 Summary

The development of the new time-domain dynamic feedback loop model represents a significant step towards a quantitative understanding of how the multi-harmonic cavity servoloops influence the longitudinal performance of the PSB. By reproducing the internal digital signal processing chain of the LLRF system — from phase-accumulator-based demodulation to CIC filtering, harmonic-dependent decimation, PI regulation, and remodulation at the DAC rate — the model provides a firmware-faithful representation of the operational system.

Benchmarking of the intermediate I/Q quantities constitutes a first level of validation. The comparison between simulated and measured baseband harmonic amplitudes demonstrates that the model accurately reproduces the demodulation and filtering behaviour of the PSB cavity servoloops. In particular, the timing of the harmonic build-up, the relative amplitudes of the first revolution harmonics, and their transient evolution are consistently captured. This confirms the correctness of the signal-processing chain implementation, including the multi-harmonic local oscillator, harmonic alignment, CIC configuration rules, and sampling-rate transitions. These results validate the first half of the model, namely the digital down-conversion and baseband representation of the cavity voltage.

A second level of validation is obtained through benchmarking of the reconstructed RF drive signal. The summed multi-harmonic drive generated by the model broadly matches the measured LLRF RF drive in both envelope evolution and modulation structure. The model correctly reproduces the onset of the drive, the linear growth towards the setpoint, and the overall amplitude profile. Residual discrepancies — such as small DC offsets and minor envelope deviations — can be attributed to analogue-chain asymmetries, acquisition offsets, and simplified modelling of amplifier and cavity dynamics. These effects lie outside the idealised digital representation and do not compromise the validity of the core signal-processing implementation. Further refinement of the per-harmonic PI gains and delay parameters is expected to improve the quantitative agreement.

The frequency-domain formulation provides an additional and complementary validation. The simulated closed-loop transfer functions reproduce the measured rejection notch at the revolution harmonics with good accuracy in both depth and bandwidth. This confirms that the analytical representation of the CIC filters, PI controller, and loop delay correctly captures the dynamic harmonic selectivity of the system. The agreement between simulation and network-analyser measurements strengthens confidence in the loop-gain modelling and establishes the model as a reliable

tool for stability and parameter studies.

Beyond benchmarking, the model has proven instrumental in clarifying the internal behaviour of the feedback loops. It provides direct access to quantities that are not easily observable in operation, such as per-harmonic baseband envelopes, controller-domain sampling rates, and intermediate filtered signals.

Finally, the architecture of the model — particularly its time-domain implementation — makes it naturally compatible with macro-particle tracking environments. The model is therefore well positioned to be integrated into particle tracking simulations such as BLoND [51], enabling self-consistent studies of beam–cavity–feedback interaction. This integration will allow investigation of intensity-dependent effects, longitudinal stability limits, and optimisation of multi-harmonic feedback settings under realistic beam conditions.

Chapter 6

Quadrupole synchrotron frequency shifts as a probe of the reactive longitudinal impedance

The longitudinal interaction between a bunched beam and its environment is described through the longitudinal coupling impedance $Z_{\parallel}(\omega)$, which relates the beam current spectrum to the induced voltage in frequency domain via $\tilde{V}_{\text{ind}}(\omega) = -Z_{\parallel}(\omega)\tilde{I}(\omega)$ [52, 53]. The impact on beam dynamics can be studied by separating $Z_{\parallel}$ into its real (resistive) and imaginary (reactive) parts.

In time domain, the beam current of a single bunch can be written in terms of the line density $\lambda(\tau)$, defined such that $\lambda(\tau)d\tau$ is the fraction of particles contained in the interval $[\tau, \tau + d\tau]$ and normalised to unity,

$$\int_{-\infty}^{+\infty} \lambda(\tau) d\tau = 1. \quad (6.1)$$

For a bunch of intensity N_p and particle charge q , the corresponding current in the bunch frame is

$$I(\tau) = q N_p f_{\text{rev}} \lambda(\tau), \quad (6.2)$$

where f_{rev} is the revolution frequency [52, 53]. In this convention, λ has units of 1/s and $I(\tau)$ is expressed in Amperes.

The real part of the longitudinal impedance represents the dissipative interaction between beam and environment. It corresponds to an induced voltage component in phase with the beam current, enabling a net transfer of energy and therefore power loss in surrounding structures. As a consequence, $\text{Re}\{Z_{\parallel}\}$ produces an average energy loss per turn that must be

compensated by the RF system, leading to a shift of the synchronous phase and a modification of the effective accelerating voltage [52, 53]. In the broadband resistive regime, and to first order in a local expansion around the bunch, the induced voltage tends to follow the bunch line density (and hence the current profile), i.e. $V_{\text{ind}}(\tau) \propto \lambda(\tau)$, which is particularly relevant for the PSB Finemet cavities operating with $h = 1$ [54].

The reactive part of the longitudinal impedance corresponds to energy being stored and returned rather than dissipated. Its main effect is to modify the effective longitudinal focusing seen by the bunch, thereby distorting the stationary RF potential, changing the equilibrium bunch shape and length, and shifting the synchrotron frequency. In the PSB, which operates below transition, the longitudinal space charge contribution is capacitive and, hence, defocusing; in the commonly used approximation where $Z_{\text{sc}}(\omega) \propto -i(\omega/\omega_0)$, the associated induced voltage is proportional to the derivative of the line density, $V_{\text{ind}}(\tau) \propto d\lambda/d\tau$ [52, 54]. This highlights that reactive effects are sensitive to the shape of the distribution rather than to its average energy loss.

The reactive impedance can therefore be probed from quadrupole synchrotron frequency shifts. This approach has been used, for example, to monitor the evolution of the reactive impedance in the CERN SPS [55]. In this chapter, the PSB reactive impedance is studied with the same method by analysing bunch length oscillations at flat bottom energy; the comparison with beam dynamics simulations is used in addition to validate the impedance model adopted for the PSB.

6.1 Principle of the method

The synchrotron frequency depends on the voltage detected by the beam, which includes the voltage induced by the impedance sources. Let V_{RF} denote the applied RF voltage and V_{ind} the induced voltage. For bunches undergoing coherent oscillations, V_{ind} comprises a contribution related to the stationary part of the bunch distribution and one related to the mismatched part. The first contribution is the incoherent shift, while the second is the coherent shift, associated with the spectral line generated by coherent oscillations. In general, for oscillations of azimuthal mode m (with $m = 1$ dipole and $m = 2$ quadrupole), the oscillation frequency can be written in the following form [56]

$$f_{s,m}(N_p) \simeq m f_{s0} + m \Delta f_{\text{inc}}(N_p) + \Delta f_{\text{coh},m}(N_p), \quad (6.3)$$

where N_p is the bunch intensity, f_{s0} is the small-amplitude synchrotron frequency, and Δf_{inc} and $\Delta f_{\text{coh},m}$ denote the incoherent and coherent shifts, respectively.

For common bunch shapes, such as the parabolic distribution, dipole oscillations exhibit a cancellation between the incoherent and the coherent contributions, which reduces the sensitivity to the impedance. Quadrupole oscillations do not generally present this cancellation and thus provide a robust probe of the reactive impedance.

6.1.1 Quadrupole oscillations and large-amplitude effects

Quadrupole oscillations correspond to a periodic exchange between bunch length and momentum spread. They arise when a geometrical mismatch between the bunch distribution and the RF bucket shape is present, which is the case at PSB injection. They can also be excited voluntarily by applying a controlled mismatch in the RF voltage, so that the distribution filaments in phase space while the bunch length exhibits oscillations.

Figure 6.1 shows the bunch motion during quadrupole oscillations in the longitudinal phase space $(\tau, \Delta E)$ for a synchrotron operating below transition. In the distribution two parts can be distinguished: the black core corresponds to the stationary (matched) component, while the violet component represents the quadrupolar perturbation (mismatch). Only particles belonging to the mismatched portion contribute to the observed coherent bunch length oscillation. The quadrupole frequency is approximately twice the synchrotron frequency and, for a stable bunch, is governed by particles at relatively large oscillation amplitudes, for which RF bucket non-linearities are relevant. In this study, the synchrotron frequency of particles with large oscillation amplitude is therefore considered.

6.1.2 Synchrotron frequency for particles with large oscillation amplitude

For a particle oscillating with a maximum amplitude τ_{max} , the non-linear synchrotron frequency is [41]

$$f_s^{(0)}(\tau_{\text{max}}) = f_{s0} \frac{\pi}{2K\left[\sin\left(\frac{\omega_{\text{RF}}\tau_{\text{max}}}{2}\right)\right]} \simeq f_{s0} \left(1 - \frac{(\omega_{\text{RF}}\tau_{\text{max}})^2}{16}\right), \quad (6.4)$$

where $K(\cdot)$ is the complete elliptic integral of the first kind, ω_{RF} is the RF angular frequency and f_{s0} is the small-amplitude synchrotron frequency. The amplitude dependence is quadratic and, by applying the substitution




Figure 6.1: Schematic representation of quadrupole (bunch length) oscillations in longitudinal phase space $(\tau, \Delta E)$. The mismatched portion of the distribution, relative to the part contributing to quadrupole motion, is shown in violet, while the matched area is shown in black.

$\tau_{max} \approx \tau_b/2$ where τ_b is the full bunch length, the quadrupole synchrotron frequency without intensity effects can be approximated as

$$f_{s2}(\tau_b) \approx 2f_{s0} \left(1 - \frac{(\omega_{RF}\tau_b)^2}{64} \right). \quad (6.5)$$

6.1.3 Incoherent synchrotron frequency shift

This section expands the derivation of the incoherent synchrotron-frequency shift Δf_{inc} arising from the induced voltage generated by the stationary bunch distribution. The key point is that the stationary induced voltage modifies the local restoring slope of the effective voltage around the synchronous particle, thereby shifting the single-particle synchrotron-frequency distribution (incoherent effect).

Induced voltage from the stationary bunch spectrum

Let $\lambda(\tau)$ be the stationary line density (normalised to unity),

$$\int_{-\infty}^{+\infty} \lambda(\tau) d\tau = 1, \quad (6.6)$$

and let the corresponding bunch current in the bunch frame be written as

$$I(\tau) = qN_p f_{\text{rev}} \lambda(\tau), \quad (6.7)$$

where q is the particle charge and f_{rev} is the revolution frequency. Introducing the Fourier transform (stationary bunch spectrum)

$$S(f) = \int_{-\infty}^{+\infty} \lambda(\tau) e^{-j2\pi f\tau} d\tau, \quad (6.8)$$

the induced voltage is expressed in the frequency domain as

$$\tilde{V}_{\text{ind}}(f) = -Z_{\parallel}(f) \tilde{I}(f) = -qN_p f_{\text{rev}} Z_{\parallel}(f) S(f), \quad (6.9)$$

and therefore in the time domain as

$$V_{\text{ind}}(\tau) = -qN_p f_{\text{rev}} \int_{-\infty}^{+\infty} S(f) Z_{\parallel}(f) e^{j2\pi f\tau} df. \quad (6.10)$$

Small-amplitude expansion and effective impedances

For particles oscillating with small amplitude around the bunch centre, the synchrotron frequency is determined by the linear component of the restoring force. We therefore expand the exponential term in Eq. (6.10) for small τ ,

$$e^{j2\pi f\tau} \simeq 1 + j2\pi f\tau + \mathcal{O}(\tau^2), \quad (6.11)$$

and keep constant and linear terms:

$$V_{\text{ind}}(\tau) \simeq -qN_p f_{\text{rev}} \left[\int_{-\infty}^{+\infty} S(f) Z_{\parallel}(f) df + \tau \int_{-\infty}^{+\infty} (j2\pi f) S(f) Z_{\parallel}(f) df \right]. \quad (6.12)$$

Separating the resistive and reactive contributions, it is convenient to define two effective impedances sampled by the stationary bunch spectrum:

$$Z_0 \equiv \int_{-\infty}^{+\infty} S(f) \text{Re}\{Z_{\parallel}(f)\} df, \quad (6.13)$$

$$Z_1 \equiv -\frac{2\pi}{f_{\text{rev}}} \int_{-\infty}^{+\infty} f S(f) \text{Im}\{Z_{\parallel}(f)\} df. \quad (6.14)$$

With these definitions, the induced voltage can be written to first order as

$$V_{\text{ind}}(\tau) \simeq -qN_p f_{\text{rev}} (Z_0 + \tau f_{\text{rev}} Z_1). \quad (6.15)$$

The term Z_0 mainly produces an asymmetry of the potential well and a synchronous-phase shift. In the following, the direct impact of Z_0 on the synchrotron frequency shift is neglected, and the focus is placed on the reactive contribution Z_1 , which modifies the local focusing slope.

From induced-voltage slope to synchrotron-frequency shift

The synchrotron frequency is proportional to the square root of the effective longitudinal focusing. In the small-amplitude limit, the induced voltage modifies the focusing through the slope at the bunch centre:

$$\left. \frac{dV_{\text{ind}}}{d\tau} \right|_{\tau \rightarrow 0} \simeq -qN_p f_{\text{rev}}^2 Z_1. \quad (6.16)$$

Treating this as a small perturbation to the RF focusing, the linear synchrotron frequency including the stationary induced voltage can be written as

$$f_{s0,\text{inc}} \simeq f_{s0} \left[1 + \frac{q f_{\text{rev}}^2}{\omega_{\text{RF}} V_{\text{RF}}} Z_1 N_p \right]^{1/2}, \quad (6.17)$$

where ω_{RF} and V_{RF} are the RF angular frequency and voltage evaluated at the synchronous phase. For weak perturbations, one may expand the square root to obtain $\Delta f_{\text{inc}} = f_{s0,\text{inc}} - f_{s0} \propto Z_1 N_p$.

Analytical evaluation for constant $\text{Im } Z/n$ and $1/\tau_b^3$ scaling

To make the dependence explicit, one can adopt a stationary bunch model and a simplified reactive impedance. For many synchrotrons, a useful approximation is to represent the reactive impedance in the frequency band sampled by the bunch spectrum by a constant $\text{Im } Z/n$,

$$\text{Im } Z_{\parallel}(f) \simeq \left(\frac{\text{Im } Z}{n} \right) \frac{f}{f_{\text{rev}}}, \quad (6.18)$$

where $n = f/f_{\text{rev}}$. Substituting Eq. (6.18) into Eq. (6.14) yields

$$Z_1 \simeq -\frac{2\pi}{f_{\text{rev}}^2} \left(\frac{\text{Im } Z}{n} \right) \int_{-\infty}^{+\infty} f^2 S(f) df. \quad (6.19)$$

For a broad class of stationary distributions (including the parabolic distribution), the integral $\int f^2 S(f) df$ scales as τ_b^{-3} , where τ_b is a characteristic bunch length (e.g. the full length for distributions with finite support, or a multiple of σ_{rms} for Gaussian-like distributions). As a consequence, the incoherent synchrotron-frequency shift exhibits the well-known cubic scaling

$$\Delta f_{\text{inc}} \propto - \left(\frac{\text{Im}Z}{n} \right) \frac{N_p}{\tau_b^3}. \quad (6.20)$$

The sign of the shift depends on the sign of the reactive impedance and on whether the machine operates above or below transition.

In the specific case of a parabolic bunch (corresponding to a binomial distribution with $\mu = 1$), and under the small-perturbation assumption, one obtains the commonly used expression for the small-amplitude incoherent shift:

$$\Delta f_{\text{inc}}(\tau \rightarrow 0) \simeq - \frac{6q}{f_{s0}} \frac{(\text{Im}Z/n) N_p}{h V_{\text{RF}} \omega_{\text{rev}}^2 \tau_b^3}, \quad (6.21)$$

where $h = \omega_{\text{RF}}/\omega_{\text{rev}}$ is the RF harmonic number and $\omega_{\text{rev}} = 2\pi f_{\text{rev}}$.

Equivalent impedance

In measurements, it is useful to define an equivalent reactive impedance $(\text{Im}Z/n)_{\text{eq}}$ as the constant $\text{Im}Z/n$ value that reproduces the measured incoherent shift for a given bunch length:

$$\left(\frac{\text{Im}Z}{n} \right)_{\text{eq}} \equiv C \Delta f_{\text{inc}} \tau_b^3 N_p^{-1}, \quad (6.22)$$

where C groups machine and RF parameters (slip factor, voltage, and revolution frequency). This definition turns the experimentally determined slope of frequency shift versus intensity into an effective impedance that can be compared with impedance models and simulations.

6.2 PSB longitudinal reactive impedance

Longitudinal space charge and the Finemet cavity longitudinal coupling impedance are the dominant contributions to the PSB longitudinal impedance model. Therefore, they are the two components considered in the quadrupole synchrotron frequency shift study.

Longitudinal space charge Z_{SC} can be modelled by a purely, constant imaginary impedance $\text{Im}Z/n$ [41] and it provides a defocusing effect on the bunch for synchrotrons operating below transition energy.

In the PSB the longitudinal space charge after-upgrade was accurately derived with the use of the Longitudinal Space Charge (LSC) code developed at SLAC [54]. The final expression obtained to compute the longitudinal space charge impedance is

$$\left| \frac{Z_{sc}}{n} \right| (E_{\text{kin}}) = \frac{\beta_0 \gamma_0^2}{\beta \gamma^2} \left| \frac{Z_{sc}}{n} \right|_0 + \frac{Z_0}{2 \beta \gamma^2} \log \left(\frac{\beta \gamma}{E_{\text{kin}} \beta_0 \gamma_0} \right), \quad (6.23)$$

where $Z_0 \simeq 377 \, \Omega$ is the impedance of free space. For this study, the reference point is chosen at $E_{\text{kin},0} = 170 \, \text{MeV}$, with $\beta_0 \simeq 0.53$, $\gamma_0 \simeq 1.18$ and $|Z_{sc}/n|_0 \simeq 608.3 \, \Omega$. More details for the derivation of the longitudinal space charge impedance as used in the PSB can be found in [54]. The longitudinal space charge impedance computed using Eq. 6.23 is shown in Fig. 6.2. The negative sign reflects the capacitive (defocusing) nature of the space charge contribution for operation below transition.




Figure 6.2: Estimation of the PSB longitudinal space charge impedance in the cycle-time C275-C805 for a 2 GeV cycle.

The precise evaluation and integration of the longitudinal space charge impedance for macroparticle simulations is necessary to avoid biases in

the comparison with measurements. An incorrect implementation can also affect the equilibrium distribution and therefore propagate to the predicted instability thresholds.

While the longitudinal space charge impedance varies strongly with energy through its $1/(\beta\gamma^2)$ scaling, the Finemet cavity impedance is essentially energy independent and broadband; the resulting interplay between the two contributions is therefore dictated by their relative magnitude in the frequency range sampled by the bunch spectrum. In particular, the slightly positive reactive component of the Finemet impedance provides a partial compensation of the capacitive space charge term. In open-loop operation, where the Finemet impedance is not reduced by the LLRF feedback, this partial compensation is stronger; when the cavities are operated in closed-loop, the feedback lowers the cavity impedance (and thus its reactive part), weakening the compensation and leaving a net impedance that is more space charge dominated.

6.3 Quadrupole synchrotron frequency shifts at PSB flat bottom

This measurement campaign aimed to probe the reactive longitudinal impedance sampled by PSB bunches at flat bottom by studying the intensity and bunch length dependence of the quadrupole synchrotron oscillation frequency f_{s2} . The results were compared with macroparticle simulations to assess the validity of the longitudinal impedance model.

6.3.1 Measurement and simulation setup

Measurements were performed at flat bottom, corresponding to a kinetic energy of $E_{\text{kin}} = 160$ MeV, with RF operation at the first harmonic ($h = 1$). A voltage of $V = 3$ kV was used, corresponding to a small-amplitude synchrotron frequency of $f_{s0} = 1.678$ kHz. Therefore, the associated reference quadrupole frequency is $f_{s2,0} \approx 3.357$ kHz. The experimental procedure consisted of the following steps:

- generation of a flat bottom energy plateau of the desired duration, in order to obtain sufficient frequency resolution,
- excitation of quadrupole motion through a non-adiabatic voltage change,

- acquisition of longitudinal profiles for multiple intensities and bunch lengths.

The FWHM bunch length was computed from the acquired bunch profiles and the frequency of the bunch length oscillations f_{s2} was obtained by identifying the maximum component of the Fast Fourier Transform.


Figure 6.3 shows two bunch length oscillation evolutions corresponding to an injected bunch length of 600 ns with the cavity control loops disabled and enabled. A faster damping is observed when the loops are off, consistent with a stronger compensation of space charge by the cavity reactance. When the loops are enabled, the effective cavity impedance is reduced and the oscillation behaviour becomes more space charge dominated.

To interpret the measured intensity dependence and to probe the impedance model, macroparticle simulations were performed with BLoND [51]. Two complementary impedance hypotheses were compared:


- Pure longitudinal space charge: the capacitive space charge contribution is included as an effective impedance at 160 MeV.
- Space charge and Finemet cavity impedance: the measured broadband Finemet impedance is added to space charge.

In Fig. 6.4 the simulations obtained with the two impedance hypotheses are compared for two injected bunch lengths. The pure space charge model is representative of the experimental closed-loop condition: when the LLRF beam-loading compensation is enabled, the beam-induced voltage is strongly reduced at the revolution-harmonic lines and the effective $\Im\{Z_{\parallel}\}$ contribution of the Finemet cavities is largely cancelled. In this regime, the cavities no longer provide a significant inductive reactance to counteract the capacitive space charge term, so that the longitudinal focusing is primarily modified by space charge alone. As a result, the net reactive impedance is more strongly defocusing and the predicted quadrupole synchrotron frequency shifts are larger than in the case where the Finemet impedance is included.

When the Finemet impedance is added in the simulations, the cavity reactance introduces an inductive contribution that partially compensates the capacitive space charge impedance. This reduces the magnitude of the net reactive impedance seen by the bunch, weakens the effective defocusing, and therefore reduces the intensity-dependent detuning of f_{s2} . This trend is consistent with the open-loop measurements, where the beam samples the full cavity impedance: the quadrupole oscillations damp faster and the detuning with intensity is smaller, indicating partial compensation of space charge by the cavity reactance.



(a) Closed-loop condition.




(b) Open-loop condition.


Figure 6.3: Comparison of measured quadrupole oscillations in closed-loop and open-loop conditions and for different intensities. The defocusing and intensity-dependent effect due to the strong longitudinal space charge effect is clearly visible when the Finemet cavity reactive impedance is compensated by the enabled action of the LLRF loops.

Measurement-simulation comparison and extracted slopes

The measured quadrupole frequency was fitted with respect to a scan of intensity with a linear model $f_{s2}(N_p) = a + bN_p$ for each bunch length con-




(a) Injected bunch length $\tau_b \approx 200$ ns.




(b) Injected bunch length $\tau_b \approx 600$ ns.

Figure 6.4: BLoND simulations of the quadrupole synchrotron frequency f_{s2} at flat bottom as a function of bunch intensity, comparing a pure space-charge model with a combined space-charge + Finemet impedance model for two injected bunch lengths.

figuration, and compared to BLoND predictions based on the space charge plus Finemet impedance model. Figure 6.5 summarises the comparison and fit results for three bunch lengths. A good agreement is obtained for the



(a) Measured quadrupole synchrotron frequency shifts in open-loop.



(b) Simulated quadrupole synchrotron frequency shifts considering impedance model equivalent to the open-loop condition.

Figure 6.5: Comparison between measured and simulated quadrupole synchrotron frequency shifts in the open-loop condition. The comparison is performed for different bunch lengths.

longer bunch cases (e.g. ~ 500 ns and ~ 350 ns), supporting the adequacy of the reactive impedance model for those operating points. For the short-

est bunch case (e.g. ~ 200 ns), the discrepancy suggests that the broadened bunch spectrum samples a wider and more structured impedance range, increasing sensitivity to the detailed frequency dependence of the cavity impedance and to narrow-band features. Overall, these flat bottom results provide an experimental benchmark for the reactive impedance model and motivate the refinement of the closed-loop impedance representation when the cavity servo loops are active.

6.4 Summary

In this chapter, quadrupole synchrotron frequency shifts were used to probe the reactive part of the PSB longitudinal impedance. The approach exploits the sensitivity of the quadrupole mode to modifications of the effective longitudinal focusing induced by the stationary induced voltage, while avoiding the incoherent-coherent cancellation that can affect dipole oscillations. The measured dependence of f_{s2} on intensity and bunch length therefore provides a practical observable to benchmark reactive-impedance models.

At PSB flat bottom, the net reactive impedance is dominated by the interplay between longitudinal space charge and the broadband Finemet cavity impedance. Space charge is capacitive and defocusing below transition and exhibits a strong dependence on energy and on bunch length, whereas the Finemet impedance is essentially energy independent and can provide partial compensation through its reactive component. This compensation is configuration dependent: in open-loop operation the cavity reactance is larger and the compensation is stronger, while in closed loop the LLRF feedback reduces the effective cavity impedance (including its reactive part), weakening the compensation and leaving a net impedance that is more space charge dominated.

Quadrupole oscillations were excited by a controlled non-adiabatic voltage perturbation and f_{s2} was extracted from the FFT of turn-by-turn bunch length measurements. The experimental results were compared with BLonD macroparticle simulations using two impedance hypotheses (space charge only, and space charge plus Finemet). The simulations reproduce the expected qualitative trend: including the Finemet impedance reduces the magnitude of the intensity-dependent frequency shift compared to pure space charge, consistent with partial compensation. Quantitatively, good agreement is obtained for the longer-bunch cases (typically ~ 500 – 350 ns), supporting the adequacy of the impedance model in that regime. The deviations observed for the shortest bunch length (~ 200 ns) indicate

increased sensitivity to impedance frequency structure and to measurement/simulation details (e.g. filamentation during the acquisition window and narrow-band features associated with RF feedback). These results provide a benchmark for the PSB reactive impedance model at flat bottom and motivate further refinements of the closed-loop Finemet impedance representation for quantitative predictions at short bunch length and high intensity.

Chapter 7

High-intensity studies in the PSB and possible future applications

Run 3 has been a period of extensive consolidation and optimisation, aimed at fully exploiting the benefits brought by the PSB hardware upgrades [57, 58]. In parallel, a large number of studies and dedicated MD sessions have been carried out, providing regular updates on the maximum achievable intensity and on the operational margins. Figure 7.1 displays the considerable impact of the PSB injection upgrade by comparing the beam intensity evolution along an ISOLDE cycle in 2018 and in 2023. In 2018 (red), a pronounced drop immediately after injection leads to a transmission of $\sim 77\%$. After the LIU upgrades (green), the same extracted intensity is obtained with a transmission of $\sim 99\%$. The figure also reports the highest intensity achieved in 2023 (blue), reaching $\sim 1.6 \times 10^{13}$ p/r while preserving a similarly high transmission, thereby highlighting the increased performance margin enabled by the upgraded injection and optimisations.

In view of present and future user requirements, in particular the beam properties demanded by current and upcoming fixed-target experiments, and notably the approval of the Search for Hidden Particles (SHiP) facility [59], the PSB and the injector chain are entering a new regime in which higher intensities are foreseen. This increase must be achieved while continuing to satisfy the requirements of the existing facilities.

Accordingly, dedicated high-intensity tests were performed, both in standard conditions and in special configurations exploiting Linac4 with dedicated injection settings, with the aim of pushing the machine performance and identifying the main limitations to be addressed. In particular, these studies were designed to:

- establish the intensity headroom beyond routine operation and iden-




Figure 7.1: Single-ring ISOLDE beam intensity in 2018, before the LIU upgrades, and in 2023, before and after performance optimisations. The red trace shows pre-LIU performance, while the green and blue traces show post-LIU performance, with the blue one corresponding to the highest intensity achieved in 2023. *Courtesy of Tirsi Prebibaj.*

tify the dominant bottlenecks;

- quantify the evolution of transmission and emittance as intensity is increased;
- validate operational scenarios for higher total extracted intensity;
- assess hardware constraints that become critical at high intensity, such as the demands placed on the RF system.

The present chapter reports the results of the high-intensity studies performed over the last few years. It outlines the adopted strategy to identify and assess the limiting factors, through the coordinated work of accelerator physicists with complementary expertise, to which the author of the present thesis contributed significantly. It also discusses the measures identified and the upgrade paths suggested by these results, providing guidance for future improvements.

7.1 High-intensity operation and strategy to push the intensity

The ISOLDE facility is presently the most demanding PSB physics user in terms of intensity. Its physics production would benefit from increased flexibility in the extracted beam parameters, especially intensity, energy, average current, and repetition rate, which motivates dedicated studies to extend the operational margins of the PSB.

ISOLDE is currently supplied by two PSB operational cycles, both with an extraction kinetic energy of 1.4 GeV. Under nominal conditions, the PSB delivers 3.2×10^{13} protons per pulse to ISOLDE, with a transmission above 99%. The same intensity is delivered also through a test cycle operated at a higher extraction energy of 1.7 GeV.

Building on this baseline, a structured strategy was adopted to push the machine performance: the bunch population was increased stepwise while maintaining acceptable beam quality and controlled loss levels. This required acting simultaneously on the injection process and the peak line density, i.e. the main drivers of space charge effects at low energy, as well as on longitudinal and transverse stability throughout the cycle.

In the PSB, the intensity increase was therefore pursued through the following main measures:

- Transverse injection painting: slow kickers enable controlled transverse phase space painting via a decaying closed-orbit bump during injection. This redistributes particles in phase space, reduces the peak line density and therefore the space charge tune spread, ultimately improving transmission at high intensity.
- Longitudinal bunch shaping with multi-harmonic RF: the Finemet RF systems allow multi-harmonic configurations to tailor the longitudinal line density. The use of two or three harmonics reduces the peak current compared to single-harmonic operation, alleviating space charge effects and increasing the tolerance to resonance crossing and associated losses.
- Resonance compensation: resonance compensation is essential to reduce losses and sustain high transmission under high-intensity and high-brightness conditions. Although multipole correctors provide a substantial improvement, residual resonance driving terms cannot be fully suppressed, especially in ramping cycles where corrector current limits apply. In this context, reducing the peak line density

through longitudinal shaping provides an effective complementary mitigation.

With these optimisations in place, high-intensity MD campaigns were carried out in 2023 and 2024 at 1.4 GeV extraction energy, reaching record intensities each year [58, 60]. Measurements in individual rings showed that intensities could be pushed up, firstly, to approximately 1.6×10^{13} p/r, with acceleration demonstrated up to about 1.7×10^{13} p/r. However, it is important to distinguish the single-ring intensity reach demonstrated in MD conditions from the operationally qualified intensity. In fact, intensities above roughly 1.25×10^{13} p/r were not considered operational at the time of the study, as they could not be achieved while remaining within the nominal interlock limits of the DC power supply of the RF amplifiers. In standard operation, the interlock is conservatively set to 118 A, well within the RF power stage specifications and sufficient for routine running at the presently delivered intensities; however, it becomes limiting when exploring the accelerator ultimate intensity reach. During high-intensity studies, the threshold can be temporarily raised to quantify the achievable gain in intensity. For a given intensity and RF voltage program, it is therefore essential to monitor the DC current request throughout the cycle and ensure that the configured limit is not exceeded, since crossing the threshold can trigger trips of the cavity cells. In addition, the available RF voltage becomes increasingly critical at high intensity when multiple harmonics are used. In particular, the triple-harmonic configuration in the first part of the cycle is mandatory for the reasons discussed above, and it can significantly affect the peak voltage margin. These considerations motivated a dedicated characterisation of the RF power and voltage demands, and naturally extended to the study of a 2 GeV cycle. Such a cycle is known to be more stringent in terms of RF requirements and is of direct interest in view of a potential upgrade of the extraction energy to 2 GeV for the ISOLDE facility. In 2025 a comprehensive and systematic campaign was carried out to investigate these RF-driven limitations for a 2 GeV cycle.

7.2 DC current demand of the RF power supply as an observable

A key request from ISOLDE is the availability of higher energy beams; therefore, specific studies have been focused on assessing the intensity and performance reach of a 2 GeV PSB cycle. In this regime, the RF power margin becomes increasingly critical, and a practical way to quantify the

proximity to RF power limitations is to monitor the power supply DC current requested by the Finemet RF power stage.

7.2.1 Engineering interpretation of I_{DC}

The measured I_{DC} corresponds to the current drawn from the DC bus feeding the RF power stage. It therefore provides a direct operational proxy for the instantaneous electrical power taken from the DC supply,

$$P_{\text{DC}} = V_{\text{DC}} I_{\text{DC}}, \quad (7.1)$$

where V_{DC} is the DC bus voltage. The RF amplifier converts this electrical power into RF power delivered to the cavity with an efficiency η that depends on the operating point and frequency,

$$P_{\text{RF}} \simeq \eta P_{\text{DC}} \quad \Rightarrow \quad I_{\text{DC}} \simeq \frac{P_{\text{RF}}}{\eta V_{\text{DC}}}. \quad (7.2)$$

As a consequence, increases in the instantaneous RF power demand translate directly into an increased I_{DC} request and, eventually, into proximity to the corresponding interlock thresholds.

A useful decomposition of the RF power request is

$$P_{\text{RF}} \approx P_{\text{beam}} + P_{\text{cav}} + P_{\text{ctrl}}, \quad (7.3)$$

where P_{beam} is the real power transferred to the beam (beam loading), P_{cav} represents the RF power dissipated in the cavity and RF station to sustain the requested gap voltage (cavity and station losses), and P_{ctrl} accounts for additional power required by regulation, transients and imperfect matching. At high intensity, P_{beam} increases with the fundamental component of the beam current and with the accelerating component of the RF voltage, while P_{cav} depends on the voltage programme and on the effective shunt impedance of the cavity system.

7.2.2 Frequency dependence and multi-harmonic operation

The Finemet cavities are broadband and their effective shunt impedance is strongly frequency dependent. Consequently, the dissipated RF power needed to sustain a given harmonic voltage component depends on the harmonic frequency f_h . For a sinusoidal voltage at harmonic h with peak gap voltage V_h , the cavity-loss term can be expressed as

$$P_{\text{cav},h} \approx \frac{V_h^2}{2 R_{\text{sh}}(f_h)}. \quad (7.4)$$




Figure 7.2: Example of frequency dependence of the I_{DC} observable when exciting individual harmonics of the Finemet RF system. The shaded regions indicate the frequency spans covered by successive harmonics ($h \times [f_{inj}, f_{ext}]$), highlighting that the RF power demand (and therefore I_{DC}) depends on harmonic frequency through the frequency-dependent cavity/station impedance.

In multi-harmonic operation, the total cavity-loss contribution can be approximated as a sum over harmonics,

$$P_{cav} \approx \sum_h \frac{V_h^2}{2 R_{sh}(f_h)}, \quad (7.5)$$

where the frequency dependence of $R_{sh}(f_h)$ implies that different allocations of voltage across harmonics can lead to significantly different RF power (and therefore I_{DC}) demands, even for comparable total voltage programmes.

To illustrate this behaviour, Fig. 7.2 shows the measured dependence of the I_{DC} observable on frequency when exciting individual harmonics, highlighting the frequency-dependent “power cost” of sustaining RF voltage in a Finemet cavity system. This frequency dependence is a relevant ingredient when designing multi-harmonic programmes: while multi-harmonic bunch shaping reduces peak line density and can improve beam dynamics, it can also increase the RF power demand if significant voltage is applied in frequency regions where $R_{sh}(f)$ is low.




Figure 7.3: DC current requested by the PSB RF Finemet power supply during the 2 GeV cycle. The signal is acquired via an oscilloscope in the CERN control room and is used as a reference for the peak RF power demand as the beam intensity is increased. The yellow marker indicates the current value at peak acceleration demand used for the intensity scan.


7.2.3 Sampling at peak acceleration demand

Figure 7.3 illustrates the DC current observable as it was used during the studies. During the intensity scan, the DC current I_{DC} measured in one cell was taken as a reference and recorded for each tested intensity. The quantity of interest is I_{DC} at the point of maximum acceleration demand (yellow marker), i.e. where the magnetic field ramp rate dB/dt is largest. At this instant the beam energy must increase at the highest rate to follow the momentum ramp imposed by the dipole field, and the RF system must therefore deliver its maximum real power to the beam. Since I_{DC} tracks the instantaneous electrical power request of the RF chain, it reaches its peak in the vicinity of this maximum- dB/dt point, making it a robust single-point observable of proximity to RF power limitations and associated interlocks.


As stated above, the nominal interlock threshold is set to 118 A. For the purpose of these studies, the limit could be temporarily increased by approximately 25%, corresponding to a value of 143 A. It is important to underline that these thresholds are not simply “peak capability” values: the operational threshold is set conservatively to ensure reliable continuous running under worst-case conditions (thermal margins, mismatch, module spread and measurement uncertainty), while the temporary increase

was authorised only for dedicated MD studies to quantify the achievable intensity headroom.

Figure 7.4 summarises the dependence of the Finemet power stage DC current at peak acceleration on beam intensity for the 2 GeV studies. The upper plot shows the measurement for one of the rings, while the lower one overlays data acquired in two different rings to demonstrate that the observed behaviour is reproducible across rings. The black dashed line indicates the nominal interlock threshold, whereas the red dashed line corresponds to the temporarily increased threshold used during the measurement. The intensity reach for the 2 GeV cycle corresponds to 1.1×10^{13} p/r in nominal conditions and 1.4×10^{13} p/r with the higher interlock setting. The star markers therefore represent an extrapolation to the target intensity, used to estimate the expected current demand and to assess whether an increase of the available RF power (e.g. via an upgrade of the RF amplifiers) would be technically feasible and operationally justified given the potential gain in intensity reach. In these studies, the available peak RF voltage did not constitute a limiting factor; the dominant constraint was the DC current demand of the power stage.



(a) Single-ring measurement of I_{DC} versus intensity for the 2 GeV cycle. The dashed lines indicate the nominal and increased interlock thresholds; star markers correspond to extrapolated values at the target intensity.



(b) Reproducibility across rings: I_{DC} versus intensity measured in two rings under comparable conditions. The consistent trend confirms that the current-demand behaviour is not ring-specific; star markers again indicate extrapolated values at the target intensity.

Figure 7.4: DC current I_{DC} requested by the PSB Finemet RF power stage at the point of peak acceleration as a function of beam intensity for the 2 GeV cycle. The top panel shows a single-ring scan, while the bottom panel overlays measurements from two rings to demonstrate reproducibility. The black dashed line indicates the nominal interlock level and the red dashed line the increased threshold temporarily authorised for the MD; data acquisition stops once the interlock is reached. Star markers show extrapolated current values at the target intensity, used to estimate the over request relative to the interlock and to evaluate the feasibility and justification of an RF power stage upgrade. In this campaign, the peak RF voltage availability did not represent a bottleneck.

7.3 Impact of the magnetic cycle ramp-rate on the DC current demand




Figure 7.5: Comparison of the magnetic field program $B(t)$ (top) and its time derivative $\dot{B}(t)$ (bottom) for the nominal magnetic cycle and an early-extraction configuration (extraction advanced by 100 ms).

The design of the magnetic cycle has a direct impact on accelerator performance because it determines the required acceleration rate. In recent studies, particular attention has been devoted to adapting the cycle to reduce, for instance, space charge effects [61]. One possibility would be to operate the magnetic field with a faster ramp, which would minimise the time spent at low energy, hence in the region of high space charge regime. On the other hand, in the high energy part of the cycle, reducing the ramp rate lowers the energy gain per turn and thus the needed accelerating component $V_{RF} \sin \phi_s$. This translates to a lower real RF power transferred to the beam and, consequently, to a reduced I_{DC} demand from the RF power stage, allowing higher intensity to be reached before encountering the RF interlock threshold.

Figure 7.5 compares the nominal PSB magnetic cycle configuration with a modified version designed to reflect these considerations.




Figure 7.6: DC current I_{DC} requested by the PSB Finemet RF power stage at the point of peak acceleration as a function of beam intensity for the 2 GeV cycle, comparing the nominal magnetic cycle (triangles) with a modified configuration (circles). The modified cycle reduces the current demand starting for intermediate intensities. Star markers indicate the extrapolated I_{DC} values at the target intensity of 1.6×10^{13} p/b for the two configurations, highlighting the expected difference in required headroom.

To assess the impact of the modified ramp on RF related limitations, the DC current demand was characterised following the same procedure as in the previous scans, for both magnetic cycle configurations. The measurements show that the modified cycle lowers the peak DC current request at higher intensities, as shown in Fig. 7.6. This result motivates further optimisation studies to identify an optimal compromise in the magnetic cycle design, with the aim of combining the benefits of reduced time in the space charge dominated regime with increased RF power margin.

7.4 High-intensity studies at 1.4 GeV: demonstrated single-ring reach

In 2025, on the basis of the results and observations described in the previous sections, dedicated MDs were carried out using the standard 1.4 GeV cycle with the explicit goal of probing the ultimate single-ring intensity reach of the PSB. These tests led to record extracted intensities of about 1.76×10^{13} protons in Rings 1 and Ring 2 and up to 1.8×10^{13} protons in Ring 3, which represents a significant step beyond the levels routinely achieved in operation and further proves the available performance potential of the upgraded machine.

For the reach of these results, the key enabling factor was the availability of the highest source current, with an estimated maximum injectible intensity approaching 2×10^{13} protons. It is also noteworthy that tests resulted in lower performance for Ring 4, and the difference will be a matter of study during 2026. In these conditions, operation in bunch-lengthening mode at extraction was observed to reduce the amplifier current signal, providing an additional knob for the optimisation.

Table 7.1 summarises the main intensity levels highlighted in this chapter, contrasting the typical operational values with the record intensities demonstrated in dedicated single-ring MD studies during Run 3.

Table 7.1: Summary of high-intensity performance levels for the 1.4 GeV and 2 GeV cycles. The four-ring equivalent is obtained by multiplying the single-ring highest value by four.

Cycle	Case	Per ring [p/r]	4 rings [p/p]
1.4 GeV	Operational	0.8×10^{13}	3.2×10^{13}
	Future operation	1.25×10^{13}	5×10^{13}
	single-ring MD 2023	1.6×10^{13}	6.4×10^{13}
	single-ring MD 2024	1.7×10^{13}	6.8×10^{13}
	single-ring MD 2025	1.8×10^{13}	7.2×10^{13}
2 GeV	Nominal interlock current	1.1×10^{13}	4.4×10^{13}
	Increased interlock current	1.4×10^{13}	5.6×10^{13}

7.5 Summary

This chapter summarises the high-intensity studies carried out in the PSB over the last years and places them in the context of present operation and future performance needs. The results demonstrate that the LIU upgrades have not only enabled the routine delivery of higher-brightness beams, but have also substantially increased the performance potential of the machine, as evidenced by the strong improvement in injection-to-extraction transmission for ISOLDE-type cycles. The comparison of Run 3 performance with pre-LIU operation illustrates that intensities previously associated with significant post-injection losses can now be delivered with near-loss-free acceleration, confirming both the effectiveness of the hardware consolidation and the impact of the subsequent operational optimisation.

A structured strategy was adopted to push intensity while preserving beam quality and ensuring controlled loss levels. The intensity reach could only be extended through a combined optimisation of transverse and longitudinal beam parameters, reflecting the fact that different limitations dominate at different stages of the cycle. Transverse injection painting, enabled by the upgraded injection system and slow bump kickers, proved essential to mitigate space-charge-driven losses by distributing particles in a larger transverse phase-space area and reducing the peak density. In parallel, longitudinal bunch shaping using multi-harmonic RF configurations allowed control of the peak line density, reducing the maximum bunch current and alleviating both space-charge effects and susceptibility to resonance crossing. These measures, complemented by resonance compensation, constitute a robust operational toolbox for high-intensity running and provide a clear framework for further performance gains.

Within this framework, the chapter highlights a key practical outcome: at the highest intensities the dominant operational limitation is increasingly set by RF power-stage constraints rather than by a lack of RF voltage. The DC current drawn by the Finemet RF amplifiers was identified as a sensitive and operationally convenient observable of proximity to RF power limitations. Its interpretation can be directly linked to the RF power budget and to the frequency-dependent cavity/station losses in multi-harmonic operation, which affect the overall RF headroom. By sampling I_{DC} at the point of maximum acceleration demand—coincident with the maximum dB/dt of the magnetic cycle—a robust single-point indicator of peak RF power demand can be obtained.

The systematic scans performed for the 2 GeV studies showed a clear and reproducible dependence of I_{DC} on beam intensity across rings, enabling a quantitative assessment of intensity reach in relation to the inter-

lock thresholds. Under nominal interlock settings, the attainable intensity for the 2 GeV cycle was limited to $\sim 1.1 \times 10^{13}$ p/r, while a temporarily authorised increase of the threshold extended the demonstrated reach to $\sim 1.4 \times 10^{13}$ p/r. Extrapolation to higher target intensities indicates that further progress for this cycle would require increased RF power capability or a reduction of peak power demand, confirming that RF power headroom is a critical element for any operational extension towards higher-energy, higher-intensity beams.

The chapter also demonstrates that magnetic-cycle design offers an effective lever to redistribute RF power demand along the ramp and to mitigate peak RF constraints. The comparison between the nominal cycle and an early-extraction configuration shows that shaping the ramp can reduce the peak DC current request at intermediate and high intensities. This observation underlines the need for an optimisation that simultaneously considers (i) the benefits of a faster ramp at low energy, which reduces the time spent in the space-charge-dominated regime, and (ii) the benefits of a reduced ramp rate at higher energy, which lowers the accelerating component required from the RF system and therefore relaxes the peak RF power demand. Such cycle-shaping approaches are attractive because they can provide a performance gain without immediate hardware changes, although they must be evaluated against operational constraints such as cycle length, repetition rate and overall throughput for users.

Finally, dedicated high-intensity machine-development studies at 1.4 GeV demonstrated a significant single-ring intensity reach beyond routine operation, with extracted intensities approaching $\sim 1.8 \times 10^{13}$ p/r in several rings under favourable injector conditions. These results confirm that the upgraded PSB retains a substantial performance margin and that the ultimate reach is increasingly determined by a combination of hardware interlocks, cycle design, and operational optimisation rather than by a single fundamental beam-dynamics limit. The observed ring-to-ring differences (notably the reduced performance of Ring 4 in the highest-intensity studies) identify a clear priority for future work, as removing such discrepancies is essential to translate the demonstrated single-ring performance into robust multi-ring operational gains.

Overall, the studies presented in this chapter provide a quantitative basis to guide future developments of the PSB towards higher intensity and higher extraction energy. On the operational side, they confirm the effectiveness of injection and longitudinal shaping techniques for controlling space charge and maintaining high transmission. On the hardware and system side, they show that RF power-stage headroom, its monitoring via I_{DC} , and the frequency-dependent cost of multi-harmonic operation

are central to extending performance, particularly for 2 GeV scenarios. Together with magnetic-cycle optimisation, these results outline a coherent set of upgrade paths and mitigation measures that can support future user requirements while preserving the reliability expected from the first injector of the CERN complex.

Chapter 8

Conclusions and outlook

8.1 Summary of results

The upgrade of the CERN Proton Synchrotron Booster (PSB) during Long Shutdown 2, within the framework of the LHC Injectors Upgrade (LIU) programme, has enabled a new high-intensity and high-brightness operating regime required for HL-LHC beam production. A central element of this upgrade is the renewal of the RF system, with the replacement of the former ferrite cavities by wideband magnetic-alloy (Finemet) cavities and the deployment of multi-harmonic digital cavity feedback loops. While this architecture provides strong operational flexibility for bunch manipulations, it also introduces a dominant broadband longitudinal impedance that makes beam loading and RF-loop dynamics a key driver of longitudinal performance and stability. The work presented in this thesis was motivated by the need to characterise and understand these mechanisms quantitatively, and to translate them into experimentally based limits and operational guidance for PSB high-intensity operation.

The thesis combined dedicated beam measurements, longitudinal diagnostics, and modelling of the RF feedback system to address the longitudinal limitations of the upgraded PSB. The main outcomes can be summarised as follows:

- A consistent framework for interpreting PSB longitudinal dynamics was established, linking the RF configuration (single RF versus double RF; bunch lengthening mode versus bunch shortening mode), the loop conditions (open-loop versus closed-loop), and the induced-voltage mechanisms that govern stability margins.
- A systematic experimental characterisation of single-bunch longitudinal instability thresholds after LS2 demonstrated that the stability

domain depends critically on the combined use of RF configuration and cavity feedback. In particular, pure single-RF operation was found to be detrimental to stability already at intermediate intensities, and the application of the beam-loading compensation loops alone was not sufficient to ensure stable operation in the explored conditions. Conversely, double-RF operation in bunch-lengthening mode provides a robust stabilising configuration for routine high-intensity production.

- The instability studies clarified the distinct roles of the Finemet impedance and of the cavity feedback loops. In open-loop, clear fast blow-up thresholds were observed and their dependence on RF voltage was consistent with a beam–cavity interaction dominated by the uncompensated cavity impedance. When the cavity servoloops are enabled, the sharp threshold is suppressed and the reachable intensity range is significantly extended, indicating that the feedback effectively cancels most of the beam-induced voltage at the controlled harmonics. The limitation is then shifted to a gradual longitudinal emittance growth with intensity, consistent with residual closed-loop effects (finite bandwidth, phase lag, and noise) that become more pronounced as the beam-induced signal increases.
- A hardware-faithful numerical model of the multi-harmonic cavity feedback loops was developed and validated against measurements. The model reproduces the internal digital signal-processing chain (demodulation, CIC filtering, decimation rules, PI regulation, re-modulation) and matches measured intermediate I/Q quantities, reconstructed RF drive waveforms, and closed-loop transfer functions. This provides a quantitative tool to study how loop settings influence longitudinal performance and forms a natural basis for future self-consistent integration with macroparticle tracking.
- Quadrupole synchrotron frequency shifts were established as a robust experimental probe of the reactive longitudinal impedance at PSB flat bottom. Measurements of quadrupole oscillations with cavity loops enabled and disabled showed that open-loop conditions exhibit faster damping, consistent with stronger compensation of space charge by the cavity reactance, while closed-loop operation reduces the effective cavity impedance and makes the observed behaviour more space charge dominated. Comparisons with macroparticle simulations under different impedance hypotheses (space charge only

versus space charge plus Finemet) provided a direct cross-check of the impedance picture relevant for injection conditions.

- High-intensity studies in Run 3 demonstrated the substantial performance headroom of the upgraded PSB. The results confirmed that intensities previously associated with significant post-injection losses can now be delivered with nearly lossless acceleration, and record single-ring intensities approaching 1.8×10^{13} protons were demonstrated in dedicated studies. These achievements required a structured strategy combining transverse and longitudinal optimisation, reflecting the fact that different limitations dominate at different stages of the cycle.

Taken together, these results show that the upgraded PSB can reach the LIU-driven intensity and brightness targets while maintaining operational reliability, and that the longitudinal performance envelope is strongly shaped by the interplay between the wideband Finemet impedance, multi-harmonic RF configuration, and cavity feedback loop dynamics. The measurement and modelling methods developed in this thesis provide a coherent basis to interpret present limitations and to guide future optimisation.

8.2 Outlook

The upgraded PSB will continue to serve as the first intensity and brightness gateway of the CERN injector chain, and further performance increases will require continued progress in understanding and controlling beam–cavity–feedback interaction. Several directions follow from this work:

- Self-consistent beam-cavity-loop simulations. The validated time-domain model of the multi-harmonic cavity loops is well positioned to be coupled to macroparticle tracking to study intensity-dependent dynamics with a realistic representation of loop delays, bandwidth, and noise. Such integration would enable predictive scans of stability margins versus loop settings and RF programmes and would provide a model-based framework for loop optimisation.
- Quantitative attribution of closed-loop emittance growth. Since the dominant limitation in closed-loop operation shifts from sharp instabilities to gradual longitudinal emittance growth, future work should aim at separating the contributions of finite loop bandwidth, phase

lag, and measurement/actuation noise, and at identifying operational settings that minimise dilution while preserving stability.

- Refined impedance identification at injection. The quadrupole-frequency-shift method provides a complementary observable to validate reactive impedance models. Extending the technique to a wider range of bunch shapes and RF configurations (including multi-harmonic shaping) would strengthen the link between measured dynamics and impedance hypotheses and help quantify the relative roles of space charge and cavity reactance under operational conditions.
- Operational diagnostics and observables for intensity reach. The high-intensity studies indicate that reliable intensity pushing requires observables that capture proximity to limits in real time. The proposed use of RF power-demand indicators (such as DC current demand) and their correlation with longitudinal settings can be further developed into practical operational tools for cycle validation and performance qualification.

Appendix A

Beam-based RF voltage calibrations during PSB commissioning

RF voltage calibrations were performed during beam commissioning of the PSB and its injector, the linear accelerator Linac4. The calibrations were computed using longitudinal phase space tomography, by reconstructing the longitudinal phase space distribution from a set of measured bunch profiles and by adjusting the RF voltage used in the reconstruction until the simulated synchrotron motion reproduces the measured evolution.

Voltage calibrations are required because the effective RF voltage experienced by the beam can differ significantly from the programmed one. This discrepancy is due to uncertainties in the electrical gap and collective effects.

A.1 Principle of tomography-based voltage calibration

Longitudinal tomography is a beam-based technique used for the reconstruction of the bunch distribution in longitudinal phase space [62]. The reconstruction is obtained by combining a series of measured longitudinal bunch profiles with the accelerator parameters (including the RF voltage) and performing particle tracking under a longitudinal beam dynamics model. The accuracy of the reconstruction depends on the grade of knowledge of the accelerator parameters.

This appendix presents the beam-based RF voltage calibration factors

obtained for the PSB Finemet cavities in all four rings. The factors are derived by varying the RF voltage in the longitudinal tracking model and selecting the value that minimises the discrepancy between measured bunch profiles and the profiles re-projected from the reconstructed phase space.

The calibration is expressed as a multiplicative factor

$$V_{\text{eff}}(t) = k_V V_{\text{ref}}(t), \quad (\text{A.1})$$

where $V_{\text{ref}}(t)$ denotes the programmed or detected voltage waveform used as reference, and $V_{\text{eff}}(t)$ is the effective voltage seen by the beam.

Figure A.1 shows a simplified schematic which summarises the principle of the longitudinal tomography technique.




Figure A.1: Principle of longitudinal phase space tomography used for voltage calibration. Measured bunch profiles are combined with a tracking model (parameterised by machine settings, including the RF voltage) to reconstruct the phase space. The RF voltage is computed by minimising the discrepancy between measured profiles and the projections obtained from the reconstructed distribution.

During Run 3 this methodology was systematically applied in the PSB at injection. A key requirement for these measurements is operation at low beam intensity, so that collective effects remain negligible and do not interfere with the reconstruction.

A.2 Convergence study and acquisition requirements

Before applying the longitudinal tomography technique for the computations of the calibration factors in the PSB, a convergence study was carried out to determine the minimum acquisition length required for a stable and reproducible voltage calibration. In this study, the voltage error percentage is evaluated as a function of the number of synchrotron periods included in the acquired set of longitudinal profiles. Figure A.2 summarises the convergence study performed at several RF voltage setpoints. For each setting, the calibration factor was computed using profile sets spanning increasing acquisition lengths, expressed in units of synchrotron periods. The results show that stable and reproducible calibration factors are obtained once the acquisition covers approximately one synchrotron period; shorter acquisitions exhibit larger scatter and can introduce a systematic bias in the calibration.




Figure A.2: Convergence study of the tomography-based voltage calibration. The obtained voltage error stabilises once the acquisition includes approximately one synchrotron period, which was adopted as a criterion for the calibration measurements.


A.3 PSB Finemet RF voltage calibration

The PSB voltage calibration was performed on a dedicated flat cycle at injection energy in Single-RF condition. Measurements were repeated for each sector (5LI, 7LI and 13LI) as well as for the summed RF voltage in all the three sectors, and for all the four rings. In order to isolate the cavity response and to avoid additional longitudinal perturbations during the reconstruction, the phase and radial loops were disabled.


The effective voltage evaluated from tomography, which corresponds to the voltage seen by the beam, is systematically lower than the programmed voltage. The discrepancy is sector dependent, motivating the extraction of individual calibration factors and their application as a re-scaling stage in the voltage chain. In practice, the correction is applied to the programmed voltage such that, after propagation through the cavity servoloop, the detected voltage matches the beam-based scale. In Fig. A.3 the results for all the combinations are reported. Here a comparison between the voltage calibration factors computed during 2023 and 2024 commissioning is also illustrated.

A.4 Summary


RF voltage calibrations of the PSB Finemet cavities, based on longitudinal tomography, were successfully implemented during the most recent beam commissioning campaigns. To compute the calibration factors both reliably and efficiently, a dedicated convergence study was conducted to determine the minimum data acquisition length required for the tomographic reconstruction. The study showed that acquiring profiles over approximately one synchrotron period is sufficient to achieve stable and reproducible calibration results. Applying the resulting voltage scale is essential to recover the effective RF voltage detected by the beam and avoid bias in the reconstruction of longitudinal beam parameters.




(a) Ring 1



(b) Ring 2



(c) Ring 3



(d) Ring 4

Figure A.3: Tomography-based RF voltage calibration results for the four PSB rings. For each ring, the voltage error is shown for sectors 5LI, 7LI and 13LI and for the summed voltage (SUM), comparing the 2023 and 2024 measurement campaigns.

Bibliography

- [1] Mary K. Gaillard, Paul D. Grannis, and Frank J. Sciulli. “The Standard Model of Particle Physics”. In: *Reviews of Modern Physics* 71.2 (1999), S96–S111. doi: 10.1103/RevModPhys.71.S96.
- [2] M. Oliphant. *The Acceleration of Particles to Very High Energies*. Classified memorandum submitted to the Department of Scientific and Industrial Research (DSIR), United Kingdom. University of Birmingham Archive. Sept. 1943.
- [3] Helmut Wiedemann. *Particle Accelerator Physics*. Springer, 2015.
- [4] S. Y. Lee. *Accelerator Physics*. 2nd ed. London: World Scientific, 2004.
- [5] *CERN Home*. CERN. URL: <https://home.cern/> (visited on 01/08/2026).
- [6] *CERN’s Accelerator Complex*. CERN. URL: <https://home.cern/science/accelerators/accelerator-complex> (visited on 04/15/2026).
- [7] D. Manglunki et al. “LEIR Operations for the LHC and Future Plans”. In: *Proceedings of COOL 2013*. 2013. URL: <https://www.jacow.org/proceedings/COOL2013/papers/mopm2ha01.pdf> (visited on 04/15/2026).
- [8] M. Chanel et al. *LEIR: The Low Energy Ion Ring at CERN*. CERN PS Report CERN-PS-2002-015-AE. CERN, 2002. URL: <https://cds.cern.ch/record/2832141/files/CERN-PS-2002-015-AE.pdf> (visited on 04/15/2026).
- [9] M. Bodendorfer et al. *Ions: Baseline, Studies Plan and Strategy for HL-LHC*. Tech. rep. CERN, 2014. URL: <https://indico.cern.ch/event/315665/contributions/1690120/attachments/605740/833614/LIU-ION-Chamonix2014-MB.pdf> (visited on 04/15/2026).
- [10] W. Herr and B. Muratori. “Concept of Luminosity”. In: (2006). URL: <https://cds.cern.ch/record/941318/files/p361.pdf> (visited on 03/24/2026).

- [11] *The HL-LHC Project*. CERN. URL: <https://hilumilhc.web.cern.ch/content/hl-lhc-project> (visited on 03/24/2026).
- [12] J. Coupard et al. *LHC Injectors Upgrade, Technical Design Report, Vol. I: Protons*. Report CERN-ACC-2014-0337. Geneva, Switzerland: CERN, 2014. URL: <https://cds.cern.ch/record/1976692> (visited on 04/16/2026).
- [13] H. Damerau et al. *LHC Injectors Upgrade, Technical Design Report, Vol. I: Protons*. Technical Design Report CERN-ACC-2014-0337. Geneva, Switzerland: CERN, Dec. 2014.
- [14] M. Vretenar et al. *Linac4 Design Report*. CERN Yellow Reports: Monographs CERN-2020-006. Geneva, Switzerland: CERN, 2020. URL: <https://cds.cern.ch/record/2736208> (visited on 04/16/2026).
- [15] C. Roderick et al. "Accelerator Fault Tracking at CERN". In: *Proceedings of ICALEPCS 2017*. Geneva, Switzerland: JACoW, 2017, pp. 397–400. DOI: 10.18429/JACoW-ICALEPCS2017-TUPHA013. URL: <https://accelconf.web.cern.ch/icalepcs2017/papers/tupha013.pdf> (visited on 04/16/2026).
- [16] L. Landau. "On the Vibrations of the Electronic Plasma". In: *J. Phys. USSR* 10 (1946), p. 483.
- [17] E. N. Shaposhnikova et al. "Longitudinal Instabilities in the SPS and Beam Dynamics Issues with High Harmonic RF Systems". In: *Proc. HB'12*. Beijing, China: JACoW Publishing, Geneva, Switzerland, 2012, pp. 358–362. URL: <https://jacow.org/HB2012/papers/WE01A04.pdf>.
- [18] Chihiro Ohmori and Mauro Paoluzzi. "Magnetic Alloy Loaded Cavities in J-PARC and CERN". In: *Proceedings of the 68th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB'23)*. HB'23: 68th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (Oct. 9–13, 2023). Geneva, Switzerland: JACoW Publishing, 2024, WEC4C1. DOI: 10.18429/JACoW-HB2023-WEC4C1. URL: <https://cds.cern.ch/record/2900875/files/document.pdf> (visited on 02/17/2026).
- [19] Hua Shi and Yoshiro Irie. "Study of a Magnetic Alloy-Loaded Cavity for Compact Synchrotrons". In: *Chinese Physics C* 40.2, 027002 (2016). DOI: 10.1088/1674-1137/40/2/027002. URL: <https://cpc.ihep.ac.cn/fileZGWLC/journal/article/zgwlc/2016/2/PDF/2016-02-13-027002.pdf> (visited on 02/17/2026).

- [20] Hitachi Metals, Ltd. *Nanocrystalline Soft Magnetic Material, Finemet*. URL: <http://www.hitachi-metals.co.jp/products/electel/pdf/hl-fm9-h.pdf> (visited on 04/07/2026).
- [21] M. M. Paoluzzi et al. "Studies on a Wideband, Solid-State Driven RF System for the CERN PS Booster". In: *Proc. IPAC'12*. New Orleans, LA, USA: JACoW Publishing, Geneva, Switzerland, 2012, pp. 3749–3751. URL: <https://jacow.org/IPAC2012/papers/THPPP011.pdf>.
- [22] M. M. Paoluzzi et al. "Design of the New Wideband RF System for the CERN PS Booster". In: *Proc. IPAC'16*. International Particle Accelerator Conference 7. Busan, Korea: JACoW Publishing, Geneva, Switzerland, June 2016, pp. 441–443. doi: 10.18429/JACoW-IPAC2016-MOPMW024. URL: <https://jacow.org/ipac2016/papers/MOPMW024.pdf>.
- [23] M. M. Paoluzzi. *Design of the PSB Wideband RF System*. Tech. rep. CERN-ACC-NOTE-2013-0030. Geneva, Switzerland: CERN, Oct. 29, 2013.
- [24] J. C. Molendijk. "Digital Receiver and Modulator Architecture for Multi-Harmonic RF Finemet Operation". In: *Proc. LLRF'15*. Shanghai, China, 2015.
- [25] J. C. Molendijk. "Introducing Fixed Frequency Clock Operation on the CERN VXS LLRF Platform". In: *Proc. LLRF'17*. Barcelona, Spain, 2017.
- [26] D. Barrientos et al. "A New Beam Loading Compensation and Blow-Up Control System Using Multi-Harmonic Digital Feedback Loops in the CERN Proton Synchrotron Booster". In: *Proc. 13th Int. Particle Accelerator Conf. (IPAC'22)*. FPGA cavity controller; 16 harmonic Cartesian loops; embedded NWA. Bangkok, Thailand: JACoW, 2022.
- [27] J. C. Molendijk. *Multi-Harmonic Servo-Loops Implementation*. Internal note. VXS–DSP–FMC carrier implementation; 16 parallel I/Q loops; fixed 122.7 MSPS clock. CERN, Nov. 2018.
- [28] M. E. Angoletta et al. "Control and Operation of a Wideband RF System in CERN's PS Booster". In: *Proc. IPAC'17*. 2017, pp. 4050–4053. doi: 10.18429/JACoW-IPAC2017-THPAB141.
- [29] T. Schilcher. "RF Applications in Digital Signal Processing". In: *CERN Accelerator School, Digital Signal Processing (2007)*, pp. 249–270.
- [30] Frank J. Sacherer. "Bunch Lengthening and Microwave Instability". In: *IEEE Transactions on Nuclear Science* 24.3 (1977), pp. 1393–1395. doi: 10.1109/TNS.1977.4328955.

- [31] G. Besnier. "Stabilité des oscillations longitudinales d'un faisceau groupé se propageant dans une chambre à vide d'impédance réactive". In: *Nuclear Instruments and Methods* 164 (1979), p. 235. DOI: 10.1016/0029-554X(79)90241-6.
- [32] Y. H. Chin, K. Satoh, and K. Yokoya. "Instability of a Bunched Beam with Synchrotron Frequency Spread". In: *Particle Accelerators* 13 (1983), p. 45.
- [33] J.-L. Laclare. *Bunched Beam Coherent Instabilities*. Report CERN-1987-003-V-1. Geneva, Switzerland: CERN, 1987.
- [34] J.-P. Garnier. "Instabilités cohérentes dans les accélérateurs circulaires". PhD thesis. Grenoble, France: Institut National Polytechnique de Grenoble, 1987.
- [35] Katsunobu Oide and K. Yokoya. *Longitudinal Single Bunch Instability in Electron Storage Rings*. Preprint KEK-Preprint-90-10. Tsukuba, Japan: KEK, 1990.
- [36] Katsunobu Oide. "A Mechanism of Longitudinal Single-Bunch Instability in Storage Rings". In: *Particle Accelerators* 51 (1995), p. 43.
- [37] A. Chao, B. Chen, and K. Oide. "A Weak Microwave Instability with Potential Well Distortion and Radial Mode Coupling". In: *Proceedings of the 16th Particle Accelerator Conference (PAC'95)*. Dallas, TX, USA, 1995, pp. 3040–3042. DOI: 10.1109/PAC.1995.505777.
- [38] K. Y. Ng. "Mode-Coupling Instability and Bunch Lengthening in Proton Machines". In: *Proceedings of the 16th Particle Accelerator Conference (PAC'95)*. Dallas, TX, USA, 1995, pp. 2977–2979. DOI: 10.1109/PAC.1995.505756.
- [39] M. D'yachkov and R. Baartman. "Longitudinal Single Bunch Stability". In: *Particle Accelerators* 50 (1995), p. 105.
- [40] A. Mosnier. "Microwave Instability in Electron Storage Rings". In: *Nuclear Instruments and Methods in Physics Research Section A* 438 (1999), p. 225. DOI: 10.1016/S0168-9002(99)00832-3.
- [41] Alexandre Samir Lasheen. "Beam Measurements of the Longitudinal Impedance of the CERN Super Proton Synchrotron". Prepared at CERN and Université Paris-Sud. Doctoral thesis defended in Orsay on 13 January 2017. NNT: 2017SACLS005. PhD thesis. Orsay, France: Université Paris-Saclay, Jan. 13, 2017.

- [42] I. Karpov and M. Gadioux. “Mechanism of Longitudinal Single-Bunch Instability in the CERN SPS”. In: *Proceedings of IPAC’21*. 12th International Particle Accelerator Conference (IPAC’21). Campinas, SP, Brazil: JACoW Publishing, May 2021, pp. 3161–3164. ISBN: 978-3-95450-214-1. DOI: 10.18429/JACoW-IPAC2021-WEPAB227. URL: <https://jacow.org/ipac2021/doi/JACoW-IPAC2021-WEPAB227.html>.
- [43] M. Marchi et al. “Cavity Loops Influence on the Single-Bunch Instability Thresholds of the CERN PS Booster”. In: *Proc. IPAC’25*. International Particle Accelerator Conference 16. Taipei, Taiwan: JACoW Publishing, Geneva, Switzerland, Nov. 2025, pp. 2168–2171. ISBN: 978-3-95-450248-6. DOI: 10.18429/JACoW-IPAC2025-WEPM082.
- [44] Danilo Quartullo. “Simulations of RF Beam Manipulations Including Intensity Effects for CERN PSB and SPS Upgrades”. Published in open access form (CC BY-NC-ND 3.0 IT). PhD thesis. Rome, Italy: Sapienza University of Rome, Nov. 2022. ISBN: 978-88-9377-243-3. DOI: 10.13133/9788893772433.
- [45] D. Quartullo et al. “CERN PS Booster Longitudinal Dynamics Simulations for the Post-LS2 Scenario”. In: *Proc. HB’16*. Malmö, Sweden: JACoW Publishing, Geneva, Switzerland, July 2016, pp. 140–145. DOI: 10.18429/JACoW-HB2016-MOPR028. URL: <https://jacow.org/hb2016/papers/MOPR028.pdf>.
- [46] S. C. P. Albright et al. “New Longitudinal Beam Production Methods in the CERN Proton Synchrotron Booster”. In: *Proc. IPAC’21*. Campinas, Brazil: JACoW Publishing, Geneva, Switzerland, May 2021, pp. 4130–4133. DOI: 10.18429/JACoW-IPAC2021-THPAB183. URL: <https://jacow.org/ipac2021/papers/THPAB183.pdf>.
- [47] S. C. P. Albright et al. “Commissioning the New LLRF System of the CERN PS Booster”. In: *Proc. IPAC’22*. International Particle Accelerator Conference 13. Bangkok, Thailand: JACoW Publishing, Geneva, Switzerland, July 2022, pp. 2060–2063. ISBN: 978-3-95-450227-1. DOI: 10.18429/JACoW-IPAC2022-WEPOTK012. URL: <https://jacow.org/IPAC2022/papers/WEPOTK012.pdf>.
- [48] E. Métral. *Stability Criterion for the Longitudinal Mode-Coupling Instability in the Presence of Both Space-Charge and Resonator Impedances*. CERN PS Division Report CERN-PS-2001-063-AE. Geneva, Switzerland: CERN, Oct. 24, 2001. URL: <https://cds.cern.ch/record/524139/files/CERN-PS-2001-063-AE.pdf>.

- [49] C. Antuono et al. “Understanding of the New Horizontal Instability at the PS Booster after LIU”. In: *Proc. IPAC’24*. International Particle Accelerator Conference 15. Nashville, TN, USA: JACoW Publishing, Geneva, Switzerland, July 2024, pp. 2594–2597. ISBN: 978-3-95-450247-9. DOI: 10.18429/JACoW-IPAC2024-WEPR44.
- [50] M. Marchi et al. “Numerical Modeling of the CERN PS Booster Cavity Feedback Loops”. In: *Proc. IPAC’25*. Paper from which the test-stand characterisation and loop block diagram are taken. 2025.
- [51] H. Timko et al. “Beam Longitudinal Dynamics Simulation Studies”. In: *Physical Review Accelerators and Beams* 26.11 (Nov. 2023), p. 114602. DOI: 10.1103/PhysRevAccelBeams.26.114602. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.26.114602>.
- [52] Alexander W. Chao. *Physics of Collective Beam Instabilities in High Energy Accelerators*. Wiley Series in Beam Physics and Accelerator Technology. Wiley, 1993.
- [53] Bruno Zotter and Sigfried Kheifets. *Impedances and Wakes in High-Energy Particle Accelerators*. World Scientific, 1998.
- [54] Danilo Quartullo. “Studies of Longitudinal Beam Dynamics and Impedance Effects in the CERN Proton Synchrotron Booster”. PhD thesis. CERN / University of Rome La Sapienza, 2017.
- [55] Alexandre Lasheen. “Beam measurements of the longitudinal impedance of the CERN Super Proton Synchrotron”. PhD thesis. CERN / Université Paris-Saclay, 2017.
- [56] J.-L. Laclare. “Bunched Beam Coherent Instabilities”. In: *Proceedings of the CERN Accelerator School: Fifth General Accelerator Physics Course*. CAS 1985, Vol. I. Geneva, Switzerland, 1987, pp. 264–326.
- [57] F. Asvesta et al. “High Intensity Studies in the CERN Proton Synchrotron Booster”. In: *Proc. IPAC’22*. International Particle Accelerator Conference 13. Bangkok, Thailand: JACoW Publishing, Geneva, Switzerland, July 2022, pp. 2056–2059. ISBN: 978-3-95-450227-1. DOI: 10.18429/JACoW-IPAC2022-WEPOTK011. URL: <https://jacow.org/IPAC2022/papers/WEPOTK011.pdf>.
- [58] M. Marchi et al. “CERN Proton Synchrotron Booster High-Intensity Operation”. In: *Proc. HB’25*. 71st ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (Oct. 19–24, 2025). ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams 71. Huizhou,

- China: JACoW Publishing, Geneva, Switzerland, Oct. 2025, pp. 73–78. ISBN: 978-3-95450-266-0. DOI: 10.18429/JACoW-HB2025-TUIDC01. URL: <https://indico.jacow.org/event/83/contributions/11962>.
- [59] Sergey Alekhin et al. “A Facility to Search for Hidden Particles at the CERN SPS: The SHiP Physics Case”. In: *Reports on Progress in Physics* 79.12 (2016), p. 124201. DOI: 10.1088/0034-4885/79/12/124201. URL: <https://doi.org/10.1088/0034-4885/79/12/124201>.
- [60] T. Prebibaj et al. “Intensity Reach in the CERN PSB with the High-Current LINAC4 Source”. In: *Proc. IPAC’24*. International Particle Accelerator Conference 15. Nashville, TN, USA: JACoW Publishing, Geneva, Switzerland, July 2024, pp. 1706–1709. ISBN: 978-3-95-450247-9. DOI: 10.18429/JACoW-IPAC2024-TUPS31.
- [61] S. Albright et al. “Magnetic Cycle Optimisation in the CERN PS Booster”. In: *Proc. IPAC’25*. International Particle Accelerator Conference 16. Taipei, Taiwan: JACoW Publishing, Geneva, Switzerland, Nov. 2025, pp. 1467–1470. ISBN: 978-3-95-450248-6. DOI: 10.18429/JACoW-IPAC2025-TUPS026.
- [62] D. Quartullo, S. C. P. Albright, and H. Damerau. “Frequency-Dependent RF Voltage Calibration Using Longitudinal Tomography in the CERN PSB”. In: *Proc. IPAC’22*. International Particle Accelerator Conference 13. Bangkok, Thailand: JACoW Publishing, Geneva, Switzerland, July 2022, pp. 845–848. ISBN: 978-3-95-450227-1. DOI: 10.18429/JACoW-IPAC2022-TUPOST006. URL: <https://jacow.org/IPAC2022/papers/TUPOST006.pdf>.