

Advanced RF techniques for CERN's future slow-extracted beams



Pablo A. Arrutia Sota
Keble College
University of Oxford

A thesis submitted for the degree of
DPhil Particle Physics

2024

In loving memory of José Sota, my granddad

Acknowledgements

I would like to start by thanking my supervisors Philip Burrows and Matthew Fraser. Phil gave me the opportunity to join the programme at Oxford and was onboard with the project I proposed. I particularly appreciate his trust during my stay at CERN, while still efficiently giving me the necessary resources and feedback. Matt essentially taught me all I know about doing research. I am extremely grateful for his guidance, which always struck the perfect balance between encouraging crazy ideas and converging on impactful applications.

I would also like to thank my colleagues from CERN. From the ABT-BTP section, Francesco Velotti and Yann Dutheil helped me a lot in setting up simulations and measurements, my office-mates Elliott Johnson and Rebecca Ramjiawan always had time for discussions (often enhanced by unreasonably strong coffee), and Elisabeth Renner gave me tons of practical advice about work, life and how to escape from pumas. Additionally, I'm grateful for all the people from operations and the equipment groups, who supported me during the long shifts and operational tests. In particular, I thank Giulia Papotti for her readiness to troubleshoot issues related to RF, and Roberto Piandani for his patience to collect and provide data.

Both in Geneva and Oxford I've had incredible flatmates: Annabel, Jess, Hannah, Alex, Karl, Joelle and Paula. Thank you all for taking such good care of me. I am also grateful for my friends back in Algorta and around the world, who welcomed me with open arms even if we don't see each other too often. Last but not least, I thank Victor, Sof and Ali for their wisdom and warmth, which made this journey worth all the trouble.

Finally, I would like to thank my family. I am blessed with loving grandparents, parents and brothers. I do not take that for granted. I am very lucky to have their support, which I always felt close to my heart even when far away from them. *Ekin eta Jarrai.*

Abstract

Resonant slow extraction is a technique employed to provide long pulses of particles from synchrotrons. It is exploited in a wide variety of applications, such as experimental physics, hadron therapy and irradiation testing. At the CERN Super Proton Synchrotron (SPS), the technique is exploited to deliver beam to the fixed-target experiments in the North Area. In order to improve current operation and prepare for future requests, the slow-extraction system in the SPS would benefit from further optimisation and flexibility. More specifically, millisecond-scale modulations in the extracted flux are known to compromise experimental data taking. This is an issue common to all facilities that perform slow extraction and mitigation techniques need to be pursued. Additionally, a future experimental request known as Search for Hidden Particles (SHiP) will require a particle spill with a customised time structure in the nanosecond-scale. It is important to propose strategies to deliver such a beam within the constraints of the SPS system. This thesis will study the exploitation of an advanced RF technique known as empty-bucket channelling to address these challenges. In particular, it will employ a simulation model to systematically characterise the manipulation. The results will then be benchmarked with measurement. Finally, an operational implementation will be proposed.

Contents

1	Introduction	1
1.1	History of accelerators	1
1.2	Slow extraction	3
1.2.1	Inception and early history	3
1.2.2	CERN's journey	4
1.2.3	The future	8
1.3	Thesis scope	8
2	Beam dynamics	10
2.1	Reference system	10
2.2	Transverse dynamics	11
2.2.1	Hill's equation	11
2.2.2	Transverse phase space	12
2.2.3	Normalised phase space	13
2.2.4	Chromaticity	15
2.2.5	Emittance	15
2.2.6	Transverse resonance	15
2.3	1/3-integer resonance	16
2.3.1	Sextupole magnet	16
2.3.2	Three-turn formalism	18
2.3.3	The stable and unstable regions	19
2.3.4	Extraction with a septum magnet	21
2.4	Longitudinal dynamics	21
2.4.1	RF cavity	22
2.4.2	Time slippage	23
2.4.3	Longitudinal phase space	23
2.4.4	Accelerating bucket	24
2.5	Phase displacement	26

2.5.1	Reflection from a potential well	28
2.5.2	Effect on a particle ensemble	29
3	Empty-bucket channelling in the CERN SPS	31
3.1	SPS Fixed-Target cycle	32
3.1.1	Transverse aspects	32
3.1.2	Longitudinal aspects	34
3.1.3	Constant-Optics Slow Extraction	35
3.1.4	Momentum stop-band	37
3.2	Empty-bucket channelling	38
3.2.1	‘Accelerating bucket’ with COSE	39
3.2.2	Concept	40
3.2.3	Parametrisation	40
3.2.4	Setup in the SPS	42
3.3	Simulation model: <i>henontrack</i>	43
3.3.1	Map models	44
3.3.2	SPS model	45
3.3.3	Further computational speed-up	46
3.4	End-to-end SPS simulation	48
3.4.1	Modelling procedure	48
3.4.2	Analysis procedure	49
	3.4.2.1 Spill structure	49
	3.4.2.2 Longitudinal phase space	49
	3.4.2.3 Normalised phase space	50
3.5	Conclusion	51
4	Bunched beams for BDF/SHiP	53
4.1	Motivation	53
4.2	Conventional bunched-extraction methods	56
4.3	Empty-bucket channelling for bunched beam	58
4.4	Simulations	60
4.4.1	Extracted bunch length	61
4.4.2	Extracted bunch shape	62
4.4.3	Integrated intensity	63
4.4.4	Transverse effects	66
4.5	Measurements in the SPS	68
4.5.1	Fast versus slow time-structure	69

4.5.2	Extracted bunch length	69
4.5.3	Extracted bunch shape	70
4.5.4	Integrated intensity	73
4.5.5	Transverse effects	74
4.6	Conclusion	75
5	Spill-quality improvements at CERN SPS	76
5.1	Motivation	76
5.2	Spill-quality characterisation	78
5.2.1	The duty factor	78
5.2.2	From current to intensity	79
5.2.3	The tune: the main macroparameter	81
5.2.4	From tune to intensity: analytical model	81
5.3	Transfer function	84
5.3.1	Low-pass filter	85
5.3.2	Linearity and time invariance of the spill structure	86
5.3.2.1	Linearity	86
5.3.2.2	Time invariance	88
5.3.3	Nominal SPS transfer function	88
5.4	Empty-bucket channelling for ripple reduction	89
5.4.1	Concept	89
5.4.2	Semi-analytical estimate of ripple-reduction coefficient	92
5.4.3	Ripple reduction for large bucket offsets	94
5.5	Ripple-reduction simulations	96
5.5.1	Dependence on RF harmonic	96
5.5.2	Dependence on RF voltage	98
5.5.3	Dependence on ripple frequency	100
5.6	Low-intensity implementation with the 200 MHz system	100
5.7	High-intensity implementation with the 800 MHz system	103
5.7.1	Narrow-band ripple	104
5.7.2	Broad-band ripple	104
5.7.3	Long-term behaviour	106
5.8	Additional considerations	108
5.8.1	RF structure	108
5.8.2	Integrated Intensity	110
5.8.3	Transverse effects & beam loss	113

5.9	Operational implementation	115
5.10	Conclusion	116
6	Conclusion	118
A	Pulsed extraction for FLASH radiotherapy	122
A.1	RF phase displacement for pulsed extraction	122
A.1.1	Concept	123
A.1.2	Relevant parameters	124
A.1.3	Longitudinal and transverse transit times	127
A.1.4	Simulation studies	128
A.1.4.1	Single-burst extraction	128
A.1.4.2	Multiple-burst extraction	130
A.1.5	Measurements in the PS	134
A.2	Conclusion and outlook	138
B	Empty-bucket channelling with double-harmonic systems	140
B.0.1	Theory	140
B.0.2	Simulation: grid-scan	142
B.0.3	Simulation: Optimiser	145
C	Empty-bucket channelling with broadband cavities	148
C.1	Measurements in the PS	148
D	Some non-linear phenomena in empty-bucket channelling	151
D.1	Weak bunching	151
D.2	Narrowing of extracted separatrix	152
D.3	Ripple suppression: two minima	154
	Bibliography	159

List of Figures

1.1	Two synchrotrons, past and present	2
1.2	Original extraction-region sketch for a synchro-cyclotron	4
1.3	Schematic of the CERN accelerator complex	5
1.4	Original extraction-region sketch for the CERN PS	5
1.5	Layout of MedAustron facility	7
1.6	Two examples of facilities that benefit from CERN's slow extraction programme	7
2.1	Local coordinate system used for synchrotrons	11
2.2	Transverse phase space	13
2.3	Normalised phase space	14
2.4	Turn-by-turn evolution of a particle in normalised phase space when the lattice contains a single dipole error	16
2.5	Turn-by-turn evolution of a particle in normalised phase space when the lattice contains a single sextupole error	18
2.6	Kobayashi-Hamiltonian contours	20
2.7	Steinbach diagram	20
2.8	Normalised phase space at septum location	22
2.9	Stationary buckets, Hamiltonian contours in longitudinal phase space	25
2.10	Accelerating buckets, Hamiltonian contours in longitudinal phase space with acceleration	26
2.11	Accelerating buckets, Hamiltonian contours in normalised longitudinal space.	27
2.12	Energy gain (black) and frequency difference (red) vs. turn number, for a particle undergoing phase-displacement acceleration.	28
2.13	Potential-energy diagram for accelerating bucket.	29
2.14	Phase displacement in longitudinal phase space for a particle ensemble	30
3.1	Schematic of the CERN SPS and the North Area.	31

3.2	Horizontal beta-function along Long Straight Section 2	33
3.3	Normalised-sextupole strengths for all magnets in the SPS ring	33
3.4	Normalised phase space at the electrostatic septum during the SPS slow extraction	34
3.5	SPS bunch rotation in longitudinal phase space	36
3.6	Momentum programme during flattop for SFT cycle.	37
3.7	Steinbach-diagram illustration of COSE	38
3.8	Bunched beam on the 1/3-integer resonance	39
3.9	Two particles undergoing empty-bucket channelling	41
3.10	Parametrisation of empty-bucket channelling illustrated in longitudinal phase space	42
3.11	Particles in longitudinal phase space during empty-bucket channelling in the SPS	44
3.12	Schematic of SPS simulation model used in [49]	45
3.13	Schematic of SPS simulation model employed in this thesis	46
3.14	Simulation pipeline for CERN's HTCondor cluster	47
3.15	Simulation run time, CPU vs. GPU	47
3.16	Simulated slow-extracted spill	50
3.17	Example of slow-extracted particle distribution in longitudinal phase space during empty-bucket channelling	51
3.18	Initial and extracted transverse distributions in normalised coordinates	52
4.1	Schematic of the SHiP experiment's proton target, veto system and scattering detector	53
4.2	Schematic of the spill time structure for the SHiP experiment at the second and nanosecond timescales	54
4.3	Two consecutive neutrino bunches and a HS particle at the proton target and the scattering detector	54
4.4	SHiP's mass-momentum search space for HS particles	55
4.5	Schematics of conventional methods to provide bunched slow extraction	56
4.6	Normalised phase space at the electrostatic septum and extracted spill for a quadrupole-sweep extraction, with RF off and RF on	58
4.7	Conventional bunched extraction v. empty-bucket channelling in lon- gitudinal phase space and normalised phase space	60
4.8	Bunch length vs. bucket offset and stable-phase sine	61
4.9	Extracted bunch length vs. bucket offset or stable-phase sine	62

4.10	Slow-extracted particle distribution in longitudinal phase space vs. bucket offset or stable-phase sine	64
4.11	Steinbach-diagram illustration of how empty-bucket channelling affects the extracted intensity	65
4.12	Integrated intensity vs. stable-phase sine and bucket offset	65
4.13	Integrated intensity vs. bucket offset or stable-phase sine	66
4.15	Separatrix length and width vs. stable-phase sine and bucket offset	68
4.16	Separatrix width and extracted momentum spread vs. bucket offset or stable-phase sine	69
4.17	Extracted time structure with nominal settings and with empty-bucket channelling	70
4.18	Bunch length vs. bucket offset for three different 200 MHz-cavity voltages, measured and simulated.	71
4.19	Average bunch profile vs. bucket offset, measured and simulated	72
4.20	Bunch profile for SHiP candidate, measured and simulated	73
4.21	Integrated intensity vs. bucket offset, measured and simulated	74
4.22	Bunch length vs. integrated intensity, measured and simulated	75
5.1	SPS spill, Poisson spill and crystalline spill	77
5.2	Frequency-spectrum magnitudes of the spill and focusing-quad current	78
5.3	Measurement of duty factors vs. bin-width for an SPS measurement and the corresponding Poisson limit	79
5.4	Diagram of current-to-intensity perturbation propagation	80
5.5	Simplified model of chromatic slow extraction	82
5.6	Steinbach diagram of particles being pushed towards the unstable region under the influence of ripple	83
5.7	Block diagram of slow extraction for timescales slower than the transit time	85
5.8	Block diagram of slow extraction	85
5.9	Illustration of three perturbed spills with varying perturbation sizes	87
5.10	Spill-ripple amplitude vs. tune-ripple amplitude for different ripple frequencies	88
5.11	Nominal SPS transfer function	90
5.12	Empty-bucket channelling for ripple suppression, represented in longitudinal phase space and Steinbach diagram	91
5.13	Block diagram of slow extraction with empty-bucket channelling	92

5.14	Tune speed-up from empty-bucket channelling	93
5.15	Channel-width and speed-up for a representative case of empty-bucket channelling.	94
5.16	Illustration of empty-bucket rippling for a large bucket offset	95
5.17	Illustration of RF slippage along the RF waveform during one ripple period	97
5.18	50 Hz-ripple-reduction coefficient vs. RF harmonic and bucket offset .	98
5.19	50 Hz-ripple-reduction coefficient vs. bucket offset	98
5.20	50 Hz-ripple-reduction coefficient vs. stable-phase sine and bucket offset	99
5.21	50 Hz-ripple-reduction coefficient vs. bucket offset	99
5.22	Ripple-reduction coefficient vs. ripple frequency and bucket offset . .	100
5.23	Ripple-reduction coefficient vs. bucket offset	101
5.24	Ripple-reduction coefficient vs. bucket offset for three different 200 MHz- system voltages, measured and simulated	102
5.25	Illustration of the conflict between the desired working points for beam loading vs. empty-bucket channelling.	104
5.26	Ripple-reduction coefficient vs. bucket offset for 800 MHz system, mea- sured and simulated	105
5.27	Spill structure in time and frequency, measured with the NA62 GTK detector	106
5.28	Broad-band ripple-reduction coefficient for the operational test conditions	107
5.29	Measured spill-by-spill ripple-reduction coefficient before, during and after operational tests	109
5.30	Bunch length vs. RF harmonic and bucket offset	110
5.31	Slow-extracted particle distribution in longitudinal phase space for bucket offset $\Delta_{RF} = 0$ and stable-phase sine $\Gamma = 0.02$ (simulated in <i>henontrack</i>) for the two available harmonics in the SPS: $h = 4620$ (shown in Subfig. 5.31a) and $h = 18480$ shown in Subfig. 5.31b). . . .	111
5.32	Extracted 800 MHz time structure with (on) and without (off) empty- bucket channelling	111
5.33	Integrated intensity and ripple-reduction coefficient vs. bucket offset .	112
5.34	Internally-dumped intensity at the end of the cycle	113
5.35	Measured horizontal beam size and simulated separatrix length/width vs. bucket offset	114
5.36	Measurement of normalised beam loss	114

5.37	Measured ripple-reduction coefficient as a function of RF frequency-offset setting	115
5.38	Histograms of measured spill-ripple amplitudes at three different ripple frequencies, accumulated over six days running with empty-bucket channelling off and on	116
5.39	Measurement of duty factors vs. bin-width for an SPS measurement with empty-bucket channelling and without it	117
A.1	Longitudinal (left) and transverse (right) phase spaces during a single frequency sweep of RF phase displacement extraction.	125
A.2	Example of voltage and frequency programmes for a single burst . . .	126
A.3	Extracted particles n_{out} for single burst scheme with different Γ	129
A.4	Extracted particles n_{out} v. $\bar{\mu}$. The dashed lines show linear fits to the simulation data	130
A.5	Extraction time v. Γ . The dashed lines show τ_L as a function of Γ . .	131
A.6	Dependence of V on n and N for the PS, $h = 8$	131
A.7	Dependence of $\bar{\sigma}$ on n and N for the PS, $h = 8$	132
A.8	Instantaneous intensity profiles for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$	133
A.9	Voltage (top) and frequency (bottom) programmes for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$	133
A.10	Extracted intensity per burst (top) and burst turns (bottom) for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$. The dashed lines show target values.	134
A.11	Longitudinal beam distribution as a function of cycle time during a 25 ms chirp. The heat-map is reconstructed from the ring current monitor.	135
A.12	Initial and optimised programmes followed by the PS 10 MHz system during burst extraction test. Both the RF frequency (top) and the RF voltage (bottom) are varied as a function of time.	135
A.13	Measured extracted intensity $I(t)$ as a function of time, for the initial and optimised burst extractions. The integrated intensity $\int_t I(t')dt'$ is also shown in dotted lines, and the three decimal-point numbers indicate its value after each individual burst.	136
A.14	Measured extracted intensity per burst (top) and burst duration (bottom) for optimised burst scheme.	137

A.15	Extracted intensity I as a function of relative-momentum offset δ . The left subplot shows the measurement, while the right subplot shows the samples generated by using the eCDF procedure from this section.	138
A.16	Comparison of measurement and simulation for optimised multiple-burst scheme. The top subplot shows the extracted intensity as a function of time. The bottom subplot shows the integrated intensity per pulse	139
B.1	Semi-analytically computed potential and longitudinal phase-spaces for different combinations of voltage ratio and relative phase	141
B.2	ϕ -coordinate of unstable fixed points vs. voltage ratio α for two different relative phases.	142
B.3	Open and closed channels illustrated in the potential (top) and longitudinal (bottom) plots.	143
B.4	ϕ -coordinate of unstable fixed points vs. voltage ratio α for a fixed relative phase, showing whether the adjacent RF channel is open or closed.	143
B.5	Slow-extracted particle distribution in longitudinal phase space vs. voltage ratio and relative phase	144
B.6	Bunch length and integrated intensity vs. voltage ratio and relative phase	145
B.7	Regions in α/ψ -space where the double-harmonic setup outperforms the single-harmonic setup	145
B.8	Bunch length vs. integrated intensity for the SPS double-harmonic setup	147
B.9	Slow-extracted particle distribution in longitudinal (left) and normalised transverse (right) spaces for the best SPS double-harmonic configuration	147
C.1	Example of channelling scheme with isolated sine pulses	149
C.2	Evolution of beam time profile during phase displacement using barrier buckets, with the corresponding channel geometry overlaid (black). Particles arrive to flattop bunched in four tight buckets, they debunch for 200 ms and are finally pushed through the barrier bucket channels by ramping the bending field.	150
D.1	Slow-extracted bunch length vs. bucket offset, highlighting ‘strong bunching’, ‘partial debunching’ and ‘weak bunching’	151

D.2	Slow-extracted particle distribution in longitudinal phase space for three bucket offsets: strong bunching, partial debunching and weak bunching	152
D.4	Slow-extracted particle distribution in longitudinal phase space and normalised phase space, illustrating narrowing of extracted separatrix	155
D.6	50 Hz-ripple-reduction coefficient vs. RF harmonic and bucket offset .	157
D.7	Schematic of longitudinal phase space during empty-bucket channelling, including trajectories that enter the resonant region and slip with respect to the RF bucket	158

Chapter 1

Introduction


Particle accelerators – machines designed to accelerate charged particles to high energies – have become a crucial piece of technology since their inception, with widespread applications across a multitude of disciplines. From assisting in the discovery of almost all fundamental particles to advancing material science and nanotechnology, accelerators have revolutionized our understanding of matter and its interactions. Moreover, their impact extends to medicine, where they play a pivotal role in cancer treatment and medical imaging. In industry, accelerators aid in materials analysis, quality control, and non-destructive testing, ensuring the reliability and safety of critical components. Even spacecraft benefit from accelerator technology, which assists in studying the behaviour of electronics when irradiated by high-energy particle showers.

1.1 History of accelerators

The journey of accelerators began in the early 20th century when physicists sought ways to probe the innermost structure of matter. In 1920, Rolf Widerøe proposed the first design for a linear accelerator (or linac) [1], which accelerated particles in a straight trajectory by using a series of alternating electric fields. In 1930, Ernest O. Lawrence pioneered the cyclotron [2], which combined the accelerating electric fields with a fixed magnetic field, constraining particles into spiral trajectories of increasing radii. Around the same time, key technologies for high-voltage electrostatic acceleration were developed, namely the ‘Van de Graaff’ [3] and the ‘Cockcroft-Walton’ [4] generators.

In the following decades, accelerators evolved rapidly, giving birth to synchrotrons in the 1940s. Synchrotrons represented a leap forward in accelerator technology: by

synchronously ramping the bending magnetic field with the particle energy, trajectories were constrained to a constant radius. The first high-energy proton synchrotron was the Cosmotron at Brookhaven National Laboratory (BNL), which reached energies of 3.3 GeV in 1952. In 1949, Nicholas Christofilos patented the concept of Alternating Gradient focusing, which was independently discovered by Ernest Courant, M. Stanley Livingston and Hartland Snyder in 1952, allowing for tight control of particle paths. The first synchrotrons to incorporate this concept were BNL's Alternating Gradient Synchrotron and CERN's Proton Synchrotron (PS) in 1959. Since then, the quest for synchrotrons of ever increasing energy continues, with CERN's Large Hadron Collider (LHC) holding the current record at 6.8 TeV. Figure 1.1 shows both the Cosmotron and the LHC, visually illustrating the progress in the field over roughly 60 years.



(a) Cosmotron (image from BNL's website)



(b) LHC (image from CERN's website)

Figure 1.1: Two synchrotrons, past and present: the Cosmotron (1953) and the LHC (2008).

Nowadays, synchrotrons are particularly crucial for high-energy hadron beams

and electron light sources. Indeed, particles can be accelerated in a compact periodic orbit by repeatedly using the same accelerating electric fields (unlike in linacs), only limited by the maximum magnetic field that can be provided to bend the beam ¹. Furthermore, synchrotrons can easily provide a wide range of particle energies by scaling their bending fields.

1.2 Slow extraction

An important aspect of proton synchrotron operation is beam extraction. When the desired energy and beam parameters are reached inside the ring, the beam is often ejected to be transported to the users. This task has occupied physicists and engineers for several decades, producing a range of solutions depending on the specific application.

The most common extraction method is dubbed ‘fast extraction’. It typically involves employing a specialized kicker magnet, which quickly switches on and steers the beam out of the ring over one revolution period. However, for applications where the beam must be extracted over timescales much longer than the revolution period, ‘slow extraction’ is employed. By exciting resonant behaviour in a controlled fashion, particle trajectories grow in radial amplitude and the whole beam can be continuously extracted over a few milliseconds to several hours.

1.2.1 Inception and early history

The idea to provoke unstable particle motion for beam extraction dates back to the 1950s. In 1950, J. L. Tuck and L. C. Teng proposed a magnet arrangement (shown in Fig. 1.2) that would aid extraction from a synchro-cyclotron by exciting radial amplitude growth over the last revolution period [5], which was later implemented in the Liverpool synchro-cyclotron by J. K. Le Couteur in 1955 [6]. In 1961, H.G. Hereward described the possibility of employing a similar approach at the CERN PS by exploiting the radial integer resonance [7]. In the same year, C.L. Hammer and L.J. Laslett extended the idea to using the 1/2-integer resonance as a general method for extraction from alternating-gradient synchrotrons [8]. Finally, in 1966, S. Strolin studied the use of the radial 1/3-integer resonance, and developed a basic theoretical framework to understand the critical aspects of such methods [9].

¹For electrons, the emission of synchrotron radiation due to bending becomes the key limiting factor, as the power necessary to compensate this loss increases with the fourth power of the beam energy. Therefore, linacs become an attractive alternative.




Figure 1.2: Original extraction-region sketch for a synchro-cyclotron (from [6]). The peeler and regenerator magnetic fields H_s change linearly with the orbit radius r . H_s can be tuned to make particles' radial motion unstable while keeping the vertical motion stable, directing trajectories into the magnetic channel.

Nowadays, 1/3-integer resonant extraction has become the work-horse of fixed-target experimental physics, radiation-effects testing and hadron radiotherapy, as it can easily provide particle spills of adjustable duration, energy and beam size. Unlike the integer and 1/2-integer resonances, the 1/3-integer resonance is not excited directly by the linear lattice (i.e. dipoles and quadrupoles), which allows for fine control of the resonance strength by employing dedicated sextupole magnets. All higher order resonances (e.g. 1/4-integer) share the same advantage, but the setup becomes increasingly complicated without providing any obvious benefit.

1.2.2 CERN's journey

The CERN accelerator complex, situated near Geneva, Switzerland, stands as the world's largest facility for particle-physics research. Figure 1.3 shows a schematic of the complex, which consists of a series of interconnected accelerators, each designed to perform specific tasks in the journey to achieve unprecedented energies and intensities for particle beams.

CERN has a rich history of designing and hosting slow-extraction systems. The aforementioned 1961 internal report from Hereward [7] triggered an initiative to implement integer-resonance slow extraction in the CERN PS. Following a similar strategy to the synchro-cyclotron arrangement from Tuck and Teng, a special ejection magnet was designed, which would kick the resonant beam out, while leaving the circulating beam unaffected. Such a system was dubbed a 'septum' and the original 1962 drawing is shown in Fig. 1.4. By 1964, the PS was extracting proton pulses of several




Figure 1.3: Schematic of the CERN accelerator complex (image adapted from CERN’s website).

milliseconds using such method. Already in the first update report [10], Hereward identified what would later become one of the critical technical challenges in the field of slow extraction: *The instantaneous spill rate is very sensitive to magnet-voltage ripple, and with the 600 Hz ripple of our unfiltered ‘flat-top’ the particles spill in short bursts at 1.7 ms intervals.*




Figure 1.4: Original extraction-region sketch for the CERN PS from [11]. The extraction magnet (*Aimant d'éjection*) sits between main units 60 and 61 and separates the beam pipe into a field-free region for the equilibrium orbit (*Orbite d'équilibre*) and a field region that kicks large-amplitude particles out of the machine.

In the 1970s, the proposal to build the Super Proton Synchrotron (SPS) opened up a new realm of possibilities for slow-extraction implementations. Unlike the PS, which combined its bending and focusing magnetic fields into a single unit, the SPS would have separate magnets for each function, increasing the flexibility of the system.

Moreover, the substantial increase in beam energy (initially up to 300 GeV) posed new challenges in terms of beam loss and activation. In 1970, all such aspects and concerns were detailed and theoretically addressed [12], and by 1977 the tests in the SPS [13] demonstrated successful resonant extraction with the integer, the 1/2-integer and the 1/3-integer resonance, providing a variety of spill lengths from 2 ms to 2 s.

Another important chapter in the development of slow-extraction systems is related to the Low-Energy Antiproton Ring (LEAR). The experimental request called for a particle spill lasting approx. 1000 s. With a particle count of roughly 1×10^9 and a revolution period of 1 μ s this equates to extracting one antiproton per turn. W. Hardt's 1979 report [14] readily identified the main technical challenge: *to reduce the sensitivity of the spill modulation versus magnetic ripple*. A solution was found by exploiting the stochastic extraction proposed by S. van der Meer in 1976 [15], which used the Radio-Frequency (RF) cavities conventionally used for acceleration to impart random kicks onto the circulating beam until particles wandered into resonant motion. The LEAR period is notable for the invention of innovative RF-based solutions to deal with slow-extraction problems [16] [17] [18].

In the late 1990s, CERN's theoretical and technical expertise on slow-extraction systems was channelled towards building a medical synchrotron for cancer treatment. In 1999, the Proton-Ion Medical Machine Study (PIMMS) project delivered a two-volume technical report [19] [20] that detailed the accelerator-physics aspects of a facility capable of performing hadron therapy. A substantial portion of the document is dedicated to characterising and improving spill quality. The PIMMS became the blueprint for two major hadron-therapy centres in Europe, namely CNAO (Centro Nazionale di Adroterapia Oncologica) in Italy and MedAustron in Austria. These centers followed the pioneering efforts of GSI's Heidelberg Ion-beam Therapy (HIT) Centre by providing carbon-ion therapy within dedicated facilities [21]. Figure 1.5 shows MedAustron's facility layout: a synchrotron of approximately 20 m in diameter delivers proton and carbon-ion beams to four irradiation rooms via slow extraction. Each spill lasts between 1 s and 10 s.

Currently, a sizeable portion of CERN's users still rely on beam delivered via slow extraction. Specifically, the PS and the SPS are still used for providing slow-extracted proton and ion beams to the East and North Areas, respectively. On top of developing and operating slow-extraction systems for the PS and the SPS, CERN is also actively engaged in collaborations with medical applications that employ the same technique. For example, the Next Ion Medical Machine Study (NIMMS) [23] aims at developing *a new generation of compact and cost-effective ion-therapy facilities*. As




Figure 1.5: Layout of MedAustron facility (from [22]), showing the synchrotron hall (SH), injector hall (IH) and the four irradiation rooms IR1-IR4. (The vertical beam line in IR2 is a projection).

a consequence of such a wide portfolio of accelerators, CERN's slow-extraction R&D programme is rich and varied, addressing powerful and flexible machines as well as compact and reliable ones. This breath of scope is illustrated in Fig. 1.6 by showing two vastly different destinations for slow-extracted beams, namely MedAustron's irradiation room and the North Area 62's (NA62) experimental hall.



(a) Irradiation room (image from MedAustron's website)

(b) NA62 experimental hall (image from NA62's website)

Figure 1.6: Two examples of facilities that benefit from CERN's slow-extraction programme: the MedAustron hadron-therapy centre and the North Area 62 (NA62) kaon-physics experiment.

1.2.3 The future

Slow-extraction challenges still remain, as user requests become increasingly demanding in terms of energy, intensity, flexibility and precision. For instance,

- in fundamental physics, experimental users request an increasing number of protons-on-target (p.o.t.) per year. This will require higher extracted intensities while ensuring excellent beam quality and low beam loss. At CERN, the possibility of a high-intensity 400 GeV facility in the North Area is being discussed [24], capable of delivering 5×10^{19} p.o.t. per year to the ECN3 experimental hall.
- on the irradiation front, the increasing number of applications necessitates irradiation-to-electronics facilities that can provide easily variable beam intensities and energies, to simulate the wide spectrum of particle showers. For example, CERN has recently engaged in the HEARTS/CHIMERA project, which will investigate *the feasibility of delivering high energy ion beams [...] for the study of radiation effects to electronics components engineered to operate in harsh radiation environments, such as space or high-energy accelerators* [25].
- in medicine, novel therapies are looking to utilise short pulses of particles [26] (FLASH therapy) and beams with mixed particle species [27], as well as aiming to perform conventional hadron therapy faster and at a lower cost. CERN directly contributes to the R&D of such initiatives through a variety of collaborations and work packages.

In short, CERN's slow-extraction programme aims to play a role in a wide variety of topics and must therefore study innovative yet robust technical solutions to address the challenges ahead.

1.3 Thesis scope

This thesis studies the potential of RF-based techniques for CERN's future slow-extracted beams. We start by introducing the necessary theory for resonant slow extraction and RF manipulations in Chap. 2. In Chap. 3, this theory is then applied to describe the implementation of an RF technique known as empty-bucket channelling in the SPS, as well as to construct the modelling tools that will be utilised throughout the rest of the thesis. By exploiting those tools, empty-bucket channelling is studied

both to satisfy a Physics-Beyond-Colliders request (Chap. 4) and to improve the spill quality of operational systems at CERN (Chap. 5). Finally, Chap. 6 provides a conclusion and points towards possible future work.

Chapter 2

Beam dynamics

This chapter summarises the beam-dynamics concepts and equations required for this thesis. The information has been extracted and adapted from various sources. For a comprehensive treatment of beam dynamics, the reader is referred to [28] [29] [30] [31].

Throughout this chapter, simulation outcomes are used to illustrate various theoretical concepts. The simulations have been performed with a custom *python* package called *henontrack*. For more details, see Chap. 3.

2.1 Reference system

In order to mathematically describe proton trajectories in a particle accelerator, a reference system needs to be chosen. First, a reference orbit is defined, which is given by the path followed by a proton that has the accelerator's central design momentum p_0 . Then, a local coordinate system is adopted, where positions are described with respect to the reference orbit as shown in Fig. 2.1. In this system, a particle's position is uniquely determined by the 3-tuple (x, y, s) , where s is the distance along the reference orbit, x is the horizontal displacement and y is the vertical displacement. For a circular machine with bending magnetic field B and bending radius ρ , the reference momentum p_0 is given by,

$$\frac{p_0}{e} = B\rho, \tag{2.1}$$

where e is the electron charge. $B\rho$ is often referred to as the 'beam rigidity'. Positive x points towards the outside of the ring, positive y points upwards and positive s traverses the ring clockwise. With these conventions, the unit vectors $\vec{x}, \vec{y}, \vec{s}$ span an orthonormal right-handed coordinate system.




Figure 2.1: Local coordinate system used for synchrotrons (image taken from [32]). The unit vectors \vec{x} , \vec{y} , \vec{s} are the basis vectors of the coordinate system, ρ is the bending radius and θ is the azimuthal displacement.

2.2 Transverse dynamics

2.2.1 Hill's equation

Using s as the independent variable ('time coordinate'), the evolution of the transverse coordinate x is now studied for a particle with reference momentum. An equivalent treatment can be followed for y . Under the influence of linear focusing forces provided by quadrupoles, the evolution of $x(s)$ can be modelled by Hill's equation:

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

where $K(s) = \frac{1}{\rho^2(s)} + K_1(s)$ is the focusing term, with $K_1 = \frac{1}{B\rho} \frac{\partial B}{\partial y}|_{x=0, y=0}$ being the quadrupole component. The solution to this equation can be written as:

$$x(s) = A_x \sqrt{\beta_x(s)} \cos[\mu_x(s) + \mu_0], \quad (2.3)$$

where $\beta_x(s)$ is the 'beta-function', $\mu_x(s) = \int_0^s d\lambda / \beta_x(\lambda)$ is the phase advance, and A_x and μ_0 are the normalised amplitude and initial phase of a given particle, respectively. Typically, β_x and μ_x are numerically computed for a given magnetic lattice using an accelerator code such as the Methodical Accelerator Design X (MADX) [33]. It is customary to define two additional functions for convenience:

$$\alpha_x = -\frac{1}{2} \frac{d\beta_x}{ds}, \quad \gamma_x = \frac{1 + \alpha_x^2}{\beta_x}. \quad (2.4)$$

Note that $\beta_x, \alpha_x, \gamma_x$ must satisfy the periodicity condition $f(s+C) = f(s)$, where C is the ring circumference.

2.2.2 Transverse phase space

In addition to the position coordinate $x(s)$, it is useful to define the angle coordinate $x'(s) = \frac{dx}{ds}$. Then, Hill's equation can be re-written as a set of two first-order differential equations that can be expressed in matrix form:

$$\frac{d}{ds} \begin{bmatrix} x \\ x' \end{bmatrix} + \begin{bmatrix} 0 & -1 \\ K(s) & 0 \end{bmatrix} \begin{bmatrix} x \\ x' \end{bmatrix} = 0. \quad (2.5)$$

The (x, x') 2-tuple forms a phase space and, in similar fashion to Eq. 2.3, the solution $\vec{x}(s) = [x(s), x'(s)]^T$ for a particular initial condition $\vec{x}(0)$ can be expressed as,

$$\vec{x}(s) = M_{0,s} \cdot \vec{x}(0), \quad (2.6)$$

where the transport matrix $M_{0,s}$ is given by

$$M_{0,s} = \begin{bmatrix} \sqrt{\frac{\beta_x(s)}{\beta_x(0)}} [\cos \mu_x(s) + \alpha_x(0) \sin \mu_x(s)] & \sqrt{\beta_x(s)\beta_x(0)} \sin \mu_x(s) \\ \frac{[\alpha_x(0) - \alpha_x(s)] \cos \mu_x(s) - [1 + \alpha_x(0)\alpha_x(s)] \sin \mu_x(s)}{\sqrt{\beta_x(s)\beta_x(0)}} & \sqrt{\frac{\beta_x(0)}{\beta_x(s)}} [\cos \mu_x(s) - \alpha_x(s) \sin \mu_x(s)]. \end{bmatrix} \quad (2.7)$$

In a periodic lattice, the one-turn map $M_{s,s+C}$ can be computed, which maps a given initial condition at a specific location s to the final condition after one turn around the ring:

$$M_{s,s+C} = \begin{bmatrix} \cos 2\pi Q_x + \alpha_x(s) \sin 2\pi Q_x & \sin 2\pi Q_x \\ -\gamma_x(s) \sin 2\pi Q_x & \cos 2\pi Q_x - \alpha_x(s) \sin 2\pi Q_x \end{bmatrix}, \quad (2.8)$$

where $Q_x = \frac{1}{2\pi} \int_s^{s+C} \frac{d\lambda}{\beta_x(\lambda)}$ is known as the (horizontal) betatron tune and corresponds to the number of transverse oscillations performed by a particle in one turn. Q_x is a global quantity that does not depend on the particular choice of s . Figure 2.2 shows the resulting surface of section at a fixed s after repeatedly applying the one-turn map. For a given initial condition with amplitude A_x all points lie on an ellipse given by,

$$A_x^2 = \gamma_x x^2 + \beta_x x'^2 + 2\alpha_x x x', \quad (2.9)$$

which is known as the Courant-Snyder invariant.




Figure 2.2: Turn-by-turn evolution (sky-blue) of an initial condition in transverse phase space at two different SPS locations. The numbers indicate the turn number.

2.2.3 Normalised phase space

The solution from Eq. 2.3 has the form of a ‘generalised harmonic oscillator’, whose amplitude and frequency vary with $\sqrt{\beta_x}$ and $1/\beta_x$, respectively. These dependencies can be removed by choosing μ_x as the independent variable (instead of s) and by defining the normalised position X as,

$$X(\mu_x) = \frac{1}{\sqrt{\beta_x}}x = A_x \cos(\mu_x + \mu_0). \quad (2.10)$$

Then, the corresponding normalised-angle variable X' is,

$$X' = \frac{dX}{d\mu} = -A_x \sin(\mu_x + \mu_0) \quad (2.11)$$

$$X' = \frac{dX}{ds} \frac{ds}{d\mu_x} = \beta_x \frac{dX}{ds} = \sqrt{\beta_x}x' + \frac{\alpha_x}{\sqrt{\beta_x}}x. \quad (2.12)$$

The 2-tuple (X, X') forms the (horizontal) normalised phase space. Using these transformations, the matrix $M_{s,s+C}$ can be decomposed as follows:

$$M_{s,s+C} = L(s)R(2\pi Q_x)L^{-1}(s), \quad (2.13)$$

where $R(2\pi Q_x)$ is the 2D-rotation matrix by an angle $2\pi Q_x$:

$$R(2\pi Q_x) = \begin{bmatrix} \cos 2\pi Q_x & \sin 2\pi Q_x \\ -\sin 2\pi Q_x & \cos 2\pi Q_x \end{bmatrix}, \quad (2.14)$$

and $L^{-1}(s)$ and $L(s)$ are known as the normalisation matrices:

$$L^{-1}(s) = \begin{bmatrix} \frac{1}{\sqrt{\beta_x(s)}} & 0 \\ \frac{\alpha_x(s)}{\sqrt{\beta_x(s)}} & \sqrt{\beta_x(s)} \end{bmatrix}, \quad L(s) = \begin{bmatrix} \sqrt{\beta_x(s)} & 0 \\ -\frac{\alpha_x(s)}{\sqrt{\beta_x(s)}} & \frac{1}{\sqrt{\beta_x(s)}} \end{bmatrix}. \quad (2.15)$$

In the normalised phase space, the elliptical surfaces of section from Fig. 2.2 become circles as shown in Fig. 2.3. Furthermore, the Courant-Snyder invariant is now given by,

$$A_x^2 = X^2 + X'^2. \quad (2.16)$$




Figure 2.3: Turn-by-turn evolution (skyblue) of an initial condition in normalised phase space at two different SPS locations. The numbers indicate the turn number.

Finally, the difference equation for $\vec{X} = [X, X']^T$ after k turns can be easily computed using *de Moivre's* formula:

$$\vec{X}_{n+k} = R^k(2\pi Q_x) \vec{X}_n = R(2\pi Q_x k) \vec{X}_n. \quad (2.17)$$

2.2.4 Chromaticity

So far, our discussion has been restricted to particles with reference momentum p_0 . In reality, a given particle i has momentum $p_i = p_0(1 + \delta_i)$, where δ (or $\Delta p/p$) is the relative momentum offset. A particle with $\delta_i \neq 0$ has a different magnetic rigidity and experiences a different kick when passing through a magnet in the lattice. Among other effects, the different kicks result in a different betatron tune $Q_{x,i}$, which is given by,

$$Q_{x,i} = Q_x + Q'_x \cdot \delta + O(\delta^2), \text{ with } Q'_x = \left. \frac{\partial Q_x}{\partial \delta} \right|_{\delta=0}, \quad (2.18)$$

where Q'_x is the chromaticity of the lattice, and can be computed numerically with MADX.

2.2.5 Emittance

For a particle ensemble, it is useful to define the emittance ϵ_x as the area occupied by the ensemble in phase space. Under conservative forces and no acceleration, the emittance is a conserved quantity. In practice, the emittance is typically characterised by its root-mean-squared (rms) estimation $\epsilon_{x,rms}$, which is given by,


$$\epsilon_{x,rms} = \sqrt{\det\{Cov(x, x')\}} = \sqrt{Var(X)Var(X')}, \quad (2.19)$$

where $Cov(a, b)$ and $Var(a)$ are the covariance matrix of a, b and the variance of a , respectively.

2.2.6 Transverse resonance

Stable motion can be disrupted by magnetic perturbations that excite resonant behaviour. Resonances are a characteristic phenomenon of oscillating systems: when a periodic perturbation's frequency matches the system's eigenfrequency, a coherent excitation builds up over time. In a synchrotron, the distribution of magnetic perturbations can be in harmony with a particle's transverse oscillations, causing kicks to add up over many turns. For example, Fig 2.4 illustrates the excitation of an integer resonance due to a dipole error. Since the particle's tune $Q_{x,i}$ is close to an integer, the perturbation results in a turn-by-turn increase of A_x .

Typically, the lattice is designed to minimise resonant behaviour, keeping the beam stably circulating inside the beam pipe.



(a) Resonant, $Q_{x,i} = n$, with $n \in \mathbb{Z}$ (b) Non-resonant, $Q_{x,i} \neq n$, with $n \in \mathbb{Z}$

Figure 2.4: Turn-by-turn evolution of a particle in normalised phase space when the lattice contains a single dipole error of normalised strength $d = 1 \text{ m}^{1/2}$ (simulated in *python*). The arrows show the turn-by-turn kick $\Delta X' = d$. The particle's initial condition is $\vec{X}_0 = [1, 0]^T$.

2.3 1/3-integer resonance

We now explore the theory of the 1/3-integer (third-order) resonance, which is productively employed in resonant slow extraction to eject the beam from the synchrotron over hundreds-to-thousands of turns. This theory is extensively described in [34] [19].

2.3.1 Sextupole magnet

The 1/3-integer resonance can be driven by sextupolar fields. A horizontal sextupole magnet generates transverse magnetic fields B_x, B_y given by,

$$B_x = \frac{\partial^2 B_y}{\partial x^2} \Big|_{x=0, y=0} xy, \quad (2.20)$$

$$B_y = \frac{1}{2} \frac{\partial^2 B_y}{\partial x^2} \Big|_{x=0, y=0} (x^2 - y^2). \quad (2.21)$$

Under the thin-lens approximation, the effect of a sextupole can be expressed as follows in normalised coordinates:

$$\Delta X = 0, \quad \Delta X' = S \left(X^2 - \frac{\beta_y}{\beta_x} Y^2 \right), \quad (2.22)$$

$$\Delta Y = 0, \quad \Delta Y' = -2S \frac{\beta_y}{\beta_x} XY, \quad (2.23)$$

where S is the normalised sextupole strength:

$$S = \frac{1}{2} \beta_x^{3/2} K_2 L, \quad (2.24)$$

where L is the sextupole length and K_2 is the sextupole component given by,

$$K_2 = \frac{1}{B\rho} \left. \frac{\partial^2 B_y}{\partial x^2} \right|_{x=0, y=0}. \quad (2.25)$$

Equation 2.23 shows that a sextupole couples the X, Y motions. However, for horizontal extraction systems, the sextupole is located in a region where β_y/β_x is small and the vertical contribution to the extraction dynamics can typically be ignored. Then, the difference equation for \vec{X} in a linear lattice with a single thin sextupole may be written as,

$$\vec{X}_{n+1} = R(2\pi Q_{x,i}) \left\{ \vec{X}_n + S \begin{bmatrix} 0 \\ X_n^2 \end{bmatrix} \right\} \quad (2.26)$$

where $Q_{x,i}$ is computed using Eq. 2.18 to first order.


When a lattice has more than one sextupole, the magnitude S and phase μ of an equivalent virtual sextupole (resonance-driving term) can be computed as,

$$S^2 = \left[\sum_j S_j \sin(3\mu_{x,j}) \right]^2 + \left[\sum_j S_j \cos(3\mu_{x,j}) \right]^2 \quad (2.27)$$

$$\tan(3\mu) = \frac{\sum_j S_j \sin(3\mu_{x,j})}{\sum_j S_j \cos(3\mu_{x,j})}, \quad (2.28)$$

where the sum is performed over all sextupoles in the lattice.

Figure 2.5 shows the normalised-phase-space orbit of a particle following the map from Eq. 2.26. When the particle's tune is equal to $n/3$ with $n \in \mathbb{Z}$ the sextupolar kicks add coherently; when the particle's tune is not $n/3$, the kicks end up cancelling out when summed over several turns. Notice that the addition of kicks for the $1/3$ -integer resonance is more subtle than in the case of the integer resonance (Fig. 2.4), as the net increase in A_x is only visible after three turns. Regardless, the resonant particle will grow in amplitude until it is likely loss somewhere along the synchrotron aperture.



(a) Resonant, $Q_x = \frac{n}{3}$, with $n \in \mathbb{Z}$ (b) Non-resonant, $Q_x \neq \frac{n}{3}$, with $n \in \mathbb{Z}$

Figure 2.5: Turn-by-turn evolution of a particle in normalised phase space when the lattice contains a single sextupole error of normalised strength $S = 0.5 \text{ m}^{-1/2}$ (simulated in *python*). The arrows show the turn-by-turn kick $\Delta X' = SX^2$.

2.3.2 Three-turn formalism

Consider a particle with betatron tune given by,

$$Q_{x,i} = m \pm \frac{1}{3} + \Delta Q, \quad \text{for } m \in \mathbb{N}^+, \quad (2.29)$$

where $|\Delta Q| = |Q_x - \frac{n}{3} + Q' \cdot \delta_i| \ll 1$ for some $n \in \mathbb{Z}$. Then, the fractional part of $Q_{x,i}$ is close to $1/3$ (or $2/3$) and $3 \cdot Q_{x,i}$ is close to an integer. For a purely linear lattice, the change in \vec{X} after three turns $\Delta \vec{X}_3$ can be computed using Eq. 2.17 and is given by,

$$\vec{X}_3 = R(6\pi Q_{x,i}) \vec{X}_0 = I \cdot R(6\pi \Delta Q) \vec{X}_0 = \left(I + \begin{bmatrix} 0 & 6\pi \Delta Q \\ -6\pi \Delta Q & 0 \end{bmatrix} \right) \vec{X}_0 + O(\Delta Q^2), \quad (2.30)$$

$$\Delta \vec{X}_3 = 6\pi \begin{bmatrix} 0 & \Delta Q \\ -\Delta Q & 0 \end{bmatrix} \vec{X}_0 + O(\Delta Q^2), \quad (2.31)$$

where I is the identity matrix. When a single thin-sextupole is introduced, the same procedure can be applied to Eq. 2.26, which adds a quadratic term (underlined below) to Eq. 2.31:

$$\Delta \vec{X}_3 = 6\pi \begin{bmatrix} 0 & \Delta Q \\ -\Delta Q & 0 \end{bmatrix} \vec{X}_0 + \frac{3S}{4} \begin{bmatrix} 2X_0 X'_0 \\ X_0^2 - X'^2_0 \end{bmatrix} + O(\Delta Q^2, X^2) \quad (2.32)$$

Since the slow-extraction dynamics are typically much slower than three turns, one may take the three-turn time T_{3n} as the fundamental unit of time and convert

the difference equation in Eq. 2.32 into differential equations (excluding higher-order terms):

$$\frac{dX}{d(T_{3n})} = 6\pi\Delta Q X' + \frac{3}{2}SXX' \quad (2.33)$$

$$\frac{dX'}{d(T_{3n})} = -6\pi\Delta Q X + \frac{3}{4}S(X^2 - X'^2), \quad (2.34)$$

which can be derived from the Kobayashi Hamiltonian [35] H_{3n} :

$$H_{3n} = \frac{6\pi\Delta Q}{2}(X^2 + X'^2) + \frac{S}{4}(3XX'^2 - X^3), \quad (2.35)$$

using Hamilton's equations:

$$\frac{dX}{d(T_{3n})} = \frac{\partial H_{3n}}{\partial X'} \quad (2.36)$$

$$\frac{dX'}{d(T_{3n})} = -\frac{\partial H_{3n}}{\partial X}. \quad (2.37)$$

This procedure has 'smoothed out' the equations describing the dynamical system, which makes them more amenable to analytical manipulations.

2.3.3 The stable and unstable regions

For a given tune distance ΔQ , it can be shown that the Kobayashi Hamiltonian from Eq. 2.35 factorises into three straight lines when $H_{3n} = (4\pi\Delta Q)^3/S^2$. These three straight trajectories are known as separatrices and split normalised phase space into a stable triangular region and an unstable region that surrounds it, as shown in Fig. 2.6. Particles in the stable region perform bounded periodic motion, while particles in the unstable region grow in normalised amplitude along three of the six separatrix arms.

Figure 2.6 shows that as $|\Delta Q|$ is reduced, the size of the stable region decreases. The stable area will eventually disappear altogether when $\Delta Q = 0$. More precisely, the limiting triangle can be shown to have an area in phase space that is given by

$$Area = \frac{48\sqrt{3}\pi}{S^2}(\Delta Q)^2\pi. \quad (2.38)$$

If one starts with a linear lattice and the sextupole is turned on adiabatically, normalised-phase-space trajectories for particles near the origin will deform from circles to the new configuration shown in Fig. 2.6, preserving the enclosed area. Therefore, one may equate the area of the triangle to the area enclosed by a trajectory in




Figure 2.6: Kobayashi-Hamiltonian contours, showing the normalised phase space near the $1/3$ -integer resonance. The separatrices (pink) trace out a triangular region, with the black arrows indicating the direction of motion along each arm.

a linear lattice ($Area = A_x^2\pi$) to find the limiting normalised amplitude:

$$A_x = (48\sqrt{3}\pi)^{1/2} \frac{\Delta Q}{S}. \quad (2.39)$$

This equation defines the boundary between stable and unstable motion in $(\Delta Q, A_x)$ -space and can be visualised using the Steinbach diagram as shown in Fig. 2.7.




Figure 2.7: Stability diagram near the $1/3$ -integer resonance in $(\Delta Q, A_x)$ -space, often referred to as the ‘Steinbach diagram’. Particle a is unstable, while particles b and c are stable because they have a larger $|\Delta Q|$ and smaller A_x , respectively.

2.3.4 Extraction with a septum magnet

It has been demonstrated that A_x growth can be excited using a sextupole magnet, and that the boundary between stability and instability can be described in terms of controllable parameters, namely S and ΔQ . By varying these parameters over time, stable trajectories can be made unstable, pushing particles into resonance in a predictable fashion. That is essentially how slow extraction is performed.

Still, a final element is required to successfully extract the beam out of the synchrotron: a septum magnet. A septum is a magnet with thin wires that split the beam pipe into two regions: a zero-field region where the circulating beam passes unperturbed, and a high-field region where large-amplitude particles are kicked towards the extraction channel. The septum is located in a favorable s -location in the ring, so that the wires become the first aperture restriction encountered by unstable particles.

Figure 2.8 shows the normalised phase space at the septum location for five unstable particles about to be extracted. Particles hop from the neighbourhood of one separatrix arm to another, gaining enough amplitude in their last three turns to jump over the septum wires into the high-field region. A few percent of the particles may impinge in the septum wires and get lost.

For the idealised case where the septum blade has zero thickness and zero length, the septum can be modelled as follows:

$$\begin{cases} \text{keep tracking,} & \text{if } X_i < X_{septum} \\ \text{extract,} & \text{if } X_i \geq X_{septum}, \end{cases} \quad (2.40)$$

where $X_{septum} = x_{septum}/\sqrt{\beta_{x,septum}}$ is the normalised aperture restriction, with x_{septum} representing the wires' distance from the reference orbit and $\beta_{x,septum}$ is the horizontal beta-function at the septum location.

2.4 Longitudinal dynamics

The reference particle, which follows the reference orbit and has design momentum p_0 , has a revolution frequency f_0 . Another particle i will have relative momentum offset $\delta_i = (p_i - p_0)/p_0$ and relative time-of-arrival τ with respect to the moving frame established by the reference particle.




Figure 2.8: Turn-by-turn evolution of particles in normalised phase space at the septum location (simulated in *henontrack*). Five particles (each in a different color) gain amplitude over three turns due to the sextupole kicks. In this particular setup the fractional part of Q_x is $\frac{2}{3}$, so particles rotate clockwise by 240 degrees each turn. The blue contours and the pink separatrices show the underlying Kobayashi Hamiltonian, and the vertical grey band represents the septum wires.

2.4.1 RF cavity

In synchrotrons, acceleration is performed by employing sinusoidally time-varying electric fields provided by RF cavities. The voltage provided by the accelerating mode of an RF cavity can be modelled as

$$V(t) = V_0 \sin(2\pi f_{RF}t + \theta), \quad (2.41)$$

where f_{RF} is the RF frequency and θ is the initial RF phase. For a cavity whose frequency is a multiple of f_0 , one may write the turn-by-turn evolution of δ for a given particle i as (in thin-lens approximation),

$$\delta_{n+1,i} = \delta_{n,i} + \frac{eV_0}{\beta_r^2 E} \sin(\phi_{n,i} + \theta), \quad (2.42)$$

where $\phi_{n,i} = 2\pi h f_0 \tau_{n,i}$ is the particle's phase, h is the RF-cavity harmonic number, β_r is the relativistic speed-factor and E is the reference-particle's energy.

For an RF cavity with arbitrary frequency, θ will slip with respect to the reference particle and become turn-by-turn dependent,

$$\theta_{n+1} = [\theta_n + 2\pi \frac{f_{RF}}{f_0}] \bmod 2\pi, \quad (2.43)$$

where $[a] \bmod b$ is the modulo operation, which returns the remainder of dividing a by b .

If a ring has several RF cavities operated at the same frequency (and the delay between consecutive cavities is synchronised to the reference particle), one may compute a virtual RF cavity whose voltage V_0 and initial phase ϕ are given by,

$$V_0^2 = [\sum_j V_{0,j} \sin \theta_j]^2 + [\sum_j V_{0,j} \cos \theta_j]^2 \quad (2.44)$$

$$\tan(\theta) = \frac{\sum_j V_{0,j} \sin \theta_j}{\sum_j V_{0,j} \cos \theta_j}, \quad (2.45)$$

where the sum is performed over all cavities in the ring.

2.4.2 Time slippage

An off-momentum particle ($\delta \neq 0$) follows a different-length path along the accelerator, as the kicks from the lattice magnets are different from those applied to the reference particle. Furthermore, the speed along the trajectory is also different due to the different momentum. These two factors result in a revolution frequency f_i that differs from the reference f_0 :

$$f_i = f_0[1 - \eta\delta + O(\delta^2)], \text{ with } \eta = -\frac{1}{f_0} \frac{\partial f}{\partial \delta} \Big|_{\delta=0}, \quad (2.46)$$

where η is known as the slippage factor of the lattice, and can be obtained numerically with MADX. As a consequence of this frequency difference, a particle i will slip in time-of-arrival with respect to the reference particle:

$$\tau_{n+1,i} = \tau_{n,i} - \frac{\eta}{f_0} \delta_{n,i}. \quad (2.47)$$

2.4.3 Longitudinal phase space

The 2-tuple (τ, δ) can be represented in longitudinal phase space. Since the relevant dynamics are typically slow compared to the reference period $T = 1/f_0$, one may choose T as the fundamental unit of time and convert the difference equations in Eqs. 2.47–2.42 into the following differential equations:

$$\frac{d\tau}{dT} = -\eta\delta \quad (2.48)$$

$$\frac{d\delta}{dT} = \frac{eV_0f_0}{\beta^2E} \sin(2\pi hf_0\tau + \theta), \quad (2.49)$$

which can be derived from a Hamiltonian H_T if one chooses the canonical coordinate pair $(\tau, \beta^2E\delta)$:

$$H_T = -\frac{\eta}{2\beta^2E}(\beta^2E\delta)^2 + \frac{eV_0}{2\pi h} \cos(2\pi hf_0\tau + \theta), \quad (2.50)$$

using Hamilton's equations:

$$\frac{d\tau}{dT} = \frac{\partial H_T}{\partial(\beta^2E\delta)} \quad (2.51)$$

$$\frac{d(\beta^2E\delta)}{dT} = -\frac{\partial H_T}{\partial\tau}. \quad (2.52)$$

H_T 's contour lines are shown in Fig. 2.9. The trajectory with value $H_T = -\text{sign}(\eta) \cdot eV_0/2\pi h$ traces a separatrix, which splits longitudinal phase space into a region of periodic motion (the bucket) and two regions of non-periodic motion (one above and one below the bucket).

2.4.4 Accelerating bucket

In order to represent acceleration (or deceleration), we now introduce an additional term (underlined below) to the Hamiltonian in Eq. 2.50:

$$H_T = -\frac{\eta}{2\beta^2E}(\beta^2E\delta)^2 + \frac{eV_0}{2\pi h} [\cos(2\pi hf_0\tau + \theta) + \underline{2\pi hf_0\Gamma\tau}], \quad (2.53)$$

where $\Delta E = eV_0\Gamma = eV_0 \sin \phi_s$ is the energy-change per turn and ϕ_s is the stable phase. This yields a new differential equation for the evolution of δ :

$$\frac{d\delta}{dT} = \frac{eV_0f_0}{\beta^2E} \sin(2\pi hf_0\tau + \theta) + \frac{eV_0f_0}{\beta^2E} \Gamma. \quad (2.54)$$

The corresponding difference equation can be computed to be:

$$\delta_{n+1,i} = \delta_{n,i} + \frac{eV_0}{\beta^2E} \sin(2\pi hf_0\tau_i + \theta) + \frac{\Delta E_i}{\beta^2E}, \quad \text{with } \Delta E_i = eV_0\Gamma \quad (2.55)$$

The new Hamiltonian contours are shown in Fig. 2.10. The two non-periodic regions are now connected by a 'channel' of trajectories. A particle may travel from below to above the bucket (or vice versa during deceleration) by following the channel.




Figure 2.9: Stationary buckets, Hamiltonian contours in longitudinal phase space. The pink dashed lines trace out the separatrices and the quantities in square brackets indicate the units for each axis, where [1] refers to a unitless quantity.

The half-height of the bucket $\hat{\delta}_{bucket}$ (often simply called ‘the bucket height’) is given by,

$$\hat{\delta}_{bucket} = 2\sqrt{\frac{eV_0}{2\pi\beta^2 E h |\eta|} Y(\Gamma)}, \quad (2.56)$$

where $Y(\Gamma) \in [0, 1]$ for $\Gamma \in [0, 1]$ is the monotonically-decreasing bucket-height factor that can be found tabulated in [28]. If one now performs the normalisation given by $(\tau, \delta) \rightarrow (\tau f_{RF}, \delta/\hat{\delta}_{bucket})$ the variation in bucket height and bucket length are removed and the longitudinal topology can be fully parametrised in terms of Γ , as shown in Fig. 2.11. In this normalised space, changes in Γ can be interpreted as trading bucket area for channel area. Two limiting cases can be identified: when $\Gamma = 0$, the bucket area is maximal and the channel width is zero; when $\Gamma = 1$ the bucket area is zero and the channel width is maximal.

In a typical acceleration scheme (the bucket lift), particles are located inside the accelerating bucket, while $B\rho$ is ramped synchronously. The larger the Γ the bigger the energy-gain per turn, but the smaller the bucket area.




Figure 2.10: Accelerating buckets, Hamiltonian contours in longitudinal phase space. The pink dashed lines trace out the separatrices and the quantities in square brackets indicate the units for each axis, where [1] refers to a unitless quantity. The non-periodic motion below the bucket connects to the non-periodic motion above the bucket through ‘channels’ between consecutive buckets. The black dots show two particle trajectories (simulated in *henontrack*): one executes periodic motion; the other one channels from below to above the bucket.

2.5 Phase displacement

The trajectories previously shown in Fig. 2.10 reveal an additional means of accelerating the beam: particles starting below the bucket will eventually flow through the channel into the region above the bucket, gaining energy in the process. This method is called phase-displacement acceleration and was used in the Intersecting Storage Rings (ISR) as the nominal acceleration scheme [36].

Unlike the bucket lift, phase displacement is an unstable acceleration scheme, as particles can only traverse the channel once. As they move from below to above the bucket, they experience a total energy gain equal to [37],

$$\Delta E = A_0 f_0 \alpha(\Gamma), \text{ with } A_0 = 16 \sqrt{\frac{2\pi\beta^2 E eV}{h|\eta|}}, \quad (2.57)$$

where $\alpha(\Gamma) \in [0, 1]$ for $\Gamma \in [0, 1]$ is the monotonically-decreasing bucket-area factor




Figure 2.11: Accelerating buckets, Hamiltonian contours in normalised longitudinal space $(\tau f_{RF}, \delta / \hat{\delta}_{bucket})$. As the stable-phase sine ($\Gamma = \sin \phi_s$) is varied from 0 to 1, the bucket area shrinks and the channel width between consecutive buckets increases.

that can be found tabulated in [28].

Figure 2.12 shows the turn-by-turn energy gain ΔE_i of a particle i undergoing phase displacement. In the beginning, the magnitude of the frequency difference $|\Delta f_i| = |f_i - f_{RF}|/h$ between the particle and the RF cavity is large. Therefore, ΔE_i oscillates as the particle pseudo-randomly samples different voltages along the RF waveform. Then, the particle's f_i approaches f_{RF}/h (either by changing f_i or f_{RF}) and $|\Delta f_i|$ becomes small. During this time, the particle consecutively receives kicks near the RF-waveform crest and ΔE_i rapidly increases. Finally, as f_i moves away from f_{RF}/h , the resonant behaviour ceases and acceleration stops.

In order to sustain acceleration for long periods of time, one may repeatedly apply

the process by momentarily setting $V_0 = 0$ and changing f_{RF} such that Δf_i is reset to its initial value. Because of voltage limitations and the beam's large momentum spread, the ISR was operated by performing 184 bucket traversals over a period of 20 minutes to get the beam from an injection energy of 26.5 GeV to a flattop energy of 31.4 GeV [36].




Figure 2.12: Energy gain ΔE_i and frequency difference Δf_i vs. turn number, for a particle i undergoing phase-displacement acceleration (simulated in *henontrack*). When $|\Delta f_i|$ is small, the RF kicks coherently add up, rapidly accelerating the particle. The black horizontal line shows the energy gain expected from the theory.

2.5.1 Reflection from a potential well

A helpful representation of phase displacement can be obtained by looking at trajectories in a potential-energy diagram. First, one splits the Hamiltonian H_T from Eq. 2.53 into a kinetic term K and a potential term U as follows (with $\theta = 0$ without loss of generality),

$$K = -\text{sign}(\eta) \cdot \frac{\eta\beta^2 E}{2} \delta^2 \quad (2.58)$$

$$U = \text{sign}(\eta) \cdot \frac{eV_0}{2\pi h} [\cos \phi + \Gamma \phi] \quad (2.59)$$

$$H' = K + U, \quad (2.60)$$

where $H' = \text{sign}(\eta) \cdot H_T$ represents the ‘total energy’ of a given particle, and the multiplication by $\text{sign}(\eta)$ is needed to enforce the convention that peaks/troughs in U correspond to unstable/stable fixed points both below and above transition.

Now, one may produce a potential-energy diagram like the one shown in Fig. 2.13. Periodic motion inside the bucket occurs when a trajectory's H' -value is small enough to be trapped in a potential well, reflecting back and forth between two potential walls. Non-periodic motion occurs when a trajectory's H' -value is large enough to overcome one of the potential walls after having been reflected. For the latter motion, phase displacement is precisely the process of reflecting against the potential wall, before drifting in the opposite ϕ -direction.




Figure 2.13: Schematic of potential-energy (U) diagram for accelerating bucket, showing the ‘total energy’ (H') v. relative phase (ϕ). The horizontal lines show two particle trajectories: periodic motion (red) and reflection from the potential wall (black). This reflection is a representation of phase displacement.

2.5.2 Effect on a particle ensemble

Liouville’s theorem states that for a Hamiltonian system *the density of system points in the vicinity of a given system point traveling through phase space is constant with time* [38] [39]. Informally speaking, the beam distribution behaves like an incompressible fluid in phase space. This theorem can be used to understand how the phase-displacement phenomenon affects an ensemble of particles.

For simplicity, let us assume the beam starts uniformly distributed in longitudinal phase space. Therefore, it occupies an area given by its length times its height. Figure 2.14 shows a tracking simulation of phase displacement. As the RF frequency is brought close to a harmonic of the beam’s revolution frequency, empty buckets approach the waiting particle stack. Since the separatrix encloses the bucket area,

particles that start outside the bucket can never enter the bucket. Thus, the empty buckets act as obstructions in phase space, which must be circumvented by the beam. As a consequence, particles below the bucket are forced to travel up the channel.




Figure 2.14: Phase displacement in longitudinal phase space for a particle ensemble (sky-blue), simulated in *henontrack*. The subplots show the phase space before (left), during (centre) and after (right) the manipulation. The buckets (black) are displaced down and through the beam, which cause particles to be displaced up. The process is analogous to submerging two bucket-shaped objects (black) into a fluid (sky-blue) resting in a container (straight black lines).

In short, the behaviour is analogous to what would happen if a bucket-shaped object were submerged into a pool of incompressible fluid (e.g. water): the electric forces from the RF cavity move the beam towards larger δ in phase space. From this interpretation, the expression for the energy-gain quoted in Eq. 2.57 can be intuitively understood: as the buckets traverse the beam, they cause a net displacement equal to their area.

Chapter 3

Empty-bucket channelling in the CERN SPS

The CERN SPS is a 7 km synchrotron capable of accelerating proton and ion beams up to an energy of 450 GeV. For the SPS Fixed-Target (SFT) cycle, the SPS receives beam from the PS at 14 GeV, which is accelerated to an extraction energy of 400 GeV. Then, a stable spill of 4×10^{13} protons is slow extracted over 4.8 s into Transfer Tunnel 20 (TT20). Finally, the protons are transported to the various targets located in the North Area (NA) along several transfer lines. A schematic of the path followed by the SFT beam is shown in Fig. 3.1. The ECN3 hall is the present location of the North Area 62 (NA62) experiment (Chap. 5). It is also the potential future location of the ‘Search for Hidden Particles’ (SHiP) (Chap. 4) or the HIKE/SHADOWS (Chap. 5) experiments.




Figure 3.1: Schematic of the CERN SPS and the North Area transfer lines. The dashed black line schematically shows the path followed by the SFT beam to reach the ECN3 hall.

3.1 SPS Fixed-Target cycle

In order to achieve a 4.8 s spill, the SPS utilises 1/3-integer-resonance slow extraction. The necessary theory to understand the extraction has been discussed in Chap. 2. The key machine and beam parameters of the SFT cycle are compiled in Table 3.1. Unless mentioned otherwise, this thesis will employ these parameters as a reference.

Quantity	Symbol	Value
Initial reference momentum	$p^{(0)}$	400 GeV/c
Reference betatron tune	Q_x	80/3
Chromaticity	Q'	-26.67
Virtual-sextupole strength	S	$170 \text{ m}^{-1/2}$
Reference revolution frequency	f_0	43.375 kHz
Slippage factor	η	1.84×10^{-3}
Maximum virtual-RF voltage	V_0	7 MV (1.4 MV)
Rf-cavity harmonic	h	4620 (18480)
R.m.s. emittance (Gaussian)	$\epsilon_{x,rms}$	$1.88 \times 10^{-2} \text{ mm mrad}$
Full momentum-width (Uniform)	$\hat{\delta}_{beam}$	3×10^{-3}
R.m.s. momentum stop-band	$\hat{\delta}_{stopband,rms}$	5.4×10^{-5}
R.m.s. extracted momentum	$\hat{\delta}_{rms}^{out}$	4×10^{-5}
R.m.s. separatrix (radial) length	l_{rms}^{out}	$3.1 \times 10^{-4} \text{ m}^{1/2}$
R.m.s. separatrix (azimuthal) width	w_{rms}^{out}	$0.98 \times 10^{-4} \text{ m}^{1/2}$
R.m.s. extracted emittance	$\epsilon_{x,rms}^{out}$	$3.1 \times 10^{-2} \text{ mm mrad}$

Table 3.1: SFT-cycle parameters. The properties of the main RF-system are listed first, and the properties of the auxiliary RF system are shown in parentheses. More details about these parameters can be found throughout this chapter. In particular, the r.m.s. momentum stop-band quantifies the relative momentum spread of the portion of the beam that is resonant at any given moment during slow extraction.

3.1.1 Transverse aspects

The SFT beam is extracted from the SPS with the help of the electrostatic septum (ZS), which is located in Long Straight Section 2 (LSS2). The horizontal optics in LSS2 are shown in Fig. 3.2, where the ZS location and the virtual-sextupole location have been highlighted. The virtual-sextupole magnitude and phase can be computed using Eq. 2.28. As shown in Fig. 3.3, the SPS uses four sextupoles to achieve the nominal value of $S = 170 \text{ m}^{-1/2}$. The phase advance of such sextupoles has been chosen so that their individual strengths S_j add up coherently when computing the effective term S . The rest of S_j contributions come from other magnets such as

chromatic sextupoles. Their contributions are small and quasi-uniformly distributed along the lattice, essentially vanishing when computing the total S .




Figure 3.2: Horizontal beta-function β_x along Long Straight Section 2 (LSS2) as a function of physical distance s and horizontal phase advance μ_x . The locations of the virtual sextupole and the entrance of the electrostatic septum (ZS) have been highlighted with vertical lines.




Figure 3.3: Normalised-sextupole strengths S_j for all magnets in the SPS ring, shown in a polar plot with respect to $3\mu_x$ (modulo 2π). The virtual-sextupole strength S is the phasor sum of all S_j . The location of the entrance of the electrostatic septum (ZS) has been highlighted.

Figure 3.4 shows the normalised phase space at the ZS during slow extraction, where its thin wires separate the beam pipe into two regions: a zero-field region

where the circulated beam passes unperturbed, and a high-field region where large-amplitude particles are kicked towards TT20. The ZS wires are made out of NbTi and have a thickness of $500\ \mu\text{m}$ [40]. During slow extraction, some particles ($\sim 3\%$ [41]) are lost in the thin wires, leading to equipment activation and equipment aging. It is critical to keep these losses as low as possible, which is primarily achieved by optimising the optics, the sextupole strengths and the alignment between the beam and the septum.




Figure 3.4: Normalised phase space at the electrostatic septum (ZS) during the SPS slow extraction (simulated in *henontrack*). The grey vertical line represents the wires of the ZS, the light blue dots show particle trajectories, and the fuchsia dots show the coordinates of extracted particles integrated throughout the entire spill.

3.1.2 Longitudinal aspects

The SPS is equipped with six $h = 4620$ cavities and two $h = 4 \times 4620 = 18480$ cavities [42]. The former are often referred to as ‘the 200 MHz system’ and can provide a combined V_0 of up to 7 MV, while the latter are referred to as ‘the 800 MHz system’ and can provide a combined voltage V_0 of up to 1.4 MV. During beam acceleration, the 200 MHz system is the main RF system, bunching the beam at intervals of 5 ns.

When the beam arrives to flat-top and before extraction, an RF manipulation known as ‘bunch rotation’ or ‘phase jump’ is performed with the aim of increasing the beam’s momentum spread $\hat{\delta}_{beam}$. A larger momentum spread leads to a larger tune spread ($dQ = Q' \cdot d\delta$), which makes the slow extraction process more robust to errors/perturbations on the betatron tune. Bunch rotation is illustrated in Fig. 3.5 and can be described as follows:

1. Arrival at flattop: The beam starts bunched inside the RF buckets (Fig. 3.5a). Individual particles perform stable oscillations around the stable fixed point.
2. Phase jump: The RF-waveform phase is rapidly shifted by 180 degrees, moving bunches near the unstable fixed point. Particle trajectories now follow the bucket separatrix and the beam-momentum spread is stretched (Fig. 3.5b).
3. Re-phasing & bunch rotation: The RF phase is shifted back by 180 degrees, bringing the stretched bunches back to stable motion. The bunches are not matched to the bucket anymore, so the momentum and time distributions vary over time (Fig. 3.5c).
4. De-bunching: The RF voltage is switched off when the bunches are ‘up-right’, i.e. the momentum spread is maximal and the time spread is minimal. The beam de-bunches until it covers the entire length of the synchrotron (Fig. 3.5d).

After bunch rotation has been performed and the RF has been switched off, one waits a few-hundred milliseconds for the beam to coast freely and cover the full length of the SPS ring. At this point, the SFT beam is ready to be extracted.

3.1.3 Constant-Optics Slow Extraction

The method employed to slow extract the SFT beam is known as Constant-Optics Slow Extraction (COSE) [43]. In COSE, the machine’s beam rigidity $B\rho$ is slowly varied as a function of time, so that the reference momentum p_0 changes accordingly. Since all magnetic fields in the lattice are varied with $B\rho$, all resonant particles ‘see’ the same optics as the resonant tune is swept. This eliminates the linear and non-linear chromatic perturbations that are present in a ‘conventional’ quadrupole sweep. Then, a particle’s betatron tune $Q_{x,i}$ evolves as follows:

$$Q_{x,i}(t) = \frac{80}{3} + \Delta Q(t) = \frac{80}{3} + Q' \cdot \delta(t) = \frac{80}{3} + Q' \left[\frac{p_i}{p_0(t)} - 1 \right], \quad (3.1)$$

and will eventually become resonant when $|\Delta Q|$ becomes small enough. Notice that the tune distance is reduced by changing the reference momentum p_0 and not the particle’s physical momentum p_i , so no actual acceleration occurs. In the SPS, the initial-beam distribution has $\delta > 0$ and p_0 is increased (magnetic fields are ramped up) to extract the beam, but one may equally start with $\delta < 0$ and decrease p_0 (magnetic fields are ramped down).




Figure 3.5: SPS bunch rotation in longitudinal phase space (simulated in *henontrack*). The different stages of the RF manipulation are shown in different subplots. For each subplot the pink dashed lines illustrate the underlying RF bucket, the light-blue dots show particles at the start of the time period, and the fuchsia dots show the particles at the end of the time period. The marginal probability distributions are projected onto the $\tau \cdot f_{RF}$ and δ axes.

In order to perform COSE, multipolar strengths K_n remain unchanged by ramping all magnetic components B_n according to the following programme:

$$B_n(t) = B\rho(t)K_n = \frac{p_0(t)}{e}K_n, \text{ with } B_n = \left. \frac{\partial^n B_y}{\partial x^n} \right|_{x=0,y=0}. \quad (3.2)$$

The momentum programme employed operationally before this thesis is shown in Fig. 3.6 and can be expressed as,

$$p_0(t) = p^{(0)} + \lambda t + \sum_j a_j \sin(2\pi f_j t + \phi_j), \text{ for } t \in [0, T_s], \quad (3.3)$$

where $p^{(0)}$ is the initial reference momentum, λ is the linear ramp rate, T_s is the spill

length, and a_j, f_j, ϕ_j are the j -th higher-order-correction amplitude, frequency and phase, respectively. For the SFT cycle, the relationship $p^{(0)} \gg \lambda t \gg a_j$ holds for $t \in [0, T_s]$, as the relative beam-momentum spread is small ($\sim 0.3\%$). The linear term is the ‘main contribution’, while the sinusoidal terms compensate for slowly varying non-uniformities in the beam-momentum distribution. In the SPS, the feed-forward program known as AutoSpill [44] is used to calculate the higher-order corrections, up to a maximum frequency of 25 Hz. Only the reproducible non-uniformities can be corrected, since cycle-to-cycle variations cannot be addressed using a feed-forward system.




Figure 3.6: Momentum programme during flattop for SFT cycle.

3.1.4 Momentum stop-band

COSE can be illustrated in a Steinbach diagram as shown in Fig. 3.7. For a given normalised amplitude A_x , the unstable region has a given width $\hat{\delta}_{stopband}$, called the momentum stop-band width. This width can be related to A_x in terms of the machine parameters by means of Eq. 2.39 as follows:

$$\hat{\delta}_{stopband} = \frac{1}{(48\sqrt{3}\pi)^{1/2}} \frac{S}{Q'} A_x. \quad (3.4)$$

If one chooses A_x such that the area enclosed by the single-particle phase-space trajectory is equated to the beam’s $\epsilon_{x,rms}$ (i.e. $A_x^2\pi = \epsilon_{x,rms}\pi$), the ‘root-mean-squared’ stop-band width $\hat{\delta}_{stopband,rms}$ may be computed for the SPS:

$$\hat{\delta}_{stopband,rms} = \frac{1}{(48\sqrt{3}\pi)^{1/2}} \frac{S}{Q'} \sqrt{\epsilon_{x,rms}} = 5.4 \times 10^{-5}. \quad (3.5)$$






Figure 3.7: Steinbach-diagram illustration of COSE. The horizontal arrow indicates that the beam (brown) is pushed from the stable region (blue) into the unstable region (fuchsia). The red band indicates the strip of particles that is resonant at a given instant. The momentum spread of such a strip is known as the r.m.s. momentum stop-band.

This r.m.s. momentum stop-band width corresponds to the momentum spread of particles that are simultaneously resonant during the slow extraction process (highlighted in red in Fig. 3.7). In other words, if one were to measure the momentum spread of the extracted beam at any given moment, the value would approximately be given by the r.m.s. momentum stop-band width. Since $\hat{\delta}_{stopband,rms} \ll \hat{\delta}_{beam}$ for the SPS, the extraction is almost purely chromatic (i.e. along the δ direction).

The quantity $\hat{\delta}_{stopband,rms}$ provides a useful visual tool to indicate the coupling between transverse and longitudinal beam dynamics. As shown in Fig. 3.8, one may draw a band of width $\hat{\delta}_{stopband,rms}$ and centre δ_{res} (such that $Q_x(\delta_{res}) = \frac{80}{3}$) to indicate the extent and location of the unstable region in longitudinal phase space, implicitly encapsulating the transverse dynamics.

3.2 Empty-bucket channelling

Empty-bucket channelling [16] is a phase-space manipulation that combines resonant extraction with phase displacement. These two processes have been detailed independently in Chap. 2. This section shows that the two manipulations can be applied together for the current operational SPS system.



(c) Normalised phase space

Figure 3.8: Bunched beam on the 1/3-integer resonance (simulated in *henontrack*). The different subplots show how the unstable transverse region can be represented in the longitudinal phase space. The Steinbach diagram connects the root-mean-squared normalised amplitude $A_x = \sqrt{\epsilon_{x,rms}}$ to the root-mean-squared momentum stop-band width $\hat{\delta}_{stopband,rms}$.

3.2.1 ‘Accelerating bucket’ with COSE

Even though COSE does not produce a ‘real acceleration’ of particles, it does change the reference energy. This results in a relative energy change per turn with respect to the RF, which is enough for phase displacement to occur. Using Eq. 3.3 and the differential relationship $dE = \beta c \cdot dp$, the acceleration term in the longitudinal equations of motion can be computed as follows:

$$dE = \beta c \cdot dp = \beta c \cdot d(p - p_0) = -\beta c \cdot dp_0 \quad (3.6)$$

$$\Delta E = \int_0^{1/f_0} \frac{dE}{dt} dt = -\frac{\beta c}{f_0} \lambda + O(\lambda^2) + O(a_j), \quad (3.7)$$

where $O(a_j)$ refers to the high-order corrections provided by AutoSpill, listed in Eq. 3.3. By comparing Eq. 3.7 to Eq. 2.55, one obtains the following relationship between λ and Γ :

$$\Gamma = -\frac{\beta c}{eV_0 f_0} \lambda + O(\lambda^2), \quad (3.8)$$

which shows that, for a given ramp-rate λ , Γ may be adjusted by changing the RF-cavity voltage V_0 . The minus sign indicates that when the magnetic fields are ramped up/down, particles are effectively decelerated/accelerated with respect to the RF.

3.2.2 Concept

The goal in empty-bucket channelling is to align the momentum stop-band with the phase-displacement channel, so that particles are extracted from the machine in the middle of the phase-displacement procedure, as shown in Fig. 3.9.

This technique has two main effects:

1. Beam bunching: as particles gain amplitude, they occupy a limited region of longitudinal phase space, constrained between empty buckets. When a particle reaches the septum and leaves the ring, its longitudinal structure ‘freezes’, as it no longer receives kicks from the RF¹. This results in an extracted spill with a strong bunched structure at the RF frequency. This phenomenon is studied in detail in Chap. 4 for the provision of bunched beams.
2. Tune speed-up: during phase displacement, a particle’s revolution frequency is resonant with f_{RF}/h (see Sec. 2.5) and particles receive coherent voltage kicks near the RF crest. Therefore, in empty-bucket channelling, particles about to be extracted rapidly change their δ_i value. In a ring with chromaticity like the SPS, this results in a rapid change in $Q_{x,i}$ since $dQ_{x,i} = Q' \cdot d\delta_i$. This phenomenon is studied in detail in Chap. 5 for spill-quality improvement.

3.2.3 Parametrisation

Three key normalised parameters can be identified when implementing empty-bucket channelling (shown in Fig. 3.10):

¹To be precise, the time structure changes along the transfer line, as particles with different momenta follow different path lengths. However, this effect is negligible compared to the nanosecond-scale bunch length produced by the RF.




Figure 3.9: Two particles undergoing empty-bucket channelling (simulated in *henon-track*). The colours (blue to fuchsia) indicate the turn number. The transverse resonance (red) and the RF (black) are represented in different subplots for additional visual aid.

- Bucket offset, Δ_{RF} : It is given by $\Delta_{RF} = (\delta_{RF} - \delta_{res})/\hat{\delta}_{bucket}$, i.e. the offset between the tune resonance δ_{res} and the bucket centre δ_{RF} , normalised to the bucket height $\hat{\delta}_{bucket}$. $|\Delta_{RF}| < 1$ ensures that the transverse resonance overlaps with the RF channel. As a convention, positive/negative Δ_{RF} indicates that the bucket is offset towards/away from the waiting beam.
- Channel width, Γ : It can be shown that the channel width at the bucket's midpoint can be approximated by $\sqrt{\Gamma/\pi}$ when $\Gamma \ll 1$ [16] (see Fig. 3.10). In other words, Γ controls the channel width, with larger/smaller Γ producing wider/narrower channels.
- Resonance width, d_{rms} : It is given by $d_{rms} = \hat{\delta}_{stopband,rms}/\hat{\delta}_{bucket}$, i.e. the ratio of the root-mean-squared momentum stop-band to the bucket height. If $d_{rms} < 1$,

resonant particles fit within the RF channel.




Figure 3.10: Parametrisation of empty-bucket channelling illustrated in longitudinal phase space, including a trajectory (grey) that reaches the resonant momentum (red). The extracted bunch has a length σ . The position and momentum coordinates have been normalised by the RF bucket height $\hat{\delta}_{bucket}$ and the RF frequency f_{RF} , respectively.

The three parameters Δ_{RF} , Γ and d_{rms} have been defined so that they are unitless and somewhat insensitive to the specific characteristics of each synchrotron (e.g. revolution frequency, RF voltage, sextupole strength...). In this way, if two different synchrotrons were operated with the same Δ_{RF} , Γ and d_{rms} one would expect the behaviour of empty-bucket channelling to be comparable between the two machines. Ultimately, however, simulations and measurements are necessary to study the full dynamics of empty-bucket channelling for a given synchrotron.

3.2.4 Setup in the SPS

Figure 3.11 shows the longitudinal phase space during the setup and execution of empty-bucket channelling in the SPS. The entire process can be summarised as follows:

1. RF Gymnastics: the beam arrives at flattop bunched. An additional frequency offset is programmed on the RF frequency, ensuring that there is a gap between

the beam and the resonance to fit an empty bucket. For a given set of empty-bucket channelling parameters $(\Gamma, \Delta_{RF}, d_{rms})$, one can compute the necessary momentum-gap δ_{gap} to be,


$$\delta_{gap} = \max\{0, \hat{\delta}_{bucket} \cdot (1 + \Delta_{RF})\}. \quad (3.9)$$

In other words, the beam is accelerated to a momentum that ensures that $\delta_0 \geq \delta_{gap} + \hat{\delta}_{beam}$. Then, the RF gymnastics (Sec. 3.1.2) are performed and the RF voltage is switched off.


2. Jump to channelling: As the beam starts to debunch, the RF is moved to the empty-bucket channelling frequency and voltage, with V_0 computed from Eq. 3.8 for a given Γ and λ . No particles are captured in the process, thanks to the gap opened in the previous step.
3. Start channelling: The reference momentum is ramped and the beam approaches the resonance (in transverse phase space) and the RF channels (in longitudinal phase space) simultaneously, with particles closer to the resonance performing empty-bucket channelling first.
4. End channelling: p_0 is swept through the beam momentum over the course of $T_s = 4.8$ s, and all particles gradually undergo empty-bucket channelling until the ring is emptied.

3.3 Simulation model: *henontrack*

Single-particle beam-dynamics simulations are crucial for slow-extraction studies. For the SPS in particular, these tools have been extensively exploited to understand the extracted phase space as well as the spill time structure [45] [46] [47]. For the studies in this thesis, a unique challenge appears in the modelling front. On the one hand, the SPS slow extraction takes $O(10^5)$ turns, with millisecond-timescale phenomena such as power-converter ripple being a core subject of study. On the other hand, longitudinal manipulations such as empty-bucket channelling occur in the sub-turn scale, requiring knowledge of the beam structure in the sub-nanosecond timescale. This section develops a modelling methodology to address this challenge efficiently, describing the custom-made *python* code *henontrack* [48] and showcasing its capabilities.



(a) RF gymnastics: full bucket is on (pink) and empty bucket is off (blue). (b) Jump to channelling: full bucket is off (blue) and empty bucket is on (pink).



(c) Start channelling: empty bucket is on (pink) and particles are pushed towards resonance.

(d) End channelling.

Figure 3.11: Particles (grey) in longitudinal phase space during empty-bucket channelling in the SPS (simulated in *henontrack*). The different subplots show different stages of the manipulation during the SFT cycle. Since the phase space is periodic in f_{RF} , particles' time coordinate τ has been aliased into the domain $\tau \in [-\frac{1}{f_{RF}}, \frac{1}{f_{RF}}]$ for visual purposes. The resonance stop-band is shown in red and the extracted particles are plotted with fuchsia squares.

3.3.1 Map models

One can build a one-turn-map of a synchrotron using a difference equation of the form:

$$\vec{Z}_{n+1,i} = \prod_j \mathcal{M}_j \circ \vec{Z}_{n,i}, \quad (3.10)$$

where \mathcal{M}_j is a map that depends on the lattice parameters and $Z_{n,i}$ is the phase-space state vector at turn n for particle i . In element-by-element tracking codes such as MADX [33], each map \mathcal{M}_j usually corresponds to a single magnet/RF cavity

(or a slice thereof). However, it is also possible to construct maps that correspond to ‘effective elements’, which lump several physical elements together. The latter approach can help minimise computational time while still preserving the relevant dynamics one is interested in modelling.

For example, previous studies in the SPS have demonstrated that an effective-map approach with $Z = [X, X', \delta]^T$ can be successfully used to model the effect of power-converter ripple on the slow-extracted intensity [49]. The model is schematically shown in Fig. 3.12, where all parameters can be varied with the turn number n . The effect of all resonant sextupoles is encapsulated in a single virtual sextupole with strength S using Eq. 2.28, and the linear one-turn mapping is characterised by the global quantities Q, Q' , which are calculated numerically with MADX.




Figure 3.12: Schematic of SPS simulation model used in [49]

Using this approach, the 1500 elements of the SPS lattice are collapsed into four effective elements. Most importantly, the results show good agreement with both measurement and MADX simulations [49]. Thus, the model provides a computationally light-weight procedure to assess spill quality.

3.3.2 SPS model

Taking the model from the previous section as a starting point, an extended model was developed exclusively for the studies in this thesis. The model has been packaged into a python module named *henontrack* [48].

Henontrack performs single-particle-dynamics simulations in 2 degrees of freedom: normalised phase space and longitudinal phase space. The single-particle state vector for particle i at turn n is given by $\vec{Z}_{n,i} = [X, X', \tau, \delta]^T$. Fig. 3.13 shows the extended model, where the effect of all RF cavities is encapsulated in a single virtual RF cavity of effective voltage V_0 using Eq. 2.45. The linear one-turn mapping now also includes the time slippage characterised by η , which is calculated numerically with MADX.

By incorporating the longitudinal motion, the model can be used to study MHz-GHz bunch structure, as well as the Hz-kHz spill structure.




Figure 3.13: Schematic of SPS simulation model employed in this thesis

3.3.3 Further computational speed-up

The model provided above already provides substantial computational speed-up with respect to element-by-element tracking (at the cost of accuracy, of course). Still, two additional architectural features were implemented to further reduce the simulation time.

The first strategy was to develop a generic framework that allows *python* scripts to be run in CERN’s parallel computing cluster, HTCondor. A schematic of the dataflow is shown in Fig. 3.14. Essentially, one defines the main script, which can take a variety of parameters as inputs. Then, each node in the cluster executes the same script with different parameters, and the outputs are collected and analysed to produce the final result. Such a setup substantially facilitates performing parameter scans, which will be repeatedly used in later studies.

The second strategy was to make *henontrack* compatible with Graphical Processing Units (GPUs), which outperform Central Processing Units (CPUs) in linear-algebra operations. Figure 3.15 compares the performance between CPU and GPU for a representative simulation job. When the number of particles in the simulation is low, the CPU outperforms the GPU, as the latter has a much larger overhead to set up the computations. However, as the particle number is increased, the GPU run time stays practically constant, as operations are fully parallelised. Meanwhile, the CPU run time increases linearly. For this specific machine, the GPU starts to outperform the CPU at around 6×10^4 particles per simulation. At some point, the parallelisation capacities of the GPU saturate and its run time increases linearly as well, but it remains significantly faster than the CPU’s. Unfortunately, most nodes




Figure 3.14: Schematic of simulation pipeline exploiting CERN’s HTCondor cluster. The user provides a *run.py* script and a list of parameter arrays (p_0, p_1, \dots, p_n). The framework generates the necessary infrastructure and executes a *job.sub* file on the submit machine, which distributes the tasks among different worker nodes. The results are collected back by the submit machine and analysed by the *analysis.py* script to provide output back to the user.

in the HTCondor cluster lack a GPU and, therefore, the benefit of using them diminishes when one needs to run large parameter scans. Still, when the job requires only a handful of nodes, the GPU approach remains superior.




Figure 3.15: Simulation run time as a function of number of particles for a representative *henontrack* simulation. For this particular machine, the GPU implementation outperforms the CPU implementation when the particle count exceeds 6×10^4 .

3.4 End-to-end SPS simulation

This chapter has detailed the nominal SPS slow-extraction system, as well as the implementation of empty-bucket channelling. Furthermore, an efficient modelling tool for slow-extraction simulations has been described, which aims at characterising the extracted time structure at timescales ranging from supra-milliseconds to sub-nanoseconds. In particular, this thesis will exploit these capabilities to understand the impact of empty-bucket channelling on the spill. This section combines everything mentioned so far to produce a simulated SPS spill under empty-bucket channelling conditions. It then describes the analysis applied to the resulting data. The upcoming chapters (Chaps. 4 and 5) will rely on this workflow to characterise different aspects of empty-bucket channelling.

3.4.1 Modelling procedure

1. A set of empty-bucket channelling parameters (Γ, Δ_{RF}) is chosen. The rest of the parameters are taken from Table 3.1.
2. All non-linear components are set to zero ($S = 0$ and $V_0 = 0$). The initial-beam coordinates are generated using the following random variables:

$$X \sim \mathcal{N}(\mu = 0, \sigma^2 = \epsilon_{x,rms}), \quad X' \sim \mathcal{N}(\mu = 0, \sigma^2 = \epsilon_{x,rms}) \quad (3.11)$$

$$\tau \sim \mathcal{U}(a = -\frac{1}{2f_0}, b = \frac{1}{2f_0}), \quad \delta \sim \mathcal{U}(a = \delta_0 - \hat{\delta}_{beam}, b = \delta_0 + \hat{\delta}_{beam}), \quad (3.12)$$

where \mathcal{N} is a normal random variable of mean μ and variance σ^2 , and \mathcal{U} is a uniform random variable sampled over the interval $[a, b]$. δ_0 is determined using Eq. 3.9 in order to ensure no particles are captured in the empty bucket.

3. The non-linear components are slowly ramped to their nominal simulation value, quasi-adiabatically deforming particle trajectories onto the new Hamiltonian contours.
4. A time-dependent tune perturbation may be introduced to model the effect of power-converter ripple on the extracted spill (See Chap. 5 for details), which can have a variety of spectral components. For the SPS, the spill ripple is dominated by a few narrow-band components at harmonics of 50 Hz. Each of these components can be modelled as a sinusoidal ripple on the tune at turn n :

$$Q_x[n] = \frac{80}{3} + 2q_i \sin(2\pi f_i n), \quad (3.13)$$

where $2q_i, f_i$ are the ripple amplitude and frequency, respectively.

5. The model in Fig. 3.13 is run for N turns. For each turn n , the phase-space coordinates of the extracted particles are stored. The relative time-of-arrival τ for a particle extracted at turn n can be converted to absolute time t using the relationship,

$$t = \tau + \sum_{j=0}^{j=n} \frac{1}{f_{0,j}}. \quad (3.14)$$

3.4.2 Analysis procedure

A representative simulation was executed following the modelling procedure from above, with the relevant parameters listed in Table 3.2.

Quantity	Symbol	Value
Rf-cavity harmonic	h	4620
Virtual-RF voltage	V_0	50 kV
Stable-phase sine	Γ	0.1
Normalised RF offset	Δ_{RF}	0.8
Normalised resonance width	d_{rms}	0.5
Tune-ripple amplitude	$2q_i$	1×10^{-5}
Tune-ripple frequency	f_i	$2.3 \times 10^{-3} \cdot f_0 = 100$ Hz

Table 3.2: Simulation parameters.

3.4.2.1 Spill structure

For time-structure analysis, two time domains are identified: super-turn and sub-turn. The former includes timescales longer than the one-turn time, while the latter is concerned with timescales shorter than the one-turn time. In the case of the SPS, the one-turn time is around 23 μ s, corresponding to a revolution frequency of 43 kHz.

The super-turn (or informally ‘slow’) spill structure of a representative simulation is shown in Fig. 3.16. Both the time structure $I(t)$ and the frequency structure $|I(f)|$ are shown, which are connected via the Discrete Fourier Transform.

3.4.2.2 Longitudinal phase space

The sub-turn (or informally ‘fast’) spill structure of a representative simulation can be analysed by looking at the longitudinal phase space shown in Fig. 3.17. The time-of-arrival τ for a particle extracted at turn n has been aliased by the RF frequency using the following relationship:




Figure 3.16: Simulated slow-extracted spill in time domain (left) and frequency domain (right). The frequency-domain representation shows a red x at the frequency of the injected tune perturbation.

$$\tau \rightarrow \tau + \Delta\tau_n - \frac{1}{f_{RF}} \lfloor f_{RF} \cdot (\tau + \Delta\tau_n) \rfloor, \quad (3.15)$$

where $\lfloor a \rfloor$ rounds a to the nearest integer and $d\tau$ is the cumulative time slippage of the RF cavity,

$$\Delta\tau_n = \sum_{j=0}^{j=n} \left(\frac{1}{f_{0,j}} - \frac{1}{f_{RF,j}} \right). \quad (3.16)$$

This transformation aliases the extraction time of all particles within a single RF period, removing the delay that particles accumulate by virtue of being extracted in different RF periods. In other words, this procedure yields the ‘average’ RF structure by integrating the bunch-by-bunch counts throughout the entire spill to obtain enough statistics. The fast spill structure will be studied in detail in Chap. 4.

Additionally, one may compute the instantaneous r.m.s. extracted momentum width $\hat{\delta}_{rms}^{out}$ by analysing the δ distribution in Fig. 3.17. If the extraction is performed slowly, one would expect $\hat{\delta}_{rms}^{out} \approx \hat{\delta}_{stopband,rms}$. In the case of the SPS, they differ by $\sim 20\%$ for the nominal extraction.

3.4.2.3 Normalised phase space

The simple model developed in this section can also provide information on the extracted normalised phase space in the transverse plane. As shown in Fig. 3.18, the initial and final X, X' distributions are vastly different, as the transformation between them is determined by the non-linear extraction process. The extracted separatrix will be characterised by three r.m.s quantities: its emittance ϵ_x^{out} , its radial length l_{rms}^{out} and its azimuthal width w_{rms}^{out} . The first is computed by applying Eq. 2.19 on the




Figure 3.17: Example of slow-extracted particle distribution in longitudinal phase space (τ, δ) during empty-bucket channelling (simulated in *henontrack*). The secondary axes show the normalised longitudinal space, using the RF frequency (f_{RF}) and the bucket height ($\hat{\delta}_{bucket}$). The marginal probability distributions are projected onto the τ and δ axes. The time-delay τ of particles has been aliased to fit within one RF period to provide the distribution of the ‘average’ bunch. The contour of the RF bucket (fuchsia) and the transverse stop-band (red) are shown for additional visual aid.

distribution in Fig. 3.18b, while the other two correspond to the square roots of the eigenvalues of the covariance matrix between X and X' .

3.5 Conclusion

This chapter has described the longitudinal and transverse beam-dynamics aspects of the slow-extraction scheme at the SPS. The beam is accelerated, manipulated longitudinally to increase its momentum spread, and extracted by synchronously ramping all magnets in the lattice to keep the optics constant. It has been shown that such extraction procedure is compatible with the RF technique known as empty-bucket channelling, and a basic parametrisation of the manipulation has been presented. Empty-bucket channelling will be exploited in the next chapters.

Furthermore, in order to study slow extraction in a computationally efficient manner, a simplified simulation model has been developed (*henontrack*), which reduces




Figure 3.18: Initial (left) and extracted (right) transverse distributions in normalised coordinates (simulated in *henontrack*). The black lines in the right subplot indicate the eigenvectors of the covariance matrix, scaled by their corresponding eigenvalues. These can be understood as the radial length and the azimuthal width of the extracted separatrix. The grey vertical line in Subfig. 3.18b represents the wires of the electrostatic septum (ZS).

the entire lattice to a few effective elements and can be executed both on CERN’s computing cluster and on GPUs. This model was then applied to simulate the SPS slow extraction under the influence of a sinusoidally modulated betatron tune. Finally, an analysis procedure was described, which computed the extracted beam’s time structure at the super-turn and sub-turn timescales, as well as its main transverse phase-space features.

All in all, the strategy outlined in this chapter could be easily adapted to implement and characterise empty-bucket channelling (or other RF manipulations) on other synchrotrons that perform slow extraction, as the parametrisation and modelling tools introduced remain, for the most part, machine-independent. Moreover, as the code is written in *python*, there are endless options for customising it and combining it with other routines. Future developments will aim at incorporating more simulation features, with a focus on integrating the implementation within the recently released *xsuite* code, which is becoming CERN’s main beam-dynamics simulation tool.

Chapter 4

Bunched beams for BDF/SHiP

4.1 Motivation

During nominal operation, the future SHiP experiment will require a ~ 1 s spill of 400 GeV protons from the SPS, which is as uniform as possible at all time scales. However, this mode of operation is limited by the neutrino background. As shown in Fig. 4.1, hidden sector (HS) particles may have an identical signature to neutrino events in the scattering detector (SD). Indeed, both particles would not be intercepted by the veto system and would produce an electromagnetic shower at the SD.




Figure 4.1: Schematic of the SHiP experiment's proton target, veto system and scattering detector (SD). From left to right: protons ($p+$) produce neutrinos (ν) and hidden-sector particles (HS). Both HS and ν pass the veto systems (hadron absorber, muon shield, etc.) undetected and can eventually produce electromagnetic showers in the SD. Their signatures are practically identical.

For this reason, an alternative operational mode has also been requested [50], which would be employed if an event in the SD required further studying. In this new mode the spill is tightly bunched at the nanosecond scale, while leaving the spill macro-structure largely unaffected. A sketch of both operational modes is shown in Fig. 4.2.

The alternative mode would allow discrimination between HS candidates and neutrinos by using time-of-flight information, further reducing the neutrino background.




Figure 4.2: Schematic of the spill time structure for the SHiP experiment at the second and nanosecond timescales. Two operational modes are requested: a de-bunched nominal mode and a bunched alternative mode.

Unlike in the de-bunched operational mode, the alternative bunched mode provides information on the time of production of both the neutrinos and the HS particle at the proton target. If the speed between the neutrinos and the HS particle were different enough, either due to a difference in mass or a difference in momentum, their arrival times at the SD detector would differ, since there is a 40 m drift between proton target and SD detector. Figure 4.3 provides a schematic of the scenario where time-of-flight discrimination can be successfully applied.




Figure 4.3: Illustration of two consecutive neutrino bunches and a HS particle at the proton target (red) and the scattering detector (blue). The first neutrino bunch arrives at the scattering detector at $t = t_1$, while the second neutrino bunch arrives at $t = t_1 + dt$, where dt is the separation between bunches. After the 40 m drift between the target and the detector, the HS particle sits between the two consecutive neutrino bunches as it arrives at the scattering detector at $t = t_2$, with $t_1 < t_2 < t_1 + dt$. Therefore, it can be identified using time-of-flight information.

Two conditions must be satisfied for successful identification of the HS particle: (i) the drift time from the proton target to the SD detector must be large enough for the HS particle to lag behind the neutrino bunch, and (ii) the time difference between two consecutive neutrino bunches must be large enough so that the HS particle is not overtaken by the next neutrino bunch. If one assumes that the bunch has a total time

width of 4σ , that neutrinos travel at the speed of light c and that the proton target is much shorter than the distance d from the target to the SD, the accumulated time slippage δt between the neutrino bunch and the HS particle must satisfy the inequalities,

$$4\sigma \leq \delta t \leq \frac{1}{f_{RF}} - 4\sigma, \text{ with } \delta t = \frac{d}{c(1 - \beta_{HS})}, \quad (4.1)$$

where $\frac{1}{f_{RF}}$ is the time spacing between consecutive bunches and β_{HS} is the HS particle's relativistic speed-factor¹. Therefore, one ideally wants to extract short bunches with a large inter-bunch separation, while still satisfying the number of protons-on-target requested (4×10^{19} p.o.t per year [51]) by SHiP. If β_{HS} is expressed in terms of the HS particle's rest mass m and momentum p , Eq. 4.1 can be recast as:

$$a_1 mc^2 \leq pc \leq a_2 mc^2, \text{ with } a_1 = \frac{1 - \frac{d}{4\sigma c}}{1 - (1 - \frac{d}{4\sigma c})^2}, \quad a_2 = \frac{1 - \frac{d}{c(1/f_{RF} - 4\sigma)}}{1 - [1 - \frac{d}{c(1/f_{RF} - 4\sigma)}]^2}, \quad (4.2)$$

which yields a detection region bounded by two straight lines, as shown in Fig. 4.4 for the SHiP configuration. Guided by different HS models, the SHiP experiment converged on a request of $4\sigma = 1.5$ ns and $\frac{1}{f_{RF}} = 5$ ns.




Figure 4.4: Schematic of SHiP's mass-momentum search space for HS particles. In the grey regions (discrimination regions), the SHiP experiment can successfully differentiate between neutrinos and HS particles using time-of-flight discrimination. In the white region, the neutrino background cannot be eliminated. The dashed arrows illustrate that the discrimination region can be expanded upwards by providing shorter proton bunches, and downwards by providing longer inter-bunch spacing. (Image adapted from [50]).


This chapter focuses on the use of the 200 MHz RF system of the SPS to provide slow-extracted bunched beams with a $5 \text{ ns} = \frac{1}{200 \text{ MHz}}$ bunch spacing and bunch length

¹To be more precise, one can also detect HS particles that lag enough to arrive within the search window of posterior bunches. To account for this one would make the substitution $\delta t \rightarrow \delta t \bmod \frac{1}{f_{RF}}$. This has been ignored.

of $4\sigma = 1.5$ ns to satisfy SHiP’s bunched-mode request. First, an overview of the conventional strategy to provide bunched beams is presented, explaining its limitations in the SPS. Then, empty-bucket channelling is introduced as an alternative approach for delivering bunched beams, a technique dubbed ‘micro-bunching’ when tested in Brookhaven National Lab [52] [53] [54]. By using simulation and measurement, it is shown that the latter technique can overcome the limitations of the conventional approach. Finally, the best candidate for SHiP’s request is identified.

4.2 Conventional bunched-extraction methods

Conventional bunched-extraction methods typically rely on keeping the beam bunched inside an RF bucket during the slow-extraction process. Instead of switching off the RF system at flat-top to de-bunch the beam, the RF voltage is kept on until the end of the cycle. Typically, one uses the ring quadrupoles to sweep the unstable-tune region through the momentum stack and/or a transverse exciter to blow-up the betatron amplitudes into the static unstable region. These procedures are illustrated in Fig. 4.5.



(a) Transverse exciter: The red arrow indicates that particles are pushed into the unstable region by transverse RF kicks.

(b) Quadrupole sweep: The red arrow indicates that the quadrupoles are ramped, sweeping the unstable region through the waiting beam.

Figure 4.5: Schematics of conventional methods to provide bunched slow extraction. Particles (brown) enter the unstable region either by a change in their normalised amplitude A_x (Subfig. 4.5a) or in their tune offset ΔQ (Subfig. 4.5b). Throughout the extraction process, particles remain bunched inside the RF bucket in longitudinal phase space (τ, δ) . Note that $\Delta Q = Q'\delta$ for a machine with non-zero chromaticity Q' .

Such schemes are already employed, for example, at the Fermilab Mu2e experiment [55] and in several medical synchrotrons [56, 57]. The former uses the bunched

structure for background discrimination (in similar fashion to the SHiP request), while the latter keep the beam bunched to easily change treatment energy during the same synchrotron fill, significantly reducing treatment time. However, several features make these schemes unattractive for the SPS:

1. Beam loss on septum wires: since high voltages (7MV) are needed to keep the whole beam bunched, particles execute large momentum oscillations due to synchrotron motion. This significantly perturbs their extraction trajectories in transverse phase space, which blows up the emittance of the extracted beam, as shown in Fig. 4.6. Low-intensity bunched-beam irradiation tests performed in the past at the SPS [58] showed that this effect increases beam losses by a more than a factor of 10 if a chromatic quadrupole-based extraction is performed [45]. An exciter-aided low-chromaticity scheme [59] could be envisioned to mitigate synchrotron coupling, but this would entail a complete makeover of the current SPS extraction scheme. In particular, the impact of the low chromaticity on beam instabilities and the increased beam loss [60] would have to be addressed, and the poor spill quality that is characteristic of such schemes [61] mitigated. The latter stems from the fact that low chromaticity makes the beam's tune spread small and thus, sensitive to small errors in tune.
2. Spill quality: Not only does synchrotron motion affect beam loss, but it also modulates the time structure of the integrated intensity. In essence, this introduces a strong contribution at the SPS synchrotron frequency $\nu_s \approx 1/210$ turns, as particles enter and exit the unstable region several times before getting extracted. This is particularly dangerous in the SPS because the high harmonic number $h = 4620$ leads to a high $\nu_s \propto \sqrt{h}$, which aggravates the effect due to higher crossing speed [62]. Moreover, synchrotron motion cuts the spill time in half: particles from the top-half of the bucket descend into the lower-half and get a chance to get extracted as the tune sweep reaches the mid-point of the bucket [45]. The latter effect can be mitigated by a slower quadrupole ramp, but the former is a more fundamental issue. Both phenomena can be seen in the spill shown in Fig. 4.6.
3. Longitudinal emittance: As SHiP will require a large number of protons on target, the SPS will aim to accelerate high-intensity beams for each spill. For such beams, several factors such as the upstream beam handling or the SPS acceleration process (especially at critical points like transition crossing) can

increase the longitudinal emittance or add jitter from cycle to cycle. This will directly translate into a lengthening of the extracted bunch, with no obvious way to mitigate the problem.

4. Incompatibility: The current SPS extraction scheme, known as Constant-Optics Slow Extraction (COSE, see Sec. 3.1.3), relies on ramping all multipoles in the ring synchronously, effectively changing the reference momentum of the machine. All operational aspects have been optimised for COSE, while conventional bunched extraction was tested in the SPS with a quadrupole-only ramp. This could compromise the reductions in beam loss obtained from silicon-crystal channelling [63], among other operational procedures.




Figure 4.6: (Left) Normalised phase space at the electrostatic septum (wires shown in black), and (right) extracted spill for a quadrupole-sweep extraction. Both subplots show RF off (de-bunched, sky-blue) and RF on (bunched, fuchsia) at operational voltage $V = 7$ MV (simulated in *henontrack*).

Given these limitations, an alternative bunching scheme has been pursued at CERN.

4.3 Empty-bucket channelling for bunched beam


Empty-bucket channelling was described in detail in Chap. 3. In this chapter, the technique is tuned for the delivery of long spills with a bunched structure at a 200 MHz repetition rate, which constitutes the first empty-bucket channelling implementation in the SPS. Unlike conventional bunched extraction, empty-bucket channelling is performed on an initially de-bunched beam. If the empty bucket and the transverse-resonance stop-band are well-aligned, particles will be bunched in narrow RF channels just before exiting the ring [53].

For a chromatic extraction like the one in the SPS, the momentum strip that needs to be bunched at any given moment $\hat{\delta}_{stopband,rms}$ is much narrower than the beam's total momentum spread $\hat{\delta}_{beam}$, i.e. $\hat{\delta}_{stopband,rms}/\hat{\delta}_{beam} \sim 1/50$. By employing empty-bucket channelling, one can use a small empty bucket to only produce bunching² within $\hat{\delta}_{stopband,rms}$, independently of $\hat{\delta}_{beam}$. On the other hand, conventional bunched extraction starts with a bunched beam and must therefore be performed with enough voltage to bunch the entire beam with momentum spread $\hat{\delta}_{beam}$, independently of $\hat{\delta}_{stopband,rms}$. This key difference is shown in Fig. 4.7, where conventional bunched extraction produces an emittance blow-up of 7x compared to the modest blow-up of 1.2x produced by empty-bucket channelling. From this simple argument it seems reasonable that a $O(50)$ -times smaller voltage could be used for empty-bucket channelling to obtain a similar bunch structure to full-voltage conventional bunched extraction.


In principle, the new method addresses all the conventional limitations listed in the previous section:

1. Beam loss on septum wires: Since only a small portion of the beam with narrow momentum spread needs to be bunched at any given time (the portion that is resonant), one can perform empty-bucket channelling with much lower voltages than conventional bunched extraction and still get short bunches. This significantly reduces the perturbation of the transverse dynamics.
2. Spill quality: empty-bucket channelling lacks periodic synchrotron modulation, although the impact of RF kicks on the transverse motion must still be kept in mind. In fact, this can be exploited to improve spill quality as explored in Chap. 5.
3. Longitudinal emittance: the bunch length provided by empty-bucket channelling does not depend on the accelerated longitudinal emittance, since the process is performed on a de-bunched beam.
4. Incompatibility: COSE is fully compatible with empty-bucket channelling, as outlined in Chap. 3.


²To be precise, the circulating beam is not stably bunched in this configuration, but the time structure of the extracted beam will be distributed in short bunches.




(a) Conventional bunched extraction, longitudinal phase space: large voltage must be used because initial beam is bunched.



(b) Empty-bucket channelling, longitudinal phase space: small voltage can be used because only resonant beam is bunched.



(c) Conventional bunched extraction, normalised phase space: the extracted emittance is 7x the nominal value due to the big synchrotron oscillations.



(d) Empty-bucket channelling, normalised phase space: the extracted emittance can be kept at 1.2x the nominal value thanks to the small RF bucket.

Figure 4.7: Bunched slow extraction for two different methods (simulated in *henon-track*): conventional chromatic (left) and empty-bucket channelling (right). The grey and sky-blue dots shows the initial and extracted particle distributions, respectively. The dashed pink lines indicate the underlying bucket separatrices for each case.

4.4 Simulations

A variety of simulation studies were performed to verify the technique's capabilities and limitations. These were performed with the SPS nominal parameters (as listed in Table 3.1), employing the $h = 4620$ system to obtain a 5 ns bunch spacing. Both the bucket offset (Δ_{RF}) and the stable-phase sine (Γ) were scanned.

4.4.1 Extracted bunch length

The goal of this subsection is to identify the (Γ, Δ_{RF}) region where the SHiP specification of $4\sigma = 1.5$ ns can be achieved. This chapter quotes the standard deviation σ of the extracted bunch as the figure of merit, normalised to the RF period/bunch-spacing. The SHiP specification σ_{SHiP} is then given by,

$$\sigma_{SHiP} = \frac{4\sigma}{4} \cdot f_{RF} = \frac{1.5 \text{ ns}}{4} \cdot \frac{1}{5 \text{ ns}} = 0.075. \quad (4.3)$$

Figure 4.8 shows the results of a (Γ, Δ_{RF}) grid-scan within the SPS-RF parameter limits. It can be seen that the bunch length requested by SHiP is achievable with the SPS 200 MHz system. The ‘SHiP specification’ region is concentrated around small Γ (high voltage) and $\Delta_{RF} = 0$ (perfect alignment between transverse resonance and RF bucket). The minimum voltage required to achieve the SHiP specification is approximately $V_0 = 0.1$ MV, which corresponds to $\Gamma = 0.05$. Notice that this voltage is $70\times$ smaller than the voltage employed for conventional bunched extraction ($V_0 = 7$ MV). The qualitative estimate in the previous section predicted a $O(50)\times$ reduction in voltage, which is in good agreement with this result.




Figure 4.8: Bunch length (σ , normalised to the RF period $1/f_{RF}$) vs. stable-phase sine (Γ) and bucket offset (Δ_{RF}), simulated in *henontrack*. The dashed red contour shows the parameter domain where the ‘SHiP specification’ ($4\sigma = 1.5$ ns) is achieved.

In order to isolate the dependence on Δ_{RF} and Γ , Fig. 4.9 shows a 1-dimensional cross-section of the grid-scan from Fig. 4.8. On the one hand, Fig 4.9a shows that when the voltage is kept fixed ($\Gamma = 0.05$), the minimum σ is indeed achieved when $\Delta_{RF} = 0$. Interestingly, a secondary ‘weak-bunching’ region appears near $\Delta_{RF} = 2.5$. Appendix D.1 explains the mechanism for this phenomenon in more detail; but in this chapter, only the ‘strong bunching’ region near $\Delta_{RF} = 0$ will be exploited. On

the other hand, Fig. 4.9b shows that when the RF frequency is kept fixed and well-aligned with the resonance ($\Delta_{RF} = 0$), larger voltages (smaller Γ -s) result in shorter bunches. Both observations make sense intuitively: smaller Γ leads to narrower channels between buckets, and $\Delta_{RF} = 0$ ensures that the beam is bunched precisely when it becomes resonant and exits the synchrotron.




Figure 4.9: Extracted bunch length (σ , normalised to the RF period $1/f_{RF}$) vs. bucket offset Δ_{RF} (Subfig. 4.9a) or stable-phase sine Γ (Subfig. 4.9b), when the other parameter is kept fixed at the shown value (simulated in *henontrack*).

4.4.2 Extracted bunch shape

Due to the non-linear nature of this bunching method, it is important to study not only the length, but also the shape of the extracted bunch. This is because the shape of the incident bunches will determine the time distribution of the produced neutrinos and hidden-sector particles. SHiP would like to include this information in their analysis pipeline.

Figure 4.10a shows the extracted beam distribution in longitudinal phase space, for simulations of varying Δ_{RF} and holding Γ constant at $\Gamma = 0.05$. As expected, the bunch is short when $\Delta_{RF} = 0$ and becomes longer as $|\Delta_{RF}|$ increases. At some point, the tail of the bunch overlaps with the head of the next bunch, eliminating the search window altogether. Something to keep in mind is that the extracted-bunch shape will in general be asymmetrical. This is no surprise: since all particles approach the empty bucket from the top, they will follow trajectories along the tilted RF channels. Nevertheless, when Δ_{RF} is close to zero this asymmetry is small, which coincides with σ being at its minimum.


Figure 4.10b shows the extracted beam distribution in longitudinal phase space, for simulations of varying Γ and Δ_{RF} constant at $\Delta_{RF} = 0$. When Γ is small, the RF provides large voltage kicks and the bucket height is much larger than the stop-band width. In this scenario, the extracted longitudinal distribution is strongly deformed along the bucket's separatrix arms. As Γ is increased, the effect from the RF becomes weaker and the extracted distribution 'rests' along the bucket edge without being heavily perturbed. This latter case is desirable, since it will have a smaller impact on the transverse beam dynamics and, consequently, on beam loss.

4.4.3 Integrated intensity


In conventional bunched extraction, particles execute phase-momentum oscillations around the RF-bucket's stable fixed point. Therefore, if a particle is not extracted during its first pass through the resonance, it will get many subsequent chances to do so as its Q_x oscillates back and forth due to chromatic coupling. In empty-bucket channelling this is no longer the case: particles are phase displaced from above to below the RF bucket, being bunched only during this crossing. Figure 4.11 shows that, if the RF kicks are large enough, this phenomenon will limit the intensity that can be extracted with empty-bucket channelling. This section quantifies such a limit.

Figure 4.12 shows the results of a (Γ, Δ_{RF}) grid-scan within the SPS-RF parameter limits. It is clear that the extracted integrated intensity I_0 will be reduced if one wants to achieve the SHiP-specification bunch length. For the largest Γ within the 'SHiP specification' region ($\Gamma = 0.05$), 93% of the beam can be extracted, while the nominal de-bunched extraction can approximately extract 99% of the beam. Still, the non-extracted beam is not lost at the ZS and one could envision several strategies to deal with this:

1. Dispose of the remaining beam on the SPS internal dump. The protons-on-target for SHiP would go down slightly, but the concerns regarding ZS activation would be avoided.
2. Employ a transverse damper to slightly blow-up the transverse emittance before and/or during extraction. This technique helps push particles out of the beam core and into the resonance. It is already used successfully in the CERN PS [64]. An implementation in the SPS would require a study of the impact on beam loss due to the transverse blow-up from the exciter.



(a) $\Gamma = 0.05$



(b) $\Delta_{RF} = 0$

Figure 4.10: Slow-extracted particle distribution (simulated in *henontrack*) in longitudinal phase space vs. bucket offset Δ_{RF} (Subfig. 4.10a) or stable-phase sine Γ (Subfig. 4.10b). The density plots on the τ and δ axes show the projections of the beam distribution onto the τ and δ coordinates, respectively. The relative time-of-arrival τ has been aliased to fit all particles within one RF period to provide the distribution of the ‘average’ bunch. The contour of the RF bucket (pink) and the transverse stop-band (red) are shown for additional visual aid.




Figure 4.11: Steinbach-diagram illustration of how empty-bucket channelling affects the extracted intensity. An additional time-of-arrival (τ) dimension has been attached to the conventional normalised-amplitude (A_x) and tune-distance (δQ) dimensions to sketch the alignment of the empty bucket. If particles (brown ellipses) with low- A_x receive large RF kicks, they will exit the unstable region on the other side and never get extracted (shown with a red cross).

3. Ramp the magnetic fields in the opposite direction, sweeping the beam through the resonance for a second time. It is likely that the relative beam loss will be large compared to the first sweep, as the beam remaining in the ring will have degraded quality after its first pass through the resonance. However, the loss should be small in absolute terms, as little beam is left to extract.




Figure 4.12: Integrated intensity (I_0 , normalised to total intensity) vs. stable-phase sine (Γ) and bucket offset (Δ_{RF}), simulated in *henontrack*. The dashed red contour shows the parameter domain where the ‘SHiP specification’ ($4\sigma = 1.5$ ns) is achieved.

The dependence of the integrated intensity on Γ and Δ_{RF} separately is shown in Fig 4.13. Figure 4.13a shows that the largest reduction in integrated intensity happens when $\Delta_{RF} = 0$ and it quickly disappears as $|\Delta_{RF}|$ is increased. Figure 4.13a demonstrates that smaller Γ leads to less integrated intensity, as the voltage kicks become larger.





Figure 4.13: Integrated intensity (I_0 , normalised to total intensity) vs. bucket offset Δ_{RF} (Subfig. 4.13a) or stable-phase sine Γ (Subfig. 4.13b), when the other parameter is kept fixed at the shown value (simulated in *henontrack*).

4.4.4 Transverse effects

As discussed in Chap. 2, particles with different A_x -s become unstable at different δ -s. In other words, the resonant beam has a linear correlation of the form $\delta \propto A_x$. On the other hand, during phase displacement, particles with different δ receive different RF kicks $V(\delta)$ (a semi-analytical approximation of $V(\delta)$ is provided in Chap. 5). These two effects are combined during empty-bucket channelling as is illustrated in Fig. 4.14a. The RF kicks may stretch or shrink the effective momentum stop-band width $\hat{\delta}_{stopband,rms}$ which will in turn affect the transverse distribution as the beam reaches the septum.

This coupling will deform the outgoing transverse distribution, changing the presentation of the separatrix on the septum. To first order, the deformation of the separatrix can be characterised by the change on its radial length l_{rms}^{out} and azimuthal width w_{rms}^{out} (See Sec. 3.4.2.3), which have been calculated numerically and are shown in Fig. 4.15 as a function of Δ_{RF} and Γ . Large changes in either can result in an increase in losses at the ZS. The change in length can typically be compensated by adjusting the spiral step, while the change in width remains a more fundamental



(a) Steinbach-diagram illustration of the impact of empty-bucket channelling on the transverse motion. An additional time-of-arrival (τ) dimension has been attached to the conventional normalised-amplitude (A_x) and tune-distance (δQ) dimensions to sketch the alignment of the empty bucket. Particles (brown ellipses) with different δ sample different regions of the channel, which results on A_x -dependent acceleration due to the slope of the resonance boundary.

challenge. This simulation study shows that empty-bucket channelling mainly has an impact on w_{rms}^{out} . Interestingly, a region exists in the Δ_{RF}, Γ -space where w_{rms}^{out} can be reduced. Appendix D.2 explores the mechanism for this phenomenon, as it could possibly be used to reduce slow-extraction losses.

Finally, Fig. 4.16 shows the dependence of w_{rms}^{out} on Δ_{RF} when Γ is kept fixed, and vice versa. The extracted momentum spread $\hat{\delta}_{rms}^{out}$ is shown in the same plot, which covaries almost exactly with w_{rms}^{out} . This confirms the mechanism through which empty-bucket channelling affects the extracted transverse distribution: in essence, empty-bucket channelling modifies the resonance stop-band width by increasing or decreasing the correlation between A_x and δ (i.e. ‘changing’ the slope of the resonance boundary). The Γ -dependence is easy to understand: as the voltage becomes large, the relative acceleration between the low- A_x and high- A_x particles becomes big, and the distribution is stretched out. The Δ_{RF} -dependence, on the other hand, is highly non-linear (See Appendix D.2). For the ‘SHiP specification’ case, where good alignment between resonance and bucket ($\Delta_{RF} = 0$) is needed, moderate blow-up of w_{rms}^{out} is expected ($\sim 30\%$) if one operates at $\Gamma = 0.05$.




Figure 4.15: Separatrix length l_{rms}^{out} (left, normalised to nominal extraction) and width w_r^{out} (right, normalised to nominal extraction) vs. stable-phase sine (Γ) and bucket offset (Δ_{RF}), simulated in *henontrack*. The dashed red contour shows the parameter domain where the ‘SHiP specification’ of $4\sigma = 1.5$ ns is achieved.

4.5 Measurements in the SPS

A machine-development test was conducted to demonstrate bunching with empty-bucket channelling in the SPS, running at low-intensity (4×10^{11} particles) to avoid equipment damage from beam loss. Informed by the simulations discussed in the previous section and guided by live feedback during the test, the 200 MHz cavity voltage and frequency were varied each cycle. The three tested voltages were 1 MV, 0.1 MV and 0.05 MV and are referred to as high-V, mid-V and low-V, respectively. These voltages correspond to $\Gamma = 0.005, 0.05, 0.1$. The aim of the test was to conduct a proof-of-principle implementation of the technique in the SPS. The measurement data also allowed one to benchmark the simulation model developed so far.




Figure 4.16: Separatrix width w_{rms}^{out} and extracted momentum spread $\hat{\delta}_{rms}^{out}$ (normalised to their corresponding nominal values) vs. bucket offset Δ_{RF} (Subfig. 4.16a) or stable-phase sine Γ (Subfig. 4.16b), when the other parameter is kept fixed at the shown value (simulated in *henontrack*).


4.5.1 Fast versus slow time-structure

The North Area 62 (NA62) GigaTracKer (GTK) [65] was used for this measurement campaign, as it could provide timing information at a resolution of 100 ps. Figure 4.17 shows slow and fast time-structure measurements acquired with the GTK for a shot with and without empty-bucket channelling. As requested by SHiP, the macrostructure (Fig. 4.17a) is mostly unaffected, modulo a slight difference in start time that could be easily compensated in an operational implementation. On the other hand, the beam structure in the nanosecond scale is vastly different between the two cases, with a strong time structure being present for the empty-bucket-channelling shot (Fig. 4.17b). This time structure will be referred to as a ‘bunch’ throughout this chapter.


4.5.2 Extracted bunch length

In order to characterise the length of the extracted bunches, the bunch standard deviation was computed over integration periods of 1 ms and compared to the simulation predictions as shown in Fig. 4.18. It can be seen that the data is in good agreement with the model. Even the global σ minimum (dubbed ‘strong bunching’ before) and the local σ minimum (dubbed ‘weak bunching’ before) are clearly visible in measurement, especially for the mid-V case.

The high-V setting of 1 MV (1/7th of total voltage) causes a clear increase in beam loss (20% or more), but, for the low-V and mid-V cases, the losses become



(a) Slow, binned at 25 Hz



(b) Fast, binned at 10 GHz and averaged over 4 s


Figure 4.17: Extracted time structure with nominal settings and with empty-bucket channelling (bunched). Note the different timescales between Subfig. 4.17a and Subfig. 4.17b. During the test, the magnetic ramp was not fine-tuned to make the spill macro-structure rectangular or to adjust the starting time of the spill, as seen in Subfig. 4.17a. In an operational implementation, this additional step could be performed with the current feed-forward approach, i.e. ‘AutoSpill’ [44].

close to nominal. Amongst the measured settings, the most promising configuration is highlighted in Fig. 4.18b (mid-V, $\Delta_{RF} \approx 0$) and will be referred to as the ‘SHiP candidate’. Its bunch length ($4\sigma = 1.75 \text{ ns} \pm 0.25 \text{ ns}$) satisfies the request with negligible increase in beam loss.


4.5.3 Extracted bunch shape

To further characterise the extracted bunches, as well as to benchmark the model, the bunch profiles were explored next. For each of the three tested voltages, a ‘bunch-grid’ is shown in Fig. 4.19, where the bunch distributions (seen ‘from above’) are plotted as a function of Δ_{RF} . Both from the model and the measurements, one can intuitively see how the bunch shape becomes tighter as Δ_{RF} is varied towards zero, but the dependence is not entirely symmetric with respect to $\Delta_{RF} = 0$, as already predicted in Sec. 4.4.


As a final step, the bunch profile for the ‘SHiP candidate’ was characterised in detail. Its bunch profile is shown again in Fig. 4.20, with focus on observing the distribution tails. It can be seen that the data agrees both with its corresponding Gaussian fit and simulation, up to the ‘Poisson limit’ coming from finite particle counts. In fact, to gain more information about the distribution tails, a higher-statistics run would be needed.



(a) High-V




(b) Mid-V




(c) Low-V


Figure 4.18: Bunch length (σ , normalised to RF period) vs. bucket offset (Δ_{RF}) for three different 200 MHz-cavity voltages (each shown separately in Sub-figs. 4.18a, 4.18b, 4.18c), measured and simulated in *henontrack*. For the measurement, the error-bars show the total millisecond-to-millisecond variation in σ . Total beam loss for each spill is included in the right subplot. The black circle highlights the identified ‘SHiP candidate’.



(a) High-V



(b) Mid-V



(c) Low-V

Figure 4.19: Average bunch profile (particle counts vs. normalised time-of-arrival $\tau \cdot f_{RF}$) vs. bucket offset Δ_{RF} , measured (left) and simulated in *henontrack* (right). Each row represents a single spill, where the black dots show the data-points used for interpolation. Subfigures 4.19a, 4.19b, 4.19c show the outcomes for each of the three different voltages tested with the 200 MHz system.




Figure 4.20: (Left) Bunch profile (particle counts vs. normalised time-of-arrival $\tau \cdot f_{RF}$) for SHiP candidate, and (right) absolute difference between measurement, *henontrack* simulation and Gaussian. The ‘ 3σ limit’ line shows the 3σ significance limit due to finite detector counts: differences above the limit are said to be significant.

All in all, this section demonstrates that (i) the simulation model from Chap 3 makes powerful predictions and can provide insights on the underlying beam dynamics, (ii) empty-bucket channelling can be successfully exploited for beam bunching and that (iii) higher-statistics measurements are ultimately necessary to provide a detailed description of the extracted-distribution tails.

4.5.4 Integrated intensity

To quantify the amount of integrated intensity, the ring’s DC Beam Current Monitor (DC-BCT) was used, subtracting the ring intensity before and after slow extraction. This value is then plotted for all measurement settings in Fig. 4.21. The measurements are in good agreement with the simulation results. For the ‘SHiP candidate’, the integrated intensity is around 94 %.

Comparing the results from the last two sections, there seems to be a fundamental trade-off between the attainable integrated intensity and bunch length. This is not surprising from the discussion above, as the parameters that maximise integrated intensity (low-voltage and large $|\Delta_{RF}|$) are precisely the ones that harm beam bunching. Figure 4.22 demonstrates this trade off by plotting σ v. I_0 for both measurements and simulations. Using the simulation data, a ‘Pareto front’ has been traced, which consists of all points for which one cannot improve σ without worsening I_0 , and vice versa. The points that are not part of the Pareto front are said to be dominated. To find the desired working point, one would choose a compromise between the two metrics along the Pareto front.

4.5.5 Transverse effects

Unfortunately, the test in this chapter was performed at an intensity too low to obtain reliable beam-size measurements of the extracted beam. Suffice to say that the ‘SHiP candidate’ setting did not result in a substantial increase in beam loss, as already shown in Fig. 4.18. Still, the transverse effects of empty-bucket channelling are measured (and compared to simulation) in Sec. 5.8.3 (next chapter), where high-intensity tests were performed.




Figure 4.21: Integrated intensity (I_0 , normalised to nominal case) vs. bucket offset (Δ_{RF}), measured and simulated in *henontrack*. The measurement error-bars show the total shot-to-shot variation. Subfigures 4.21a, 4.21b, 4.21c show the outcomes for each of the three different voltages tested with the 200 MHz system.

4.6 Conclusion

This chapter demonstrated the use of empty-bucket channelling for providing low-loss slow-extracted bunched beams from the SPS. It was shown that the proposed scheme can overcome the limitations of the conventional method to satisfy the SHiP request. By using the 200 MHz system at 0.1 MV, bunches of $4\sigma = 1.75 \text{ ns} \pm 25 \text{ ns}$ at a repetition rate of 5 ns were extracted with minimal increase in beam loss. However, it was shown that the new scheme would reduce the integrated intensity from 99 % to 94 %. This could be mitigated with the help of the transverse exciter or by re-ramping the magnetic lattice. If further optimisation were needed, Appendices B C include additional strategies to add more flexibility to empty-bucket channelling, namely employing double-harmonic RF systems and broad-band RF cavities, respectively.




Figure 4.22: Bunch length (σ , normalised to RF period) vs. integrated intensity (I_0 , normalised to nominal), measured and simulated in *henontrack*. The plot includes the upper limit (‘Ideal’), only achievable if no trade-off existed between σ and I_0 , as well as the Pareto-front line, which consists of all points for which one cannot improve σ without worsening I_0 , and vice versa.

Chapter 5

Spill-quality improvements at CERN SPS

5.1 Motivation

The R&D performed for the slow extraction of bunched beams (described in Chap. 4) was useful to acquire knowledge about RF-cavity control and instrumentation, as well as to benchmark the simulation tools developed in Chap. 3. In this chapter, that knowledge is exploited to improve one of the key aspects of present operation: spill quality.

One can informally define spill quality as the level of uniformity of the intensity profile delivered to the user. Ideally, one would extract a perfectly rectangular spill, whose particle rate $I(t)$ would be described by:

$$I(t) = \begin{cases} I_0/T_s & \text{if } t \in [0, T_s] \\ 0 & \text{elsewhere,} \end{cases} \quad (5.1)$$

where T_s is the spill time and $I_0 = \int_0^{T_s} I(t)dt$ is the total integrated intensity. However, this is not physically attainable, as it would require a ‘crystalline’ beam, i.e. a particle lattice with predictable spacing. In a realistic ‘hot’ beam, the scenario with the smallest rate variations will be given by a Poisson process of mean rate λ :

$$I(t) = \begin{cases} Pois(\lambda = I_0 \cdot \Delta t/T_s) & \text{if } t \in [0, T_s] \\ 0 & \text{elsewhere,} \end{cases} \quad (5.2)$$

where Δt is the chosen time-binning. However, even such a spill is practically impossible to deliver, as slow extraction is sensitive to a variety of undesired perturbations. As an example, Fig. 5.1 shows an SPS spill measurement, where the crystalline and

Poisson spills are also plotted for comparison. It becomes clear that the real variations far exceed those expected from the statistical nature of a non-crystalline distribution.




Figure 5.1: SPS spill (measurement data), Poisson spill and crystalline spill sampled at $\Delta t = 250 \mu\text{s}$.

One of the most dangerous sources of non-uniformity comes from the ring power converters. When the main grid 50 Hz AC current is rectified into a DC current, some small amount of AC power remains. A typical SPS power converter will contain narrow-band peaks at 50 Hz harmonics (i.e. $n \cdot 50 \text{ Hz}$, $n \in \mathbb{Z}$) over a broad-band noise ‘slab’. This spectrum will result in a modulation on the magnetic field and, ultimately, perturb the extracted particle rate. To illustrate this connection between power converter and spill spectra, Fig. 5.2 shows the frequency-domain representations of the SPS spill and the focusing-quadrupole current (the major contributor to horizontal tune perturbations [49]). This comparison also captures the sensitivity of the resonant slow-extraction process: power-converter perturbations of relative magnitude $O(10^{-6})$ result in spill modulations of relative magnitude $O(10^{-1})$, i.e. an amplification of ~ 5 orders of magnitude.

Amongst the North-Area users, the North Area 62 (NA62) experiment has expressed concerns with the large particle-rate variations in the millisecond-scale, which regularly lead to the overload of their First-In-First-Out (FIFO) buffer and loss of valuable data [66]. In fact, concerns about rate variation are common amongst slow-extracted beam users and a variety of mitigation techniques have been developed




Figure 5.2: Measurement of frequency-spectrum magnitudes of the spill (left) and focusing-quad current (right). Both spectra share the same qualitative features: a broad-band floor (dashed grey) and narrow-band spikes (black) at the 50 Hz harmonics.

over the years. These include feedback/feed-forward controllers [67, 44], smoothing via stochastic noise [15, 68] and adjusting the resonant-sextupole strength [46], to name a few. In this chapter, the problem is addressed by employing empty-bucket channelling, leveraging the data and expertise gathered in Chap. 4.

Even though the technique has been deployed before in other machines [16, 69], the first ever implementation in the SPS is presented, which operates at significantly higher beam energies and intensities. Furthermore, a systematic comparison study between measurement and simulation is provided, which compliments the mostly empirical approaches realised in the past. This study starts by presenting a framework for spill-quality modelling and characterisation. Then, the potential of empty-bucket channelling is outlined. After that, an initial low-intensity SPS implementation is discussed. Finally, SPS tests at nominal intensity are presented, which lead to an operational implementation of the technique.

5.2 Spill-quality characterisation

5.2.1 The duty factor

It is customary to quantify the spill quality by the figure of merit known as the duty factor \mathcal{F} [70], which is given by

$$0 \leq \mathcal{F} = \frac{\text{Mean}\{I(t)\}^2}{\text{Mean}\{I^2(t)\}} = \frac{\text{DC power}}{\text{Total power}} \leq 1, \quad (5.3)$$

where $\mathcal{F} = 1$ represents a crystalline spill. Using the relationship $Mean\{I^2\} = Var\{I\} + Mean\{I\}^2$, Eq. 5.3 can be re-written in the form:

$$0 \leq \mathcal{F} = \left[1 + \frac{Var\{I(t)\}}{Mean\{I(t)\}^2} \right]^{-1} = 1 - \frac{Var\{I(t)\}}{Mean\{I(t)\}^2} + (\dots) \leq 1, \quad (5.4)$$

where (...) includes higher-order terms.

Furthermore, one can introduce the Poisson-process statistics $Mean\{I\} = Var\{I\} = I_0 \cdot \Delta t / T$ into Eq. 5.4 to obtain \mathcal{F} 's practical upper limit $\mathcal{F}_{Poisson}$ for a given time-binning Δt :

$$0 \leq \mathcal{F} \leq \mathcal{F}_{Poisson} = \left[1 + \frac{T}{I_0 \cdot \Delta t} \right]^{-1} = 1 - \frac{T}{I_0 \cdot \Delta t} + (\dots). \quad (5.5)$$

This formalism can be used to characterise a spill (or set of spills) as shown in Fig. 5.3 for the SPS. To produce this plot, a representative spill has been binned at various timescales (Δt) and its corresponding Poisson distribution ($\lambda = I_0 \cdot \Delta t / T_s$) generated via a random-number generator. As expected, the real spill has a lower duty factor than the Poisson spill at all timescales. This comparison captures the net effect of all the perturbations present during the slow-extraction process.




Figure 5.3: Measurement of duty factors (\mathcal{F}) vs. bin-width (Δt) for an SPS measurement (data) and the corresponding Poisson limit (Poisson).

5.2.2 From current to intensity

The duty factor characterises intensity variations. However, the goal of this chapter is to mitigate them. To do so, it is important to understand the causal chain

that connects power-converter current ripple to extracted-intensity modulations. For example, let current j have the form,

$$j(t) = j_{ref} + dj, \text{ with } j_{ref} \gg dj, \quad (5.6)$$

where j_{ref} is the requested current and dj is a (small) ripple. We'd like to compute the perturbation dI caused on the extracted intensity I :

$$I(t) = I_{ref} + dI, \quad (5.7)$$

where I_{ref} is the requested intensity.

The propagation process from dj to dI can be broken down into different steps for better conceptual understanding, as shown in Fig. 5.4. First, the current ripple causes a magnetic-strength ripple. Then, that magnetic ripple affects the beam-dynamics macroparameters, such as the tune or the chromaticity. Finally, the beam-dynamics perturbations disrupt the slow-extraction process, resulting in an intensity perturbation.




Figure 5.4: Diagram of current-to-intensity ($dj \rightarrow dI$) perturbation propagation. The perturbation propagation is broken down as follows: current (dj) to j -th magnetic strength (dK_j), to beam-dynamics macroparameters ($dQ_x, dQ', dS\dots$), to extracted intensity (dI). The red box shows the piece of the propagation modelled in *henontrack*.

A full characterisation of the current-to-intensity propagation would require an extremely sophisticated setup, including part-per-million measurement of all magnet currents (and their relative delays), as well as the consideration of physical phenomena such as frequency-dependent vacuum-chamber shielding [71]. However, in this work, a simplified approach will be followed, focusing on the macroparameters-to-intensity part of the propagation. This mapping will be studied numerically using the *henontrack* model described in Chap. 3.

5.2.3 The tune: the main macroparameter

A given magnetic perturbation can affect several macroparameters simultaneously. For example, a quadrupolar perturbation dK_1 will influence both the betatron tune and the chromaticity (amongst others) as follows,

$$dQ_x = \frac{1}{4\pi}\beta_x dK_1 L \quad (5.8)$$

$$dQ' = -\frac{1}{4\pi}\beta_x dK_1 L, \quad (5.9)$$

where β_x is the beta function at the quadrupole location and L is the quadrupole length. It can be seen that both perturbations have the same magnitude.

However, not all macroparameter perturbations have the same size of effect on the extracted intensity. In fact, the size of effect of a perturbation can be quantified by the change it effects on the transverse stable area. Using Eq. 2.39, one obtains:

$$Area \propto \frac{\Delta Q^2}{S^2} \rightarrow \frac{dArea}{Area} = \frac{2}{\Delta Q} d\Delta Q - \frac{2}{S} dS, \quad (5.10)$$

where $\Delta Q/S \sim O(10^{-4} - 10^{-5})$ for a chromatic extraction. It can already be seen that S -perturbations may be neglected when compared to ΔQ -perturbations. Then, one may further expand $d\Delta Q$ as follows,

$$\delta Q = Q_x + Q'\delta - \frac{80}{3} \rightarrow d\Delta Q = dQ_x + \delta dQ', \quad (5.11)$$

where $\delta \sim O(10^{-4} - 10^{-3})$ for a typical particle near the unstable region. Therefore, it can be concluded that a magnetic perturbation dK_1 can be modelled as a tune perturbation dQ_x to first approximation. Past studies in the SPS [49] have confirmed that only modulating the tune produces comparable intensity modulations to the ones produced by rippling K_1 directly.

5.2.4 From tune to intensity: analytical model

Tune perturbations have been identified as the main contributors to the modulation of the extracted intensity. The mapping from dQ to dI will now be approximated following an analytical approach to obtain some critical insights. This can only be done for perturbations slower than the SPS revolution frequency $f_0 = 43.3$ kHz, since the tune is a meaningless construct for faster timescales. As we are mostly concerned with the millisecond-scale, this constraint is not limiting in any significant way. For an initial derivation, two additional simplifying assumptions are made:

- Particles are immediately extracted upon touching the resonant region, i.e. the model focuses on timescales which are longer than the transit time (around 1-10ms for the SPS, See Sec. 5.3.1).
- The resonance boundary does not depend on betatron amplitude, i.e. the extraction is fully chromatic. In the SPS this is a reasonable simplification, since the momentum stop-band width is much smaller than the stack's total momentum spread (See Chap. 2).

In this situation the resonance boundary at the tune $Q_x = 80/3$ acts as a separation between inside and outside the ring, and slow extraction can be modelled as a one-dimensional continuity equation as shown in Fig. 5.5. Then, the extracted intensity $I(t)$ can be expressed in terms of the tune-control program $\underline{Q}_0(t)$ and the undesired time-varying ripple $q(t)$ as follows:




Figure 5.5: Simplified model of chromatic slow extraction. The tune (Q_x) of the beam distribution ($n(Q_x)$, brown) is pushed towards $Q_x = 80/3$, which acts as a boundary between inside (cyan) to outside (magenta) the ring. During a time interval dt , the distribution is displaced towards the resonance by the drift term $\underline{Q}_0 dt$, while the oscillatory term of size $\dot{q} dt$ perturbs this process.

$$I(t) = \left. \frac{\partial n}{\partial t} \right|_{Q_x=80/3} = - \left. \frac{\partial}{\partial Q_x} [n \cdot (\underline{Q}_0 + \dot{q})] \right|_{Q_x=80/3} = \left. \left\{ - \frac{\partial n}{\partial Q_x} (\underline{Q}_0 + \dot{q}) - n \frac{\partial (\underline{Q}_0 + \dot{q})}{\partial Q_x} \right\} \right|_{Q_x=80/3}, \quad (5.12)$$

where n denotes the number of particles. Since \underline{Q}_0 and q are created by the magnets, no dependence on the Q_x coordinate is expected, arriving at the following expression:

$$I(t) = - \rho(Q_x) \left[\dot{\underline{Q}}_0 - \frac{d}{dt} \sum_i 2q_i \cos(2\pi f_i t + \phi_i) \right] \Big|_{Q_x=80/3}, \quad (5.13)$$

where $\rho(Q_x) = \frac{\partial n}{\partial Q_x}(Q_x)$ is the local particle density and the spectrum of q has been expanded, with $2q_i, f_i, \phi_i$ representing the ripple amplitude, frequency and phase components, respectively. Typically, $\dot{Q}_0(t)$ is adjusted to compensate for the slow variations in $\rho(Q_x)$, which allows us to simplify $\rho(Q) \rightarrow \rho$ and $Q_0(t) \rightarrow \dot{Q}_0 t$ without much loss of generality, where \dot{Q}_0 is the average speed at which particles change their betatron tune. This finally leads to

$$I(t) = -\rho[\dot{Q}_0 + \sum_i (2\pi f_i) \cdot 2q_i \sin(2\pi f_i t + \phi_i)]. \quad (5.14)$$

The factor $(2\pi f_i)$ indicates that slow extraction acts as a differentiation operator on the tune. As a consequence, the faster the tune ripple is, the stronger its impact on the extracted intensity.

Finally, one can substitute the obtained expression for $I(t)$ into Eq. 5.3 and, by assuming no correlation between spectral components, the duty factor is given by

$$\mathcal{F} = [1 + x]^{-1} = 1 - x + (\dots), \quad x = \frac{\frac{1}{2} \sum_i (2\pi f_i)^2 (2q_i)^2}{\dot{Q}_0^2}. \quad (5.15)$$

Since \dot{Q}_0 is in the denominator of x , it can be concluded that a faster tune speed leads to better spill quality. This can be intuitively grasped through Fig. 5.6. As the resonance boundary oscillates back and forth, the number of extracted particles will vary for a given $\dot{Q}_0 dt$, depending on the instant in time. But, if \dot{Q}_0 becomes sufficiently large, the oscillatory region will be avoided altogether.




Figure 5.6: Steinbach diagram of particles (brown) being pushed towards the unstable region by an amount $\dot{Q}_0 dt$ under the influence of ripple. Due to ripple, the unstable wedge oscillates back (fuchsia) and forth (light pink). When \dot{Q}_0 is small (solid black), the number of extracted particles varies between one and three. When \dot{Q}_0 is large (dashed blue) all three particles get extracted, independently of the wedge location.

However, increasing \dot{Q}_0 naively by a faster magnetic ramp would also lead to a shorter spill. Still, there are clever ways of counteracting this, such as increasing the beam momentum spread (since $dQ_x = Q' \cdot d\delta$), which is why most chromatic slow extractions are preceded by ‘momentum-stretching’ RF gymnastics. In fact, it will later be shown that empty-bucket channelling also exploits the increase of \dot{Q}_0 in a productive way.

5.3 Transfer function

The framework described in the previous section can be naturally expanded with the formalism of control theory, making a direct connection between time domain $q(t), I(t)$ and frequency domain $q(f), I(f)$ representations of ripple and intensity over the spill time T_s via the Fourier transform ¹,

$$q(f) = \int_0^{T_s} q(t) \exp(-i2\pi ft) dt, \quad I(f) = \int_0^{T_s} I(t) \exp(-i2\pi ft) dt, \quad (5.16)$$

and finding the transfer function \mathcal{T} that connects them:

$$\underline{I}(f) = \frac{I(f)}{I_0} = \mathcal{T}(f; \bar{p}) \circ q(f), \quad (5.17)$$

where \circ is the function-composition operation, \bar{p} is a set of machine parameters such as chromaticity and sextupole strength, and $\underline{I}(f)$ is the intensity normalised by the DC component $I_0 = I(f=0) = \int_{T_s} I(t) dt = \rho \dot{Q}_0 T_s$. I_0 is established by the experimental request and, thus, perturbations are only meaningful in their size relative to this baseline.

By substituting $I(t)$ from Eq. 5.14 into Eq. 5.17, one can already identify a first expression for \mathcal{T} :

$$\mathcal{T} = 2\pi f \cdot \frac{1}{\dot{Q}_0}, \quad (5.18)$$

which can be conceptualised as a differentiator circuit followed by an amplifier of gain $G = 1/\dot{Q}_0$, as represented in the block diagram in Fig. 5.7. As expected, the implications are identical to the ones highlighted in the time domain: (i) fast ripple is amplified, (ii) large \dot{Q}_0 reduces the impact of ripple.

¹Control theory works with the more general Laplace transform with basis $\exp(-st) = \exp[-(\lambda + i2\pi f)t]$, but in this chapter we set $\lambda = 0$ since we are only concerned with oscillatory terms.



Figure 5.7: Block diagram of slow extraction for timescales slower than the transit time

5.3.1 Low-pass filter

An important feature of slow extraction, which has been ignored so far by restricting the analysis to slow timescales, is the fact that particles take a non-zero transit time $T_{transit}$ to reach the septum. Moreover, the time depends on the initial phase-space coordinates, producing a spread $\Delta T_{transit}$. These characteristics define the ‘building block’ of the spill, formally called the impulse response or Green’s function [72], i.e. the system output to a Dirac-delta input. The transit time and its spread add a delay $\tau \propto T_{transit}$ and a low-pass filter with cutoff frequency $f_c \propto 1/\Delta T_{transit}$ to the transfer function, respectively. In fact, one can actively try to maximise $\Delta T_{transit}$ by adjusting the machine parameters [46], thus decreasing the threshold f_c above which perturbations are suppressed. It is beyond the scope of this thesis to study the detailed dependence of $T_{transit}$ and $\Delta T_{transit}$ on machine parameters, which is comprehensively explored in [73, 19]. In so far as this chapter is concerned, it is relevant to note that the SPS transfer function exhibits a low-pass filter whose cutoff is $f_c \approx 100$ Hz [46]. For this reason, studies in this chapter will focus on ripple frequencies in the range $\sim 10 - 200$ Hz, as larger frequencies are already heavily suppressed by the low-pass filter.

Fig. 5.8 compiles all the observations made so far, providing a block-diagram model of slow extraction. This model establishes a useful framework to discuss spill-quality phenomena.




Figure 5.8: Block diagram of slow extraction

5.3.2 Linearity and time invariance of the spill structure

In order to obtain the transfer function \mathcal{T} from data, one needs to have knowledge of both the input q and the output \underline{I} . In simulation it is easy to obtain both, but in the real machine it is practically very challenging to measure q , as the perturbations are typically in the part-per-million range with respect to the nominal-tune band.

However, this thesis is only interested in relative improvements in spill quality, and not in an absolute characterisation of the SPS transfer function. For the latter, the reader is referred to [49]. For the former, it is possible to ‘normalise away’ the specifics of the input q if one assumes linearity and time invariance. These two assumptions are addressed below.

5.3.2.1 Linearity

So far, the particular size of the input perturbation q_i has not been addressed, other than loosely specifying that it should be small. In fact, the ‘smallness’ of q_i can be made precise. If the inequality

$$|\dot{Q}_0| \geq \left| \sum_i (2\pi f_i) \cdot 2q_i \sin(2\pi f_i t + \phi_i) \right| \quad (5.19)$$

holds, the ripple q_i is said to be small. On the contrary, if the inequality is violated, the contributions from the modulations q_i become larger than the mean speed \dot{Q}_0 . At this point, there will be certain moments when the beam distribution recedes from the resonance, instead of moving towards it. This breaks down the ‘continuity equation’ model described by Eq. 5.12, as it would predict that particles are ‘un-extracted’ when $-(\dot{Q}_0 + \dot{q}) < 0$. This is obviously unphysical.

In reality, if the ripple becomes large enough, the spill breaks up into discrete intensity bursts as illustrated in Fig. 5.9. In the frequency-domain interpretation, an input modulation at frequency f_i will not only excite an output modulation at frequency f_i , but also at all higher harmonics required to ‘compose’ discrete bursts ².

At this point, \mathcal{T} will lose the linearity properties given by,

$$\mathcal{T} \circ (\alpha q_1) = \alpha \mathcal{T} \circ q_1, \quad \mathcal{T} \circ (q_1 + q_2) = \mathcal{T} \circ q_1 + \mathcal{T} \circ q_2, \quad (5.20)$$

with $\alpha \in \mathbb{C}$, and ripple inputs q_1, q_2 . When this threshold has been crossed the concept of a ‘transfer function’ becomes much less useful as most of its generality is lost.

²This phenomenon is analogous to that caused by half-wave rectifiers in AC-to-DC conversion.




Figure 5.9: Illustration of three perturbed spills with sinusoidal tune perturbations of varying perturbation sizes. For a linear perturbation (blue), the flux oscillates around the mean in time domain. In frequency domain, a single spectral component is excited. For a non-linear perturbation (red), the spill breaks up into discrete bursts in time domain, and several harmonics are excited in the frequency domain. For an infinite/asymptotic perturbation (grey), the spill consists of periodic delta functions both in time and frequency domain.

For a single sinusoidal ripple of amplitude $2q_i$, the inequality from Eq. 5.19 can be simply expressed as,

$$|\dot{Q}_0| \geq 2\pi f_i \cdot 2q_i. \quad (5.21)$$

The linearity threshold can be verified in simulation, by introducing a sinusoidal tune ripple with increasingly large $2q_i$. The result of such a simulation study is shown in Fig. 5.10, where the procedure has been applied for several input-ripple frequencies f_i . As soon as the linearity threshold is passed, the transfer function at f_i starts to saturate. The ‘missing’ power is used to excite high harmonics of f_i , i.e. $n \cdot f_i, n \in \mathbb{Z}^+$ (not shown), as was already illustrated in Fig. 5.9.

Supported by previous characterisations of the SPS spill [49], this chapter assumes that condition 5.19 holds, which enables us to work in the linear regime. Under linearity, \mathcal{T} does not depend on the magnitude of q and one can compute the relative ripple-reduction coefficient $G(f)$ between an old and a new machine configuration as follows:

$$G(f) = \left| \frac{\mathcal{T}_{new}(f)}{\mathcal{T}_{old}(f)} \right| = \left| \frac{\underline{I}_{new}(f)}{\underline{I}_{old}(f)} \right|, \quad (5.22)$$

where \underline{I}_{new} and \underline{I}_{old} are the new and old extracted-intensity spectra, respectively, which can be directly obtained from measurement/simulation. Notice that there is no need to have knowledge about the specifics of $q(f)$. Even if linearity were not satisfied exactly, the error would grow slowly with the magnitude $|q|$, still making linear analysis a useful first-order approximation.




Figure 5.10: Spill-ripple amplitude ($|I(f_i)|$, normalised to integrated intensity I_0) vs. tune-ripple amplitude ($q(f_i)$) for different ripple frequencies (f_i), simulated in *henontrack*. $q(f_i)$ has been normalised so that the linearity boundary is reached when the input equals 1. The physical limit is reached asymptotically, as the input amplitude grows towards infinity.

5.3.2.2 Time invariance

Linearity is crucial in order to ‘normalise away’ the specific input q , but it is not sufficient. If one wants to compare two different instants in time (either within the same spill or across different spills), time invariance of q is also required.

Under time invariance, one can compute the relative ripple-reduction coefficient $G(t)$ between an old and a new machine configuration as follows:

$$G(t) = \sqrt{\frac{1 - \mathcal{F}_{new}(t)}{1 - \mathcal{F}_{old}(t)}}, \quad (5.23)$$

where \mathcal{F}_{new} and \mathcal{F}_{old} are the new and old duty factors, respectively, which can directly be obtained from simulation/measurement.

In the real machine, time invariance is not exactly satisfied. For example, the voltage and frequency of the main grid are known to have modulations over time, which will inevitably affect the power-converter AC-to-DC conversion. Still, time invariance is approximately satisfied and its violation will simply be taken as a source of statistical error.

5.3.3 Nominal SPS transfer function

Finally, the nominal SPS slow-extraction transfer function was obtained from simulation. This transfer function will serve as a baseline to quantify the ripple-reduction

obtained via empty-bucket channelling. It also demonstrates the features from the block-diagram framework outlined in this chapter. Two different methods were used to obtain the transfer function:

- Sinusoidal-input scan: A single sinusoidal perturbation is injected for each spill. The frequency f_i is scanned across the relevant range from spill to spill. The tune perturbation is applied each turn k as follows:

$$q[k] = 2q_i \sin(2\pi f_i k). \quad (5.24)$$

- Noise input: a wide-band perturbation is injected for each spill, with a spectral-density profile that covers the entire frequency range of interest. Several spills are averaged to reduce the intrinsic variations of the noise. A Gaussian-noise tune perturbation is applied each turn k as follows:

$$q[k] \sim \mathcal{N}(\mu = 0, \sigma). \quad (5.25)$$

Since both inputs and outputs are known in simulation, one can then extract the transfer function from either method. The former method can be conceptualised as a point-by-point approach, while the latter is a statistical approach. Figure 5.11 shows the outcome of the study, highlighting all the characteristics discussed so far. The result is in good quantitative agreement with previous tune-to-intensity characterisations of the SPS transfer function [46], showing a cutoff frequency $f_c \approx 100$ Hz and a low-pass filter of approximately second order. In conclusion, the framework of control theory allows us to formalise the problem of spill quality. Moreover, simulation tools can be exploited to extract quantitative insights about \mathcal{T} . The same approach will be used throughout this chapter.

5.4 Empty-bucket channelling for ripple reduction

5.4.1 Concept

The concept and basic parametrisation of empty-bucket channelling were described in detail in Chap. 3. It was shown that one can restrict longitudinal trajectories to tight channels between buckets as transverse trajectories gain amplitude and reach the septum. Furthermore, due to chromaticity Q' , the longitudinal kicks from the RF have an effect on the transverse motion, as a voltage kick V over one revolution period $1/f_0$ effectively becomes a transverse-tune kick $\int_0^{1/f_0} dQ_x$ via the relationship




Figure 5.11: Nominal SPS transfer function (\mathcal{T}), simulated in *henontrack*. Both the magnitude (top) and the phase (bottom) of \mathcal{T} are shown.

$$\int_0^{1/f_0} dQ_x = Q' \int_0^{1/f_0} d\delta = Q' \cdot \frac{eV}{\beta_r^2 E_0}. \quad (5.26)$$

This is, in essence, the coupling one can exploit to improve spill quality, as originally realised in the 1980s by Cappi and Steinbach [16]. As described in the section above, increasing the tune speed \dot{Q}_0 across the resonance boundary is equivalent to reducing the gain of the transfer function that connects tune ripple to spill. If one does so naively by increasing the magnetic-ramp speed by a factor K , the spill length will be shorter by that same factor. However, by using an empty RF bucket near the resonance boundary, one can create that same speed-up K only locally, as illustrated in Fig. 5.12. This splits momentum/tune space into two regions: (i) a highly populated region far from the RF frequency, where the tune speed is low and determined by the magnetic ramp, (ii) a low-population region with a tune speed higher by a factor K , where RF and beam are close in frequency, causing voltage kicks to add up

coherently.




Figure 5.12: Empty-bucket channelling for ripple suppression, represented in longitudinal phase space (bottom) and Steinbach diagram (top), simulated in *henontrack*. The RF generates empty buckets (fuchsia) aligned with the resonance (red). Particles (sky-blue) channel between consecutive buckets and cross the resonance boundary at high tune-speeds.

In fact, it is easy to derive K 's upper bound for a given peak voltage V_0 . At most, particles will have a net tune speed $\dot{Q}_{0,max}$ given by

$$\dot{Q}_{0,max} = \dot{Q}_0 + f_0 \int_0^{1/f_0} dQ_x = \dot{Q}_0 + f_0 Q' \frac{eV_0}{\beta_r^2 E_0}, \quad (5.27)$$

i.e. the nominal tune speed plus the additional contribution from the peak of the RF wave (from Eq. 5.26). This leads to a maximum speed-up K_{max} that can be expressed as

$$K_{max} = \frac{\dot{Q}_{0,max}}{\dot{Q}_0} = 1 + \frac{1}{\Gamma}, \quad (5.28)$$

where $\Gamma = \sin \phi_s$ is the stable-phase sine. In reality, particles with different arrival times will sample a different voltage along the sinusoidal RF waveform, yielding a smaller average speed-up than K_{max} . Regardless, the net speed-up can be made quite

large as particles are not uniformly distributed along the ring during channelling. In short, the improvement in spill quality can be understood as a consequence of the average speed-up K across the resonance boundary, which will depend on the empty-bucket channelling parameters as well as the beam distribution.

We can conclude that empty-bucket channelling introduces a ripple-reduction coefficient $G_{ebc} = 1/K$, which corresponds to adding an attenuator of gain $G_{ebc} < 1$ to the slow-extraction block diagram, as shown in Fig. 5.13. The appeal of empty-bucket channelling becomes apparent in comparison to actual ‘control theory’-based (e.g. feedback, feed-forward) solutions: a broad-band ripple suppression can be achieved through the beam dynamics directly, circumventing typical concerns such as loop stability and measurement noise. G_{ebc} can be extracted from simulations/measurements using Eqs. 5.23–5.22.




Figure 5.13: Block diagram of slow extraction with empty-bucket channelling

5.4.2 Semi-analytical estimate of ripple-reduction coefficient

A semi-analytical formula for K can be derived in terms of the bucket offset $\Delta_{RF} = (\delta_{RF} - \delta_{res})/\hat{\delta}_{bucket}$, which controls the bucket alignment; and the stable-phase sine Γ , which controls the maximum speed-up. Figure 5.14 provides a visual sketch of the derivation (details available in [74]), which is performed as follows:

1. For a given Γ and Δ_{RF} , the longitudinal Hamiltonian from Eq. 2.53 is used to identify the channel boundary points τ_{left} and τ_{right} at the resonant momentum.
2. The RF voltage is averaged over the channel domain, which, if aligned properly for ripple suppression, only samples points where the speed-up is large.
3. The average voltage is compared to the energy change per turn provided by the magnetic ramp (Eq. 3.7), arriving at the expression:

$$K = -\frac{1}{\Gamma \Delta \phi} \int_{\phi_{left}}^{\phi_{right}} \sin \phi d\phi = 1 + \frac{1}{\Gamma} \cdot \alpha, \quad \alpha = \left[\frac{\cos \phi_{right} - \cos \phi_{left}}{\Delta \phi} \right], \quad (5.29)$$




Figure 5.14: Illustration of tune speed-up due to empty-bucket channelling, shown in longitudinal phase space. Particles crossing the resonance sample the RF voltage along the channel domain ($\tau \in (\tau_{left}, \tau_{right})$, solid black), which provides additional acceleration (quantified by the RF voltage divided by the energy-change per turn $\frac{eV}{\Delta E}$).

where $\phi = -2\pi f_{RF}\tau$ is the phase w.r.t the RF voltage and $\Delta\phi = \phi_{right} - \phi_{left}$ the channel-width. Figure 5.15 shows K for $\Gamma = 0.1$ as a function of Δ_{RF} . Several observations can be made both from Eq. 5.29 and Fig. 5.15:

- K is directly proportional to the RF voltage, i.e. inversely proportional to Γ .
- K can be interpreted as K_{max} from Eq. 5.28 weighted by a geometric factor $\alpha \leq 1$, which depends on the alignment Δ_{RF} ³.
- K is not maximised when resonance and bucket are perfectly aligned ($\Delta_{RF} = 0$), but when Δ_{RF} is offset towards the waiting beam ($\Delta_{RF} > 0$).

³One can verify that one recovers K_{max} (i.e. $\alpha \rightarrow 1$) by re-writing $\phi_{right} = \phi_{left} + \Delta\phi$ and taking the limit $\phi_{left} \rightarrow \pi/2$, $\Delta\phi \rightarrow 0$, i.e. moving towards the crest and reducing the channel-width

- The speed-up is not symmetrical w.r.t. $\Delta_{RF} = 0$ since the accessible phase-space region is different when moving into ($\Delta_{RF} < 0$) or away ($\Delta_{RF} > 0$) from the empty bucket.

Intuitively, we may picture the beam as an incompressible fluid travelling through the RF channel (See Sec. 2.5.2), where its vertical velocity determines the speed-up K . Therefore, the narrower and the more up-right the channel is, the larger K becomes.




Figure 5.15: Channel-width and tune speed-up (K) for a representative case of empty-bucket channelling, computed semi-analytically. In this picture, the beam comes from above and encounters the resonance at some point in the traversal through the channel, which determines its K . This is why the bucket offset (Δ_{RF}) goes from negative to positive values.

The semi-analytical estimate of the ripple-reduction coefficient ($G_{ebc} = 1/K$) will be useful to explain the depth of the ripple-suppression region, as a function of relevant parameters.

5.4.3 Ripple reduction for large bucket offsets

The semi-analytical approach above assumed that particles were extracted immediately upon crossing the resonance, which allowed us to equate K to the acceleration across the boundary. Moreover, K was computed by averaging over all phases available within the RF channel, which yielded $K = 1$ as soon as $|\Delta_{RF}| > 1$ (since $\int_0^{2\pi} \sin(\phi) d\phi = 0$).

In practice, as discussed in Sec. 5.3.1, unstable particles take $O(100-1000)$ turns to get extracted and, therefore, stay near the stability boundary for non-negligible spans of time. Furthermore, particles that are kicked towards the resonance (positive K) are

more likely to become unstable than those kicked away from it (negative K), which means that the empty bucket may still produce a speed-up even if $|\Delta_{RF}| > 1$. Note that in this latter case the name empty-bucket ‘channelling’ is somewhat misleading, since particles are not really channelled between empty buckets but rather oscillate in momentum as shown in Fig. 5.16. Therefore, empty-bucket rippling will be used for configurations where $|\Delta_{RF}| > 1$.




Figure 5.16: Illustration of empty-bucket rippling for a large bucket offset $\Delta_{RF} = -1.5$. Even though trajectories (black) are not channelled between buckets, they still preferentially point into the resonance (solid), as supposed to away (dashed), because the beam approaches the resonance from above.

As a consequence of these two phenomena, the beam experiences an ‘effective’ K different from the instantaneous K from Eq. 5.29. In fact, one can think of the simulation studies in this chapter as a tool to numerically compute this effective K . Still, a simple heuristic study is presented in this section, which will help intuitively explain the dependence of G_{ebc} on various parameters. Let’s suppose that a positive voltage kick accelerates a particle towards the resonance and a negative RF kick decelerates it away from it⁴. Then, for a strip of particles with momenta slightly smaller than the resonant momentum, particles near the RF crest are the ones most likely to be pushed into the resonance by the RF kicks. If one focuses on a particle starting precisely at the RF crest, two observations can be made:

⁴In the SPS the opposite is true, since particles start with momenta larger than the resonant momentum. In any case, the same logic applies.

- During a ripple period $T_i = \frac{1}{f_i}$, the particle slips away from the crest, subsequently receiving smaller RF kicks. If the period is short (f_i is large), this effect is small. On the contrary, if the period is long (f_i is small), the particle may slip enough to reach the trough of the RF wave, resulting in a reduction of the effective K .
- As the particle drifts in phase with respect to the RF, the time needed to slip from the peak to the trough will determine the effective voltage integrated along the trajectory for a fixed time T_i . If the time to slip by one full period T_{slip} is short, the slippage effect will be significant and will reduce the effective K by averaging out all RF phases.

A graphical illustration of the two criteria is shown in Fig. 5.17. These observations may be summarised by computing the ratio r as follows:

$$r = \frac{T_i}{T_{slip}} \approx \frac{h f_0 \eta \Delta_{RF} \hat{\delta}_{bucket}}{f_i}, \quad (5.30)$$

where $\hat{\delta}_{bucket}$ is the bucket height from Eq. 2.56, which depends on Γ and h (amongst other lattice parameters). It has been assumed that $|\Delta_{RF}| \gg 1$ so that trajectories have practically constant δ and T_{slip} can be taken to be proportional to the relative momentum offset.

When r is small/large, the effective K is large/small. This simple metric r will be useful to explain the width of the ripple-suppression region with respect to Δ_{RF} , as a function of other relevant parameters.

5.5 Ripple-reduction simulations

This section explores the dependence of G_{ebc} on the bucket offset Δ_{RF} , the stable-phase sine Γ , the ripple frequency f_i and the RF harmonic h . These simulation studies are crucial to optimise the ripple suppression effect of empty-bucket channelling and select a good working point.

5.5.1 Dependence on RF harmonic

The first choice to be made in the SPS is whether to use the 200 MHz ($h = 4620$) or the 800 MHz ($h = 18480$) system to produce empty buckets. Section 5.6 uses the former, as data was already available from the bunched-beams test described in




Figure 5.17: Illustration of RF slippage (dashed and dotted lines) along the RF waveform during one ripple period T_i . The effective voltage (integral along the slippage) will depend on the slippage time T_{slip} and the ripple period T_i .

Chap. 4. Section 5.7 uses the latter, since it will turn out to be more convenient for an operational implementation.

In order to explore the functional dependence on the RF harmonic, a simulation grid-scan was performed, where h and Δ_{RF} were varied while keeping $\Gamma = 0.02$ and $f_i = 50$ Hz. This particular Γ has been chosen because it will be the one used later in Sec. 5.7. In any case, the dependence on Γ will be studied in detail in the following section.

Unlike in the real machine, any h can be chosen in the simulation study, providing a more complete picture of the functional dependence of G_{ebc} , as shown in Fig. 5.18. It can be seen that lower h provides a wider ripple-reduction valley, which would make the implementation more robust to jitter and errors in general. This dependence is well captured by r : a smaller h leads to a longer RF period; therefore, a particle at the RF crest takes a longer time to slip out of phase and ‘lose’ its initial speed-up K .

The functional dependences for the SPS RF systems ($h = 4620, 18480$) are shown again in Fig. 5.19, where h has been fixed and the variation of G_{ebc} with respect to Δ_{RF} is shown. The same pattern as above is clear: the higher-harmonic cavity exhibits a narrower valley of ripple suppression. In this regard, the 200 MHz system is more desirable for an operational implementation.




Figure 5.18: 50 Hz-ripple-reduction coefficient ($G_{abc}(50 \text{ Hz})$, Subfig. 5.18a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.18b) vs. RF harmonic h and bucket offset Δ_{RF} (for stable-phase sine $\Gamma = 0.02$), simulated in *henontrack*.




Figure 5.19: 50 Hz-ripple-reduction coefficient ($G_{abc}(50 \text{ Hz})$, Subfig. 5.19a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.19b) vs. bucket offset Δ_{RF} (for RF harmonics $h = 4620, 18480$ and stable-phase sine $\Gamma = 0.02$), simulated in *henontrack*. The ‘Poisson limit’ line shows the minimum G achievable due to the finite statistics of the simulation.

5.5.2 Dependence on RF voltage

Once the RF harmonic has been specified, its voltage must be chosen, which will determine Γ for a given momentum-ramp rate through Eq. 3.8. Figure 5.20 shows the dependence of $G_{abc}(50 \text{ Hz})$ on Γ and Δ_{RF} for $h = 18480$. This harmonic has been chosen instead of $h = 4620$ because it will be the one used in Sec. 5.7. In any case, the suppression pattern for other harmonics has already been discussed in the previous section. The contour plot verifies the prediction from the semi-analytical estimation of K : optimal suppression is obtained when the empty bucket is offset towards the beam $\Delta_{RF} > 0$.

Figure 5.21 shows $G_{abc}(50 \text{ Hz})$ for three different Γ -s, demonstrating that higher




Figure 5.20: 50 Hz-ripple-reduction coefficient ($G_{abc}(50 \text{ Hz})$, Subfig. 5.20a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.20b) vs. stable-phase sine Γ and bucket offset Δ_{RF} (for RF harmonic $h = 18480$), simulated in *henontrack*.

voltage (smaller Γ) leads to better ripple suppression. The explanation is simple: the larger the voltage, the larger the tune speed. Interestingly, $G_{abc}(50 \text{ Hz})$ exhibits two minima as a function of Δ_{RF} , one for $\Delta_{RF} > 0$ (global min.) and one for $\Delta_{RF} < 0$ (local min.). This is a real phenomenon coming from the non-linear nature of empty-bucket channelling and is later shown in measurement in Sec. 5.7. The mechanism is explained conceptually in Appendix D.3.




Figure 5.21: 50 Hz-ripple-reduction coefficient ($G_{abc}(50 \text{ Hz})$, Subfig. 5.21a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.21b) vs. bucket offset Δ_{RF} (for stable-phase sines $\Gamma = 0.1, 0.02, 0.005$ and RF harmonic $h = 18480$), simulated in *henontrack*. The ‘Poisson limit’ line shows the minimum G achievable due to the finite statistics of the simulation.

5.5.3 Dependence on ripple frequency

Once h and Γ have been chosen, the last step is to scan Δ_{RF} and find a working point where ripple suppression is optimised. To do so, G_{ebc} is computed for a few outstanding frequencies f_i in the spill as Δ_{RF} is varied. Figure 5.22 shows the G_{ebc} landscape as a function of f_i and Δ_{RF} , with $\Gamma = 0.02$ and $h = 18480$. It can be observed that slower f_i show narrower suppression valleys. This is consistent with the qualitative metric r introduced in Sec. 5.4.3, Eq. 5.30: the slower f_i , the longer the time T_i particles have to slip away from the RF crest.




Figure 5.22: Ripple-reduction coefficient ($G_{ebc}(f_i)$, Subfig. 5.22a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.22b) vs. ripple frequency f_i and bucket offset Δ_{RF} (for Rf harmonic $h = 18480$ and stable-phase sine $\Gamma = 0.02$), simulated in *henontrack*.

In the case of the SPS, the 50 Hz and 100 Hz ripples are the most prominent, and therefore the best observables to optimise spill quality. Figure 5.23 shows the dependence of G_{ebc} on Δ_{RF} for these two frequencies. The 100 Hz valley is wider than the 50 Hz valley, which is also encapsulated in r . This means that it is easier to minimise $G_{ebc}(f_i = 100 \text{ Hz})$ for an initial coarse search, which must be later complemented with a minimisation of $G_{ebc}(f_i = 50 \text{ Hz})$ during a finer scan.

5.6 Low-intensity implementation with the 200 MHz system

The data used below were extracted from the acquisitions made for bunched beams in the previous chapter, allowing us to demonstrate the ripple-suppression effect of empty-bucket channelling without the need to request additional beam test time. These data include three different 200 MHz RF-system voltages (high-V: 1 MV, mid-V: 0.1 MV and low-V: 0.05 MV) at a variety of different frequency offsets Δ_{RF} . They





Figure 5.23: Ripple-reduction coefficient ($G_{abc}(f_i)$, Subfig. 5.23a) and metric from Sec. 5.4.3, Eq. 5.30 (r , Subfig. 5.23b) vs. bucket offset Δ_{RF} (for ripple frequency $f_i = 50$ Hz, 100 Hz, RF harmonic $h = 18480$ and stable-phase sine $\Gamma = 0.02$), simulated in *henontrack*. The ‘Poisson limit’ line shows the minimum G achievable due to the finite statistics of the simulation.


correspond to $\Gamma = 0.005, 0.05, 0.1$, respectively. It is important to note that the data were taken at $\sim 1/100$ th of nominal intensity, i.e. 4×10^{11} protons-per-pulse.

Via spectral analysis of the spill time structure, $G_{abc}(f_i = 50, 100$ Hz) was computed for each Γ, Δ_{RF} pair, by using Eq. 5.22. Figure 5.24 shows the outcome of such analysis, as well as the corresponding simulation output and the semi-analytical estimate $G_{abc} = 1/K$ (where K is computed from Eq. 5.29). The semi-analytical estimate can predict the maximum achievable suppression (smallest G_{abc}) for a given voltage, but fails to capture the dependence on Δ_{RF} (as expected from Sec. 5.4.3, Eq. 5.30). This is already useful, since an approximate best G_{abc} can be evaluated at practically no cost for a given available RF voltage. On the other hand, the simulation output tracks the measurements reliably for the low-V and mid-V cases. This suggests that it could be employed to explore a larger parameter space offline or even construct a surrogate model of the $(\Gamma, \Delta_{RF}) \rightarrow G_{abc}$ mapping, which has been done and exploited for other processes in the SPS [63]. The model even reproduces the distinct behaviours of $G_{abc}(f_i = 50$ Hz) and $G_{abc}(f_i = 100$ Hz), whose dependencies on Δ_{RF} are not exactly the same. This is not surprising, since the ‘effective’ G_{abc} is not expected to be identical as the ratio r from Sec. 5.4.3, Eq. 5.30 will be different. For the high-V case, however, a larger disagreement can be observed between data and simulation. For this configuration, the perturbation from the RF kicks is very large, and a more complete model of the SPS beam dynamics would be needed to better capture the impact on the spill ripple.


Overall, the measurements demonstrate that significant ripple suppression of the



(a) High-V



(b) Mid-V



(c) Low-V

Figure 5.24: Ripple-reduction coefficient G_{ebc} vs. bucket offset Δ_{RF} for three different 200 MHz-system voltages (each shown separately in Subfigs. 5.24a, 5.24b, 5.24c). The plots include a semi-analytical estimate (formula), measurement data and simulation in *henontrack*. The measurement error-bars show total shot-to-shot variations. Horizontal black lines have been drawn at suppression factors of 1, 2 and 5x for additional visual aid.

50 Hz and 100 Hz ripples can be achieved via empty-bucket channelling. It is important to mention, however, that the higher the voltage the larger the perturbation on the extracted beam will be. In particular, as studied in Chap. 4, the high-V setting produces a significant increase in beam loss due to substantial emittance blow-up of the extracted separatrix. This makes it unsuitable for an operational implementation as already discussed in Chap. 4. Still, both the mid-V and low-V setups include settings with negligible loss increase and spill-quality improvements of $\sim 5x$. In general, one should choose the minimum voltage that still leads to acceptable ripple suppression.

5.7 High-intensity implementation with the 800 MHz system

Having demonstrated the potential of empty-bucket channeling with the 200 MHz RF system at low intensity, additional time was secured to perform tests at operational intensity (4×10^{13} protons-per-pulse). In this case, the 800 MHz RF system was chosen for the implementation. There are two critical reasons why the latter system is preferable:

1. Experimental users in the North Area are less affected by $1.25 \text{ ns} = \frac{1}{800 \text{ MHz}}$ structures compared to $5 \text{ ns} = \frac{1}{200 \text{ MHz}}$. On the one hand, 1.25 ns is close to their detectors' time-resolution; on the other hand, for some experiments the particle rate is too low to measure the structure altogether.
2. The 200 MHz system is already used in operation for beam-loading compensation during the first 800 ms of flattop [75]. It would be challenging to overload the program with the additional task of empty-bucket channelling, not least because these two manipulations need to be operated at different frequency offsets from the beam, as illustrated in Fig. 5.25.

The next step was to choose the RF voltage of the 800 MHz system. For the tests, a single cavity was available, which limited the voltage in the range $V_0 \in [0.2 \text{ MV}, 0.7 \text{ MV}]$, which corresponds to $\Gamma \in [0.006, 0.02]$. The lower limit of V_0 is imposed to ensure the cavity's closed-loop stability. In order to minimise the unwanted perturbations on the extracted beam, the minimum available voltage was chosen for the high-intensity implementation ($V_0 = 0.2 \text{ MV}$, $\Gamma = 0.02$). Based on the simulation results from Sec. 5.5.2 and the measurements with the 200 MHz system




Figure 5.25: Illustration of the conflict between the desired working points for beam loading vs. empty-bucket channelling. The vertical lines show the fundamental-harmonic frequency $f_{rev} = f_{RF}/h$ for each of the RF systems.

(Sec. 5.6), such voltage was predicted to be more than enough to obtain significant ripple reduction.

5.7.1 Narrow-band ripple

In a similar fashion to the 200 MHz tests, the first step was to scan Δ_{RF} while observing $G_{ebc}(50\text{ Hz})$ and $G_{ebc}(100\text{ Hz})$. Figure 5.26 shows the result of the scan, including both measurement data and simulation outcome. It can be concluded that the 800 MHz system can also be exploited to successfully suppress ripple. Moreover, the simulation outcome captures the different functional behaviours of G_{ebc} for the 50 Hz and 100 Hz ripples. The latter shows a wider ‘suppression valley’, which is expected from the discussion in Sec. 5.4.3, Eq. 5.30. In addition, the ‘two-minima’ structure of G_{ebc} alluded to earlier (See Sec. 5.5.2) can be observed in both simulation and measurement.

5.7.2 Broad-band ripple

So far, the discussion of ripple has focused on particular narrow-band components, namely 50 Hz and 100 Hz. This makes sense from a practical standpoint, since they are the main contributors to spill non-uniformities and, therefore, the priority when it comes to mitigation. From an academic standpoint, it must not be forgotten that





Figure 5.26: Ripple-reduction coefficient G_{abc} vs. bucket offset Δ_{RF} for 800 MHz system (for stable-phase sine $\Gamma = 0.02$, RF harmonic $h = 18480$). Subfigures 5.26a, 5.26b show measurement, simulation in *henontrack* and semi-analytical estimate (formula) for a ripple frequency of 50 Hz and 100 Hz, respectively. The measurement error-bars show the total shot-to-shot variations.


empty-bucket channelling should deliver broad-band suppression of ripple by acting on the beam dynamics directly. This subsection demonstrates that this is indeed the case.

In order to acquire a statistically significant sample of spills, empty-bucket channelling was carried out with the fixed settings $\Delta_{RF} = 0.8$ and $\Gamma = 0.02$. The broad-band spectrum of the spill was compared to the spills delivered during nominal operation. Figure 5.27 shows the comparison of the two scenarios. As a caveat, the 50 Hz magnitude is the only one that cannot be directly compared, since during operation a feed-forward controller is deployed on the focusing quadrupole current to minimise it [76]. Although empty-bucket channelling is fully compatible with it, the controller was not active during the test to study G_{abc} in isolation. However, this is not a problem for this study since the focus lies on understanding the effect of empty-bucket channelling on the overall spectrum distribution and not on particular narrow-band components. In other words, we want to characterise how the statistical properties of the spill perturbations are affected. When looking at Fig. 5.27, the effect of empty-bucket channelling is clear: in the time domain, the particle count variation is significantly reduced; in the frequency domain, all AC magnitudes are suppressed.

The spill data with empty-bucket channelling on and off can be exploited to compute G_{abc} , both as a function of time (with Eq. 5.23) and frequency (with Eq. 5.22). These observables can be compared to those obtained via simulation as shown in Fig. 5.28. The simulation study was performed by injecting broad-band white noise on the betatron tune Q_x , in similar fashion to the procedure utilised in Sec. 5.3.3



(a) Nominal operation




(b) Empty-bucket-channelling test (RF harmonic $h = 18480$, stable-phase sine $\Gamma = 0.02$ and bucket offset $\Delta_{RF} = 0.8$).

Figure 5.27: Spill structure in time (left) and frequency (right) measured with the NA62 GTK detector. Ten representative cycles have been used to obtain the mean and standard deviation. The ‘Poisson limit’ line indicates the approximate noise floor of the measurement due to finite statistics.


to extract the SPS transfer function. Good agreement of the average G_{ebc} can be observed both in the time and frequency domains. For the latter, the effect of the low-pass filter at $f_c \approx 100$ Hz is clearly visible, as $|I(f)|$ is suppressed down towards the Poisson limit for both configurations.

5.7.3 Long-term behaviour

Having characterised both the narrow-band and broad-band effects, the SPS North-Area physics experiments agreed to a series of three operational tests. The goal was to study the long-term behaviour of $G_{ebc}(50$ Hz) and $G_{ebc}(100$ Hz), as well as to assess the compatibility with nominal operation. This period also allowed the physics



(a) Time domain



(b) Frequency domain

Figure 5.28: Broad-band ripple-reduction coefficient G_{ebc} for the operational test conditions (RF harmonic $h = 18480$, stable-phase sine $\Gamma = 0.02$ and bucket offset $\Delta_{RF} = 0.8$), measured and simulated in *henontrack*. In both simulation and measurement, ten spills have been used to compute G_{ebc} , which leads to some noise due to the finite statistics per bin. The output is shown both in time domain (Subfig. 5.28a) and in frequency domain (Subfig. 5.28b). Horizontal black lines have been drawn at suppression factors of 1, 2 and 5x for additional visual aid.

experiments to acquire data under the new conditions, which was later used to assess whether the implementation affected their analysis. The empty-bucket channelling setup was fixed at $\Delta_{RF} = 0.8$, $\Gamma = 0.02$ and $h = 18480$. For all tests, the feed-forward controllers that correct the 50 Hz and 100 Hz ripples were locked at their last pre-test values. By keeping the controllers in a steady state, the effect of empty-bucket

channelling could be decoupled from that of the controllers.

Figure 5.29 shows the spill-by-spill $G_{ebc}(f_i = 50, 100 \text{ Hz})$ throughout all three operational tests, with data before and after the tests serving as a nominal baseline for comparison. By looking at the histograms in the right-hand side subplots, it is clear that empty-bucket channelling provides a significant improvement. Both the mean and spread of G_{ebc} are reduced, with the average improvement factor ranging between 3 and 4x, depending on the day and the frequency.

It is important to mention that empty-bucket channelling provides a relative ripple suppression G_{ebc} , so drifts or variations of the ripple baseline $q(f)$ (i.e. violation of time invariance, see Sec. 5.3.2) will still be seen on the spill. For example, Fig. 5.29a clearly demonstrates a factor ~ 2 drift of this kind over the course of test (a time length of $\sim 1\text{h}$). This is caused by the slow variations in the 50 Hz AC current provided by the main grid.


5.8 Additional considerations

On top of achieving ripple suppression, one must ensure that empty-bucket channelling does not perturb the extraction significantly. As discussed in Sec. 4.4, the RF kicks affect the transverse and longitudinal structure of the extracted beam, which can lead to beam loss and a reduction of integrated intensity. Unlike in Chap. 4, where RF structure was deliberately imprinted on the beam, an implementation for nominal operation would have to be as transparent as possible. This section examines these aspects using both simulation and the results from the tests described above.


5.8.1 RF structure

It is inevitable that empty-bucket channelling will introduce some longitudinal structure on the extracted beam, as particles must be channeled between consecutive RF buckets. Unlike in Chap. 4, the goal in this chapter is to minimise this RF structure while maintaining substantial ripple suppression. In fact, strong RF structures are undesirable for fixed-target experiments, as they can spoil data acquisition.


Firstly, if one aims at minimising RF structure, it turns out to be beneficial to use a high-harmonic cavity. In the case of the SPS, using the $h = 18480$ system instead of the $h = 4620$ system turns out to be the correct choice. To verify this point, Fig. 5.30 shows the outcome of a simulation study, where the standard deviation of the extracted bunch σ (normalised to the RF period) has been computed as a function of h and Δ_{RF} (for $\Gamma = 0.02$). Larger σ implies less RF structure, with



(a) Test on 23/9/2022.



(b) Test on 13/10/2022.



(c) Test on 19/10/2022.

Figure 5.29: Measured spill-by-spill ripple-reduction coefficient G_{abc} before, during (grey) and after operational tests. Subfigures 5.29a, 5.29b, 5.29c show the outcomes for a different test day. For each test, the right subplots show the G_{abc} histograms for each ripple frequency f_i . Horizontal black lines have been drawn at suppression factors of 1, 2 and 5x for additional visual aid.

$\sigma = 1/\sqrt{12} = 0.289\dots$ being the standard deviation of a uniform distribution. Indeed σ becomes larger as h is increased, and the advantage of the $h = 18480$ system with respect to the $h = 4620$ system is clear.




Figure 5.30: Bunch length (σ , normalised to RF period) vs. RF harmonic h and bucket offset Δ_{RF} (for $\Gamma = 0.02$), simulated in *henontrack*. Subfigure 5.30a shows a grid scan over a wide range of harmonics, while Subfig. 5.30b shows the cross sections for the two available harmonics in the SPS.

The fact that higher h leads to larger σ can be understood when one looks at the empty-bucket channelling phenomenon in longitudinal phase space. Figure 5.31 shows the extracted particle distribution in longitudinal phase space for identical choices of $\Gamma = 0.02$ and $\Delta_{RF} = 0$ and different harmonic numbers, namely $h = 4620, 18480$. The $h = 18480$ bucket has half the height of the $h = 4620$ bucket which results in a 2x increase in d_{rms} (resonance stop-band width divided by bucket height). Inevitably, this leads to a larger σ , as some particles become resonant where the RF channel is still wide. It is important to emphasise that the channels have identical normalised widths (since only Γ controls the bucket geometry as shown in Sec. 2.4.4), but the ‘integration window’ d_{rms} is larger for the $h = 18480$ case.

NA62 performed their nominal analysis routine on the data acquired during the empty-bucket-channelling tests. They confirmed that the manipulation was transparent to their analysis [77]. Still, the mild 800 MHz structure could be observed on the acquired data, as shown in Fig. 5.32, which indirectly verified that empty-bucket channelling was properly aligned.

5.8.2 Integrated Intensity

The mechanism for the reduction of integrated intensity was discussed in detail in Sec. 4.4.3. In essence, low- A_x particles can be accelerated too quickly through the resonance stop-band, becoming stable again before they can gain enough amplitude




Figure 5.31: Slow-extracted particle distribution in longitudinal phase space for bucket offset $\Delta_{RF} = 0$ and stable-phase sine $\Gamma = 0.02$ (simulated in *henontrack*) for the two available harmonics in the SPS: $h = 4620$ (shown in Subfig. 5.31a) and $h = 18480$ shown in Subfig. 5.31b).




Figure 5.32: Extracted 800 MHz time structure with (on) and without (off) empty-bucket channelling, measured with NA62 Giga-TrackKer detector (GTK). The plot shows a window of 5 ns, averaged over 4 s to obtain sufficient statistics.

to reach the septum. Figure 5.33 shows the dependence of the integrated intensity on Δ_{RF} for the empty-bucket channelling settings used in the high-intensity tests (i.e. $h = 18480$, $\Gamma = 0.02$). Both simulation and measurement are included. In similar fashion to the RF structure, the region where integrated intensity is reduced does not correspond to the region where ripple suppression is maximised. Therefore, it is possible to find a working point where the integrated intensity is close to nominal, while ripple is still substantially suppressed. A grey region has been outlined in

Fig. 5.33, which indicates a ‘forbidden region’ where the integrated intensity is reduced by more than 0.5%. To obtain both ripple reduction and high integrated intensity, the manipulation must be operated outside the ‘forbidden region’, which corresponds to approximately $|\Delta_{RF}| \geq 0.8$ for the setup shown here.




Figure 5.33: (Top) Integrated intensity (I_0 , normalised to nominal extraction) vs. bucket offset Δ_{RF} , measured and simulated in *henontrack*. (Bottom) ripple-reduction coefficient G_{ebc} vs. Δ_{RF} (measurement mean).

In order to assess the long-term behaviour of I_0 , Fig. 5.34 shows the non-extracted intensity (i.e. $1 - I_0$) for the three different operational tests, as well as a reference day with nominal conditions. It can be seen that the non-extracted intensity can be slightly higher during the tests. This is most likely for two reasons: (i) the chosen working point with $\Delta_{RF} = 0.8$ is right at the edge of the region highlighted in Fig. 5.33 and (ii) changes in super-cycle as well as machine jitter can perturb Δ_{RF} . This could be mitigated by choosing a slightly larger Δ_{RF} or by actively optimising the RF frequency from cycle to cycle with an automatic feed-forward routine.




Figure 5.34: Internally-dumped intensity at the end of the cycle, as a percentage of total intensity. The violin plots correspond to the three tests and a reference day when empty-bucket channelling was off (ref).

5.8.3 Transverse effects & beam loss

The last aspect to consider is beam loss. Several R&D efforts have been made in the SPS to reduce beam loss during slow extraction [47] [63] [40]. Loss reduction is critical to ensure equipment longevity and acceptable radio-activation while fulfilling future high-intensity requests [41]. Therefore, it is important that an operational empty-bucket-channelling implementation is compatible with these efforts.

Section 4.4.4 discussed the mechanism by which empty-bucket channelling can lead to a deformation of the transverse phase space presented at the septum. It was identified that, to first order, the effect can be characterised by the change in length l_{rms}^{out} and width w_{rms}^{out} of the extracted separatrix. The simulated dependence of these two quantities as a function of Δ_{RF} (for $\Gamma = 0.02$ and $h = 18480$) can be compared to measurement using the Beam Screen Grids (BSGs) in the transfer line. This is done by choosing two grids that have a convenient phase advance with respect to the extraction septum, so that the measured beam size corresponds almost exclusively either to the radial length (spatial dimension at the septum) or to the azimuthal width (angular dimension at the septum) of the rotating separatrix. Figure 5.35 shows this comparison between simulation and measurement, where one can observe good functional agreement between the horizontal beam size and the separatrix width or length as a function of Δ_{RF} . For a more detailed prediction of the beam's transverse

features, element-by-element tracking would be required, both in the SPS ring and the TT20 transfer line.




Figure 5.35: Measurement (data): Horizontal beam size (x_{std}) vs. bucket offset Δ_{RF} at two different BSGs (Subfigs. 5.35a, 5.35b). Simulation (sim.): Separatrix length l_{rms}^{out} (Subfig. 5.35a) or width w_{rms}^{out} (Subfig. 5.35b) vs. bucket offset Δ_{RF} .

Finally, Fig. 5.36 shows the long-term behaviour of the beam loss during the three operational tests, as well as a reference day operated at nominal conditions. The beam loss during the tests is comparable to that of the reference, which confirms that the deformation of the extracted transverse profile is acceptable.




Figure 5.36: Measurement of normalised beam loss. The violin plots correspond to the three tests and a reference day when empty-bucket channelling was off (ref).

5.9 Operational implementation

In mid-2023, the North Area users and SPS operations expressed their interest in making empty-bucket channelling part of nominal operation. Initially, an attempt was made to operate the manipulation with the parameters employed during the operational tests previously detailed in Sec. 5.7.3, i.e. $\Gamma = 0.02$, $\Delta_{RF} = 0.8$. However, this configuration could not be made robust either cycle-to-cycle or supercycle-to-supercycle. The reason for this is that a configuration with positive Δ_{RF} results in an empty bucket that is near in frequency to the waiting stack. Perturbations and shot-to-shot jitter caused erratic behaviour, resulting in non-reproducible spill macrostructure and partial capture of the waiting beam inside the (no longer) empty bucket.

In view of this, an alternative configuration was implemented operationally. While Γ was kept unchanged at 0.02, Δ_{RF} was set to be negative ($\Delta_{RF} \approx -1$), ensuring enough integrated intensity and still some ripple suppression. Figure 5.37 indicates the frequency offset programmed on the SPS control software, both for the ‘global’ optimum configuration (not implemented due to fragile behaviour) and the ‘local’ optimum configuration (operationally implemented).




Figure 5.37: Measured ripple-reduction coefficient G_{etc} as a function of RF frequency-offset setting $SA.dFrevSlip.G1$. The grey regions are forbidden, as they result in reduced integrated intensity. Two valid configurations (‘global’ and ‘local’) are shown in green.

Since its commissioning, empty-bucket channelling has been running for the SPS Fixed Target beams practically uninterrupted. The system requires little to no maintenance and runs in synergy with other systems that aim at optimising the slow-extracted spill, such as the 50 Hz/100 Hz controllers and the automatic transfer-line steering. Figure 5.38 shows the histograms for the spill-ripple amplitudes at 50 Hz,

100 Hz and 150 Hz, before and after the implementation of empty-bucket channelling. A clear improvement can be observed across all three frequencies, ranging between a factor 1.5 and a factor 2.5.




Figure 5.38: Histograms of measured spill-ripple amplitudes at three different ripple frequencies (f_i), accumulated over six days running with empty-bucket channelling off and on. The legend shows the mean value μ for each histogram.

5.10 Conclusion

This chapter has detailed the ripple-reduction capabilities of empty-bucket channelling in the SPS. The technique was tested both with the 200 MHz and 800 MHz cavity systems, and the differences between them were outlined. Furthermore, the simulation model from Chap. 2 was benchmarked with measurements across a wide-variety of RF parameters, showing good agreement for both narrow-band and wide-band ripple suppression. Additionally, other considerations such as integrated intensity, bunch

structure and transverse perturbations were studied, ensuring that the implementation would remain compatible with current operational requirements. Finally, the studies culminated in an operational implementation employing the 800 MHz system at 0.2 MV, which has so far proven to be successful and easy to maintain. This implementation provides a low-frequency ripple suppression of a factor 1.5-2.5 across the most relevant spill-ripple frequencies, namely 50 Hz and 100 Hz. Figure 5.39 brings the chapter ‘full circle’ by comparing the spill with and without empty-bucket channelling, demonstrating that the duty factor has been significantly improved across a variety of timescales.




Figure 5.39: Measurement of duty factors (\mathcal{F}) vs. bin-width (Δt) for an SPS measurement with empty-bucket channelling (ebc) and without it (nominal), as well as the corresponding Poisson limit (Poisson).

Further work will aim at assessing the compatibility of empty-bucket channelling with future experimental requests in the North Area, namely HIKE/SHADOWS or BDF/SHiP. It is possible that the (mild) 800 MHz structure could perturb the future generation of high-intensity fixed-targets experiments, which would require rethinking the current implementation. In such case, other ripple-suppression strategies such as the introduction of broad-band longitudinal noise may become attractive. Alternatively, one could think of ways of ‘washing out’ the RF structure extracted from the ring throughout the transfer line. In any case, the current empty-bucket channelling implementation significantly benefits current users and it is likely that it will continue to do so for the near to mid-term future.

Chapter 6

Conclusion

This thesis has described the RF technique known as empty-bucket channelling to address two practical challenges concerning CERN's slow-extraction systems:

- A new experimental request for BDF/SHiP (Chap. 4), which required a slow-extracted beam from the SPS with a tightly bunched structure at the nanosecond-scale. The proposed solution is to exploit empty-bucket channelling using the 200 MHz RF system at a voltage of 0.1 MV, which would imprint a strong 5 ns bunched structure with minimal impact on other beam parameters.
- The mitigation of extracted-intensity ripple on the millisecond-scale (Chap. 5), which perturbs the data acquisition of current North Area users. In this case, empty-bucket channelling can be implemented using the 800 MHz RF at a voltage of 0.2 MV, which suppresses ripple effectively and is compatible with current operation.

In order to arrive at these solutions, empty-bucket channelling has been systematically studied throughout this work. In particular, Chap. 3 exploited the beam-dynamics theory outlined in Chap. 2 to develop custom-made simulation tools. The goal was to design a computationally-efficient minimal model that would capture the relevant phenomenology of empty-bucket channelling at timescales ranging from hundreds of milliseconds to sub-nanoseconds. Then, extensive comparisons between measurement and simulation were conducted in Chap. 4 and Chap. 5, which clearly demonstrate the predictive power of the model. Most importantly, these comparisons provide insights on how empty-bucket channelling affects key beam-dynamics parameters, namely extracted bunch structure, spill ripple, extracted intensity and extracted transverse profile.

Apart from yielding the applications detailed in Chap. 4 and Chap. 5, the knowledge and tools developed in this thesis have been exploited to address other challenges. These additional studies have not been included in the main body of this thesis, but are mentioned below for completeness (with their respective references available for the interested reader):

- Empty-bucket channelling in the CERN PS: The first implementation of empty-bucket channelling dates to the 1980s in the CERN PS [16], but ceased to be used operationally at some point. In 2021, the technique was brought back and is ready to be made operational again if needed. See [78] for details.
- A novel method for pulsed extraction: The technique of phase displacement can be employed to quickly push a portion of the beam into the resonance. This ‘fast’ slow extraction is promising for FLASH radiation therapy and some experimental requests at CERN. See App. A (or [79]) for details.
- A novel technique for momentum blow-up: the RF kicks from phase displacement may be used to blow up the momentum distribution of a debunched beam. This approach may be a good alternative to bunch rotation (See Sec. 3.1.2) in certain circumstances. See [80] for details.

All in all, these studies demonstrate the synergies between RF manipulations and slow extraction, which by no means have been exhausted here. In fact, the findings in this thesis are not only relevant for CERN’s machines (particularly the SPS and PS), but can find applicability in many other experimental and medical synchrotrons that perform slow extraction.

In terms of spill-quality improvement, most synchrotrons aim at reducing the impact from power-converter ripple, as degraded spill quality leads to longer treatment times. For example, the MedAustron synchrotron currently uses empty-bucket channelling during treatment to suppress ripple [81]. The CNAO synchrotron has also tested its implementation [69]. In both instances, the system was optimised empirically, but the studies shown here could lead to a deeper understanding of the manipulation and, ultimately, to further improvements in the beam quality delivered to patients. An active slow-extraction collaboration between CERN and MedAustron has been ongoing since 2019, which aims at facilitating the sharing of technical expertise. More broadly, concerns about spill quality and beam time structure extend to other synchrotrons around the world, and a variety of facilities could benefit from the findings in this thesis.

With regards to modelling, this thesis has demonstrated that simplified simulation procedures can show impressive agreement with measurement when it comes both to ripple suppression and extracted RF structure. This is important for present and future facilities that are concerned with these topics, as it demonstrates that insights about the extracted beam can be obtained in a computationally efficient manner. For example, the approach outlined in this thesis could be exploited to characterise and optimise the spill for the future medical synchrotron being studied under NIMMS. An initiative is being pursued to transfer the knowledge and support their studies.

In short, the efforts to study, tailor and improve the slow-extracted spill of the SPS are relevant not only for North Area experiments, but extend to the East Area in the PS, medical initiatives at CERN and the broader slow-extraction community. Indeed, the technical solutions in this thesis have been met with academic and practical interest within the Resonant EXtraction (REX) work-package launched by the Innovation Fostering in Accelerator Science and Technology (I.FAST) project, as well as at the Slow-Extraction mini-workshops organised by the International Committee for Future Accelerators (ICFA) bi-annually. Through these frameworks, the key results outlined here will potentially benefit a myriad of slow-extraction users around the world who often request improvements in spill quality and fine control of the RF time structure.

Further work

Recently, empty-bucket channelling has been made part of nominal operation in the SPS, with the experiments in the North Area (NA62 in particular) reporting significant improvement in the low-frequency ripple. Specifically, the slow variations in particle rate have been reduced by more than a factor 2. On the operational side, no show-stoppers have been reported either. Machine operators have been briefly trained on the main issues that can arise and how to troubleshoot them. Still, a more formal documentation and a Graphical User Interface (GUI) to monitor and control the key parameters would be a valuable addition. This could be particularly useful after super-cycle changes, where the technique needs to be re-optimised, as well as during periods of degraded spill quality, where the malfunctioning source needs to be assessed over a variety of systems.

On the modelling front, this thesis has focused on simplified and computationally-efficient approaches. However, more comprehensive tools exist and are being developed for CERN's slow-extraction systems, as certain problems require more detailed tracking of the beam dynamics. In particular, there are active efforts to implement

both the PS and SPS systems in the recently released *xsuite* code, which is a *python*-based tracking code that incorporates all major beam-dynamics simulation needs into a single framework. One could develop implementations for the RF manipulations in this thesis, as well as an interface to simulate power-converter ripple. Such an effort will allow other users to seamlessly add these effects to their own studies.

Finally, other relevant RF techniques may be studied using a similar framework to the one used in this thesis. For example, both the PS and SPS could potentially benefit from an implementation of stochastic slow-extraction [15], where longitudinal RF noise is used to transport particles into resonance. For the PS, this method would enable slow-extraction without the need to ramp any magnetic components during flattop, which introduces many complexities due to the different saturation levels of the dipole and quadrupole components in the combined-function main units. In the SPS, the RF noise could be added on top of the operational COSE scheme to further aid with ripple suppression. Similar interest exists for transverse RF noise, which is already used in the PS to slow-extract low-energy ions with the method known as RF Knock Out (RFKO).

In summary, this thesis has exploited computational tools and advancement in instrumentation to study and implement RF-based solutions for slow extraction, which aim at contributing to present and future systems by improving and manipulating the beam time structure delivered to the users. It is clear that many avenues of research remain unexplored and that further studies are needed to bring new technical solutions to light.

Appendix A

Pulsed extraction for FLASH radiotherapy

The work presented below is an adapted version of the peer-reviewed article available in [79]. Whereas the article only included simulation results, this Appendix also includes the first measurements from a brief test in the PS.

A.1 RF phase displacement for pulsed extraction

Within the Physics Beyond Colliders (PBC) Study Group at CERN, novel methods of resonant slow extraction from synchrotrons are being explored to satisfy the requirements of future fundamental physics experiments. Recent studies for Enhanced NeUtrino BEams from kaon Tagging (ENUBET) investigated the extraction of millisecond bursts of protons from the CERN Super Proton Synchrotron (SPS) by pulsing the main quadrupole circuit [82]. Furthermore, a novel radiation oncology therapy known as FLASH [83] has awoken interest in the delivery of fast bursted spills of particles in the medical community. FLASH is an ultra-high dose rate irradiation technique that relies on total dose delivery times of the order of milliseconds, often with RF time structure of several MHz. As a consequence, ongoing studies [26] are evaluating the potential of different accelerator types as candidates for delivery of FLASH therapy. In this chapter, RF phase displacement is explored to control the bursting of a resonant third-integer extraction. The same mechanism could also be exploited to feed other resonances like the half-integer, which has been employed in the SPS for fast resonant extractions [84] in the past. The flexibility of digital low-level RF control systems in synchrotrons makes this technique easy to implement and optimise. Additionally, no ramping of magnetic elements is required, avoiding magnet ramp rate limits and optics perturbations that could render certain loss reduction

techniques such as crystal shadowing [63] or octupole folding [47] inefficient. RF phase displacement also has the benefit of extracting at constant momentum, while quadrupole sweep techniques produce a time-varying momentum profile. For these reasons, the scheme could be attractive for the exploitation of existing experimental and medical synchrotrons as burst extraction facilities.

RF phase displacement was originally used in the 1960s as a beam acceleration technique in the CERN Intersecting Storage Rings (ISR) [36]. In 1999, the Proton-Ion Medical Machine Study (PIMMS) [19] considered RF phase displacement as a potential mechanism for slow-extraction spills of around 1 s, but it was ultimately discarded due to the strong modulation it introduced at the repetition period of the sweeping of the RF frequency [85]. For the burst extraction application, the method's initial weakness is turned into a strength by maximizing the modulation to provide short discrete pulses of particles. Here we study the transverse and longitudinal beam dynamics of the method via simulation; we propose both single-burst and multi-burst schemes and assess their limits as a burst extraction technique. The CERN Proton Synchrotron is employed as a case study, as the ENUBET project would be interested to use its beam for equipment testing. Nevertheless, the procedure can be generalised to any synchrotron and a comparison with SPS and PIMMS parameters is also qualitatively addressed.

A.1.1 Concept

The RF phase displacement scheme, implemented in the model described in Chap. 2, is shown in Fig. A.1 and can be summarised as follows:

1. An ensemble of particles awaits coasting stably, away from the third-integer resonance (Fig. A.1a).
2. The RF system is programmed to follow a linear frequency ramp that sweeps empty buckets through the stack of particles in longitudinal phase space (Fig. A.1b).
3. As particles are phase-displaced, they are kicked coherently by the RF cavity and accelerated, which results in a change in momentum and ultimately their tune via chromaticity.
4. Particles with resonant tune gain amplitude until they jump over the septum and are extracted.
5. After the sweep, some particles may remain in the machine (Fig. A.1c).

6. Steps 2-4 are repeated n times to provide n bursts.

The dynamics might remind the reader of the ‘Liouville interpretation’ of phase displacement described in Chap. 2. In other words, the empty buckets are ‘submerged’ into the incompressible beam, which must displace in the opposite direction through the available channels.

An example of the voltage and frequency programmes is shown in Fig. A.2 . To produce n bursts, n of these programmes were concatenated.

A.1.2 Relevant parameters

When the RF frequency is swept through the circulating beam, the distribution of $\Delta p/p$ experiences both an average coherent displacement μ and a blow-up with root mean square increase in spread σ . The effect on the $\Delta p/p$ distribution can be derived analytically [37][28] and expressed as,

$$\mu = \begin{cases} -\text{sign}(\Gamma) \frac{A_0}{2\pi} \alpha(\Gamma) & \text{if } |\Gamma| \leq 1 \\ 0 & \text{elsewhere,} \end{cases} \quad (\text{A.1})$$

$$\sigma = \begin{cases} \frac{A_0}{2\pi} |\Gamma| & \text{if } |\Gamma| \leq 1 \\ \frac{A_0}{2\pi} \left(\frac{\pi}{4}\right)^{3/2} |\Gamma|^{-1/2} & \text{if } |\Gamma| \gg 1, \end{cases} \quad (\text{A.2})$$

where $\Gamma = \sin \phi_s$, ϕ_s is the synchronous phase, $\alpha \approx \frac{1-\Gamma}{1+\Gamma}$ is the bucket area factor and A_0 is the stationary bucket area in $(\phi, \Delta p/p)$ space given by,


$$A_0 = 16 \sqrt{\frac{eV}{2\pi\beta^2 E h |\eta|}}, \quad (\text{A.3})$$

where e is the electron charge, V is the RF voltage, β is the relativistic speed factor, E is the beam energy, h is the RF harmonic number and η is the slip factor. The domain $|\Gamma| \leq 1$ indicates that some closed trajectories still survive inside the bucket.


The expressions for $|\Gamma| \leq 1$ can be normalised to the initial momentum spread of the beam $(\Delta p/p)_0$ to obtain,

$$\bar{\mu} = \frac{\mu}{(\Delta p/p)_0} = r\alpha(\Gamma), \quad \bar{\sigma} = \frac{\sigma}{(\Delta p/p)_0} = r\Gamma, \quad (\text{A.4})$$


where $r = \frac{A_0}{2\pi(\Delta p/p)_0}$ is the ratio between the stationary bucket area and the phase space area of the coasting beam. This relationship provides an intuitive interpretation: to push a fraction $1/n$ of the beam into the resonance, the sweeping bucket must have a factor $1/n$ of the stack area. However, one must additionally account for the blow-up



(a) Before frequency sweep



(b) During frequency sweep



(c) After frequency sweep

Figure A.1: Longitudinal (left) and transverse (right) phase spaces during a single frequency sweep of RF phase displacement extraction.




Figure A.2: Example of voltage and frequency programmes for a single burst: (i) the voltage ramps up to its nominal value while the frequency starts a linear sweep, (ii) the voltage stays flat while the frequency continues its sweep, (iii) the voltage ramps down to zero while the frequency ends its linear sweep and (iv) the voltage stays at zero while the frequency resets to its initial value.

contribution, which comes from the time-varying non-adiabatic nature of the sweep. In fact, A_0 and Γ vary slightly ($\sim 1\%$ for the PS) during the sweep as E changes but V stays fixed, which makes r correct only for the average parameters. Both of these effects were included in the numerical tracking. On the other hand, the variation of η as a function of E was not included (since linear longitudinal transport was assumed), but the change is small ($\sim 1\%$ for the PS) if the machine is operated far from transition energy. For a large number of sweeps, the variation in η could produce a noticeable perturbation [19].

In practice Γ will have a maximum value for a given V , since the RF frequency cannot be ramped arbitrarily fast. The ramp rate $\frac{df_{RF}}{dt}$ is given by,

$$\frac{df_{RF}}{dt} = \frac{\eta f_{RF}^2 eV}{\beta^2 E} \Gamma = F(V) \Gamma,$$

where f_{RF} is the RF frequency and $F(V)$ is the proportionality constant between Γ and ramp rate in physical units. This establishes a maximum $\Gamma = \Gamma_{max}$ given by,

$$\Gamma_{max} = \frac{1}{F(V)} \left. \frac{df_{RF}}{dt} \right|_{max},$$

where $\left. \frac{df_{RF}}{dt} \right|_{max}$ is the maximum RF ramp rate. This directly sets a minimum $r = r_{min} = r(\Gamma_{max})$.

A.1.3 Longitudinal and transverse transit times

In the longitudinal plane, a timescale for each burst can be estimated by ignoring dynamic effects; simply the time needed to move the entire bucket area across the same amount of area in the stack. In other words, the (virtual) synchronous particle inside the sweeping bucket must travel a distance $\Delta p/p = \bar{\mu}(\Delta p/p)_0 + 2H$, where H is the bucket height. Ultimately, numerical simulations are needed to account for dynamic effects. The timescale τ_L (in number of turns) is given by,

$$\tau_L = \frac{[1 + \frac{\pi}{2}rY(\Gamma)]\bar{\mu}(\Gamma)}{\Gamma} \left[\frac{(\Delta p/p)_0}{\frac{eV}{\beta^2 E}} \right], \quad (\text{A.5})$$

where $Y(\Gamma)$ is the bucket height factor [28]. The first term captures the dependence on Γ and the second term provides the time needed for the synchronous particle to cross the entire stack at a given voltage. The latter can easily be compared across machines by setting V to the maximum voltage V_{max} , as shown in Table A.1 by $\tau_{L,0}$.

In the transverse plane, particles take a finite amount of time to gain amplitude and jump over the extraction septum. This transit time varies between particles and it is dependent on their initial coordinates when entering the resonant region and their tune speed [19]. Nevertheless, a characteristic timescale τ_T (in number of turns) can be obtained [86] and expressed as,

$$\tau_T = \frac{8}{S\sqrt{\epsilon_{G,RMS}}},$$

where S is the virtual sextupole strength of the machine [19]. Dynamic effects faster than τ_T will limit the extraction process, e.g. certain particles could cross the resonant region too fast to be extracted. In RF phase displacement the coherent kicks from the RF could make the tune speed of certain particles too large, transporting them back to a stable tune before they can reach the septum. Numerical simulations quantify the severity of this effect.

Table A.1 summarises the relevant parameters for the PS (both for $h = 8$ and $h = 16$), the SPS and a PIMMS-like synchrotron. The PIMMS-like machine and the PS ($h = 8$) differ considerably in nominal parameters, but the normalised parameters τ_T , $\tau_{L,0}$ and r are within 40% of each other. Therefore, we expect results in the PS to be indicative of the PIMMS-like machines, as long as they are expressed in terms of unitless parameters. Due in part to its high RF harmonic number, the SPS parameters differ from the PS parameters.

Parameter	PS, h=8 (16)	PIMMS	SPS
Kinetic E. [GeV]	24	0.25	400
$\epsilon_{G,RMS}$ [mm mrad]	0.059	0.67	0.0018
T [μ s]	2.1	0.4	23
h [1]	8 (16)	1	4620
V_{max} [kV]	200	4	10000
η [1]	0.024	-0.37	0.0018
S [$\text{m}^{-1/2}$]	77	30	170
$F(V_{max})$ [kHz ms $^{-1}$]	0.37 (0.74)	18	0.40
$(\Delta p/p)_0$ [10^{-3}]	6	4	3
τ_T [turns]	430	330	350
r_{max} [1]	1.1 (0.78)	1.3	0.58
$\tau_{L,0}$ [turns]	720	450	120


Table A.1: Relevant machine parameters for RF phase displacement burst extraction.

A.1.4 Simulation studies


A.1.4.1 Single-burst extraction

Single burst simulations were performed for a fixed voltage $V = V_{max}$, and fixed harmonic $h = 8$ (as it provided a larger r), whilst scanning Γ . The instantaneous and cumulative intensities of the extracted spills are shown in Fig. A.3. As predicted by Eq. A.5, the larger the Γ , the shorter the bursts (shorter τ_L), but fewer particles are extracted as predicted by Eq. A.4 because of the smaller $\bar{\mu}$.

Simulations were repeated with identical V and Γ , but with $h = 16$. This changes r and τ_L while keeping all other parameters constant. Fig. A.4 shows the percentage of extracted particles n_{out} as a function of $\bar{\mu}$. In an idealised situation without blow-up or finite extraction time, the simulated data points would lie on the $\bar{\mu} = n_{out}$ line. At low $\bar{\mu}$ blow-up effects drive a substantial number of particles into or away from the resonance due to the large $\bar{\sigma}$. For all settings, a certain number of particles experience large changes in tune per turn. This fast change in tune has the potential to increase the transverse emittance of the extracted beam [45] and some particles $n_{crossed}$ even cross the resonance region too quickly to be extracted, becoming non-resonant again and remaining in the machine. This reduces the efficiency of the scheme but could be combated by shortening τ_T , as particles would then need less turns to reach the septum. These effects are included in all simulation results. Nevertheless, the simulated dependence of $\bar{\mu}(n_{out})$ remains approximately a linear trend (dashed lines), and especially for $h = 16$, where the blow-up effects are smaller due to the smaller r .



(a) Instantaneous extracted particle number.



(b) Cumulative extracted particle number.

Figure A.3: Extracted particles n_{out} for single burst scheme with different Γ .

Fig. A.5 shows the extraction time as a function of Γ . The extraction time was computed as the number of turns required to extract from 1% to 99% of the total extracted intensity, i.e. including only the central 98%. This metric was chosen to stop outliers from dominating the calculation, as they can dramatically increase the




Figure A.4: Extracted particles n_{out} v. $\bar{\mu}$. The dashed lines show linear fits to the simulation data. The horizontal lines ($\bar{\sigma}$) indicate the size of the momentum blow-up effect calculated from Eq. A.4. The grey vertical lines show the portion of particles $n_{crossed}$ that crossed the resonance but did not get extracted. They are only added in the higher n_{out} direction since they indicate the additional percentage of particles that would have been extracted if the extraction happened instantaneously when entering the resonance (i.e. they express the asymptotic limit when $\tau_T = 0$).

total extraction time (by a factor 2 or more). τ_L is within a factor 2 of the simulation output for all the data points. However, dynamic effects clearly play an important role at large Γ .

A.1.4.2 Multiple-burst extraction

One can exploit the knowledge acquired in the previous subsection to generalise the scheme to multiple bursts, which is highly relevant for FLASH therapy. For a given set of machine parameters, one can estimate an initial guess for (V, Γ) such that the extraction has n bursts with N turns per burst by solving the system of equations,

$$\begin{cases} \bar{\mu}(V, \Gamma) = \frac{1}{n} \\ \tau_L(V, \Gamma) = N, \end{cases} \quad (\text{A.6})$$

as shown in Figs. A.6 and A.7. Such an estimate does not account for blow-up and finite transit time and would have to be further studied and optimised. The scheme would be limited by V_{max} and τ_T . The latter could be reduced with higher S or larger $\epsilon_{G,RMS}$.

Fig. A.6 shows the required voltage for different combinations of n and N , which establishes the forbidden region due to voltage limitations. Furthermore, larger Γ




Figure A.5: Extraction time v. Γ . The dashed lines show τ_L as a function of Γ .




Figure A.6: Dependence of V on n and N for the PS, $h = 8$.

leads to larger $\bar{\sigma}$, which may complicate operation and optimisation of consecutive bursts. Figure A.7 shows $\bar{\sigma}$ for different combinations of n and N .

To test the validity of the approximations made in Figs. A.6 and A.7 tracking simulations were carried out for the specific combination of $n = 3$ and $N = 15000$




Figure A.7: Dependence of $\bar{\sigma}$ on n and N for the PS, $h = 8$.

(chosen as an example). The settings from the contour plots were used as initial conditions for an optimisation procedure employing the algorithm Py-BOBYQA [87]. In order to account for $n_{crossed}$, the target n_{out} was set to 90% of the total intensity. This led to an initial guess of $V = 28$ kV, $\Gamma = 0.16$ and an expected blow-up of $\bar{\sigma} = 0.06$. The optimiser was allowed to vary V and Γ for each sweep independently. The cost function was specified as,

$$C = \sqrt{\frac{\sum_i (n_{out,i} - 1/3)^2}{(1/3)^2} + \frac{\sum_i (N_i - 15000)^2}{15000^2}},$$

where $i = 1, 2, 3$ is the sweep number and $n_{out,i}$ was normalised to 90% of the total intensity. Fig. A.8 shows the instantaneous intensity profiles for both the initial guess and the optimised settings found by the optimiser: $V = (36, 36, 67)$ kV, $\Gamma = (0.21, 0.23, 0.24)$. Fig. A.9 shows the voltage and frequency programmes for both schemes.

The simulation was executed fifty times with 10^3 particles to test the robustness of the scheme and the results are shown in Fig. A.10. The initial setting produced longer extraction times than the target for all bursts, which is consistent with Fig. A.5; τ_L consistently underestimates the extraction time for $\Gamma > 0.1$. The optimised settings brought the extraction time down reducing the average relative error in the pulse duration from 37% to 12%. The average relative error in n_{out} increased from 5%




Figure A.8: Instantaneous intensity profiles for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$.




Figure A.9: Voltage (top) and frequency (bottom) programmes for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$.

to 9%, but since the cost function weighs both errors equally this solution is preferred. A different cost function could be more suitable depending on the specific operational requirements. Furthermore, the solution to this cost function is likely to be non-unique and further optimisation could be possible (e.g. by including more sophisticated control on the voltage and frequency programmes). This particular implementation is a simple showcase of the flexibility of the scheme.




Figure A.10: Extracted intensity per burst (top) and burst turns (bottom) for multiple burst scheme in the PS ($h = 8$) with $n = 3$ and $N = 15000$. The dashed lines show target values.

A.1.5 Measurements in the PS

Given the promising simulation results, a machine development test was performed in the PS to conduct a proof of concept implementation. During the test the $h = 8$ setup was used to sweep through the beam, as it provided a larger bucket area and was already the default setting of the East Area extraction configuration. Fig. A.11 shows the beam distribution along the ring during the RF sweep. As expected, the beam picks up the RF structure as the empty bucket traverses its momentum spread. This can be used to correctly align the frequency sweep range of the manipulation.

Following the results from simulation and after an initial discussion with ENUBET [82], a 3-pulse ($n = 3$) extraction was then pursued, aiming at a ~ 10 ms pulse-length ($N \approx 4700$). Starting from V, Γ estimated using Eq. A.4, the voltage and frequency programmes (shown in Fig. A.12) were empirically optimised to reach the desired configuration. Due to the 23-bit resolution of the digital control system, the frequency programme is restricted to a minimum step of ~ 8 Hz, which can be clearly observed in the measurement. An upgrade to a 32-bit system is being pursued for other RF applications [88], which would improve the fine-tuning of the technique if needed.

Fig. A.13 shows the initial and optimised spills, where the instantaneous extracted intensity was measured with the transfer line nitrogen scintillator and the integrated extracted intensity with the ring current transformer. As expected from the simulation studies, the flexibility of the voltage and frequency programmes allows to correct




Figure A.11: Longitudinal beam distribution as a function of cycle time during a 25 ms chirp. The heat-map is reconstructed from the ring current monitor.




Figure A.12: Initial and optimised programmes followed by the PS 10 MHz system during burst extraction test. Both the RF frequency (top) and the RF voltage (bottom) are varied as a function of time.

for the initial differences between pulses, which result from the perturbation each RF sweep causes on the non-extracted beam.




Figure A.13: Measured extracted intensity $I(t)$ as a function of time, for the initial and optimised burst extractions. The integrated intensity $\int_t I(t')dt'$ is also shown in dotted lines, and the three decimal-point numbers indicate its value after each individual burst.

In order to understand the stability of the technique, nine spills were recorded at the optimised setting. Their statistics concerning pulse intensity and pulse duration are shown in Fig. A.14. The pulse duration was computed by considering the time needed to extract the central 90% of the pulse intensity, i.e. $t_{95\%} - t_{5\%}$. The within-pulse variation is small with $\sigma < 4\%$ for all quantities, while the pulse-to-pulse differences are larger ($\sigma \approx 10\%$ for n_{out} and $\sigma \approx 5\%$ for duration). The latter is a systematic offset and could be further optimised by fine-tuning the frequency and voltage programmes. Moreover, the total extracted intensity is only 60% of the injected one, which might be a limiting factor if high-intensities need to be delivered. However, since the requirements of the ENUBET test beams are still being discussed, additional tuning of these quantities has not been pursued yet. Still, this novel extraction method is now available as a possible approach if pulsed beams are needed in the East Area.

As a final check, the optimised settings from obtained from the test were fed-back into the simulation model, looking to benchmark its predictive capabilities. One of




Figure A.14: Measured extracted intensity per burst (top) and burst duration (bottom) for optimised burst scheme.

the critical unknowns was the initial longitudinal momentum distribution. Although the previous simulation studies in assumed it to be uniform, it is common for it to have additional structure due to the unstable nature of the ‘momentum-stretching’ RF gymnastics performed before extraction. Since there is no device to directly measure the distribution in the ring (e.g. Schottky plates), an indirect measurement was performed:

1. Once the RF gymnastics had been performed, the particles were slow extracted by ramping the lattice mains (combined function magnets).
2. We assume the extraction is fully chromatic, which is valid since $\hat{\delta}_{stopband} \ll \hat{\delta}_{beam}$ for the PS. Then, the magnetic rigidity $B\rho$ and the relative momentum offset δ can be related via the equation:

$$p/q = B\rho \rightarrow \delta = dp/p = d(B\rho)/B\rho \quad (\text{A.7})$$

3. Since the magnetic rigidity time-dependence can be directly obtained from the extraction programme, the spill $I(t)$ can be expressed as $I(d(B\rho)/B\rho)$ and finally as $I(\delta)$.
4. $I(\delta)/\int_{\mathbb{R}} I(\delta)d\delta$ provides an empirical probability density function (ePDF), which can be integrated to obtain an empirical cumulative distribution (eCDF). Then, the eCDF is sampled using the inverse transform method. Fig. A.15 shows the outcome of this procedure.




Figure A.15: Extracted intensity I as a function of relative-momentum offset δ . The left subplot shows the measurement, while the right subplot shows the samples generated by using the eCDF procedure from this section.

Finally, this momentum distribution, along with the voltage and frequency programmes shown in Fig. A.12 was run for the optimised configuration. Fig. A.16 shows the extracted spill and the intensity per burst. It can be seen that model and measurement are in good agreement in terms of pulse duration, pulse separation and extracted intensity per pulse.

A.2 Conclusion and outlook

RF phase displacement acceleration was exploited to provide short bursts of particles from a synchrotron in a scheme that could be attractive for exploitation at existing experimental and medical synchrotrons, with FLASH therapy in mind. Employing a simplified model of the PS, it was shown that 80 - 90% of the total beam intensity could be extracted in a single burst of 20000 - 30000 turns. This corresponds to 40 - 60 ms in the PS and 8 - 12 ms in PIMMS-like machines. Moreover, a few basic parameters were computed to characterise the extraction scheme. The parameters




Figure A.16: Comparison of measurement and simulation for optimised multiple-burst scheme. The top subplot shows the extracted intensity as a function of time. The bottom subplot shows the integrated intensity per pulse

were informative, but simulation showed that dynamic effects ultimately played an important role in the extracted intensity and extraction time. A 3-burst scheme with 15000 turns per burst could be realised in a simulation of the PS with optimised machine parameters. Finally, the technique was implemented in the PS and may be exploited for irradiation tests requested by the ENUBET project.

Appendix B

Empty-bucket channelling with double-harmonic systems

The RF voltage and frequency can be varied to obtain different bunch lengths and extracted intensities during empty-bucket channelling, as studied in Chap. 4. The freedom for customisation could be further increased with additional degrees of freedom on the RF system. In the case of the SPS, the 200 MHz system can be complemented with up to 1.4 MV from the 800 MHz system, a feature employed in operation to alter the synchrotron tune distribution to tame bunch instabilities for high-intensity beams [89].

B.0.1 Theory

The additional degrees of freedom provided by the 800 MHz-system can be used to modify the channel geometry. Two parameters are defined: the ratio between 800 MHz and 200 MHz voltages α and relative phase between the two waveforms ψ . This leads to the following potential U :

$$U(\phi) = V \cdot \text{sign}(\eta) \cdot \left\{ -\Gamma\phi + \cos\phi + \alpha \cos(4\phi + \psi) \right\}, \quad (\text{B.1})$$

where the factor 4 comes from the frequency ratio between the two systems and U has the quasi-periodic property $U(\phi + 2\pi n) = U(\phi) - 2\pi n V \text{sign}(\eta) \Gamma$. Figure B.1 shows U for a few representative values of α and ψ , demonstrating how the new knobs affect the longitudinal phase space.

For such a system, it becomes non-trivial to compute the ϕ -values of the unstable fixed points ϕ_{ufp} , which determine the channel locations. The ϕ_{ufp} s can be computed by numerically finding the roots of,




Figure B.1: Semi-analytically computed potential (far left) and longitudinal phase-spaces (rest) for different combinations of voltage ratio α and relative phase ψ , with $\Gamma = 0.05$.

$$\frac{\partial U}{\partial x}(\phi_{ufp}) = 0 \rightarrow \sin \phi_{ufp} + 4\alpha \sin(4\phi_{ufp} + \psi) = \Gamma, \text{ if } \frac{\partial^2 U}{\partial x^2}(\phi_{ufp}) < 0, \quad (\text{B.2})$$

where the ‘if statement’ discards the stable fixed points and the computation has been carried out above transition energy, i.e. $sign(\eta) > 0$. Figure B.2 shows the

solutions to Eq. B.2 for two different values of ψ as α is increased. Unlike in the single-harmonic case, the number of unstable fixed points varies.




Figure B.2: ϕ -coordinate of unstable fixed points ϕ_{ufp} vs. voltage ratio α for two different relative phases (ψ).

Even if the ϕ_{ufps} have been found, there is no guarantee its corresponding channel will be available for empty-bucket channelling. As shown in Fig. B.3 one can have ‘trapped’ RF channels, inaccessible to the beam outside due to a larger bucket blocking access to it. In general, a channel will be open if

$$U(\phi_{ufp,i}) > U(\phi_{ufp,j}), \quad \forall \phi_{ufp,j} \geq \phi_{ufp,i} \quad \text{if } \Gamma \geq 0, \quad (\text{B.3})$$

where i, j enumerate all the unstable fixed points. The inequality can be understood from the potential point of view of phase displacement (see Chap. 2): trajectories coming from infinity must reflect from the potential wall adjacent to the unstable fixed point. If a preceding potential wall blocks this, the channel will be closed.

Moreover, the availability of a channel, i.e. being open or closed, may change as α/ψ are varied. Figure B.4 shows the result of applying the criterion in Eq. B.3 to the unstable fixed points computed previously for $\psi = 0.6\pi$ as a function of α . At around $\alpha = 0.6$ the open channel moves from one unstable fixed point to another.

B.0.2 Simulation: grid-scan

Given the good agreement between simulation and measurement in Chap. 4, a simulation study was performed to explore the bunch-shaping capabilities of the double-harmonic manipulation. This was done by configuring the 200 MHz system at the ‘SHiP candidate’ setting of mid-V and $\Delta_{RF} = 0$, and scanning the 800 MHz-system settings by changing α and ψ . Since the maximum voltage attainable by the 800 MHz




Figure B.3: Open and closed channels illustrated in the potential (top) and longitudinal (bottom) plots.




Figure B.4: ϕ -coordinate of unstable fixed points ϕ_{ufp} vs. voltage ratio α for a fixed relative phase ($\psi = 0.6\pi$), showing whether the adjacent channel is open or closed.

system is 1.4 MV, we have $\alpha \in [0, 1.4]$. Fig. B.5 shows nine representative bunch shapes, with their corresponding longitudinal phase spaces. It is clear that the higher-frequency cavity can have a large impact on both the bunch length and shape. For example, increasing α for $\psi = 1.4\pi$ generates an intermediate small empty bucket between the two main buckets, splitting the original channel into two and delivering a bi-modal bunch profile.




Figure B.5: Slow-extracted particle distribution in longitudinal phase space vs. voltage ratio α and relative phase ψ for the SPS double-harmonic setup, simulated in *henontrack*.

Figure B.6 summarises the results of a simulation grid-scan, showing dependences of the bunch length σ and the integrated intensity I_0 on the parameters α and ψ . The regions labelled ‘OK’ correspond to the parameter spaces where the dual-harmonic setup beats the single-harmonic setup. Approximately speaking, when both systems are in-phase/out-of-phase, σ and I_0 are decreased/increased.

The ‘OK’ regions from Fig. B.6 have been combined in Fig. B.7. It can be seen that the regions are non-overlapping except at $\alpha = 0$. In other words, it is not possible to find a solution that outperforms the single-harmonic solution on σ and I_0 simultaneously.




Figure B.6: Bunch length (σ , normalised to RF period) and integrated intensity (I_0 , normalised to total intensity) vs. voltage ratio α and relative phase ψ . The ‘OK’ regions surrounds the parameter space where the single-harmonic setting ($\alpha = 0$) is outperformed.




Figure B.7: Regions in α/ψ -space where the double-harmonic setup outperforms the single-harmonic setup in reducing σ (red) or increasing I_0 (white).

B.0.3 Simulation: Optimiser

So far, the 200 MHz-system’s voltage and frequency have been kept fixed at the ‘SHiP candidate’ settings. However, it is possible that a better double-harmonic setup could be found if Γ and Δ_{RF} were allowed to change. In order to address this question, an optimisation routine was set up. The problem was defined as follows:

1. The goal is to find a configuration (or set of configurations) in the $(\Gamma, \Delta_{RF}, \alpha, \psi)$

search space that dominates the current candidate, i.e. a configuration that outperforms the current choice both in I_0 and σ .

2. A cost function C is set up, which we aim to minimise and is given by:

$$C = (1 - I_0) + \sigma \geq 0. \quad (\text{B.4})$$

3. In order to avoid searching for setups on the tails of the Pareto front (i.e. strongly prioritise small σ or large I_0) as well as to stay within hardware constraints, the search-space is restricted within the following bounds:

$$\sqrt{\Gamma/\pi} \in [0.123/2, 2 \times 0.123], \Delta_{RF} \in [-1, 1], \alpha \in [0, 1.4], \psi \in [0, 2\pi] \quad (\text{B.5})$$

4. One-hundred random initial conditions are initialised within the chosen bounds and each of them is run through the ‘Nelder-Mead’ [90] local optimiser. This approach allows us to perform a global search while effectively exploiting cluster computing for parallelisation of each random seed.

Figure B.8 shows the output of the final configurations obtained via this method. Firstly, it becomes clear that the problem is non-convex, as several random seeds (47 out of 100) provide solutions that are outperformed by the single-harmonic Pareto front. Nevertheless, many seeds (53 out of 100) manage to produce the desirable outcome, surpassing the Pareto limits of the single-harmonic setup. Moreover, a few seeds (6 out of 53) dominate the SHiP-candidate configuration, providing better outcomes both for σ and I_0 . In short, this exercise demonstrates that the additional degrees of freedom provided by the 800 MHz system do in fact allow us to push empty-bucket channelling beyond its single-harmonic limits.

Finally, Fig. B.9 shows the transverse and longitudinal phase spaces of the best configuration found by the optimiser ($\sqrt{\Gamma/\pi} = 0.094, \Delta_{RF} = -0.3, \alpha = 0.6, \psi = 1.4\pi$), which increases I_0 from 93 % to 96 % and decreases σ from 0.073 to 0.07, with respect to the best single-harmonic case. Interestingly, the ‘strategy’ of the solution can be understood somewhat intuitively: the 800 MHz system is used to create an intermediate bucket between the two main 200 MHz buckets, providing a flat basin where the resonant beam is bunched. However, unlike in the single-harmonic case, the empty buckets are offset away from the beam ($\Delta_{RF} = -0.3$) in order to align this basin (and not the channels) with the resonance. In fact, particles are extracted before




Figure B.8: Bunch length (σ , normalised to RF period) vs. integrated intensity (I_0 , normalised to nominal) for the SPS double-harmonic setup. The plot includes the Pareto front for the 200 MHz-only scan. The optimiser outputs (Opti.) are marked with a grey cross (Opti. Pareto) if they surpassed the 200 MHz Pareto front. The solutions that dominate the SHiP candidate are those in the bottom-right quadrant of the fuchsia grid.

being fully channelled, as channelling would produce the double-humped distribution already discussed in Fig. B.5. Furthermore, the transverse phase space remains largely unaffected, which indicates that the proposed solution will not have a strong impact on beam loss.




Figure B.9: Slow-extracted particle distribution in longitudinal (left) and normalised transverse (right) spaces for the best SPS double-harmonic configuration, including the histograms for the SHiP candidate (simulated in *henontrack*).

Appendix C


Empty-bucket channelling with broadband cavities

Broad-band cavities can provide voltage signals of (almost) arbitrary shape by superimposing a linear combination of many sinusoidal signals. For example, the PS is equipped with a broad-band Finemet cavity capable of generating so-called barrier buckets [91]. These are created by concatenating isolated sine wave pulses of period T_1 repeatedly generated every time interval T_2 , which produces an elongated bucket structure in longitudinal phase space, as shown in Fig. C.1. Unlike in conventional sinusoidal pulses, the RF channel-width and the channel repetition frequency become fully decoupled, the former being controlled by T_1 and the latter by T_2 . This allows to tune the bunch length and bunch spacing independently during empty-bucket channelling.


C.1 Measurements in the PS

The flexibility of the PS Finemet cavity was briefly explored in a machine development test. Since no fast detection device was present in the transfer line, the ring fast beam current monitor (fBCT) was used to characterise the RF structure imparted on the circulating beam, as it was phase displaced inside the machine. As can be seen in Fig C.2, different barrier bucket configurations were successfully probed, varying both T_1 and T_2 . The phase displacement phenomenon can still be observed, as there is no fundamental difference between a conventional bucket and more ‘unconventional’ geometries (as already seen in the double-harmonic case).

Even though no such broad-band system is available in the SPS, past tests have shown that a similar phenomenon can be achieved by exploiting the filling time of the 200 MHz system to modulate its voltage [92], due to the fact that it is a travelling



(a) Voltage and potential. The dashed vertical lines show the τ -coordinate of the various fixed points.



(b) Longitudinal phase space. The blue to pink heatmap corresponds to different increasing turn number in a particle tracking simulation.

Figure C.1: Example of channelling scheme with isolated sine pulses

wave structure. However, this remains mostly a topic of academic interest, since the SHiP request does not foresee the need for additional customisation of the bunch structure.




Figure C.2: Evolution of beam time profile during phase displacement using barrier buckets, with the corresponding channel geometry overlaid (black). Particles arrive to flattop bunched in four tight buckets, they debunch for 200 ms and are finally pushed through the barrier bucket channels by ramping the bending field.

Appendix D

Some non-linear phenomena in empty-bucket channelling

Empty-bucket channelling is a technique that combines a non-linear transverse process (1/3-integer resonance) with a non-linear longitudinal process (phase displacement). This results in many interesting non-linear phenomena. In this Appendix, a few of them are explained and their possible applications/consequences are described.

D.1 Weak bunching

Chapter 4 studied empty-bucket channelling as a technique to provide slow-extracted beams with a 5 ns bunch structure. It was shown that, although bunch length σ is minimised when bucket and resonance are perfectly aligned, a second ‘weak bunching region’ exists, as shown in Fig. D.1. It can be seen that, for this particular setting, this region is located at approximately $\Delta_{RF} = 2$, i.e. when the bucket is offset towards the waiting beam by two bucket heights.




Figure D.1: Simulation: slow-extracted bunch length (σ , normalised to the RF period $1/f_{RF}$) vs. bucket offset (Δ_{RF}), with RF harmonic $h = 4620$ and stable-phase sine $\Gamma = 0.05$. The vertical lines highlight the regions of ‘strong bunching’, ‘partial debunching’ and ‘weak bunching’.

In order to understand this phenomenon, Fig. D.2 shows the extracted particle distribution in longitudinal phase space for three bucket offsets. ‘Strong bunching’ was discussed in detail in Chap. 4 and occurs when the outgoing beam is trapped between two consecutive buckets. As the bucket is offset towards the waiting beam (positive δ), bunching is gradually lost, as the beam populates the tilted lower-end of the RF channel. This effect is at its maximum in the ‘partial rebunching’ setting. As the offset towards the waiting beam continues, ‘partial rebunching’ occurs. In this case, the extracted beam has mostly left the tilted RF channel, but the oscillations caused by the RF are large enough to ‘clear out’ a region of longitudinal phase space directly below the unstable fixed point. In other words, as the large oscillations push particles towards the unstable fixed point, they become non-resonant and cannot be extracted.




Figure D.2: Slow-extracted particle distribution in longitudinal phase space (τ, δ) for three bucket offsets (simulated in *henontrack*): $\Delta_{RF} = 0$ (strong bunching), $\Delta_{RF} = 0.9$ (partial debunching) and $\Delta_{RF} = 2$ (weak bunching). The marginal probability distributions are projected onto the τ and δ axes. The time-delay τ of particles has been aliased to fit within one RF period to provide the distribution of the ‘effective bunch’. The contour of the RF bucket (fuchsia) and the transverse stop-band (red) are shown for additional visual aid.

‘Partial bunching’ may be an interesting alternative to strong bunching for applications where only some RF structure is needed. Unlike ‘strong bunching’, ‘partial bunching’ would not compromise the extracted intensity (See Sec. 4.4.3).

D.2 Narrowing of extracted separatrix

Empty-bucket channelling introduces a non-linear coupling between transverse and longitudinal beam dynamics. In fact, this coupling can significantly deform the trans-

verse profile of the extracted separatrix. An interesting phenomenon, which could perhaps be exploited for loss-reduction purposes, is the possibility of reducing the width of such separatrix. As shown in Fig. D.3, a parameter region can be found for the SPS settings where the separatrix width w_{rms}^{out} can be reduced by more than 20 %.




Figure D.3: Separatrix width (w_{rms}^{out} , normalised to nominal extraction) vs. stable-phase sine (Γ) and bucket offset (Δ_{RF}) for RF harmonic $h = 4620$, simulated in *henontrack*.

When the beam is pushed into the unstable region, a correlation is introduced between transverse amplitude A_x and momentum offset δ . This is simply because the unstable-region size changes linearly with δ and, therefore, the outgoing separatrix arms are displaced with respect to each other. This effects manifests itself in a widening of the extracted separatrix, as shown in Fig. D.4a for the nominal SPS extraction. In essence, a dispersive contribution has been introduced to the beam size.

There is a well-known technique which tunes the dispersion in the ring to counterbalance the aforementioned dispersive contribution at the extraction septum: the Hardt condition [93]. The Hardt condition sets up the following condition at the septum:

$$\vec{D}_x = -\vec{D}_{sloex}, \quad (\text{D.1})$$

where $\vec{D}_x = [D_x, D'_x]^T$ is the dispersion at the septum and \vec{D}_{sloex} is the dispersive contribution from the resonance boundary. The latter can be expressed in terms of machine parameters, which are then tuned to satisfy Eq. D.1. The separatrices have been aligned by adjusting (the effective) transverse amplitude A_x to cancel the A_x, δ correlation.

But there is yet another way to cancel the correlation, namely adjusting δ properly. This is exactly what the manipulation shown here does. In other words, empty-bucket channelling can be optimised to perform a similar task to the Hardt condition. Figure D.4b shows the mechanism by which this happens:

1. Particles with different A_x become resonant at different δ . Due to empty-bucket channelling, this means that different A_x slices get different average accelerations.
2. If the empty bucket is aligned so that its lower edge aligns with the resonant stop-band, particles starting at high δ are strongly decelerated down the RF channel, while low δ particles are accelerated by the librations below the channel. These two effects provide a (non-linear) focusing effect in δ .
3. As particles approach the septum, this phenomenon reduces the instantaneous δ width that is resonant at any given point. This results in a narrowing of the extracted transverse profile.


The reduction of the instantaneous δ width can be seen in Fig. D.5. This effect will also aid with higher-order chromatic effects coming from the non-linear lenses. In fact, the technique can also be seen as an attempt to extend Constant Optics Slow Extraction (COSE) to the particles within the resonant stop-band.

Of course, no real change in the total phase-space volume has occurred, since the total extracted momentum will remain unchanged and the dispersive correlation from the nominal extraction can anyway be corrected by matching the transfer-line optics. Nevertheless, this technique does create a narrowing of the separatrix inside the ring. Moreover, unlike with the Hardt condition, the narrowing is present throughout the entire ring circumference, and not just at the extraction septum.


D.3 Ripple suppression: two minima

When scanning the bucket offset Δ_{RF} and looking at the ripple-reduction coefficient G_{ebc} (for all other parameters fixed), one often finds that the ripple-suppression landscape exhibits two minima, as it is shown in Fig. D.6. However, this feature is not always there. In fact, Fig. D.6 shows that it only becomes apparent when the RF harmonic h is sufficiently large.

In order to understand this phenomenon, we must take into account the discussions from Sec. 5.4.2 and Sec. 5.4.3. In a nutshell, the ripple reduction comes from the fact



(a) Nominal extraction (no empty-bucket channelling)



(b) Empty-bucket channelling ($h = 4620$, $\Delta_{RF} = 1.3$, $\Gamma = 0.04$)

Figure D.4: Slow-extracted particle distribution in longitudinal phase space (τ, δ) (left) and normalised phase space (X, X') (centre), simulated in *henontrack*. Right: initial transverse amplitude A_x vs. final instantaneous momentum δ . The contour of the RF bucket (pink) and the transverse stop-band (red) are shown for additional visual aid.

that particles are accelerated into the resonance by the RF kicks. Moreover, one must ‘average out’ this acceleration over the period T_i during which a particle stays ‘near’ the resonance boundary. The effect size of this averaging is captured in the ratio $r = T_i/T_{slip}$, where T_{slip} is the time needed by a particle to slip from the RF crest to the RF trough. If r is small/big, the effective acceleration is big/small.

Figure D.7 shows the longitudinal phase spaces at the global minimum, local maximum and local minimum from Fig. D.6b. These phase spaces can be used to explain why a $G_{ebc}(\Delta_{RF})$ has two minima:

- At the global minimum, trajectories moving into the resonance are tightly restricted within the RF channel and all particles are transported deep into the resonance.






Figure D.5: Histogram of slow-extracted instantaneous momentum δ for nominal extraction and empty-bucket channelling (ebc, $h = 4620$, $\Delta_{RF} = 1.3$, $\Gamma = 0.04$), simulated in *henontrack*.

- At the local maximum, trajectories start by pointing towards the resonance. However, as one integrates for a period T_i , a non-negligible part of them change direction and carry particles away from the resonance. This effect is aggravated when $T_{slip} \propto 1/h$ is short, which is why the two-minima structure is only visible for high h . For low- h setups, the required T_i to slip by a full RF period will be extremely long. Thus, this ‘averaging effect’ will never become relevant.
- At the local minimum, some trajectories do change direction, but a substantial part of them do not have enough time to do so. Therefore, some net acceleration is still experienced by the beam.

It is important to keep this non-linearity in mind, as scanning the bucket offset is the main tool by which one aligns resonance and bucket in the real machine.




(a) h , Δ_{RF} scan.




(b) Δ_{RF} scan, $h = 18480$


Figure D.6: 50 Hz-ripple-reduction coefficient ($G_{abc}(50 \text{ Hz})$) vs. RF harmonic (h) and bucket offset (Δ_{RF}), simulated in *henontrack*. Subfigure D.6a shows the output over a broad range of harmonics, while Subfig. D.6b shows the cross section for $h = 18480$.



(a) Global min. Fig. D.6



(b) Local max. Fig. D.6



(c) Local min. from Fig. D.6

Figure D.7: Schematic of longitudinal phase space during empty-bucket channelling, including the bucket separatrix (pink) and the resonance stop-band (red). For a given integration time (T_i), trajectories enter the resonant region and slip with respect to the RF bucket. After enough time, some trajectories (dotted) will start to move away from the resonance. T_{slip} shows the slippage time to move from the RF peak to the RF trough. Subfigures D.7a, D.7b, D.7c show three different scenarios taken from Fig. D.6.

Bibliography

- [1] Rolf Wideröe. “Über ein neues Prinzip zur Herstellung hoher Spannungen”. In: *Arbeiten aus dem Elektrotechnischen Institut der Technischen Hochschule Aachen: Band III: 1928*. Ed. by W. Rogowski. Berlin, Heidelberg: Springer, 1929, pp. 157–176. URL: https://doi.org/10.1007/978-3-662-40440-9_14.
- [2] Ernest O. Lawrence and M. Stanley Livingston. “The Production of High Speed Light Ions Without the Use of High Voltages”. In: *Physical Review* 40.1 (1932). Publisher: American Physical Society, pp. 19–35. URL: <https://link.aps.org/doi/10.1103/PhysRev.40.19>.
- [3] R. J. Van de Graaff, K. T. Compton, and L. C. Van Atta. “The Electrostatic Production of High Voltage for Nuclear Investigations”. In: *Physical Review* 43.3 (1933). Publisher: American Physical Society, pp. 149–157. URL: <https://link.aps.org/doi/10.1103/PhysRev.43.149>.
- [4] J. D. Cockcroft and E. T. S. Walton. “Experiments with high velocity positive ions”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 129.811 (1930). Publisher: Royal Society, pp. 477–489. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1930.0169>.
- [5] J.L. Tuck and L. C. Teng. *Synchro-cyclotron Progress Report III*. 1951.
- [6] K. J. le Couteur. “The Regenerative Deflector for Synchro-cyclotrons”. In: *Proceedings of the Physical Society. Section B* 64.12 (1951), p. 1073. URL: <https://dx.doi.org/10.1088/0370-1301/64/12/308>.
- [7] H. G. Hereward. *The possibility of resonant extraction from the C.P.S.* AR-Int-GS-61-5. 1961. URL: <https://cds.cern.ch/record/297121?ln=en>.
- [8] C. L. Hammer and L. Jackson Laslett. “Resonant Beam Extraction from an A. G. Synchrotron”. In: *Review of Scientific Instruments* 32.2 (1961), pp. 144–149. URL: <https://pubs.aip.org/rsi/article/32/2/144/464136/Resonant-Beam-Extraction-from-an-A-G-Synchrotron>.
- [9] P. Strolin. *Third-Order Resonance Slow Extraction from Alternating Gradient Synchrotrons*. ISR-TH-66-40. 1966. URL: <https://cds.cern.ch/record/296978/files/CM-P00065556.pdf>.

- [10] H.G. Hereward. “The CPS Resonant Extraction System”. In: *Proceedings of the 4th International Conference on High-Energy Accelerators*. 4th International Conference on High-Energy Accelerators. 1964. URL: https://cds.cern.ch/record/2065209/files/HEACC63_II_197-199.pdf.
- [11] G. Auberson and K.H. Reich. *Etudes préliminaires du faisceau éjecté lent*. CERN-MPS-Int-DL-62-25. 1962. URL: <https://cds.cern.ch/record/298567/files/MPS%20Int.DL%2062-25.pdf>.
- [12] Y Baconnier et al. *Beam extraction from the 300 GeV synchrotron*. MPS-DL-70-26. 1970. URL: <https://cds.cern.ch/record/2854096>.
- [13] Yves Baconnier et al. “Extraction from the CERN SPS”. In: *IEEE Trans. Nucl. Sci.* 24 (1977), pp. 1434–1436. URL: <https://cds.cern.ch/record/319348>.
- [14] W. Hardt. *Slow extraction from LEAR*. pbar-LEAR-Note-64, CERN-PS-DL-Note-79-4, LEAR-Note-64. 1979. URL: <https://cds.cern.ch/record/134883>.
- [15] S. van der Meer. *Stochastic Extraction, a Low-Ripple Version of Resonant Extraction*. CERN-PS-AA-78-6. 1978. URL: <https://cds.cern.ch/record/2830887>.
- [16] R. Cappi and Ch. Steinbach. “Low Frequency Duty Factor Improvement for the CERN PS Slow Extraction Using RF Phase Displacement Techniques”. In: *IEEE Transactions on Nuclear Science* 28.3 (1981), pp. 2806–2808. URL: <http://ieeexplore.ieee.org/document/4331919/>.
- [17] W. Hardt W. *Slow ejection based on repetitive unstacking*. CERN-PS-DL-79-19. 1979. URL: <https://cds.cern.ch/record/322574/>.
- [18] G Molinari and H Mulder. *The Improved Ultra Slow Extraction Noise System at LEAR*. CERN-PS-94-22-AR. 1994, pp. 2376–2378. URL: <https://cds.cern.ch/record/265947>.
- [19] L. Badano et al. *Proton-ion Medical Machine Study (PIMMS) I*. CERN-PS-99-010-DI. 1999. URL: <https://cds.cern.ch/record/385378>.
- [20] L. Badano et al. *Proton-ion Medical Machine Study (PIMMS) II*. CERN-PS-2000-007-DR. 1999. URL: <https://cds.cern.ch/record/449577>.
- [21] Alberto Degiovanni and Ugo Amaldi. “History of hadron therapy accelerators”. In: *Physica Medica* 31.4 (2015), pp. 322–332. URL: <https://www.sciencedirect.com/science/article/pii/S1120179715000629>.
- [22] M. Benedikt et al. “Overview of the MedAustron design and technology choices”. In: *Proceedings of the 2010 International Particle Accelerator Conference*. IPAC10. 2010. URL: <https://accelconf.web.cern.ch/IPAC10/papers/mopea020.pdf>.
- [23] M. Vretenar et al. “The Next Ion Medical Machine Study at CERN: Towards a Next Generation Cancer Research and Therapy Facility with Ion Beams”. In: *Proceedings of the 12th International Particle Accelerator Conference*. IPAC2021. In collab. with Lin Liu (Ed.) et al. 2021. URL: <https://cds.cern.ch/record/2776467>.

- [24] Claudia Ahdida et al. *Findings of the Physics Beyond Colliders ECN3 Beam Delivery Task Force*. CERN-PBC-REPORT-2023-001. Place: Geneva. Geneva, 2023. URL: <https://cds.cern.ch/record/2847433>.
- [25] M. Fraser et al. “Feasibility of Slow-Extracted High-Energy Ions from the CERN Proton Synchrotron for CHARM”. In: *Proceedings of the 13th International Particle Accelerator Conference*. IPAC2022. 2022. URL: <https://accelconf.web.cern.ch/ipac2022/papers/wepost012.pdf>.
- [26] S. Jolly et al. “Technical challenges for FLASH proton therapy”. In: *Physica Medica* 78 (2020), pp. 71–82. URL: <https://doi.org/10.1016/j.ejmp.2020.08.005>.
- [27] Davide Mazzucconi et al. “Mixed particle beam for simultaneous treatment and online range verification in carbon ion therapy: Proof-of-concept study”. In: *Medical Physics* 45.11 (2018), pp. 5234–5243. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mp.13219>.
- [28] S. Y. Lee. *Accelerator Physics*. Vol. 4th edition. World Scientific, 2019.
- [29] A. W. Chao, ed. *Handbook of accelerator physics and engineering*. Second edition. World Scientific, 2013.
- [30] A. W. Chao. *Lecture Notes on Topics in Accelerator Physics*. SLAC-PUB-9574, 812598. 2002, SLAC-PUB-9574, 812598. URL: <http://www.osti.gov/servlets/purl/812598/>.
- [31] E. Forest. *From Tracking Code to Analysis*. Tokyo: Springer Japan, 2016. URL: <http://link.springer.com/10.1007/978-4-431-55803-3>.
- [32] J. Wenninger. *Linear Imperfections*. arXiv:2004.14001. type: article. 2020. arXiv: 2004.14001[physics]. URL: <http://arxiv.org/abs/2004.14001>.
- [33] L. Deniau et al. (*Methodical Accelerator Design*) *Version 5.06.0*. Geneva: CERN, 2020. URL: <http://mad.web.cern.ch/mad/releases/5.06.01/madxuguide.pdf>.
- [34] M. Benedikt. “Optical Design of a Synchrotron with Optimisation of the Slow Extraction for Hadron Therapy”. PhD thesis. 1997. URL: <https://cds.cern.ch/record/450847>.
- [35] Y. Kobayashi. “Theory of the resonant beam ejection from synchrotrons”. In: *Nuclear Instruments and Methods* 83.1 (1970), pp. 77–87. URL: <https://www.sciencedirect.com/science/article/pii/0029554X70905379>.
- [36] K. N. Henrichsen and M. J. De Jonge. “Acceleration by Phase Displacement in the ISR”. In: *9th International Conference on High-Energy Accelerators*. 1975, pp. 390–393. URL: <https://cds.cern.ch/record/312835>.
- [37] E. W. Messerschmid. *Scattering of particles by phase displacement acceleration in storage rings*. CERN-ISR-TH-73-31. 1973. URL: <https://cds.cern.ch/record/1256629/files/ISR-73-31.pdf>.

- [38] *Liouville's theorem (Hamiltonian)*. In: *Wikipedia*. Page Version ID: 1140439605. 2023. URL: [https://en.wikipedia.org/w/index.php?title=Liouville%27s_theorem_\(Hamiltonian\)&oldid=1140439605](https://en.wikipedia.org/w/index.php?title=Liouville%27s_theorem_(Hamiltonian)&oldid=1140439605).
- [39] J. Liouville. “Note sur la Théorie de la Variation des constantes arbitraires”. In: *Journal de Mathématiques Pures et Appliquées* (1874). URL: http://www.numdam.org/item/JMPA_1838_1_3__342_0.pdf.
- [40] B. Goddard et al. “Reduction of 400 GeV / c slow extraction beam loss with a wire diffuser at the CERN Super Proton Synchrotron”. In: *Physical Review Accelerators and Beams* 23.2 (2020), p. 023501. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.023501>.
- [41] M. Fraser et al. “Slow Extraction at the SPS: Extraction Efficiency and Loss Reduction Studies”. In: *Proceedings of the Injector MD Days 2017*. Injector MD Days 2017. Vol. 2. 2017, pp. 87–87. URL: <https://e-publishing.cern.ch/index.php/CP/article/view/653>.
- [42] H. Damerou et al. *LHC Injectors Upgrade, Technical Design Report*. 2014. URL: <https://cds.cern.ch/record/1976692>.
- [43] V. Kain et al. “Resonant Slow Extraction with Constant Optics for Improved Separatrix Control at the Extraction Septum”. In: *Physical Review Accelerators and Beams* 22.10 (2019). Publisher: American Physical Society, p. 101001. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.101001>.
- [44] V. Kain, K. Cornelis, and E. Effinger. “New Spill Control for the Slow Extraction in the Multi-Cycling SPS”. In: *Proceedings of IPAC2016*. IPAC2016. 2016. URL: <https://accelconf.web.cern.ch/ipac2016/papers/tupmr051.pdf>.
- [45] J. Prieto et al. “Beam Dynamics Simulations of the Effect of Power Converter Ripple on Slow Extraction at the CERN SPS”. In: *Proceedings of the 9th Int. Particle Accelerator Conf. IPAC2018*. In collab. with Satogata Todd (Ed.) and Schaa RW (Ed.) Volker. 2018. URL: <https://accelconf.web.cern.ch/ipac2018/papers/tupaf050.pdf>.
- [46] M. Pari et al. “Characterization of the slow extraction frequency response”. In: *Physical Review Accelerators and Beams* 24.8 (2021), p. 083501. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.083501>.
- [47] M. A. Fraser et al. “Demonstration of slow extraction loss reduction with the application of octupoles at the CERN Super Proton Synchrotron”. In: *Physical Review Accelerators and Beams* 22.12 (2019). Publisher: American Physical Society, p. 123501. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.123501>.
- [48] P. A. Arrutia Sota. *henontrack*. 2023. URL: <https://gitlab.cern.ch/parrutia/henontrack>.
- [49] M. Pari. “Study and development of SPS slow extraction schemes and focusing of secondary particles for the ENUBET monitored neutrino beam”. PhD thesis. 2020. 205 pp.

- [50] SHiP Collaboration. *SHiP Experiment, Progress Report*. CERN-SPSC-2019-010, SPSC-SR-248. 2019. URL: <https://cds.cern.ch/record/2654870/>.
- [51] SHiP Collaboration. *BDF/SHiP at the ECN3 high-intensity beam facility*. CERN-SPSC-2022-032, SPSC-I-258. 2022. URL: <https://cds.cern.ch/record/2839677/>.
- [52] J. W. Glenn. *Modeling of Micro-Bunching During AGS Slow Extraction*. BNL-104842-2014-IR. 1996. URL: <https://www.osti.gov/biblio/1151333>.
- [53] J. Glenn et al. “Micro-bunching of the AGS slow extracted beam for a rare kaon decay search”. In: *PACS2001. Proceedings of the 2001 Particle Accelerator Conference (Cat. No.01CH37268)*. 2001 Particle Accelerator Conference. Vol. 2. Chicago, IL, USA: IEEE, 2001, pp. 1529–1531. URL: <http://ieeexplore.ieee.org/document/986737/>.
- [54] K A Brown et al. “Mini-bunched and Micro-bunched Slow Extracted Beams from the AGS”. In: *Proceedings of EPAC 2004*. EPAC 2004. 2004. URL: <https://accelconf.web.cern.ch/e04/PAPERS/TUPLT179.PDF>.
- [55] V. Nagaslaev et al. “Third Integer Resonance Slow Extraction Scheme for a Mu2e Experiment at Fermilab”. In: *Proceedings of HB2010*. HB2010. 2010. URL: <https://lss.fnal.gov/archive/2010/conf/fermilab-conf-10-395-ad.pdf>.
- [56] S. Sorge, P. Forck, and R. Singh. “Measurements and Simulations of the Spill Quality of Slowly Extracted Beams from the SIS-18 Synchrotron”. In: *9th International Particle Accelerator Conference*. IPAC2018. Vol. 1067. 2018, p. 052003. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/1067/5/052003>.
- [57] K. Noda et al. “Advanced RF-KO slow-extraction method for the reduction of spill ripple”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 492.1 (2002), pp. 253–263. URL: [https://doi.org/10.1016/S0168-9002\(02\)01319-0](https://doi.org/10.1016/S0168-9002(02)01319-0).
- [58] M. Bozzo et al. “Observation of the Slow-Extracted 25ns Bunched Beam”. In: *Proceedings of 11th Workshop of the LHC*. 11th Workshop of the LHC. 2001. URL: <https://cds.cern.ch/record/567178>.
- [59] M. Tomizawa et al. “Slow extraction from the J-PARC main ring using a dynamic bump”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 902 (2018), pp. 51–61. URL: <https://doi.org/10.1016/j.nima.2018.06.004>.
- [60] L. Stoel et al. “Progress Toward a Dynamic Extraction Bump for Slow Extraction in the CERN SPS”. In: *Proceedings of IPAC2018*. IPAC2018. 2018. URL: <https://accelconf.web.cern.ch/ipac2018/papers/tupaf055.pdf>.

- [61] R. Muto et al. “Current status of slow extraction from J-PARC Main Ring”. In: *10th International Particle Accelerator Conference*. Vol. 1350. 2019. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/1350/1/012105>.
- [62] S. Sorge, P. Forck, and R. Singh. “Spill ripple mitigation by bunched beam extraction with high frequency synchrotron motion”. In: *Physical Review Accelerators and Beams* 26.1 (2023), p. 014402. URL: <https://doi.org/10.1103/PhysRevAccelBeams.26.014402>.
- [63] F. M. Velotti et al. “Septum shadowing by means of a bent crystal to reduce slow extraction beam loss”. In: *Physical Review Accelerators and Beams* 22.9 (2019). Publisher: American Physical Society, p. 093502. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.093502>.
- [64] Marc Delrieux. *Discussion on PS Transverse Feedback (private communication)*. 2023.
- [65] E. Cortina Gil et al. “The beam and detector of the NA62 experiment at CERN”. In: *Journal of Instrumentation* 12.5 (2017), P05025. URL: <https://dx.doi.org/10.1088/1748-0221/12/05/P05025>.
- [66] B. Dobrich. “Highlights and requests from the North Area experiments for proton beams”. Injectors and Experimental Facilities Workshop. 2021. URL: <https://indico.cern.ch/event/1063281/contributions/4468572>.
- [67] M. Gyr. *Proposal for a new Servo-Spill System: Power Requirements for different Configurations*. CERN-SL-95-103. 1995. URL: <https://cds.cern.ch/record/292415>.
- [68] D. Boussard et al. *Slow extraction at 400 GeV/c with stochastic RF noise*. CERN-SPS-AOP-MG-jf, CERN-SPS-Improvement-Report-179. 1980. URL: <https://cds.cern.ch/record/1020094>.
- [69] M. G. Pullia et al. “Betatron Core Driven Slow Extraction at CNAO and MedAustron”. In: *Proceedings of IPAC2016*. IPAC2016. 2016. URL: <https://accelconf.web.cern.ch/ipac2016/papers/tupmr037.pdf>.
- [70] D. Bloess et al. *Measurement of the CERN slow ejected beam time structure*. CERN-MPS-CO-71-3, CERN-MPS-CO-71-3. 1971. URL: <https://cds.cern.ch/record/2831259/>.
- [71] P. Burla et al. *Power Supply Ripple Study at the SPS*. CERN-SL-94-11-AP. 1994. URL: <https://cds.cern.ch/record/263879>.
- [72] J. Bechhoefer. *Control Theory for Physicists*. Edition: 1. Cambridge University Press, 2021. URL: <https://www.cambridge.org/core/product/identifier/9780511734809/type/book>.
- [73] M. Pullia. “Dynamique de l’éjection lente et son influence sur les lignes de transfert”. PhD thesis. Geneva, 1999. 145 pp.

- [74] M. Crescenti. *RF Empty Bucket Channelling Combined with a Betatron Core to Improve Slow Extraction in Medical Synchrotrons*. CERN-PS-97-068-DI. 1998. URL: <https://cds.cern.ch/record/346139/>.
- [75] R. Piandani. “Spill quality with different RF off timings and voltages”. SPS MPC #24. CERN, 2022. URL: <https://indico.cern.ch/event/1176807>.
- [76] M. A. Fraser. “Requirements and achievement at CERN PS and SPS”. iFAST-REX Collaboration Meeting. 2022. URL: <https://indico.gsi.de/event/14171/contributions/59828/>.
- [77] R. Piandani. *Discussion on NA62 analysis (private communication)*. 2023.
- [78] Pablo Andreas Arrutia Sota et al. “Implementation of RF Channeling at the CERN PS for Spill Quality Improvements”. In: *Proceedings of IPAC2022*. IPAC2022. Vol. 2022. 2022, pp. 2114–2117. URL: <https://cds.cern.ch/record/2845777>.
- [79] P. A. Arrutia Sota et al. “Millisecond burst extractions from synchrotrons using RF phase displacement acceleration”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment (2022)*, p. 167007. URL: <https://www.sciencedirect.com/science/article/pii/S0168900222004363>.
- [80] P. A. Arrutia Sota et al. “RF techniques for spill quality improvement in the SPS”. In: *Proceedings of the 14th International Particle Accelerator Conference*. IPAC23. 2023. URL: <https://doi.org/10.18429/JACoW-IPAC2023-MOPA116>.
- [81] Andrea De Franco et al. “Slow Extraction Optimization at the MedAustron Ion Therapy Center: Implementation of Front End Acceleration and RF Knock Out”. In: *Proceedings of the 9th Int. Particle Accelerator Conf*. IPAC2018 (2018). In collab. with Satogata Todd (Ed.) and Schaa RW (Ed.) Volker. Artwork Size: 4 pages, 1.036 MB ISBN: 9783954501847 Medium: PDF Publisher: JACoW Publishing, Geneva, Switzerland, 4 pages, 1.036 MB. URL: <http://jacow.org/ipac2018/doi/JACoW-IPAC2018-MOPML025.html>.
- [82] ENUBET. *The ENUBET Project*. CERN-SPSC-2018-034, SPSC-I-248. 2018. URL: <https://cds.cern.ch/record/2645532/>.
- [83] M.-C. Vozenin et al. “The Advantage of FLASH Radiotherapy Confirmed in Mini-pig and Cat-cancer Patients”. In: *Clinical Cancer Research: An Official Journal of the American Association for Cancer Research* 25.1 (2019), pp. 35–42. URL: <https://doi.org/10.1158/1078-0432.ccr-17-3375>.
- [84] K. Kissler et al. “Fast Resonant Extraction from the CERN SPS”. In: *Nuclear Science, IEEE Transactions on* 26 (1979), pp. 3228–3230. URL: <https://cds.cern.ch/record/134662>.
- [85] U. Dorda et al. “Simulation Studies of Longitudinal RF-noise and Phase Displacement Acceleration as Driving Mechanism for the MedAustron Synchrotron Slow Extraction”. In: *Proceedings of IPAC2013*. IPAC2013. 2013. URL: <https://accelconf.web.cern.ch/IPAC2013/papers/wepea001.pdf>.

- [86] K. Cornelis. *Momentum cleaning in the LHC with resonant Extraction*. LHC Project 176. 1998. URL: <https://cds.cern.ch/record/692007/files/project-note-176.pdf>.
- [87] C. Cartis et al. “Improving the Flexibility and Robustness of Model-based Derivative-free Optimization Solvers”. In: *ACM Transactions on Mathematical Software* 45.3 (2019), 32:1–32:41. URL: <https://dl.acm.org/doi/10.1145/3338517>.
- [88] M. Vadai et al. *Implementation of synchronised PS-SPS transfer with barrier buckets*. arXiv:2210.05416. type: article. 2022. arXiv: 2210.05416[physics]. URL: <http://arxiv.org/abs/2210.05416>.
- [89] E. Shaposhnikova et al. “Intensity upgrade plans in SPS with 800 MHz system”. LIU-SPS Coordination Meeting. 2012. URL: <https://indico.cern.ch/event/195397/>.
- [90] F. Gao and L. Han. “Implementing the Nelder-Mead simplex algorithm with adaptive parameters”. In: *Computational Optimization and Applications* 51.1 (2012), pp. 259–277. URL: <https://doi.org/10.1007/s10589-010-9329-3>.
- [91] M Vadai et al. “Beam manipulations with barrier buckets in the CERN PS”. In: *Journal of Physics: Conference Series* 1350 (2019), p. 012088. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/1350/1/012088>.
- [92] Heiko Damerau. “Creation and storage of long and flat bunches in the LHC”. PhD thesis. 2005.
- [93] W. Hardt. *Ultraslow extraction out of LEAR (Transverse Aspects)*. PS-DL-LEAR-81-6. 1981. URL: <https://cds.cern.ch/record/1025914/>.