Investigations of a Novel Bunch Compression Technique for Hadron Accelerators

Johan Holmberg

A thesis presented for the degree of
Master of Science

Supervised by Natalia Milas and Francesca Curbis
Division of Synchtrotron Radiation Research
Lund University, Sweden
Project duration: 10 months
June 2026

Abstract

Bunch compression is necessary for many accelerator applications, one of which being the creation of muons for the proposed muon collider. The design of a proton driver to deliver a short and very intense proton pulse to a target for the creation of muons is being developed as part of the International Muon Collider Collaboration. Some challenges with the proposed baseline design, notably the ambitious requirements for the radiofrequency system, motivate the investigation of alternative compression methods. Such a method is developed and investigated in this thesis.

The fundamental idea of this method is that a series of chirped bunch trains can be compressed in the same compressor ring as in the baseline design, removing the need for an accumulator ring and radiofrequency cavities in the compressor ring. To study this novel scheme, the chirp creation was first developed using off-frequency cavities at the end of the linac. A toy model was also developed to investigate the parameters for the compression, giving a basic set of parameters which was then improved upon. The scheme was implemented in the particle-in-cell simulation code PyORBIT, adding non-linear and collective effects to the compression. An iterative process led to an optimised scheme and corresponding set of machine and simulation parameters.

Of most interest in the investigation of the scheme was the impact of collective effects, particularly space charge. This was investigated for two compressor lattices operating above and below transition respectively. The strength of the space charge forces caused strong instabilities above transition, leading to microbunching, emittance growth and eventually the breaking up of the bunch. It was concluded that this was due to the negative mass instability. This finding prompted the investigation of a lattice below transition, but in this case the space charge repulsion was found to hinder the compression entirely.

In summary, the collective effects experienced during the compression were found to be too detrimental for realising the proposed scheme with the muon collider parameters in both cases. However, the compression scheme was also investigated at lower bunch intensities, reducing the strength of the collective effects. For the lattice above transition, this allowed for successful compression at 2 % of the original bunch intensity. Based on this, an alternative use case for the scheme is proposed for a neutrino factory, by modifying the repetition rate and bunch structure.

Abbreviations and Acronyms

- ESS – European Spallation Source
- IMCC – International Muon Collider Collaboration
- LBNF – Long-Baseline Neutrino Facility
- Fermilab – Fermi National Accelerator Laboratory
- RF – Radio Frequency
- linac – Linear Accelerator
- LHC – Large Hadron Collider
- HL-LHC – High Luminosity Large Hadron Collider
- FCC – Future Circular Collider
- CLIC – Compact Linear Collider
- ESPP – European Strategy for Particle Physics
- COM – Centre-of-mass
- RMS – Root Mean Squared
- IP – Interaction Point
- NMI – Negative Mass Instability
- DTL – Drift Tube Linac
- SCL – Superconducting Linac
- ToF – Time-of-Flight
- HEBT – High Energy Beam Transport
- HBL – High-Beta Linac
- PIC – Particle-in-Cell
- SNS – Spallation Neutron Source
- SPL – Superconducting Proton Linac
- IBS – Intra-Beam Scattering

Relevant Vocabulary

- Luminosity – The number of interactions occurring per second for a process with a cross section of one barn (1 barn $= 1 \times 10^{-28}\,\mathrm{m}^2$).
- Betatron function – The solution to the position dependent Hill's equation, giving the amplitude at a given location for a pseudo-harmonic oscillator. Provides the envelope for the oscillating particle motion in accelerators.
- Courant–Snyder parameters – The beta function and derived quantities which parametrise the phase space motion of particles in an accelerator.
- Emittance – A quantity proportional to the area covered in phase space by particles performing oscillations. Can also be used as a measure for the phase space area (or volume) covered by an entire bunch.
- Liouville's theorem – Theorem stating in the context of accelerator physics that the phase space volume (emittance) of a particle distribution is conserved under the influence of exclusively conservative forces.
- Tune – The advance of a particle in phase space after one full turn through a storage ring, either in transverse or longitudinal. Relates to the property phase advance, which can be derived from the beta function.
- Ideal orbit – The orbit of a particle through an accelerator which maintains exactly the same transverse and longitudinal position turn after turn. This is the central stability point around which all particles oscillate.
- Dispersion – Measure of the orbit excursion of a particle depending on its deviation from the momentum of a particle following the ideal orbit.
- Synchronous particle – The fictional particle which experiences exactly the energy gain required to stay on the ideal orbit every turn, around which all particles oscillate longitudinally.
- Separatrix – The boundary in phase space separating stable from unstable motion, specifically in the context of the longitudinal stability in the presence of an RF cavity.
- Bucket – The area enclosed by the separatrix, where all particles in a bunch should ideally be located.

Contents

List of Figures

List of Tables

1 Introduction

In many types of particle accelerators, longitudinal bunch compression is used to reach the desired intensities and pulse lengths. It is used to deliver short and brilliant X-ray pulses in free electron lasers, to create intense neutrino pulses in neutrino factories or long baseline neutrino experiments, to deliver exact treatments in medical applications, and to create bright neutron pulses with the desired spectral properties in spallation sources. Improving and developing the techniques used for bunch compression is therefore of interest in many fields of research, where shorter and more intense pulses are required to advance the research which can be performed.

Particularly, this can be applied to a muon collider, a proposed accelerator which is being investigated by the International Muon Collider Collaboration (IMCC). To create the required muon beam, a very intense and compressed proton beam is first required. When sent to a target, the resulting particle cascade would decay partially into muons which can be captured and cooled to create a useable muon bunch. The constraints on the resulting muon bunches determine the requirements on the proton beam, which are more ambitious than those of any existing or proposed proton accelerator. The proton accelerator has to deliver 5×10^{14} protons in 2 ns bunches with a 5 Hz repetition rate, corresponding to an average power of 2 MW. Comparing this to another high intensity proton driver, the Long-Baseline Neutrino Facility (LBNF) beamline at Fermi National Accelerator Laboratory (Fermilab), the number of protons is an order of magnitude greater with a bunch length three orders of magnitudes shorter [1]. Compared to LBNF, the proton energy is also considerably lower, with the muon collider proton driver operating at either 5 GeV or 10 GeV compared to the 60–120 GeV of the LBNF. This increases the impact of space charge forces considerably, which alongside other collective effects are the main limitations for the performance [2, p. 72]. Such effects all increase with particle density, and are thus expected to have large impacts on the quality of the beam. To avoid such effects, it is important to achieve a quick compression in a low number of turns.

The proposed proton driver achieves compression in two stages, and utilises very high field gradients in the Radio Frequency (RF) cavities of the second ring to achieve the compression. The feasibility of this design has not been demonstrated yet, inviting the possibility of alternative compression schemes. The purpose of this project was to investigate one such method which removes both the accumulator ring (see Figure 2) and the RF cavities in the compressor ring. Instead, a longitudinal energy correlation (chirp) is imprinted on the bunches from the linac, which in the compressor ring is converted to a compression by the large momentum compaction factor of the ring. By adding a pair of RF cavities in the linac, operating at frequencies slightly different than the nominal RF frequency (which determines the bunch spacing), the proton bunches will experience a superposition of fields which results in different energy gains for the different bunches. An approximately linear correlation can be imprinted this way by appropriate choice of frequency and corresponding chopping pattern of the linac pulse. The process is schematically shown in Figure 1.

This scheme has several advantages compared to the original design: the removal of the accumulator ring simplifies the design and reduces the footprint significantly, while moving the RF system away from the compressor ring relaxes the requirements on the RF. It also presents some considerable challenges. Since the accumulation is done simultaneously with the compression in the compressor ring, the total compression time is limited by the injection rate of the linac. This sets the total compression time to at least 1000 turns, which is 20 times longer than in the original design. Consequently, collective effects are a significantly larger concern with this scheme than in the original one, as the particles in the bunch have more time to interact. It is expected that these effects, especially space charge, will lead to negative effects like instabilities and particle loss [3, 4]. Limiting these effects is the key challenge in achieving the desired compression with the given parameters.

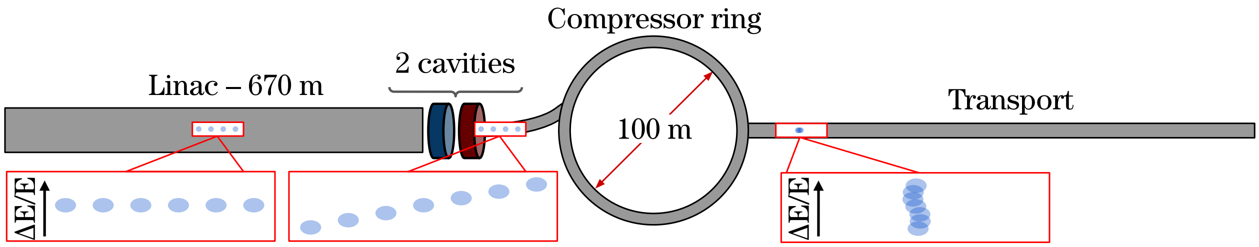

Figure 1: A bunch train is accelerated to full energy with a chirp being imprinted by two off-frequency cavities at the end of the linac, making the bunch compress in the compressor ring.

In this thesis, the study case of the muon collider will first be introduced. Then, the fundamentals required for understanding the compression are introduced, namely the transverse and longitudinal dynamics of particle motion in an accelerator, allowing for a convenient description of the compression using the concept of phase space. This is done for both linacs and storage rings, which have slightly different conventions and conditions. Following that is an explanation of different collective effects which impact the beam quality and compression, where space charge is the most critical. Thereafter is an explanation of the drawbacks of the baseline compression scheme, motivating the investigation of this novel scheme. Subsequently, this scheme and its implementation are explained in more detail. A potential proof-of-principle measurement for the chirp creation at the ESS linac is also explained. The scheme is then investigated analytically to identify key parameters, defining a lower bound on the final bunch length while using a toy model implementation of the compression scheme to define the required frequency and amplitude of the chirp. The simulations of the compressions are then explained, starting with the analytical model and moving to gradually more involved particle-in-cell simulations. These simulations were performed using PyORBIT [5]. The results of these are subsequently given and discussed, where the limitations from collective effects are investigated in detail. Different intensities are investigated, as is an alternative lattice, allowing finally for an optimal set of parameters for the machine to be given, as well as the alternative use case of a neutrino factory.

2 Muon Collider

Investigating the standard model of physics, which describes the fundamental interactions of all particles around us, is a high priority of the international research community. This is seen in the ambitious plans set out by the European Strategy for Particle Physics (ESPP), which proposes the construction of the Future Circular Collider (FCC) at CERN in Geneva, a decades long project with a proposed budget on the order of 10 BCHF [6, 7]. If this project is realised, it would define the particle physics landscape for the four decades following the end of the High-Luminosity Large Hadron Collider (HL-LHC) operations. Although many questions in fundamental physics could potentially be answered by such a machine, there are limitations to what it can deliver. The limitations mainly come from the fact that the high-energy and precision frontiers are investigated separately. At first, the discovery potential of the machine is limited by the energy of the machine (for the electron–positron collider FCC-ee), and then by the precision of its measurements (for the hadron collider FCC-hh). A muon collider would investigate these avenues simultaneously, offering precision measurements of the standard model up to energies equivalent to those of the FCC-hh [8].

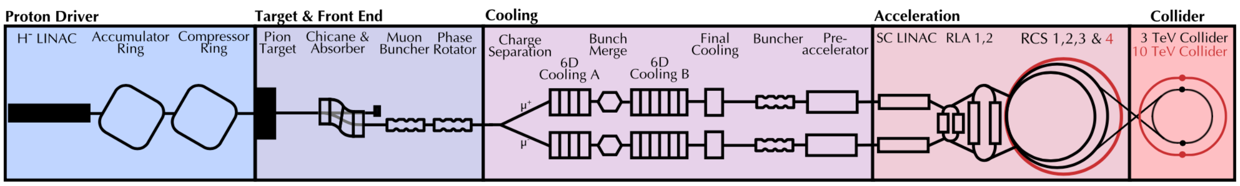

Figure 2: The proposed layout of the muon collider [9].

2.1 Physics Case

The centre-of-mass (COM) energy in a collider is a measure for how deep the experiment can probe into unknown physics. In a circular accelerator, the achievable energy is limited by the radius of the machine, the magnet strength, and the synchrotron radiation emitted by the circulating particles. A collider like the FCC is pushing the technological and budgetary restrictions to these parameters as far as possible, but it is inherently limited by the choice of particle species to accelerate. Specifically the synchrotron radiation becomes an issue, as it relates to both the machine radius and the mass of the accelerated particle. The radiation corresponds to an energy loss of the particles every turn, which must be replenished by RF cavities in the ring. The energy loss is proportional to the ratio of the particle energy to its rest energy, also known as the Lorentz factor $\gamma = E/E_0$, to the fourth power as

$$U_0 = \frac{e^2 I_2}{6\pi\varepsilon_0}\gamma^4 \approx \frac{e^2 C}{6\pi\varepsilon_0}\gamma^4, \tag{1}$$

where I_2 is the second synchrotron radiation integral around the ring, which in the simplest case is approximately equal to the circumference of the ring [10, p. 439]. Since the scaling with rest energy is with the fourth power, the energy limit varies greatly among particle species. Electrons, which like all massive leptons are ideal for precision measurements due to their point-like nature, have a very low mass and are therefore limited by the energy loss per turn due to synchrotron radiation. Protons are not limited by their mass, but are composite particles where the constituent quarks and gluons each carry only a fraction of the total momentum. This effectively limits the available COM energy to around a tenth of the total COM energy of the collision. For the muon, a heavier lepton than the electron, this limitation does not hold, meaning a 10 TeV muon collider would offer similar physics reach as the 100 TeV FCC-hh [8]. The muon mass is 105.7 MeV, 207 times the electron mass [11], meaning the energy limit is a factor of $207^4 = 1.8 \times 10^9$ greater than that of an electron machine with the same radius, comfortably enabling energies up to the proposed 10 TeV even at a smaller circumference.

2.2 Accelerator Complex

While muons have several desirable characteristics for collider experiments, the acceleration of muons is far from trivial due to their short lifetime of 2.2 µs [11]. One of the main motivations for circular colliders over linear ones (like the proposed CLIC [12]) is that the particle bunches can be reused many times, increasing the average number of collisions per unit time. This corresponds to an increase in a quantity called *luminosity*, which is the figure of merit describing how many physics events are produced per unit time per unit cross section. It is defined as

$$\mathcal{L} = \frac{N^2 n_b f_{\mathrm{rev}}}{4\pi\sigma_x\sigma_y}, \tag{2}$$

for a revolution frequency f_{rev}, where N is the number of particles in a bunch, with n_b being the number of bunches and $\sigma_{x/y}$ the beam size in the two transverse dimensions. The luminosity is typically expressed in units of $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. This is clearly dependent on the number of bunches colliding per unit time, meaning collider experiments have an inherent advantage to linear colliders which are limited by the repetition rate of the accelerator. This is exemplified in the case of CLIC and the proposed muon collider in Figure 3, where the increase in luminosity with beam power has a plateau for linear colliders when increasing the COM energy, unlike a circular one. Since the energy used by the accelerators to reach collision energy is expended only once per cycle for a circular machine (a cycle typically lasting several hours), it uses considerably less power than a linear machine which must expend this energy for every new collision event (with a repetition rate on the order of 1 Hz or higher). It is clear that in order to reach the desired luminosity and energy for a future collider at a reasonable power consumption, a circular muon collider is preferred to a linear electron collider.

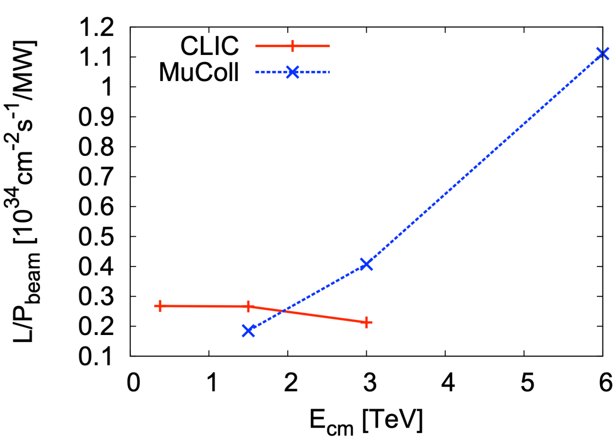

Figure 3: The achievable luminosity normalised to beam power at different COM energies for a linear collider (CLIC) and the proposed muon collider [8].

The luminosity is proportional to the square of the number of muons per bunch, meaning it is essential to create and maintain muon bunches with large particle densities. This is limited greatly by the lifetime of the muons which prevents traditional techniques to increase bunch intensity, such as accumulation in a storage ring. Hence the very stringent proton driver parameters, which ensure that the muon bunches are populated enough to deliver the required luminosity. If the number of protons was kept constant, but divided in more bunches, the total luminosity would decrease considerably considering the linear scaling with n_b but quadratic scaling with N in Equation 2. The parameters required for the proton driver to deliver a luminosity on the order of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ [9, p. 9] are given in Table 1, with the first option being the preferred one and thus the one studied most in this thesis.

Table 1: The proposed parameters for the proton driver of the muon collider, as defined in [9, p. 72], where the first option is the one primarily studied in this thesis.

Quantity	Value		Unit
	Option 1	Option 2	
Energy, E_{kin}	5	10	GeV
Average power, P_{av}	2	4	MW
Number of protons, N	5×10^{14}	5×10^{14}	protons
Repetition rate, f_{rep}	5	5	Hz
Bunch length (RMS), τ_b	2	2	ns

2.2.1 Proton Driver

To deliver this performance, a proton driver consisting of a high-power H^- linac and two storage rings is suggested, as shown in Figure 2. The linac delivers short pulses of H^- ions with low energy spread at the final energy of 5 GeV or 10 GeV to an accumulator ring, where injection occurs by stripping. The pulses are then accumulated to an intense proton bunch with an RMS bunch length of 120 ns [9]. The energy spread must be kept low so that the bunch can be compressed via a phase space rotation in the subsequent compressor ring, as shown schematically in Figure 4. This principle is explained further in Section 3.3. In principle, the accumulator ring does not require any RF cavities to manipulate the bunch, though these may be used for keeping the bunch from debunching. For the compressor ring however, very strong RF fields are required to perform the required bunch rotation within a short number of turns.

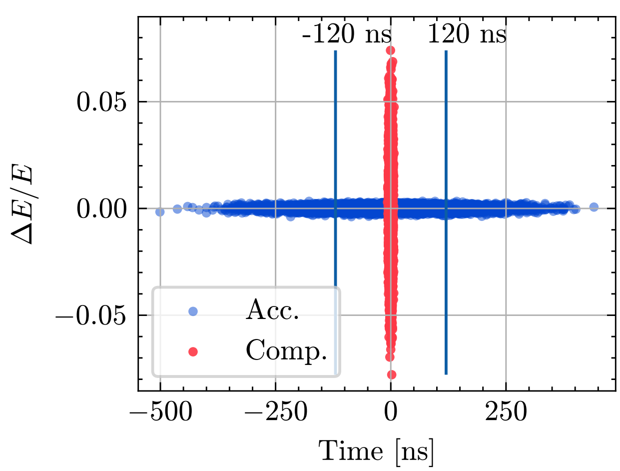

Figure 4: The bunch shape shown schematically in longitudinal phase space in the accumulator ring (blue), and compressor ring (red). The low energy spread in the accumulator ring is what allows for a short final bunch length. The vertical lines show the RMS bunch length in terms of one standard deviation of the bunch size in each direction.

Because of the unprecedented bunch intensity and bunch length, it is expected that many collective effects, some of which are explained in Section 3.4, will reduce the bunch quality. Mitigating these effects is therefore crucial to realise this proton driver design, where a short compression cycle is one key method. Another alternative is to use the higher energy option, because as is explained in Section 3.4, space charge effects scale with the Lorentz factor as γ^{-2} or γ^{-3} [13, p. 540–555]. Thus, using the 10 GeV option has a considerable advantage in reducing the impact of collective effects, while maintaining a similar muon production rate [9]. However, the increased energy poses challenges for the linac, compressor ring, and target design, which might outweigh the advantages.

3 Accelerator Dynamics

To describe bunch compression, the transverse and longitudinal behaviour of the charged particles in an accelerator must first be understood. The transverse behaviour is determined by the lattice optics, where elements such as dipole and quadrupole magnets control the beam. The longitudinal behaviour is normally defined largely by the presence of RF cavities which provide energy to the particles. In the proposed compression scheme however, no RF is used in the compressor ring, meaning some key features typically associated with particle bunches in a storage ring do not hold. The following description will nonetheless include this traditional treatment, since it is necessary for understanding the purpose of the alternative compression scheme proposed here. The following treatment of fundamental concepts of accelerator dynamics treats circular accelerators, also known as *storage rings*. Linear accelerators, or *linacs*, are given a simplified treatment in Section 3.2.4, highlighting the differences between them and storage rings.

3.1 Single Particle Dynamics—Transverse Dynamics

In a particle accelerator, the charged particles are kept in their stable orbits around the ring through the use of various electric and magnetic fields. If there were no disturbances in the lattice, one could treat the propagation of the beam as occurring entirely along an ideal orbit, defined by the dipoles which bend the particles at a given energy. In practice however, the particles do not have the exact energy or position to follow this ideal orbit. Therefore, a stable region around the ideal orbit is created by using linear restoring forces on the particles, creating a pseudo-harmonic oscillator for particles within this region. In the transverse plane, this force is created by quadrupole magnets, and in the longitudinal it can be provided by the RF cavities giving energy to the particles. In phase space[1], these oscillations manifest as ellipses, with each particle advancing along an ellipse at a certain rate as illustrated in Figure 5 (right). This description is very instructive in understanding the basic ways a beam of particles can be controlled.

[1]In this context, the term *phase space* is used even when two variables are not canonical conjugate variables of each other. Thus, a more correct term would be *trace space*, but as this difference is often neglected in literature within the field, the term *phase space* will be used with the meaning of *trace space* in this work.

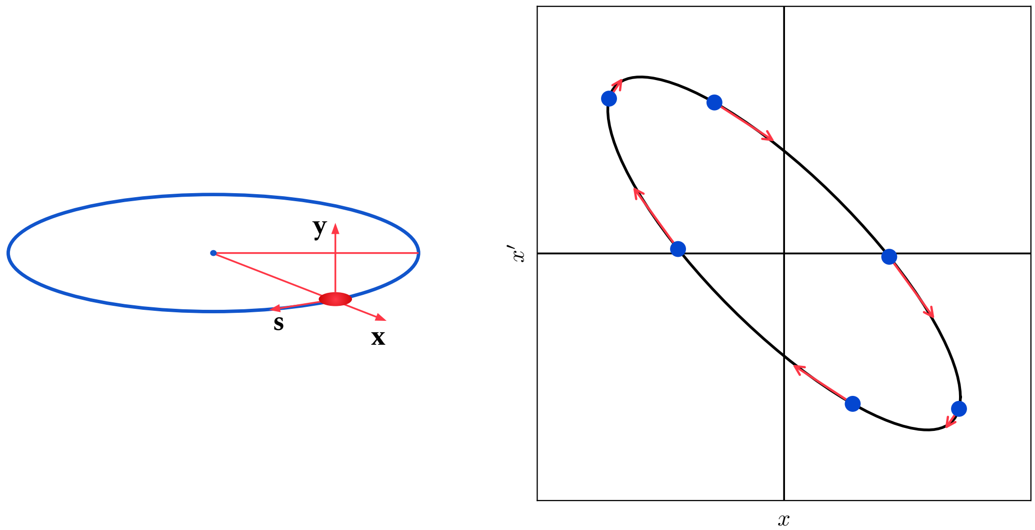

Figure 5: The coordinate system used in this work for the particles in a circular accelerator (left) and the advancing phase of a particle in phase space at the given location, with the length of the arrow corresponding to the rate of phase advance in that point (right).

Consider now a single particle traversing a circular accelerator, i. e. a storage ring. It will encounter optical elements (mainly quadrupoles) which give a kick depending on the deviation from the ideal orbit. Each such kick represents a change in the momentum of the particle, which depends on the current position of the particle, as shown in Figure 6. In phase space, this means that if a particle arrives at the quadrupole with a large transverse position, it will receive a strong kick, changing its momentum considerably. When it has low transverse displacement, the momentum will not change noticeably. This accounts for the harmonic oscillator-like behaviour of the particle, since large momenta coincide with small displacements, and vice versa.

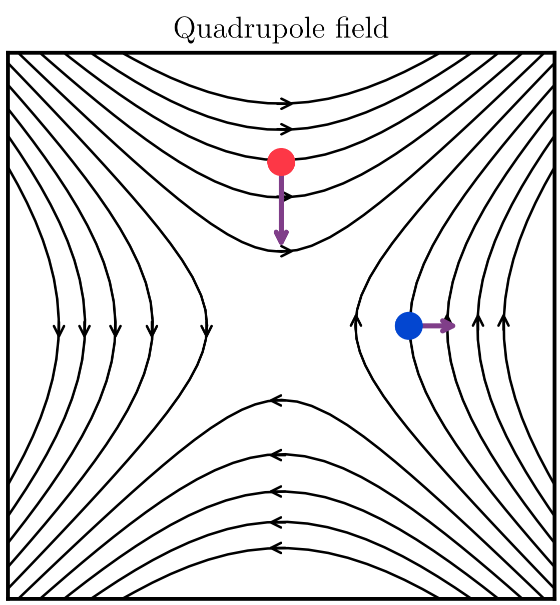

Figure 6: A schematic showing the function of a quadrupole magnet and the force it enacts on a traversing particle. Note that the force scales with the deviation from the centre, and that the force is focussing in one plane, while defocussing in the other.

However, the particle in an accelerator will pass through several quadrupoles with different strengths, and at different phases. This position dependence of the restoring force gives rise to a key behaviour of accelerators called *betatron oscillations*. The equation for a position dependent harmonic oscillator is known as Hill's equation, which for the horizontal coordinate can be expressed as

$$\frac{d^2x(s)}{ds^2} + k(s)x(s) = 0, \tag{3}$$

with $k(s)$ giving the position dependent restoring force, with s being the longitudinal coordinate around the ring, according to the coordinate system shown in Figure 5 (left). Because the lattice is periodic, the function $k(s)$ is periodic over the circumference of the ring. Solving Hill's equation gives a solution with varying amplitude, which is typically expressed in terms of *beta functions*, $\beta(s)$, which define the transverse size of the beam. These act as envelopes for the stable motion in the accelerator, and depend solely on the linear elements in the lattice. To see the impact this has on the phase space motion of the particles, the so-called *Courant–Snyder parameters* (often referred to as Twiss parameters in other literature) are introduced:

$$\begin{cases} \beta(s) = \beta(s) \\ \alpha(s) = -\frac{\beta'(s)}{2} \\ \gamma(s) = \frac{1+\alpha^2(s)}{\beta(s)} \end{cases} . \tag{4}$$

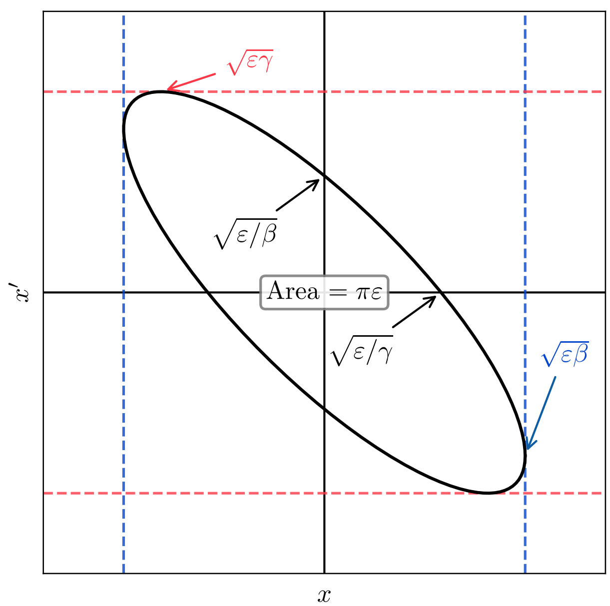

Figure 7: The phase space ellipse expressed in terms of the Courant–Snyder parameters and the emittance, ε.

When parametrising the phase space ellipse in a certain point, these three parameters can be used to determine the size and angle of the ellipse, as illustrated in Figure 7. The beta function controls the beam size directly as

$$\sigma = \sqrt{\varepsilon\beta}, \tag{5}$$

and the alpha function controls the tilt of the ellipse indirectly through γ. This means that the shape and angle of the ellipse changes around the ring. However, the total area of the ellipse is conserved, which is a consequence of *Liouville's theorem* (see Section 3.3), and this area is generally referred to in terms of the quantity *emittance*, ε. A single particle moving through the accelerator will remain on an ellipse with an area $\pi\varepsilon$ which changes shape with the position s according to Equation 4. Its location on the ellipse also changes with s, according to the quantity *phase advance*,

$$\mu(s) = \int_0^s \frac{ds}{\beta(s)}. \tag{6}$$

With this quantity, it is possible to see how much a particle has advanced in phase angle around the phase space ellipse between two given points (by setting the bounds of the integral to the corresponding values). Specifically, it is very useful to see the phase advance for one full revolution, i. e. with $s = C$, the circumference of the ring. This quantity is proportional to the tune,

$$Q = \frac{1}{2\pi} \oint \frac{ds}{\beta(s)}, \tag{7}$$

which is one of the key characteristics of a storage ring. This value shows how much the particle will have advanced in phase when returning to the same location. Here, the particle advances around a single phase space ellipse with a constant shape, due to the periodicity condition of the lattice, as in Figure 5 (right). The tune can be used to determine whether the beam is at risk of resonance behaviour, as will be described further in Section 3.4.1. Therefore, it must be controlled carefully by varying quadrupole strengths around the ring to determine the phase advance.

This holds for both transverse dimensions, and to first order there is no interaction between the two planes. This means it is possible to treat the phase spaces separately. Similar quantities can be defined in the longitudinal dimension, which will be explained further in the following section. However, there is a coupling between longitudinal and horizontal motion induced by *dispersion* which is treated in this work, though it is weak enough to treat as a perturbation. Typically the emittance and tune (or phase advance) differ between the two transverse planes.

3.2 Single Particle Dynamics—Longitudinal Dynamics

Many concepts from the transverse dynamics can be applied for the longitudinal dynamics as well. Notably, Hill's equation still determines the stable motion of the particles, and thus the notions of Courant–Snyder parameters and phase advance are important for the description. However, the restoring force in the longitudinal direction comes from the RF cavities providing energy to the particles, not from magnetic quadrupoles. To see how this occurs, the equations of motion in a storage ring are studied to find the characteristic frequency for the longitudinal motion, corresponding to the tune in the transverse dimensions. This frequency is very relevant for bunch compression, as it determines how fast compression occurs.

3.2.1 RF Cavities

RF cavities provide the energy to particles, and can be operated either so that the net energy of the particles increases every turn, or so that it is kept constant. Either way, they provide a stabilising force which is the key to longitudinal dynamics.

To understand this, the acceleration of particles in an RF cavity must first be treated. A cavity is designed to support a resonant mode at a specific frequency, and is fed with a high power electromagnetic pulse at this frequency which resonates in the cavity for a certain time. The beam is bunched at intervals with the same periodicity as the RF wave in the cavity, and the timing of the pulses is controlled to the precision of degrees within this sinusoidal waveform. This ensures that the particle bunches arrive at a phase where the cavity has a large accelerating field in the direction of the particle's motion, and exit the cavity before the standing wave has changed direction and would decelerate the particles. Since the frequencies are in the range of radio frequencies, ~0.1–1 GHz, the period of an oscillation is on the order of 1 ns. For particles travelling at a velocity approaching the speed of light, this corresponds to a distance of around 30 cm, which is a comparable distance to the size of a cavity (~1 m). It is clear that the field will change considerably during the transit of a bunch, which affects the total energy transferred by the cavity to the particles. For a wave with an amplitude of V_0, a particle will receive the energy

$$\delta E = q \cdot T_t \cdot V_0, \tag{8}$$

where q is the charge of the particle (typically $q = e$, the elementary charge) and T_t is the *transit time factor*, which takes into account the the change in the field during the transit of the particles. For a gap of length g, it is defined as

$$T_t = \frac{\int_{-g/2}^{g/2} E(z) \cos\left(\frac{\omega_{\rm rf} z}{\beta c}\right) dz}{\int_{-g/2}^{g/2} E(z)\, dz}, \tag{9}$$

for an angular frequency of the RF wave of $\omega_{\rm rf}$, with cosine being the convention for expressing the cavity field in linacs, unlike in storage rings where sine is typically used. The electric field strength $E(z)$ relates to the amplitude of the wave as

$$V_0 = \int_{-g/2}^{g/2} E(z)\, dz, \tag{10}$$

which for a constant electric field along the longitudinal axis can be simplified to $V_0 = E \cdot g$. In reality, a cavity also has a radial dependence on the electric field, in which case the electric field is expressed as $E(\varrho, z)$ for a radius ϱ. For further reading on this dependence and the various electromagnetic modes which are used for acceleration, see [14, 15]. The field modes are influenced by the shape of the cavity, with the simplest one being the so-called *pillbox cavity*, which is a cylinder with its radial axis along the beam direction. It can be shown [14] that the frequency of the cavity mode is inversely proportional to the radial size of the cavity, ϱ, as

$$f_{\rm cav} \propto \frac{1}{\varrho}. \tag{11}$$

Consider again the energy gain for a particle, but this time adding the dependence on the time, τ, it arrives with respect to the RF wave,

$$\delta E(\tau) = q \cdot T_t \cdot V_0 \sin(\omega_{\rm rf}\tau) \quad \longrightarrow \quad \delta E(\phi) = qV_{\rm rf} \sin(\phi), \tag{12}$$

where the arrival time is translated to a phase, ϕ, and the effective voltage of the cavity is $V_{\rm rf} = T_t \cdot V_0$. The energy gained by the accelerating particles now depends on their phase, i. e. location in the bunch. This dependence can be used to create stable motion akin to that described in the transverse plane.

3.2.2 Equations of Motion

The longitudinal phase space motion is often described in terms of a so-called *bucket* which encloses all particles which undergo stable motion in the accelerator. This can be seen if the phase space motion around the synchronous phase is treated for a storage ring, here at constant energy. The bunches are separated by a multiple of the RF period, meaning that a single bunch falls within one period of the RF wave, with the total number of RF periods contained in the ring being known as the *harmonic number*, $h = \frac{f_{\rm rf}}{f_{\rm rev}}$. In each such period, there are two zero-crossing of the RF wave, where a hypothetical particle would receive no energy from the field. Such a particle is by definition stable, as it stays on the ideal orbit every turn. In reality however, particles have different longitudinal positions and thus different phases with respect to the RF field, corresponding to different energy gains. It is important that a bunch of particles lies in a region of stable motion, which as can be seen intuitively is only the case for one of two zero-crossings at a time. For this to be the case, particles which are ahead or behind in phase must receive an energy which drives them towards the zero point, with a stronger force if they deviate more in phase. This behaviour would then be equivalent to that of particles in the transverse plane, driven towards the ideal orbit by a position dependent force. A particle located at the stable zero point is known as the *synchronous particle*, and its phase and energy as the *synchronous phase*, ϕ_s, and *synchronous energy*, E_s, as shown in Figure 8.

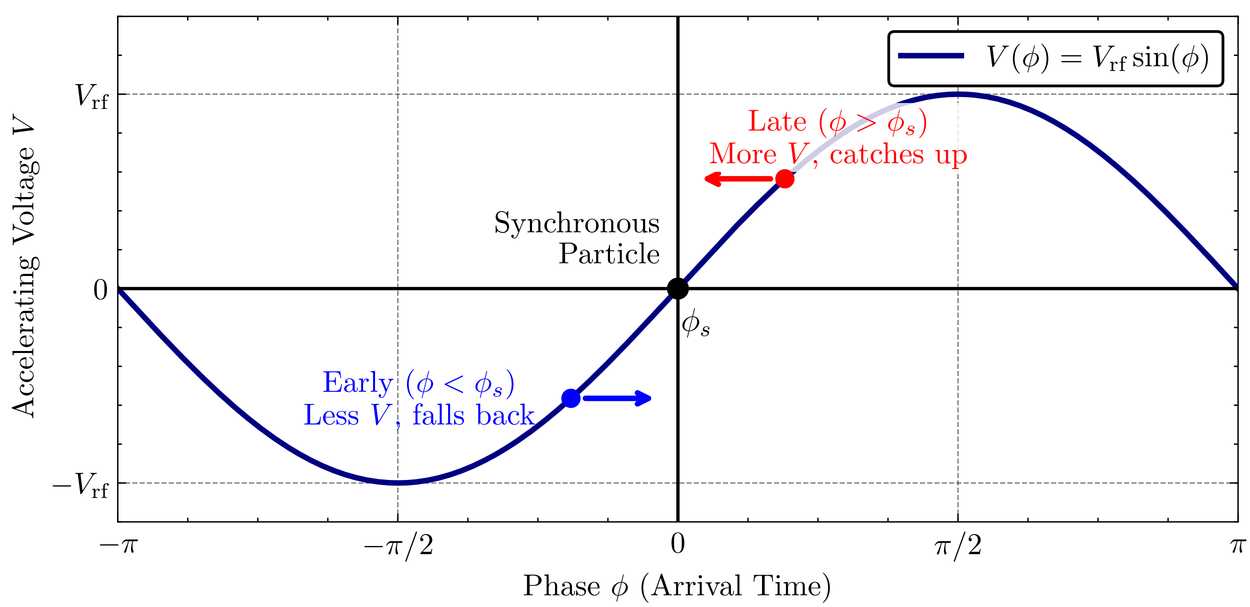

Figure 8: The stable phase of a constant energy storage ring.

To find the equations of motion which govern this motion, two quantities are first introduced, the relative energy and phase to the synchronous particle are defined as $w = E - E_s$, and $\varphi = \phi - \phi_s$. Using these, the energy gain from the RF wave of a particle at a certain phase, as described in Equation 12, can be expressed as [14]

$$\frac{dw}{dt} = \omega_{\rm rev,s}\frac{qV_{\rm rf}}{2\pi}\left(\sin(\phi_s + \varphi) - \sin(\phi_s)\right) \tag{13}$$

From the comparison with the transverse motion, it is clear that finding a differential relation describing the phase shift with time is now key to describing the full phase space motion. That such a relation exists is intuitive, particles with different energies would reasonably arrive back to the RF cavity at different times, and thus with different phases. To quantify this relation, the *slippage factor* is introduced, which connects the relative momentum deviation to the relative deviation in revolution frequency of a particle, as compared to the synchronous particle,

$$\eta = -\frac{d\omega_{\rm rev}/\omega_{\rm rev,s}}{dp/p} \iff \frac{dp}{p}\eta = -\frac{d\omega_{\rm rev}}{\omega_{\rm rev,s}}. \tag{14}$$

Although it might appear intuitive that this factor would always be negative, such that a particle with higher energy would have an increased revolution frequency, this is not always the case in an accelerator. Because the particles are relativistic, the path length excursion of higher energy particles might be more impactful than the marginally increased particle velocity, leading to a reduced revolution frequency. To use the slippage factor with the quantities w and φ above, the relative momentum deviation is converted to a relative energy deviation as

$$\frac{dp}{p} = \frac{1}{\beta^2}\frac{dE}{E} \iff \frac{dp}{p} = \frac{w}{\beta_s^2 E_s}, \tag{15}$$

for the relativistic beta factor of the synchronous particle, β_s. Equation 14 can then be expressed as

$$\frac{d\omega_{\mathrm{rev}}}{\omega_{\mathrm{rev,s}}} = -\eta\frac{w}{\beta_s^2 E_s}. \tag{16}$$

The phase relative to the RF wave can be expressed in terms of the revolution frequency of a given particle, $\omega_{\mathrm{rev}} = \omega_{\mathrm{rev,s}} + d\omega_{\mathrm{rev}}$ with relation to the synchronous particle with revolution frequency $\omega_{\mathrm{rev,s}}$

$$\varphi(t) = \int \omega_{\mathrm{rf}}\, dt - h\int \omega_{\mathrm{rev}}\, dt = \int h\omega_{\mathrm{rev,s}} - h\omega_{\mathrm{rev}}\, dt = -h\int d\omega_{\mathrm{rev}}\, dt, \tag{17}$$

which can be differentiated w.r.t. time to give

$$\frac{d\varphi}{dt} = -hd\omega_{\mathrm{rev}}, \tag{18}$$

and using Equation 16 one arrives at the differential relation connecting the change of φ to w,

$$\frac{d\varphi}{dt} = \frac{h\eta\omega_{\mathrm{rev,s}}}{E_s\beta_s^2}w. \tag{19}$$

The desired pair of differential equations for the longitudinal motion have been found:

$$\begin{cases} \frac{dw}{dt} = \omega_{\mathrm{rev,s}}\frac{qV_{\mathrm{rf}}}{2\pi}\left(\sin(\phi_s+\varphi) - \sin(\phi_s)\right) \\ \frac{d\varphi}{dt} = \frac{h\eta\omega_{\mathrm{rev,s}}}{E_s\beta_s^2}w \end{cases}. \tag{20}$$

Differentiating the phase twice with respect to time, with the approximation of small phase deviations φ and thus $\sin(\phi_s+\varphi) - \sin(\phi_s) \simeq \varphi\cos(\phi_s)$, a second order differential equation emerges:

$$\frac{d}{dt}\frac{d\varphi}{dt} = \frac{d}{dt}\frac{h\eta\omega_{\mathrm{rev,s}}}{E_s\beta_s^2}w = \frac{h\eta\omega_{\mathrm{rev,s}}}{E_s\beta_s^2}\frac{dw}{dt} = \frac{h\eta\omega_{\mathrm{rev,s}}}{E_s\beta_s^2}\frac{qV_{\mathrm{rf}}\omega_{\mathrm{rev,s}}}{2\pi}\left(\sin(\phi_s+\varphi) - \sin(\phi_s)\right) \tag{21}$$

$$\xRightarrow{\text{small angle}} \frac{d^2\varphi}{dt^2} - \frac{qV_{\mathrm{rf}}h\omega_{\mathrm{rev,s}}^2\eta\cos(\phi_s)}{2\pi E_s\beta_s^2}\varphi = 0. \tag{22}$$

A harmonic oscillator has appeared, allowing to once again connect the phase space motion, this time in the coordinates φ and $w = \Delta E$ which correspond to z and p_z, to an ellipse in phase space. The frequency of this harmonic oscillator, i.e. the frequency of the oscillation around the synchronous phase is then

$$\omega_s = \sqrt{-\frac{qV_{\mathrm{rf}}h\omega_{\mathrm{rev,s}}^2\eta\cos(\phi_s)}{2\pi E_s\beta_s^2}}. \tag{23}$$

Since this frequency must be real for a stable oscillation, the following condition must be satisfied:

$$\eta\cos(\phi_s) < 0, \tag{24}$$

where this case at constant energy means that $\cos(\phi_s) = \pm 1$. The sign of this factor then determines the allowed sign of the slippage factor η, which is positive for $\phi_s = \pi$ and negative for $\phi_s = 0$. These two cases are referred to as being above and below transition energy respectively, a concept which is explained further in Section 3.2.3. Thus, it is understood that the stable motion occurs under opposite conditions for the two cases. Below transition, a particle with positive w has a larger revolution frequency than the synchronous particle, and will thus advance in phase relative to it. To counteract this and enable stable motion, particles which are ahead of the synchronous particle in phase must consequently receive less energy than the synchronous particle. Then, the phase gradually decreases until the particle has less energy than the synchronous particle and starts moving back towards it. This gives rise to the rotation in phase space seen in Figure 9. To get this behaviour, the synchronous phase must be on the rising edge of the RF wave, whereas it would be on the falling edge for a lattice above transition.

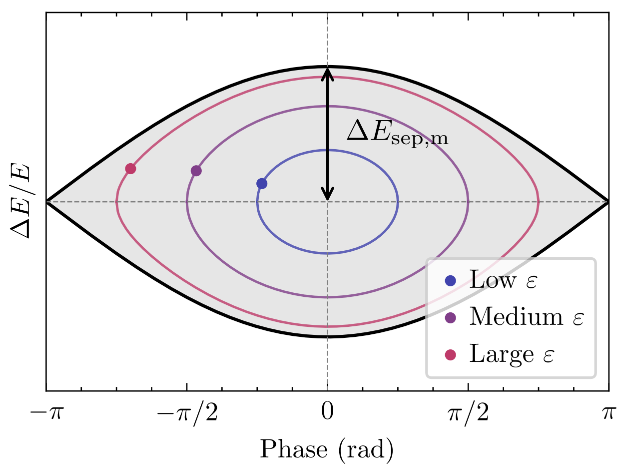

Figure 9: The bucket in longitudinal phase space, enclosed by the separatrix in black. The maximum height of the bucket, which is at the synchronous phase (here with $\phi_s = 0$) is marked, and the phase space trajectories of particles with different longitudinal emittances are shown.

There is a boundary in phase space separating regions of stable from unstable motion known as the *separatrix*. The area enclosed by this boundary is known as the *bucket*, and particles oscillate with different amplitudes within it, along phase space ellipses covering an area proportional to the *longitudinal emittance* of the particle. For an accelerator operating at constant energy (i. e. $\delta E = 0$ at $\phi = \phi_s$), this bucket covers the full RF period in phase, with a varying *bucket height* which is maximal at ϕ_s, as shown in Figure 9. This means that for a harmonic number h, there are h buckets along the circumference of the accelerator where bunches can fit. The size of the bunch in longitudinal phase space in relation to the bucket size is a key consideration in accelerator design. Especially for bunch compression, this becomes important as the rotating bunch distribution must at all times fit within the bucket. The height of the bucket can be expressed as [14]

$$\Delta E_{\mathrm{sep,m}} = \sqrt{\frac{2qV_{\mathrm{rf}}\beta_s^2 E_s}{\pi h|\eta|}}, \tag{25}$$

where the inverse dependence on the slippage factor η means that accelerators with strong compression, i. e. large η, will have a smaller bucket height. For compression, where the longitudinal spread in particle position is converted to an energy spread, this is an issue since the allowed energy spread must be as large as possible, placing constraints on the voltage of the RF system. Furthermore, since the energy is provided by a sinusoidal RF wave, there is an amplitude dependence on the energy received by the particles which was neglected in the small angle approximation above. This means the behaviour of the particle oscillation resembles that of a pendulum, where particles which oscillate with larger amplitudes have a different frequency than those with lower amplitudes. This further motivates the importance of having a sufficiently large bucket so that the particle trajectories are close to the centre and thus have reasonably linear behaviour with amplitude.

3.2.3 Dispersion and Momentum Compaction

The energy and phase of a particle in a storage ring will, as explained above in relation to the slippage factor, affect the orbit which the particle follows in the horizontal dimension. A higher energy corresponds to less bending in the dipoles, and thus a longer orbit. The difference in orbit excursion for particles at different energies (or momenta) in any given point of the ring depends on the strength of the dipole field there, and the proportionality between the excursion and momentum deviation is a quantity known as *dispersion*. It is defined as

$$D_x(s) = \frac{dx(s)/x}{dp(s)/p}, \tag{26}$$

where x and p are the orbit radius and momentum of the ideal orbit, respectively, and can be calculated around the entire ring. Normally the dispersion is always positive, but there are cases where it can take on negative values.

It is also interesting to know the average increase in the beam excursion as a function of the momentum deviation. The beam excursion is expressed in terms of the mean radius $R = \frac{C}{2\pi}$ of the orbit, with C again being the circumference. Beyond this, it is useful to define the *bending radius*, ρ, of the ring, which denotes the total radius of bending in the dipole magnets of the ring, i. e. the radius of the ring if it consisted of only dipoles and no quadrupoles, diagnostics, insertion devices, or similar. From the centripetal force on the particles in the ring, this quantity can be calculated using

$$B\rho = \frac{p}{q}, \tag{27}$$

for momentum p and charge q of the particles in the ring with a dipole field strength B. Integrating the dispersion divided by the bending radius one can calculate the *momentum compaction factor*

$$\alpha_p = \frac{1}{2\pi R}\int_0^C \frac{D_x(s)}{\rho(s)}\,ds = \frac{\langle D_x \rangle_\rho}{R}, \tag{28}$$

noting that the bending radius, $\rho(s)$, is infinite outside the dipole magnets. This relates the momentum and mean radius deviations as

$$\frac{dp}{p}\alpha_p = \frac{dR}{R}. \tag{29}$$

The momentum compaction factor also connects to the concept of slippage factor. If one combines the increased orbit length from the momentum compaction with the change in particle velocity, one finds the slippage factor and how the actual revolution time changes with momentum. This connection is easy to understand at the point where $\eta = 0$, i. e. the transition point. At a certain particle energy, the increased particle velocity exactly matches the increase in orbit length from the momentum compaction factor. In this point, known as the *transition energy* of the lattice, all particles circulate with the same revolution frequency in the linear approximation used here. The relation between slippage factor and momentum compaction can be found by expressing the revolution frequency

$$\omega_{\text{rev}} = \frac{\beta c}{R}, \tag{30}$$

which lets the differential change in revolution frequency from Equation 14 be expressed as

$$\frac{d\omega_{\text{rev}}}{\omega_{\text{rev}}} = \frac{d\beta}{\beta} - \frac{dR}{R}. \tag{31}$$

Using the known relativity relation [14]

$$\frac{dp}{p} = \gamma^2 \frac{d\beta}{\beta} \tag{32}$$

and Equation 29, this can be written as

$$\frac{d\omega_{\text{rev}}}{\omega_{\text{rev}}} = \frac{1}{\gamma^2}\frac{dp}{p} - \alpha_p \frac{dp}{p} \iff -\frac{d\omega_{\text{rev}}/\omega_{\text{rev}}}{dp/p} = \alpha_p - \frac{1}{\gamma^2}, \tag{33}$$

which using Equation 14 gives

$$\eta = \alpha_p - \frac{1}{\gamma^2}. \tag{34}$$

A lattice is operating below transition for $\eta < 0$, and above for $\eta > 0$. Thus, one can find the energy at which the transition occurs, expressed conveniently in terms of the transition gamma, γ_t, by setting $\eta = 0$:

$$0 = \alpha_p - \frac{1}{\gamma_t^2} \iff \alpha_p = \frac{1}{\gamma_t^2}, \tag{35}$$

enabling a useful alternative expression for the slippage factor:

$$\eta = \frac{1}{\gamma_t^2} - \frac{1}{\gamma^2}. \tag{36}$$

It is worth noting that these expressions are all linearised, and to take into account higher order behaviour one can develop the slippage factor and momentum compaction to higher orders, such that the slippage factor at a specific momentum p_0 is given by

$$\eta_{p_0} = \eta_0 + \eta_1 \frac{\Delta p}{p_0} + \eta_2 \left(\frac{\Delta p}{p_0}\right)^2 + \cdots \quad . \tag{37}$$

3.2.4 Linear Accelerator

Before a beam enters a storage ring, it must be accelerated by, typically, a linear accelerator (linac). Many of the concepts described in the sections above do not fully translate to linacs as the particles only traverse it once, removing the periodicity condition from storage rings. However, it is nonetheless true that the particles do not follow the ideal trajectory through the linac, so there is still a need for stable motion and thus the phase space concept is still useful, particularly for the longitudinal motion.

A linac accelerates the beam through a series of RF cavities which provide energy to the incoming particles based on their phase in relation to the standing wave in the cavity[2]. Unlike the stationary case studied for storage rings above, all particles in a linac experience a net energy gain, meaning the synchronous phase must be placed away from the zero-crossing of the RF wave, but still in a region where stable motion is possible. Consequently, the synchronous phase cannot be placed too close to the crest of the RF wave, as seen in Figure 10. It is due to this difference that some conventions differ between linacs and storage rings, such as the use of cosine for the RF wave, rather than sine.

[2] Travelling wave cavities also exist, but are only used for electrons as protons do not become ultrarelativistic until very high energies.

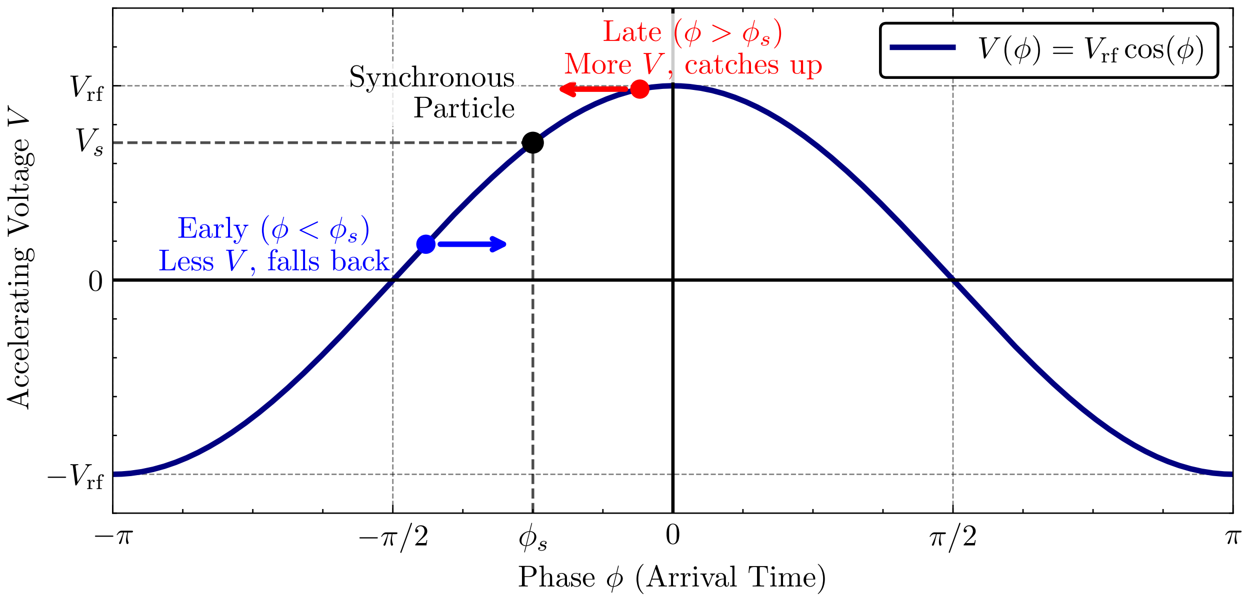

Figure 10: The phase of a particle bunch with relation to the standing RF wave in a cavity. Placing the bunch off-crest gives a stable motion like the one seen in Figure 9.

The acceleration in the cavities is repeated at increasing velocities with gaps of appropriate length between the cavities until the bunch has reached the desired energy. Due to the different requirements for the beam at different energies, various types of cavities are used for the different regimes to account for the change in velocity of the beam. For example, the ESS linac, which shares many concepts with the proposed high power linac for the muon collider, has five different types of cavities within its 400 m length [16]. At the end of the linac, the protons are relativistic and thus the velocity is almost constant, at which point the same type of cavity can be used indefinitely. Before this is the case, cavity types like *drift-tube cavities* must be built to accommodate the increasing velocity of the particles. In these, the bunches are shielded from the standing RF wave when it is decelerating, and gaps between these shielding *drift tubes* are placed so that the field is maximally accelerating when the bunches arrive there, as seen in Figure 11.

To simplify the description of the elements in the linacs, the positions are often expressed in terms of $\beta\lambda$, which is the distance the particles travel in one period of the RF field at the given part of the linac. This distance of course increases significantly from the first part of the drift tube section (DTL) at an energy of a few MeV to the end of the linac at 2 GeV. For this reason, it is also possible to change the frequency of the cavities somewhere along the linac to a multiple of the original frequency, provided the particles all fit within the bucket. This reduces the transverse cavity size considerably, c.f. Equation 11, which is particularly useful for superconducting cavities which must fit within a cryomodule. The maximum amplitude of the field which can be contained in a cavity also varies with the frequency, which can be a further reason to move to a different frequency. This is done for the superconducting linac (SCL) section of the linac at both ESS and in the muon collider linac. The superconducting cavities suffer several orders of magnitudes smaller losses than normal conducting ones, allowing considerably stronger field amplitudes in the cavity. They also have a much narrower frequency acceptance than normal conducting cavities, meaning proper tuning of the cavities is very important [17].

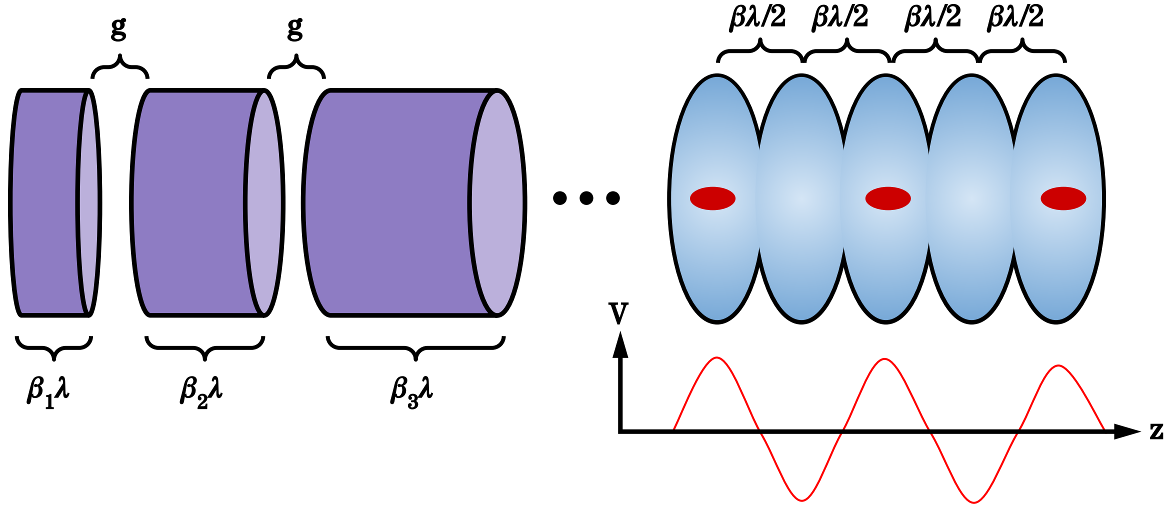

Figure 11: Schematic layout of the initial drift tube (DTL) section of a linac, followed by the final superconducting (SCL) part shown as a π-mode cavity. The tubes of increasing length shield the bunches from the accelerating field, and they are accelerated in the constant length gaps when the field is maximal. In the SCL, the standing wave switches sign every two cavities between accelerating and decelerating, but the spacing is constant as the bunches are relativistic. Note that the field also changes with time.

3.3 Multiparticle Dynamics

So far, the behaviour of individual particles in an accelerator has been treated. However, as a particle accelerator contains bunches of particles with at least 1×10^{10} particles, it is necessary to also develop a statistical intuition for how a group of particles behaves. Here, the concept of emittance is very important, as the phase space description permits a statistical distribution of particles to be handled almost as an individual particle. Thus far, the emittance has described the area of a single particle trajectory in phase space, as seen in Figure 5 (right). All particles in a given location, s, follow an elliptical trajectory with the same shape, since the Courant–Snyder parameters only depend on the magnetic lattice. However the particles have different emittances, ε, and thus different amplitudes of the oscillation. Therefore, an emittance can be defined for a hypothetical particle which traces an ellipse containing within it a certain percentage of the particles in a bunch. Typically, a Gaussian bunch shape is assumed, and it is convenient to define an emittance which contains one standard deviation of the particles in a bunch, the RMS emittance, ϵ_{RMS}. Then, the emittance does not correspond to any given particle, but to the distribution of particles, describing how confined it is in phase space. This is a key measure of the quality of the beam. Achieving a small beam emittance requires well-constructed optics and thoughtful considerations of imperfections and non-linear effects which disturb the beam. A small emittance beam confines the oscillatory trajectories of all particles to small amplitudes, leading to a more focussed and intense beam as is required in many accelerator applications such as light sources and colliders. The beam size at a given location can then be calculated analogously to the maximum position of a single particle, as seen in Figure 7, for any given choice of emittance, ϵ, as

$$\sigma(s) = \sqrt{\epsilon\beta(s)}. \tag{38}$$

Another way of achieving a tightly focussed beam in the spatial dimensions is thus by designing the optics for small beta functions in specific locations, in accordance with Equation 38, which is also done in the aforementioned applications. Notably the interaction points (IPs) of colliders utilise a very extreme version of this. However, this method is limited in its usefulness since a small beta function at one location necessarily implies a large angular divergence. This means that a tightly focussed beam in one location quickly becomes large after a short distance. This effect is clear in the phase space picture of the beam. Consider the full beam in phase space, represented by an ellipse enclosing the particles (for this purpose, one is free to choose if one standard deviation, 99% or some other percentage of the particles are enclosed). Since all particles follow phase space trajectories of the same shape, determined by the Courant–Snyder parameters at the given location, the ellipse represents the trajectory of some particle at a given emittance, which encloses the trajectories of all particles with smaller emittances. As the beam moves through the ring towards the focussing point, the phase space ellipse changes angle with the phase advance, towards a narrow ellipse which is vertically straight ($\alpha = 0$). The process is shown in Figure 12. Because the beam now has a large angular divergence, it will quickly rotate back to a state with large beam size, i. e. large $\beta(s)$. A reduced emittance here means that the beam can be squeezed further for a given allowed maximum divergence and consequently maximum beam size after the IP. The limit here often comes from the size of the beam pipe, since particles in the beam will be lost if the beam is too large in relation to the physical constraint. This process can be performed analogously in the longitudinal phase space, which is what is done for bunch compression. Minimising the longitudinal beam emittance is then also key to achieving a short final bunch length.

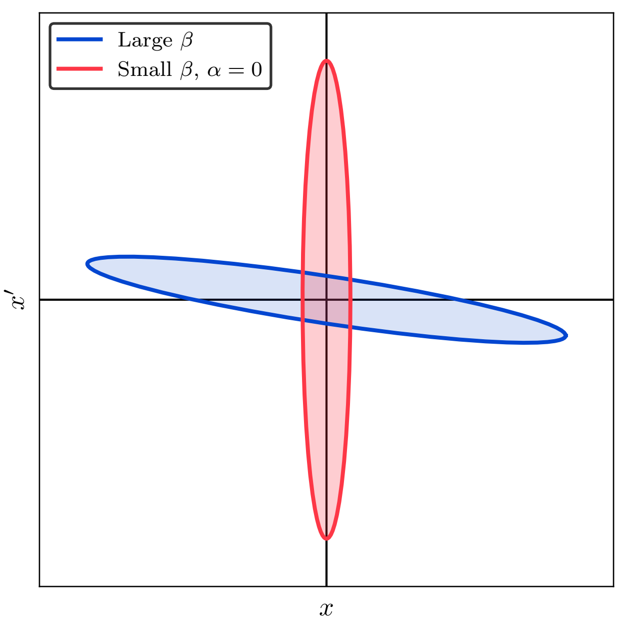

Figure 12: Visualisation of the horizontal phase space ellipse rotating to a compressed state at an interaction point (IP), with $\alpha = 0$ and minimal beam size. Note that the emittance is conserved, even though the horizontal beam size is reduced.

A key feature of the emittance is that it is conserved under the influence of conservative forces, which is a result of Liouville's theorem. In the case of a simple accelerator model under the influence of external fields, as has been discussed so far, this condition holds. For the self-interaction of the bunch due to e. g. space charge forces, non-conservative forces are introduced and Liouville's theorem no longer holds. Technically, the conservation of emittance is only true for the six-dimensional phase space volume described by the emittance in all three spatial dimensions and their corresponding momenta or divergences. However, in a well-designed accelerator there is little coupling between the spatial dimensions, and the emittance can be considered conserved for each dimension independently. This is very important for understanding the example above.

3.4 Collective Effects

Bunch compression can be described using the concepts listed above, but such a description would neglect a key aspect of the compression, which is the self-interaction of the bunch, as well as the interaction with the fields it induces in the surrounding beam pipe. This is what is referred to as *collective effects*, which introduces non-linear and non-conservative forces to the simplified description in previous sections. These additional forces disturb the stable particle motion already mentioned, with consequences such as emittance increase, tune shifts, particle loss, etc. To understand the consequences of these effects, it is necessary to first understand that the orbit stability from the oscillatory behaviour developed above has limitations, namely in the presence of resonances.

3.4.1 Tune Resonances

In any accelerator, there are imperfections in the fields of the machine. This can manifest as higher order fields from the dipoles and quadrupoles, which come naturally from the magnets, but also arise from misplaced magnets in relation to the beam orbit. Furthermore, stray fields may come from beam instrumentation, the beam pipe, other bunches, or even within the bunch itself. If such a field is periodic, the effect on the oscillations of particles in the beam becomes important. Should the periodicity be such that it agrees with the oscillation frequency of a particle, a resonance behaviour appears, driven by this field imperfection.

Depending on the type of field error, the conditions for the resonance varies. To understand this, it is useful to use the phase space view. For a dipole kick to a particle which has an integer tune, i. e. a first order resonance, the kick occurs with the same direction at the same location in phase space every turn. This leads to an increasing amplitude every turn, as shown in Figure 13, and quickly leads to particle loss. For a second order resonance however, the kick is applied at two locations separated by 180° in phase space, meaning the effect of the kick cancels out. For a quadrupole however, the direction and strength of the kick depends on the transverse position of the particle, and thus a second order resonance is also unstable, just as a first order one with a quadrupole error. The same reasoning can be used to deduce which order resonances are excited to unstable behaviour by higher order field errors.

The excitation of tune resonances is one of the main mechanisms for particle loss in accelerators. Specifically, lower order resonances cause larger issues, whereas higher order ones generally can be ignored. That particles end up on resonant frequencies (cross a resonance) is caused by the tune spread and shift from various sources. Tune spread is mainly caused by *chromaticity*, a feature of the lattice which can be corrected using sextupole magnets. Tune shifts have many causes, with the main one in the context of this work being space charge. When the amplitude of the betatron oscillations increases sufficiently, it will be lost as it encounters some physical object in the beam pipe. This would often be intentional collimation of the beam, but if the loss occurs rapidly enough, it could hit the beam pipe itself, potentially damaging it.

To identify which combinations of tunes are permitted, a *tune diagram* is used, where the horizontal and vertical *fractional tunes* of particles are plotted against each other. The fractional tune is the non-integer part of the tune, and is what determines which resonances a particle might encounter. The total tune is seldom interesting, and the non-integer part of the tune is often used interchangeably with the tune, as it is the one which determines the behaviour of the particles. The different

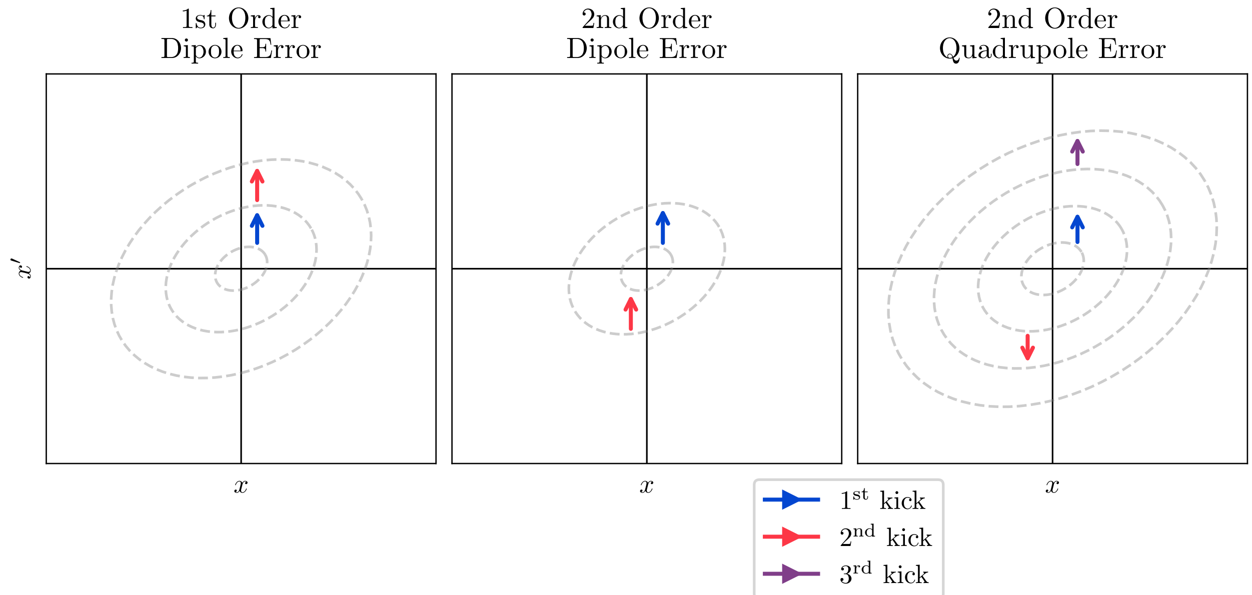

Figure 13: Illustration of the kicks from different field errors in phase space at different order resonances. A 1st order resonance corresponds to an integer tune, with kicks occurring every turn at the same location in phase space. A 2nd order resonance corresponds to half-integer tunes, and thus the kicks occur with a 180° angle in phase space, leading to cancellation for dipole errors, but not quadrupole errors.

order resonances are then represented by lines in this diagram. These lines are solutions to the equation

$$kQ_x + mQ_y = r, \qquad k, m, r \in \mathbb{Z}, \tag{39}$$

where $|k| + |m|$ is the order of the resonance. Thus, not only resonances such as $Q_x = 1/2$ are present (represented by a straight line in the tune diagram), but also diagonal resonances which are combinations of the horizontal and vertical tunes. Resonance lines up to the fourth order are shown in Figure 14.

3.4.2 Space Charge

The most relevant collective effect for this work is the interaction of the particles with the potential generated by the bunch itself, which is known as space charge. The large density of charged particles gives rise to an additional potential which is a function of the position in the bunch. It acts both longitudinally and transversally. Because the force between protons is repulsive, this force effectively acts as a defocussing force in transverse, competing with the focussing of the quadrupoles. This leads to a tune shift, as a change to $k(s)$ in Hill's equation (Equation 3) leads to a change in the phase advance according to Equation 6. The impact of space charge can be described in terms of a contribution $k_{\text{s.c.}}$ as

$$\frac{d^2x(s)}{ds^2} + [k(s) - k_{\text{s.c.}}(s)]x(s) = 0. \tag{40}$$

Because the space charge potential is position dependent, each particle experiences a different force, and thus a unique $k_{\text{s.c.}}$. This leads to unique focussing for each particle, and consequently different beta functions and phase advances, according to Equation 6. This causes the tune spread shown in Figure 14, where the tunes of the particles are shifted across a region known as the space charge necktie. Here, the horizontal and vertical tunes of individual particles are plotted, and the resonance lines at different orders are shown. If a particle falls on one of these lines, it will begin the resonating

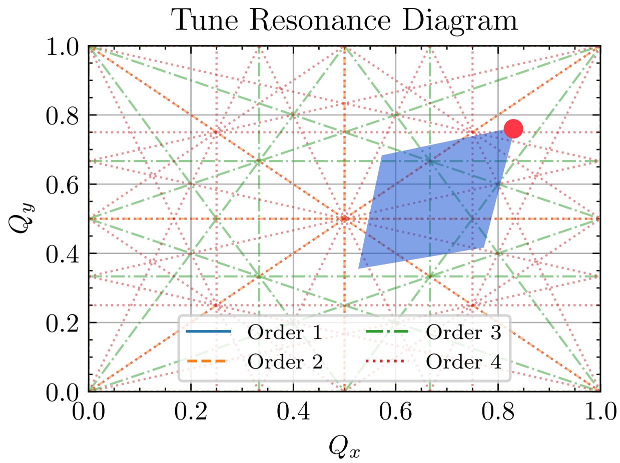

Figure 14: Tune diagram showing the horizontal and vertical tunes on the respective axes, with resonance lines up to the fourth order shown. The characteristic space charge tune spread is shown as a necktie emanating from the working point of the machine.

behaviour described in Section 3.4.1 which leads to increasing oscillation amplitude and, for lower order resonances, ultimately to loss of the particle. The magnitude of this tune shift depends on the density of the particles, the energy and many other terms. For the bunch compression studied here, the effect is expected to be very significant and one of the main mechanisms for particle loss.

It has been qualitatively shown that the space charge force affects each particle in a bunch uniquely, with a dependence on the transverse and longitudinal position of a particle. The force on a test particle at a certain location in the bunch can be calculated using the space charge potential from the whole bunch. This potential can be calculated using Maxwell's equations, either analytically or numerically for complicated distributions. Consider for simplicity an ultra-relativistic radially symmetric bunch with radius a in a perfectly conducting circular beam pipe with radius b (this boundary does not affect the transverse forces, but as shall be seen it does impact the longitudinal force). The bunch is characterised by its line density, $\lambda(s)$, which gives the longitudinal distribution of particles in the bunch. The direct transverse force can be calculated by applying Gauss' and Ampere's laws, see Figure 15, as [3]

$$F_r(r, s) = \frac{q\lambda(s)r}{2\pi\varepsilon_0\gamma^2 a^2}, \tag{41}$$

for a radial transverse position r, for a beam with radius a in vacuum with electric permittivity ε_0. Since this force has a dependence on the transverse position, it can simply be divided into the two transverse planes. Then, it can be linearised to give an expression for $k_{\text{s.c.}}$ [3] in Equation 40, with

$$k_{\text{s.c.}} = \frac{q\lambda(s)}{2\pi\varepsilon_0\beta^2\gamma^3 E_0 a^2}, \tag{42}$$

for the particle rest energy E_0 and relativistic γ and β. Notable here is the $1/\gamma^3$ dependence, showing that the space charge force decreases dramatically with increasing particle energy. From this example, the position dependent tune shift can be seen, explaining the tune spread shown in Figure 14.

In the longitudinal direction, the bunch must be non-uniform to ensure that the forces do not cancel. Therefore, a Gaussian distribution of the bunch in the longitudinal direction is used in this example, $\lambda(s) \propto e^{-s^2/\sigma_s^2}$. Then, the force on a test particle in the bunch is found to be dependent on the

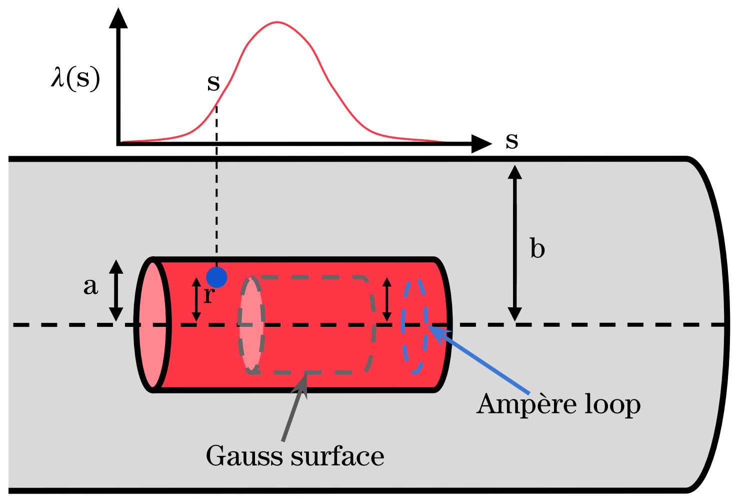

Figure 15: A radially symmetric bunched beam with line density $\lambda(s)$, contained in a perfectly conducting beam pipe of radius b.

derivative of the line density. In a perfectly conducting cylinder, the longitudinal force is

$$F_{\text{long.}}(r,s) = -\frac{q}{4\pi\varepsilon_0\gamma^2}\left(1 - \frac{r^2}{a^2} + 2\ln\left(\frac{b}{a}\right)\right)\frac{\partial\lambda(s)}{\partial s}. \tag{43}$$

The picture is however complicated by the boundary conditions which appear due to the beam pipe and other conducting devices in the accelerator. This makes a complete treatment of space charge forces fall outside the scope of this thesis, the possibility of solving the space charge potential for a given particle distribution with certain boundary conditions suffices for the understanding.

Given a longitudinal space charge potential, $V_{\text{s.c.}}(s)$, as it follows the bunch through the accelerator, it is possible to define a corresponding wakefield [3], $w_{||}(s)$. The Fourier transformation of this is an impedance, $Z_{||}(\omega)$, which has a frequency dependence. It turns out [3] that this space charge impedance is capacitive, and linearly proportional to the frequency. There are many other sources of impedance, which all can drive instabilities in the ring. However, there are also damping mechanisms, such as Landau damping [4, p. 218], which limit the impact of instabilities. One key aspect of dealing with impedances is that they contain frequency components which excite the bunch. The instabilities are often categorised by the frequency range of these driving components, for example microwave instabilities. This picture is important for understanding the behaviour of instabilities, as will be seen in the following section on the negative mass instability [4].

3.4.3 Negative Mass Instability

The longitudinal force from space charge acts repulsively between particles. Thus, a region of increased particle density gives momentum kicks to adjacent particles directed away from this region. Below transition, this leads to a debunching behaviour which counteracts bunch compression. If the density of particles is too great, this will prevent successful bunch compression [18]. Above transition however, the repulsive force tends to act attractively, since increased momentum leads to decreased revolution frequency according to Equation 14, and vice versa. This is what gives rise to the negative mass instability.

Intuitively, the negative mass instability can be understood by a simple example for an accelerator operating above transition. Consider two test particles interacting with the space charge potential of a bunch, one ahead of the bunch and one behind. They will receive momentum kicks in opposite directions from the space charge potential, being positive and negative in sign, respectively. This changes their revolution frequency according to Equation 14. The revolution frequency of the particle ahead of the bunch centre will decrease, meaning the particle moves towards the space charge potential instead of being repelled by it. For the particle behind the bunch centre, the revolution frequency instead increases, having the same effect. These test particles thus act as though they had negative mass, giving the instability its name. As a consequence, particles gather in regions of high density, causing unstable microbunching.

The physical effects of the instability can be understood by considering the impedance view of the space charge. To determine whether an impedance gives rise to unstable behaviour, the collective driving frequency, Ω, can be calculated. This frequency drives the collective motion, and depending on whether it is real or imaginary, whether this motion is simply oscillatory or leads to exponential growth, i. e. an instability. The behaviour is determined by the Vlasov equation [19], a full treatment of which falls well outside the scope of this thesis, which gives a condition for stability, as well as the growth rate of the exponential growth of the instability if this condition is not satisfied. This growth rate is [4, 20]

$$G(n,t) = \omega_0 \sqrt{\frac{\eta e I_p n Z_{||}}{2\pi\beta^2 E}}, \tag{44}$$

for a revolution frequency, ω_0, peak current, I_p, longitudinal space charge impedance, $Z_{||}$ (which is capacitive, i.e. $\mathrm{Im}(Z_{||})$), as well as particle energy, E, and corresponding relativistic β. The square root increase with η and I_p indicate that this rate will be very large in the proposed scheme, which utilises strong compression (large η) and also has a very high peak current due to the high intensity (especially towards the end of the compression).

4 Bunch Compression in the Proton Driver

Typically, bunch compression in proton accelerators is performed by making a long bunch with small energy spread perform synchrotron oscillations in a bucket from an RF system, converting the small energy spread into a small bunch length when a 90° rotation has been performed. This is done in a compressor ring, which has some particular features which maximise the compression strength for a given energy spread, that is they serve to maximise the slippage factor. In this work, two different lattices were used, one above and one below transition, shown in Figure 16.

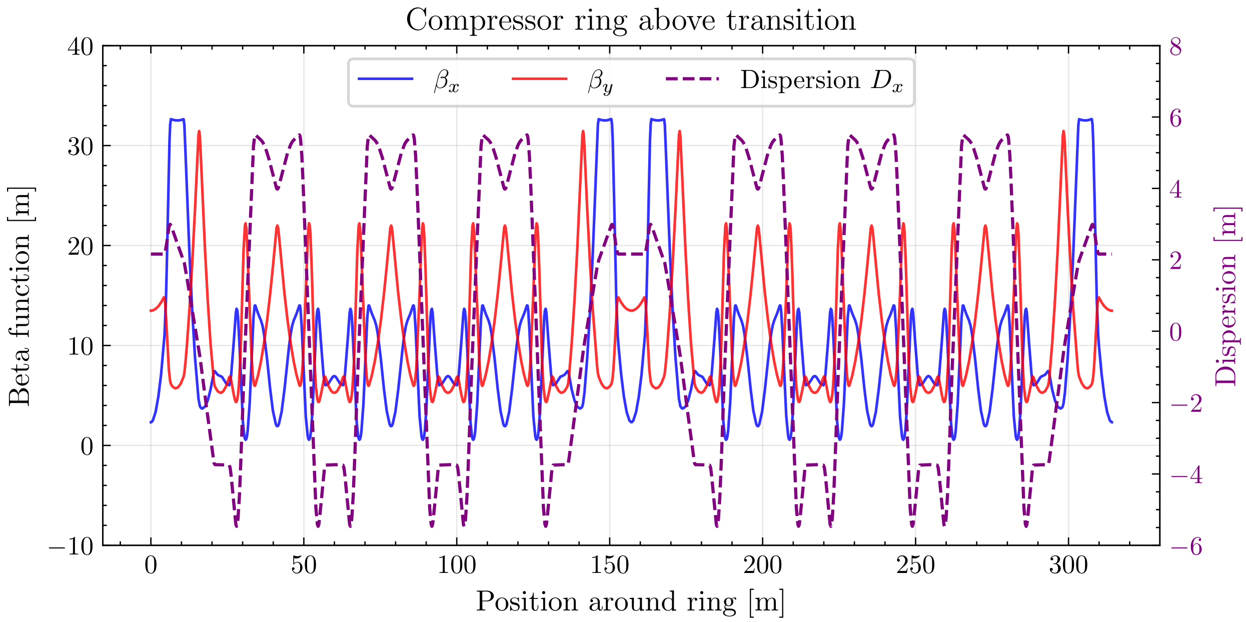

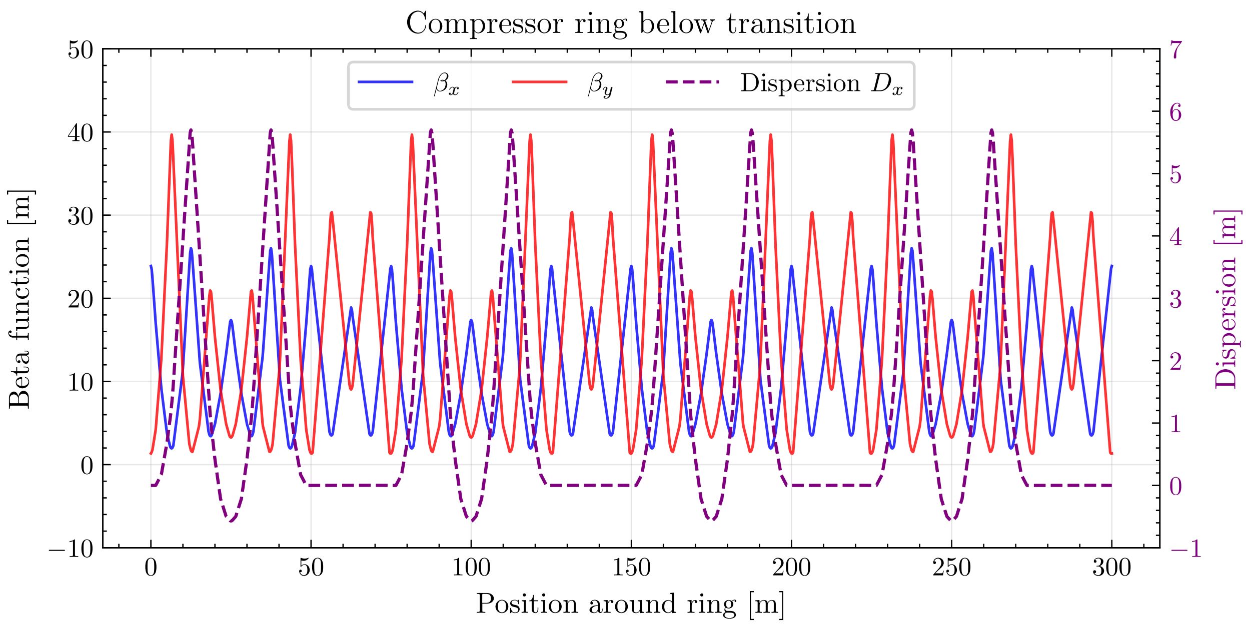

Figure 16: The lattice functions for the two compressor rings used in this work. The upper one is the main option, operating above transition, with the lower one being a four-fold symmetric lattice operating below transition. This one is inspired by the work of Shinji Machida [21]. Note that the one above transition utilises negative bends to achieve a negative dispersion in order to maximise the slippage factor.

This method is used in the baseline proton driver design for the muon collider, but has several drawbacks. Consider the bucket which has to fit the entire bunch in order to ensure stable motion during the compression. At the start, the bucket must cover the entire length of the bunch before compression, preferably with considerable margin to avoid the non-linear behaviour of the synchrotron oscillations near the edge of the bucket. This can be done by having a ring of sufficient circumference with a harmonic number of $h = 1$. However, as the bucket must also contain the bunch during the compression, particularly at the point of maximum compression and maximum energy spread, the height of the bucket must be very large. This is inversely proportional to $\sqrt{h\eta}$, as seen in Equation 25, which motivates the choice of minimising $h = 1$, and not just increasing the circumference of the ring. However, maximising η is essential to achieve fast compression, as seen in Equation 23, which also shows that the synchrotron frequency, i. e. compression rate, increases proportional to $\sqrt{h\eta}$. The RF voltage, V_{rf} is involved in both these relations, but cannot be made arbitrarily large.

4.1 Baseline Compression

Consider the bunch compression in the baseline proton driver, taking the accumulated 120 ns bunch with a small energy spread to a 2 ns bunch length with large energy spread. The rotation occurs within the bucket created by the RF system of the ring, and the bunch must at all times fit within the bucket to avoid unstable motion where particles are lost. Note that the bunch lengths and energy spreads are given as RMS values. To find the required RF voltage for the system, the other known parameters are used. The nominal kinetic energy of $E_s = 5.0\,\mathrm{GeV}$ corresponds to $\beta_s = 0.99$, which with a slippage factor $\eta = 0.17$ and harmonic number $h = 1$ (it would be advantageous to use one bunch, i. e. one bucket) allows to calculate the required RF voltage for a given bucket height. To find this, the energy spread after the compression must first be calculated. Here, a lower bound on the energy spread after the compression can be found using Liouville's theorem, i. e. neglecting space charge forces and other collective effects. Because the longitudinal emittance, proportional to the product of the energy spread, $\Delta E/E$, and the bunch length, τ, is conserved, the energy spread after the rotation can be calculated as

$$(\Delta E/E)_{\mathrm{final}} = \frac{(\Delta E/E)_{\mathrm{initial}} \cdot \tau_{\mathrm{initial}}}{\tau_{\mathrm{final}}}. \tag{45}$$

The linac can deliver an initial energy spread of around $(\Delta E/E)_{\mathrm{initial}} = 3.0 \times 10^{-4}$, which with the initial bunch length $\tau_{\mathrm{initial}} = 120\,\mathrm{ns}$ and final ditto $\tau_{\mathrm{final}} = 2\,\mathrm{ns}$, all RMS quantities, gives a final RMS energy spread

$$(\Delta E/E)_{\mathrm{final}} = \frac{3.0 \times 10^{-4} \cdot 120\,\mathrm{ns}}{2\,\mathrm{ns}} = 1.8 \times 10^{-2}. \tag{46}$$

As the bucket must contain the entire bunch, it is required that the bucket size encloses at least three standard deviations of the bunch distribution, giving a bucket height of $\Delta E_{\mathrm{sep,m}} = 5.4 \times 10^{-2} \cdot E_s = 0.27\,\mathrm{GeV}$. The required RF voltage can now be calculated using Equation 25 with the charge $q = e$, the elementary charge, as

$$V_{\mathrm{rf}} = \frac{\pi h|\eta|\Delta E^2_{\mathrm{sep,m}}}{2e\beta_s^2 E_s} = \frac{\pi \cdot 0.17 \cdot (0.27\,\mathrm{GeV})^2}{2e \cdot 0.99^2 \cdot 5.0\,\mathrm{GeV}} = 3.9\,\mathrm{MV}. \tag{47}$$

This might not seem a lot considering the superconducting cavities used in the ESS are able to deliver more than 10 MeV per cavity. However, as the RF system here must be tuned to the frequency of the ring, on the order of 1 MHz, designing a cavity with the same field strength is not possible. Moving to a higher harmonic to allow higher fields would also change the harmonic number, which as seen in Equation 25 reduces the bucket height, counteracting such a change. Additionally, if h becomes too large, the bunch from the accumulator would not fit inside the bucket. Furthermore, the transverse size of the cavity has to increase as a result of the lower frequency, roughly following the relation in Equation 11 for the pillbox cavity. Beyond being impractical and expensive, a cavity the size required for this introduces a very large impedance.

Given that bunch compression is additionally complicated by the introduction of collective effects and non-linearities in the lattice, it is clear that the above parameters are a minimum requirement. In summary, the above mentioned issues with the cavity mean that the currently suggested baseline scheme will be very challenging to implement with conventional technologies.

4.2 Novel Compression

As the baseline design for the proton driver [22] is seen to require RF cavities with currently unattainable parameter requirements, there is reason to investigate alternative methods to achieve the required bunch compression. Given the compressor lattice, which is already designed with a large slippage factor, it is possible to perform a phase space rotation on any configuration of protons having an appropriate energy distribution. In the baseline case, this energy distribution is induced by the RF cavities in the compressor ring, but this does not necessarily have to be the case. If instead the protons injected into the compressor ring arrive with a *chirp* (an appropriate energy distribution depending on the longitudinal coordinate of the particle) it is in principle possible to achieve compression by simply allowing the protons to circulate in the lattice.

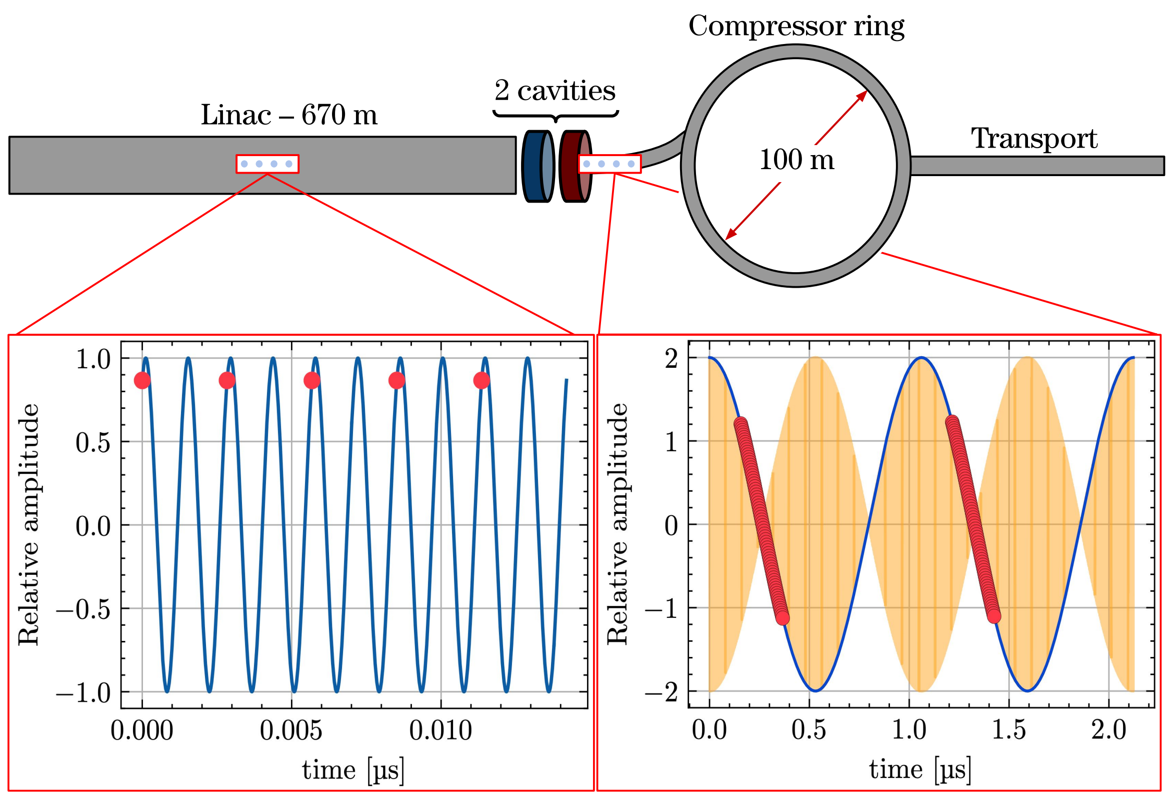

Figure 17: The spacing of bunches in relation to the individual RF fields and the superimposed envelope, showing the near-linear correlation in energy gain for a chopped bunch train from the linac.

This chirp can be created by operating two RF cavities at frequencies slightly offset from the one used by the rest of the linac. Protons arriving to these cavities will experience a superposition of the two fields, which gives a total energy gain or loss depending on their arrival time. If the correct frequency shift is chosen from the nominal RF frequency, a beating envelope can be created where the superposition of the fields creates a chirp over a sequence of many proton bunches coming from the linac, as shown schematically in Figure 17. To create this sequence, the pulse from the linac must be chopped to create bunch trains which fall on the correct region of this sinusoidal envelope. These can then be compressed in the compressor ring.

To reach the very high intensity required for the muon collider proton driver, assuming that the source can deliver an average current of 80 mA, around 3.5×10^5 proton bunches are needed. The linac cavity frequencies must be chosen to match the revolution frequency of the compressor ring, which is close to 1 µs, meaning the beating envelope should have a frequency on the order of 1 MHz. This gives a rough approximation of the injection time required to accumulate the required intensity, which is a key quantity when comparing to the baseline design. The bunch spacing from the linac is determined by the RF frequency at the start of the linac, which is $f_{\text{source}} = 352.21\,\text{MHz}$, giving a spacing of $t_{\text{bunch}} = \frac{1}{f_{\text{source}}} = 2.84\,\text{ns}$. This means that the length of a linac pulse which is not chopped would be 3.5×10^5 bunches $\cdot\, 2.84\,\text{ns/bunch} = 0.99\,\text{ms}$ to reach the required intensity. Considering a large fraction of the pulse must be chopped to fit the bunch trains on the near-linear region of the sinusoidal envelope, it is clear that the injection will take several milliseconds, corresponding to several thousand turns in the compressor ring. The baseline design instead compresses the accumulated proton bunch from the accumulator ring in around 50 turns. A longer compression time means the bunch is more exposed to collective effects which negatively affect the quality of the bunch in different ways. These effects cause various instabilities which lead to emittance increase, particle loss, and longitudinal debunching [3, 4, 19].

Beyond this, the long injection time implies that the first bunch train to be injected must circulate in the compressor ring for more than 1000 turns before the final bunch train is injected. Consequently, it must reach its maximum compression after circulating for at least 1000 turns more than the final bunch train. To achieve this, the chirp must be adapted depending on the time of injection. For a given maximum allowed chirp, $A_{\max}$, the chirp for a bunch train injected a time t after the first bunch train will be given by

$$A = A_{\max} \frac{T_{\text{compression}}}{T_{\text{tot}} - t}, \tag{48}$$

with

$$T_{\text{tot}} = T_{\text{compression}} + T_{\text{injection}} \tag{49}$$

where $T_{\text{compression}}$ is the time it takes to compress a bunch train at the maximum chirp, $A_{\max}$, and $T_{\text{injection}}$ is the total time required to inject the required number of bunch trains. Note that time here is often counted in terms of the number of turns, as seen in the chirp creation code in Appendix A.

4.2.1 Refined Design

To arrive at an optimal design for this scheme, two main factors were considered. As collective effects are expected to limit the performance of the proton driver [22], minimising the proton density in accordance with Equations 41–44 is a key consideration, as is minimising the injection time. Thus, the frequency shift should be chosen such that the injected bunch train covers as much of the ring as possible, minimising the line density, $\lambda(s)$, and peak current, I_p. This naturally has the drawback that the initial bunch length is longer, requiring more compression of the bunch and thus enforcing smaller energy spread according to Liouville's theorem. However, calculating the optimal frequency shift shows that a shift of $\Delta f = 314\,\text{kHz}$, i. e. one third of the revolution frequency, gives maximum coverage of the ring. With this frequency shift, one half-period of the beating envelope (the region where bunches can be placed to create a compressing chirp) is

$$T_{\text{rf}}/2 = \frac{1}{2 \cdot 314\,\text{kHz}} = 1.59\,\mu\text{s} \implies L_{\text{bunch train}} = cT_{\text{rf}}/2 = 477\,\text{m} \tag{50}$$

which is longer than the ring circumference of 314 m. However, as only a portion of the half-period is used for the compression to avoid the non-linearities of the sinusoidal shape, this gives the maximum spread of the initial bunch. This is quite convenient as it allows to simultaneously remedy the other main issue with the proposed scheme, the long injection time and consequently slow compression. By dividing the chopped pulse from the linac to three separate lines with off-frequency cavities, it is possible to utilise a larger portion of the pulse from the linac, shortening the total injection time. Using a fast kicker magnet, the bunch trains, separated by the rise time of the magnet, are sent to pairs of off-frequency cavities which are phased to match the corresponding bunch train, as seen in Figure 18. If the frequency shift is set to one third the revolution frequency, the length of the bunch train which fits in the linear region of the beating envelope becomes 0.76 µs, which fits in the ring, while maintaining a near-linear chirp as shown in Figure 18. This is done with a kicker magnet rise time of 300 ns, which is fast in this context. However, even 100 ns rise times are in principle attainable, depending on the exact requirements, as seen in [23] where the injection kicker magnets at the Relativistic Heavy Ion Collider are found to operate faster than the design value of 95 ns. Shortening the rise time would enable more bunches to be injected per turn, shortening the injection time even further, though the gains to be had here are relatively minimal.

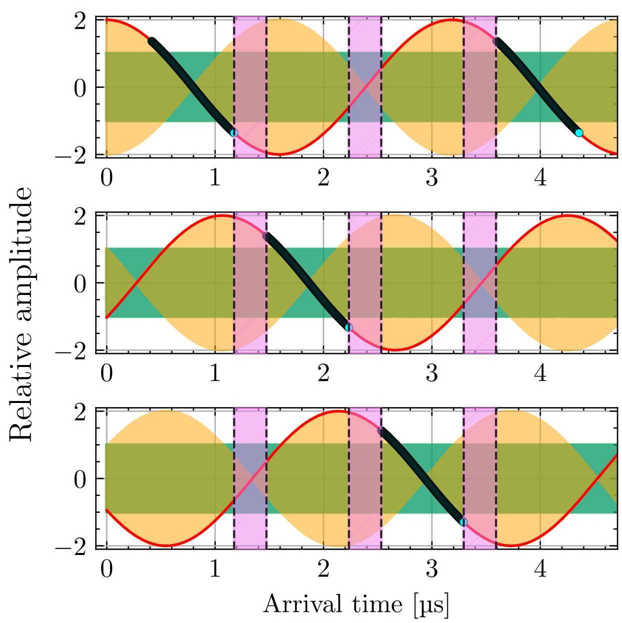

Figure 18: Fields of the three pairs of off-frequency RF cavities, with phases adapted to accept bunch trains separated by a kicker magnet rise time of 300 ns, shown as the shaded violet region between bunch trains.

4.3 Chirp Creation Measurement

In order to show that the creation of a chirp with off-frequency cavities is indeed possible, an experiment at an existing linac could be performed. However, as superconducting cavities operate within a very narrow range of frequencies, it is difficult to detune an existing cavity to a different frequency than the working point. This is why an implementation of the proposed compression scheme would have to use cavities designed to operate at shifted frequencies with respect to the rest of the linac. At an existing linac, a detuning requires two separate steps. Firstly, the cavity must be adjusted physically to resonate at a different frequency. As corrections of the resonant mode of the cavity are necessary to adjust to the real world conditions in which the cavity operates, different types of tuners are available to use during operation. Thus, this issue can be handled with relative ease.

Secondly, the RF wave fed to the cavity must be adjusted in frequency so that it feeds the cavity at its new resonant frequency. If not, the RF pulse will drive the original frequency in the cavity, but a lot of power will be lost as the cavity is not tuned to the RF pulse.

If these two adjustments are done, it would be possible to measure the chirp creation. By detuning two cavities, and then measuring the arrival time at different points downstream of the cavities, the time-of-flight (ToF) between these points can be used to calculate the energy of the particles. The difference in energy between particles at different locations on the beating envelope (specifically the near-linear region used for the chirp) can then be inferred. One can use BPMs for this purpose, as they measure the arrival time of bunches within fractions of the RF wave period (typically expressed as a phase based on this period). Depending on the amplitude of the chirp, particle energy, and the ToF in a given linac, this can be shown to suffice for measuring the desired arrival time difference. Using

$$\gamma = \frac{E + \Delta E}{E_0} \tag{51}$$

$$\beta = \sqrt{1 - \frac{1}{\gamma^2}} \tag{52}$$

$$v = \beta c \tag{53}$$

$$\implies t_{\mathrm{ToF}} = \frac{d_{\mathrm{ToF}}}{v}, \tag{54}$$

the time of flight, t_{ToF}, can be calculated for bunches travelling a distance d_{ToF}. The energy gain is ΔE and given by the cavity amplitude and bunch placement on the sinusoidal beating envelope, see Figure 19 (left).
This is done here for the ESS linac, where BPMs located immediately after the final cavities in the high-beta linac (HBL) section and at the end of the high energy beam transport (HEBT) are used. This means the distance between the two is 130 m. The particles are accelerated to an energy of $E_{\mathrm{kin}} = 870\,\mathrm{MeV}$, and the detuned cavities can provide energy up to the nominal 16 MeV [16]. However, it is reasonable to assume that this may not be achieved at detuned conditions, and thus lower amplitudes are also considered. In order to utilise the full amplitude, the detuning frequency and length of the RF pulse must align so that one full linear chirp region can be created. Although the ESS linac is not currently delivering the 3 ms pulses that it will upon completion, the RF system is capable of delivering pulses of this length. Then, individual probe bunches of 5 µs length can be delayed with respect to the long pulse RF wave in order to simulate a bunch train over the linear region of the envelope. It is crucial that the measurement is performed on the centre of the linear region in order to maximise the chirp. At a realistically achievable detuning of 100 Hz, shifting the bunch over a region of 500 µs would give an energy difference of 10 MeV at full cavity amplitude, corresponding to a difference in the time-of-flight of 1.1 ns, which is corresponds to a phase of 140° in the BPMs. This should be clearly measurable. The measured phase difference from the changing time-of-flight is shown for different cavity amplitudes in Figure 19 (right).

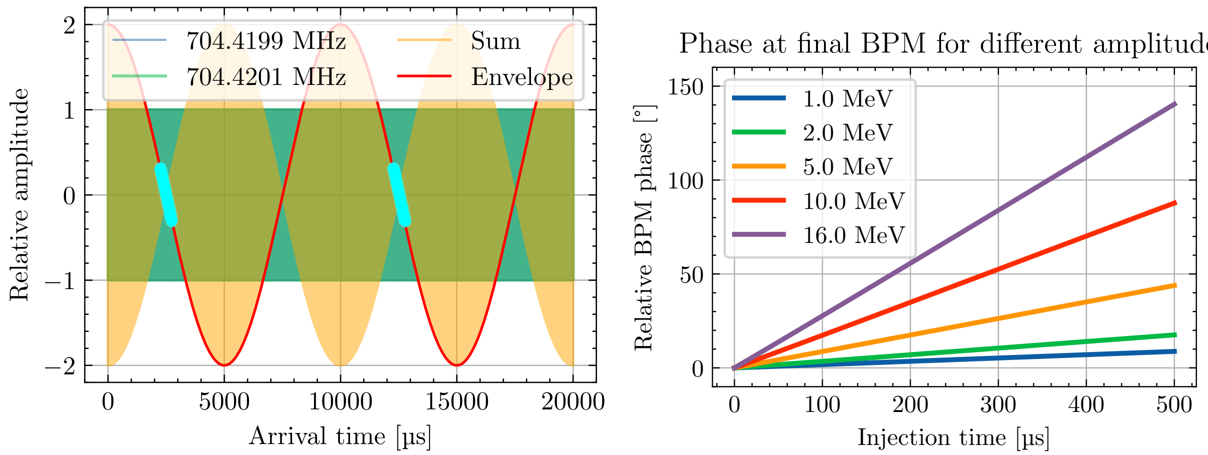

Figure 19: The placements of bunches for a simulated bunch train of length 500 µs on the sinusoidal beating envelope of two cavities detuned by 100 Hz (left), and the corresponding arrival time difference (expressed in units of degrees) between two BPMs in the ESS linac depending on the cavity amplitude in the detuned cavities (right). The injection time is given relative to the first bunch of the simulated bunch train.

5 Method

5.1 Analytical Investigations

In order to gauge the feasibility of the method, analytical investigations of the compression were initiated. This was done partially by estimations of bunch length from Liouville's theorem, as in Section 5.1.2, and partially by implementing a toy model which treats the chirp creation and compression with simple analytical expressions based on the on-axis RF waveform and slippage factor of the lattice.

5.1.1 Toy Model

As only the relative motion of the particles in a series of subsequent bunch trains are interesting for the compression, the compressor ring can in principle be modelled using only the slippage factor, η, and Equation 14, optionally including higher order terms in η, taken from the current lattice for the compressor ring. The chirp is calculated by finding the the envelope of two off-frequency sinusoidal RF waves, which are phased such that the bunches are placed on-crest to maximise the chirp for a given cavity amplitude, as shown in Figure 20. Bunches are then created with different chirps according to Equation 48 every turn at the same longitudinal location, and every bunch is then tracked turn by turn with a change in arrival time calculated using

$$\frac{\Delta T}{T_{\mathrm{rev}}} = -\frac{\Delta\omega_{\mathrm{rev}}}{\omega_{\mathrm{rev,s}}}, \tag{55}$$

in combination with Equation 14 as

$$\Delta T = -\frac{\Delta\omega_{\mathrm{rev}}}{\omega_{\mathrm{rev,s}}} T_{\mathrm{rev}} = \eta \frac{\Delta p}{p} \frac{C}{\beta c}, \tag{56}$$

for a circumference, C, with each bunch treated as a single macroparticle with a relative momentum deviation calculated from the sinusoidal RF envelope.

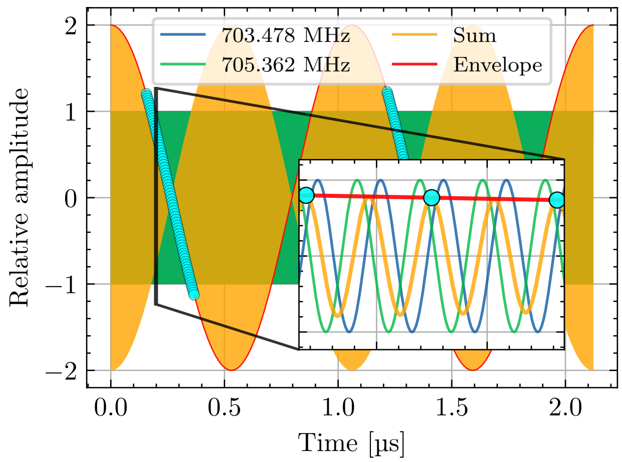

Figure 20: The bunch spacing and location of the individual bunches on the superposition of the RF waves from the two off-frequency cavities, here shown at a frequency shift of $\Delta f = 0.94\,\mathrm{MHz}$.

By injecting the number of bunch trains required to reach the desired intensity with a realistic cavity amplitude of 30 MeV for the chirp creation (corresponding to two or three pairs of off-frequency cavities, each providing a maximum of 16 MeV based on the ESS superconducting RF cavities [16]), the full compression could be studied. The results, shown in Figure 21 demonstrate simultaneous compression of all bunch trains, with a final bunch length of 5.61 ns. The same simulation was also performed at a frequency shift of $\Delta f = 0.94\,\mathrm{MHz}$, giving a final bunch length of 1.91 ns, demonstrating that it is possible to reach the desired compression with the scheme, at the cost of stronger collective effects. The parameters and results from this toy model are found in Table 2 and were used for the following particle-in-cell simulations with PyORBIT.

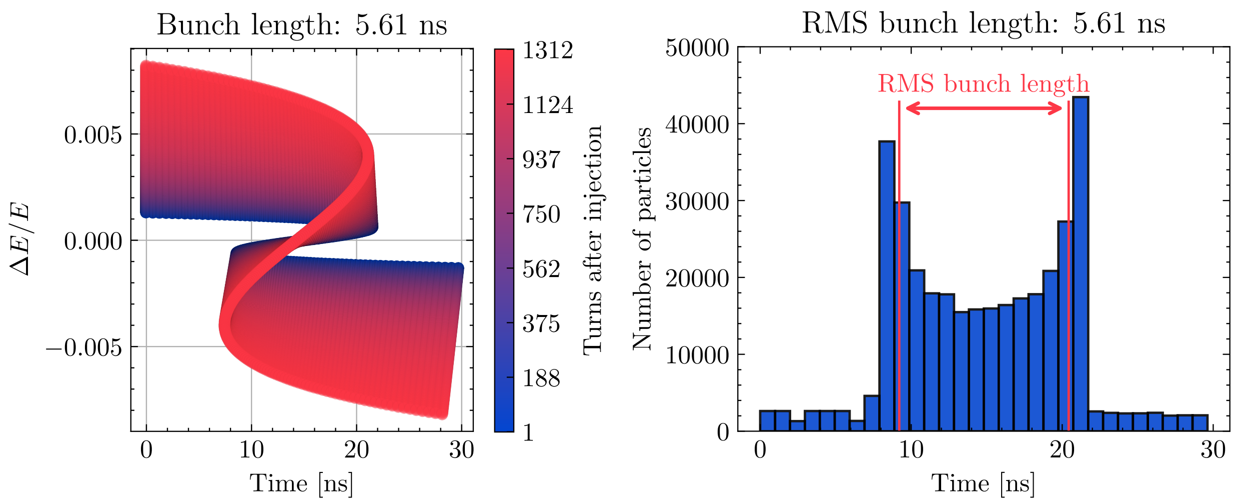

Figure 21: The results of the toy model simulation of the bunch compression scheme. Simultaneous compression of 1312 injected bunch trains with 267 bunches each to a final RMS bunch length of 5.61 ns after injection over 1312 turns and compression over 250 turns is shown to the left. Each colour represents one bunch train tracked for a certain number of turns, with a majority of the bunch trains having a low chirp amplitude. The histogram to the right shows the longitudinal distribution of these bunches after compression, with the RMS bunch length indicated. The horn-like shape is characteristic for compression with a sinusoidal chirp.

Table 2: The parameters used for the toy model, and the resulting bunch length.

Parameter	Value
Amount of pulse kept	72 %
Pulse length	1.393 ms
# turns injection	1312
# turns compression	250
Amplitude	30 MeV
Final bunch length	5.63 ns
Average power	2.01 MW

5.1.2 Minimum Achievable Bunch Length

Before investigating the proposed scheme with more detailed methods, it is interesting to know what lower bound of bunch length is permissible by Liouville's theorem. Using the same method for bunch length calculations as in Equation 45, the minimum achievable bunch length can be calculated for this novel compression method. In the case studied here, the initial RMS bunch length is $\tau_{\text{initial}} = 218\,\text{ns}$, with an RMS energy spread of $\Delta E/E = 3.0 \times 10^{-4}$, which for a nominal energy of $E_{\text{kin}} = 5\,\text{GeV}$ corresponds to $\Delta E_{\text{initial}} = 1.50\,\text{MeV}$. This gives an RMS emittance of

$$\epsilon_{\text{RMS}} = \Delta E_{\text{initial}} \cdot \tau_{\text{initial}} = 0.33\,\text{eV s}. \tag{57}$$

To find the final energy spread, which comes from the applied chirp, Equation 48 can be used. The RMS energy spread of that distribution, with a maximum t corresponding to the 1312 turns of injection required to reach the maximum intensity, is calculated. The RMS amplitude for a maximum amplitude of $A_{\max} = 30\,\text{MeV}$ from Equation 48 is $A_{\text{av}} = 16.8\,\text{MeV}$. To find the final bunch length, Equation 45 is used and the emittance is divided with the final RMS energy spread,

$$\tau_{\text{final}} = \frac{\epsilon_{\text{RMS}}}{\Delta E_{\text{final}}} = \frac{0.33\,\text{eV s}}{16.8 \times 10^6\,\text{eV}} = 19.3\,\text{ns}. \tag{58}$$

This is an order of magnitude more than the desired compression, and uses reasonable parameters in terms of initial energy spread and chirp to create the compression. In principle, it is possible to get closer to the desired bunch length with physically feasible parameters, mainly be reducing the initial bunch length by increasing the frequency shift. However, collective effects will be shown to be so detrimental to the compression that this limit set by Liouville's theorem is negligible in comparison. Nonetheless, this shows a key drawback for this scheme compared to the baseline scheme, which is not as limited by emittance conservation, provided the necessary RF voltage can be achieved.

5.2 PyORBIT Simulations

In order to properly include all relevant effects experienced by the particles in the compressor ring, particle-in-cell (PIC) simulations are needed. For these, the simulation code PyORBIT [5], developed at the Spallation Neutron Source (SNS) for high intensity proton accelerators, was used. Their facility consists of a high-intensity linac and an accumulator ring, similar to the proposed proton driver, and therefore their simulation code is benchmarked in conditions similar to those expected in this work. Other options, such as Xsuite [24] developed at CERN, offer similar capabilities, but as they are benchmarked against other types of accelerators (such as the high-energy LHC and its injector chain), it is expected that they may handle these high-intensity, space charge dominated simulations with less accuracy than PyORBIT.

As the injection from an H^- linac is in itself complicated, this aspect was not handled in the simulation. Instead, bunches were generated in the compressor ring with random distributions based on results from linac simulations [25]. The imprinted energy chirp was also created artificially, using the analytical method from the toy model to calculate the longitudinal positions and corresponding energy offsets of the bunches in the bunch train. In order to confirm the feasibility of the chirp creation in practice, the experiment described in Section 4.3 could be performed, but this was not possible to do during this thesis. Bunch trains were placed with an offset from each other within the 2.84 ns gap between bunches to avoid strong repulsion between existing bunches and injected ones. The amplitude of the chirp was varied according to Equation 48 to achieve simultaneous compression. It was investigated whether the bunch is indeed compressed in the desired manner with the added effects present in the PyORBIT simulation.

5.2.1 Simulation Parameters

The lattice model for the compressor ring developed for the muon collider is less detailed than a final lattice design would be. Certain physical parameters are yet to be specified, mostly related to the physical realisation of the accelerator elements. Thus, there remains some freedom to adapt the lattice for this compression method. Most notably, the beam pipe diameter must be chosen appropriately. This serves two purposes. Firstly, it determines the permitted orbit excursion of particles before the particles are lost. This is handled by adding apertures to the lattice which remove particles with larger transverse positions than the dimensions of the aperture at that location. Secondly, the beam pipe gives a boundary condition to the beam which is essential for accurately calculating the collective effects, as discussed in Section 3.4. Ideally, a detailed lattice with realistic beam insertions would be required to fully account for the impedance effects from the accelerator, but this cannot be done until the design has been iterated upon further.

For the purpose of this work, a relatively large beam pipe with radius 35 cm was chosen, which was done to fulfil the standard requirement that 10σ of the beam size should fit within the beam pipe. As avoiding particle loss is a key consideration, reducing the size of the beam pipe further was found to be detrimental, even though this is a very large beam pipe. This condition corresponds to a transverse emittance of $\sim$30 mm mrad at the locations with maximum beta function, but this does not take into account the dispersion contribution to the beam size. As this is expected to be considerable given the large energy spread used to achieve the compression, a smaller transverse emittance of 20 mm mrad was chosen after comparison with several combinations of emittances and beam pipe radii. Emittance blowup is also a consequence of collective effects, so it is expected that many effects contribute to increased beam sizes, and potentially particle losses.

In the preliminary simulations with space charge, it was found that fast particle losses often occurred, where a particle diverged towards infinity in the simulation between two apertures (which were spaced with roughly 30 m between them). This is a consequence of the incredibly large beam intensities used in the simulation, and demonstrates the strength of space charge potential. As PyORBIT uses an adaptive grid for calculating the space charge potential, it redefined the grid to follow these rapidly lost particles, leading to diverging grid sizes that caused the simulation to crash. This was investigated, but no other solution was found than to reduce the aperture size while adding as many apertures as could fit in the ring, such that the particles were removed before the fast loss occurred. This is not ideal, as it increases the number of lost particles since the transverse tolerance is reduced, but it was necessary for performing the simulations. For the same reason, the bunch length for the bunches from the linac was lengthened from the design value of 2 ps to 20 ps, which reduced the particle density and thus relaxed the peak space charge potential. A similar choice was made for the energy spread from the linac, after initial simulations at a design energy spread value of $\Delta E/E = 3 \times 10^{-4}$ showed larger final bunch lengths than those at larger energy spread. An optimum was found at $\Delta E/E = 8.85 \times 10^{-4}$, which was subsequently used.

Simulations in PyORBIT were performed both with and without the three-line scheme explained in Section 4.2.1, confirming the desired reduction of space charge effects. Without space charge, a larger frequency shift was preferred, whereas with space charge the final bunch length was shorter with the chosen frequency shift, and considerably fewer particles were lost.

Table 3: The key parameters used for the PyORBIT simulations.

Quantity	Value
Bunch length (linac)	20 ps
$\Delta E/E$ (linac)	8.85×10^{-4}
Transverse emittance	20 mm mrad
Beam pipe radius	35 cm
Chirp amplitude	30 MeV

5.2.2 Convergence Checks

In order to confirm that the results of the simulations are trustworthy, several simulations were launched with varying parameters for the particle-in-cell simulations. The longitudinal grid size was varied with powers of 2 from 16 to 256, while simultaneously varying the number of macroparticles per bunch from 1 to 100. The space charge node separation was also varied from 0 to 10 m (the nodes can only be placed where there are already nodes in the lattice, so a separation of 0 m corresponds to putting the space charge nodes as tightly as possible without modifying the lattice). All possible combinations of parameters in the relevant regime were tested, and the results in terms of bunch length, energy spread and lost particles at the end of the simulation was analysed to see if there was any major variation or trend in the result. It was found that the number of macroparticles per bunch had to be larger than one, where a value of four per bunch was chosen, resulting in a final number of macroparticles on the order of 1.4×10^6 after the full injection. Similarly, the longitudinal grid size was found to be optimal with 32 or 64 bins, where the latter was used for all subsequent simulations. For the transverse plane, where the dynamics are not as interesting to study, 32 bins were used. This aligns relatively well with the old wisdom that there should be roughly 10 macroparticles per grid cell in a particle-in-cell simulation [26]:

$$n_{\text{cell}} = \frac{1.4 \times 10^6 \text{ particles}}{32 \times 32 \times 64 \text{ grid cells}} = 21.4 \text{ macroparticles per cell,} \tag{59}$$

at the end of the simulation. No major variation was found apart from this, with the space charge separation of 2.5 m being used in the simulations.

Limited comparisons of the simulations with another simulation code, Xsuite [24], were also performed. This showed little qualitative difference to the results from PyORBIT.

5.2.3 Procedure

Given the specified parameters for the compression from Table 2, PyORBIT simulations were performed first without space charge with the parameters in Table 3. Initially, simulations with individual bunch trains were performed to verify the compression dynamics, also with space charge. After this, the full simulation was performed without and with space charge. New bunch trains were injected every turn, with bunch trains of 267 bunches spanning almost the full circumference of the compressor ring. These simulations were improved iteratively, fine-tuning various parameters to achieve the best compression possible given the various constraint. The most significant ones taken into account were the performance of the linac and the requirements of the target, as well as the physical constraints for designing the lattice (magnet strength, beam pipe radius, cavity amplitudes etc.). For the space charge simulations, different intensities were also tested to investigate potential other use cases for the proposed scheme.

6 Results

The results are presented in the same order that the compression method was investigated with simulations. Firstly, single bunch train compression is demonstrated for trains injected at the beginning and end of the injection, respectively. This is done without and with space charge effects. After this, the full injection is handled without space charge, demonstrating the full compression in a realistic lattice with non-linear effects, energy spread, and Liouvillian limitations. Finally, the full simulation with space charge is presented, demonstrating strong instabilities due to collective effects. Simulations with different intensities are shown, demonstrating cases where the proposed compression scheme may be feasible for other applications.

Beyond this, the results from an alternative lattice operating below transition are presented.

6.1 Single Bunch Train

Verifying the toy model parameters in Table 2 was first done by simulating a single bunch train injected with the specified chirp amplitude to verify that the bunch train was maximally compressed after the desired number of turns. This was done for the first and final bunch trains without space charge, giving the results in Figure 22. Here, the longitudinal phase space distribution of the bunch train is shown before (red) and after (blue) compression. The energy spread in the bunches is the same for both cases, but the chirp amplitude differs, as seen on the vertical axis, corresponding to the difference in compression time according to Equation 48. In both cases, the compression was found to agree well with the expectation from the toy model: the ideal number of turns for compression was found to be 1621 turns and 309 turns, respectively, within 5 % of the number predicted by the toy model. This indicates that the non-linear effects of the lattice are small. Note that the toy model prediction given in Section 5.1.1 is optimised for the full bunch train with 250 turns of compression. The optimal number of turns specifically for the final bunch train is 295 turns.

Adding space charge to these simulations shows that the dynamics are significantly affected over time, as seen in Figure 23. It is expected that the attractive space charge resulting from operating above transition leads to more rapid compression [18], which is also seen for the bunches injected early where a reduction of the compression time by around 10 % is optimal for achieving the minimal bunch length, as shown in Figure 24. For the final bunch to be injected however, the difference is not as drastic, with the final bunch length remaining largely unchanged. With the particle density of a single bunch train, containing 267 bunches, the space charge forces are strong enough to change the dynamics of the compression for bunches injected early into the compressor, with a low amplitude chirp. It follows naturally from this observation that the effects of space charge will be quite noticeable when dealing with a full injection, with more than 1000 times as many bunches. It is also notable that there is an appreciable difference in the impact of space charge when decreasing the frequency shift to reduce the particle density, as seen in Appendix C.

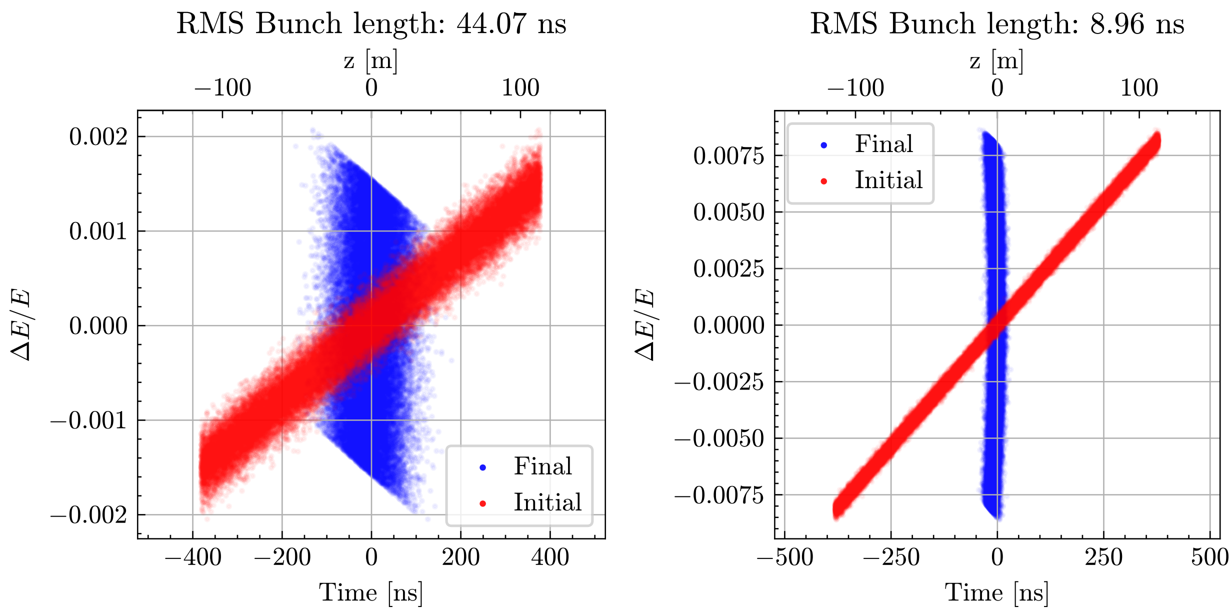

Figure 22: The compression of the first (left) and last (right) bunch trains to be injected, simulated in PyORBIT without space charge. The ideal compression time was found to be 1621 and 309 turns, respectively, increasing the compression time by less than 5 % compared to toy model predictions. The horizontal axis covers the full 314 m, or 1047 ns, circumference of the ring. Note that it shows the position around the ring converted to units of ns, rather than the arrival time of the particles as used in the plots above.

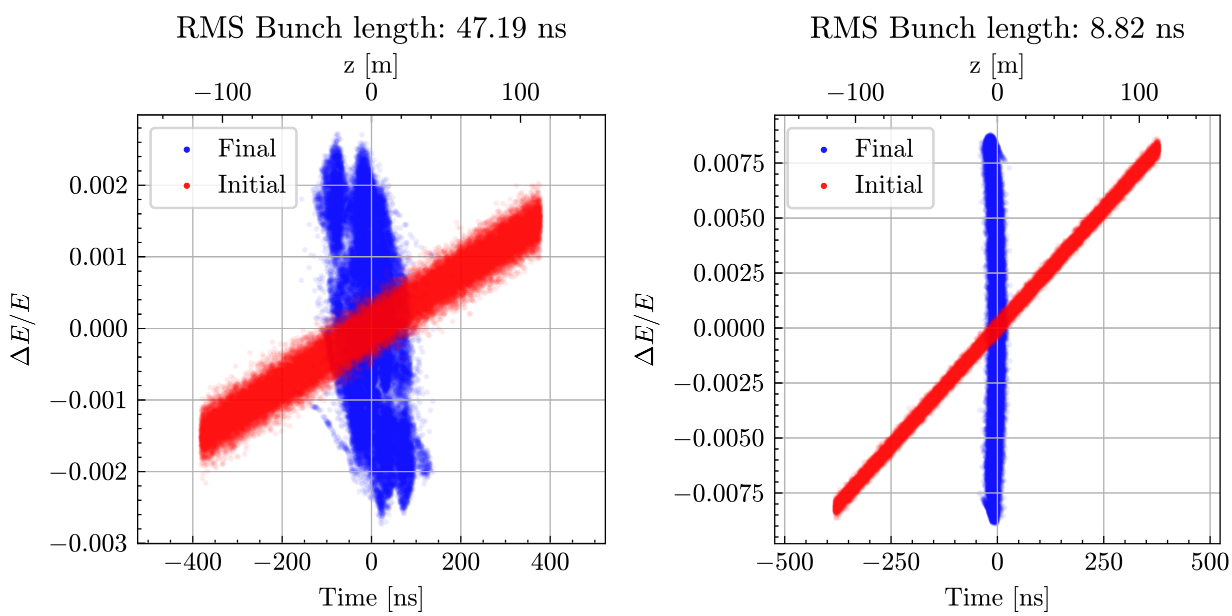

Figure 23: The compression of the first and last bunch trains, respectively, simulated in PyORBIT with space charge but otherwise identical parameters (including number of turns) to the simulations in Figure 22. Compression is almost unchanged for the last bunch train, but the dynamics of compression are clearly altered for the first bunch train. The bunch is slightly over-compressed, the shape has changed, and the energy spread increased. Without space charge, the energy of the particles after compression would have to be the same as the initial energy.

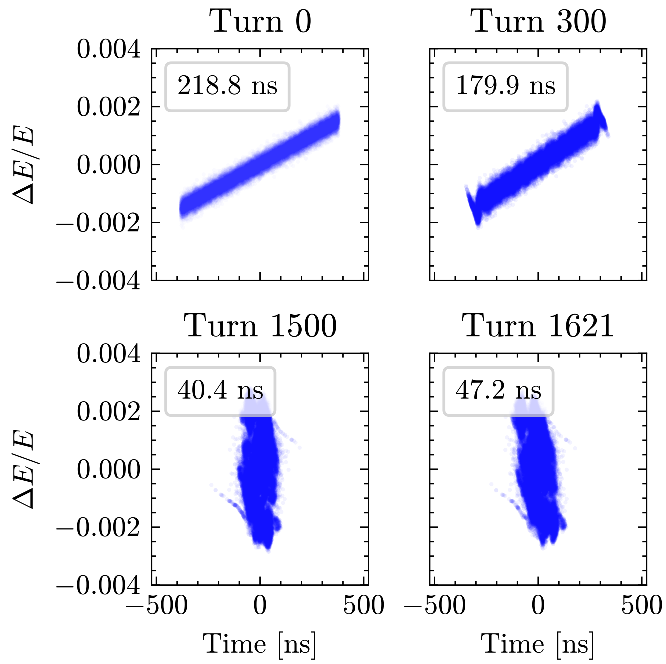

Figure 24: The longitudinal phase space distribution evolving over time for the first bunch train to be injected, as simulated in PyORBIT with space charge. The optimal compression occurs after around 1500 turns, with a bunch length of 40 ns, considerably shorter than at the number of turns predicted by the toy model. The RMS bunch length is shown in the top left inset.

6.2 Full Injection

As implied in the previous section, the increased charge density significantly affects the results of the simulation when performing the full injection of particles. If one extrapolates the results from the previous section, assuming a linear dependence of the bunch length with the compression time at the given amplitude (which was found to be a good approximation for these simulations), a total bunch length of 21 ns is found for the full bunch, consisting of 1312 bunch trains injected on top of one another and compressed. This is an order of magnitude greater than the desired bunch length, but in line with Liouvillian expectations on the final bunch length. Indeed, the predicted bunch length is only slightly longer than the 19 ns calculated as the minimum achievable one in Section 5.1.2. Injecting the full series of bunch trains however, the compression demonstrated a very different behaviour than in the case with individual bunch trains.

6.2.1 Without Space Charge

Simulations were first performed without space charge to show how the full compression would occur in the ideal case, with the parameters derived from the toy model. This gave the compression shown in Figure 25. It is clear that the energy spread limits the minimum bunch length, with increased debunching observed for bunch trains injected early. For the final bunches to be injected, compression occurs fast enough that the bunch train is short when fully compressed. It is also found that the ideal compression time varies slightly from the toy model estimate, with an optimum at 1626 turns. The RMS bunch length of 28.1 ns is slightly longer than the minimum permitted by Liouville, which is because the energy spread used in the final simulations, found in Table 3, is larger than the minimal one the linac is capable of delivering in order to reduce the impact of collective effects. In this case without space charge, this is of course not necessary, but this makes the result easier to compare to the subsequent simulations with space charge.

Interestingly, the length is greater than that estimated from the space charge simulations of single bunch trains with the same energy spread. If the full compression proceeded ideally, this should not be the case since the different bunch trains do not interact via space charge forces and should thus behave similarly to the individual bunch trains simulated above. The reason for this is that the synchronisation of the compression is handled differently between the two cases. The extrapolation was performed with the optimised compression time for all bunch trains, which was shorter than the one predicted analytically. This means that such a compression is not possible as is, but would have to be adapted for realistic conditions. Furthermore, the synchronisation was found to differ slightly between PyORBIT simulation and predictions, and the chirp amplitude is based on the analytical model. Thus, it is expected that the individual bunch trains in the full simulation are not perfectly aligned in compression, so even though the total bunch compression is optimised for the number of turns, the individual bunch trains are probably not.

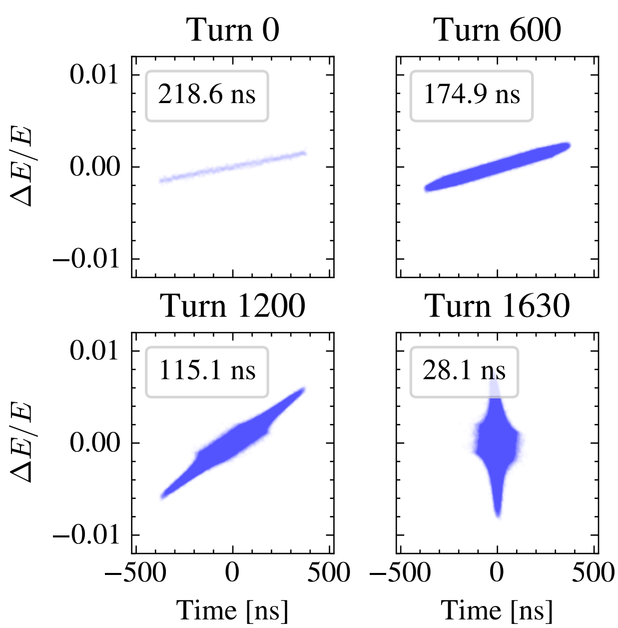

Figure 25: The simulated compression with 1312 bunch trains injected sequentially with varying amplitudes, compressing simultaneously without space charge. The RMS bunch length is shown in the top left inset.

6.2.2 With Space Charge

Deviations from the analytical expectation become much more pronounced in the case with space charge, for which the compression is shown over time in Figure 26. The beam shows clear microbunching and increased energy spread after around 150 turns, before breaking up entirely well before the compression would have finished after 1597 turns. Thus, it is not even relevant to consider the final bunch length as a figure of merit when comparing simulations with different parameters. Instead, the number of lost macroparticles in the simulation is a more relevant measure when evaluating the compression with different parameters such as beam pipe shape and radius, beam intensity, and other relevant parameters both physical and simulation related.

In the case of full intensity injection, a total of 5×10^{14} protons are injected, which corresponds to 1.4 million macroparticles. Of these, only 0.35 million or 23 % remained at the end of compression after 1597 turns. This compression time is estimated based on the toy model, and as indicated above most likely not accurate with the attractive space charge forces above transition. However, as the microbunching starts when only a tenth of the bunch trains have been injected, with considerable losses and beam breakup before the final bunch train is injected, the optimal compression time is not a concern for the full intensity injection. The increased energy spread which is observed with the microbunching naturally alters the compression dynamics, and even for the last injected bunch train, the attractive force from the space charge potential of the central particles is significant enough to ensure it compresses much differently than seen in Figure 23.

It is clear that there is a difference between this compression scheme and compression performed in an RF bucket with a certain synchrotron frequency. Here, the simultaneous accumulation and compression mean that a large spread of rotation frequencies are present simultaneously, limiting the direct parallels that can be made with such single bucket compression. However, the lack of a bucket providing stable motion becomes very clear upon the beam breakup which occurs after the full injection has been performed. Although many particles would likely have been lost even with an RF system providing some stability in the compressor ring, for example through the various fast loss mechanisms that are present, the stabilising synchrotron oscillation might have reduced some of extreme energy spread which is observed here.

The microbunching and growing energy spread are attributed to the negative mass instability (NMI), described in Section 3.4.3. The characteristic features of the negative mass instability are microbunching and exponential growth of the instability, with a growth rate determined by Equation 44. The growth rate is proportional to the square root of the bunch intensity (peak current) and slippage factor, which are both very large in this compression scheme. Additionally, because this compressor operates at a relatively low energy of 5 GeV, the strength of the instability is not reduced by the particles being ultrarelativistic. Given that the compression scheme is uniquely well-tuned to exciting the negative mass instability, it seems reasonable to connect the observed behaviour to it. The detected microbunching occurs rapidly, leading to increased energy spread, exceeding the imprinted chirp on the bunch train. The microbunches in turn begin oscillating around their own centre, due to the large energy spread. This causes the creation of small vortices within the bunch in longitudinal phase space. Eventually, some particles leave the collective motion in the bunch, leading finally to the complete breakup of the bunch, as seen in Figure 26.

In order to excite the negative mass instability, it is necessary to have local density variations around which the microbunching can start. Due to the randomly generated bunch distributions and injection here, it is not impossible that some features of the microbunching observed here derive from numerical effects. Distinguishing the physical effects from the numerical ones is key to trusting the results of the simulation. In this case, because the injection is not performed with a single uniform bunch, but many displaced trains of bunches, there will naturally be local density variations at several times even in an idealised injection. Every time a new bunch train is injected, the previous bunch trains have been compressed and are slightly shorter than the newly injected one, leading to increased density in the centre. Beyond this, the shifting relative locations of bunches in the bunch train lead to local variations in the line density. These effects should suffice to excite the instability, meaning numerical noise, although it cannot be entirely discarded as a potential source, is likely not the dominant source of the density variations.

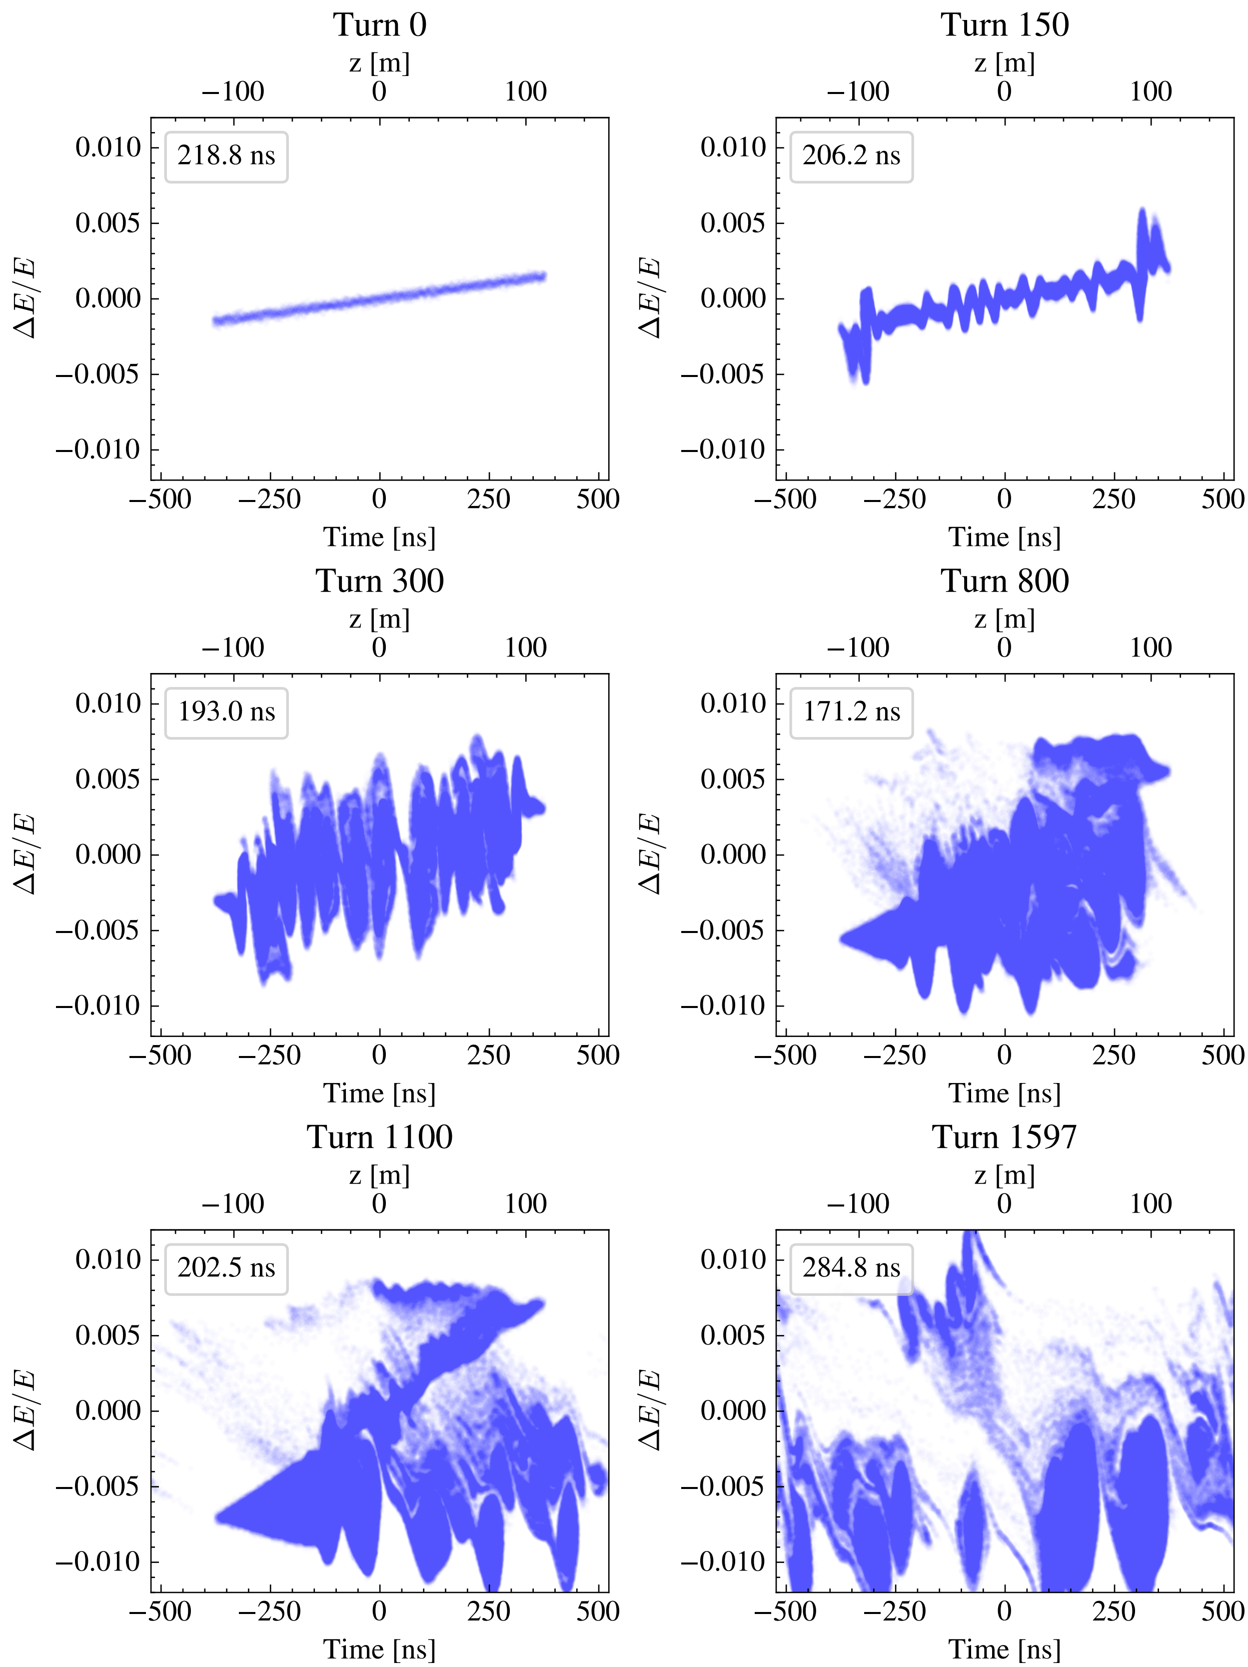

Figure 26: The compression of a full pulse injected from the linac shown over many turns. Injection of the 1312 bunch trains occurs sequentially, with the bunch trains being accumulated and compressed simultaneously. Microbunching is observed already after 150 turns, leading to increased energy spread and eventually bunch break-up. The RMS bunch length is shown in the top left inset.

6.2.3 Intensity Scan

The results with space charge clearly demonstrate that the proposed compression scheme is not feasible for the requirements of the muon collider. However, at a reduced intensity, the space charge effects are expected to be reduced. Furthermore, the injection time is shortened significantly by reducing the intensity, meaning the bunch trains are compressed for a much shorter time. This further reduces the impact from collective effects. Therefore, the compression method was attempted with intensities ranging from 1×10^{13} to the nominal 5×10^{14} protons in identical conditions, with the key results summarised in Table 4. It is seen in Figure 27 that compression is much more successful at lower intensities.

Table 4: The results from the intensity scan performed with the nominal chirp amplitude of 30 MeV at the optimal number of turns for compression. Note that the number of turns is only resolved in multiples of 50 turns due to how the data was saved from the simulation, leading to a large uncertainty in the other quantities.

Protons/10^{13}	Turns to compression	Remaining particles [%]	RMS bunch length
1	300	100 %	25 ns
2.5	350	100 %	37 ns
5	350	97 %	52 ns
10	450	71 %	72.4 ns
25	750	37 %	130 ns
50	1597	19 %	N/A

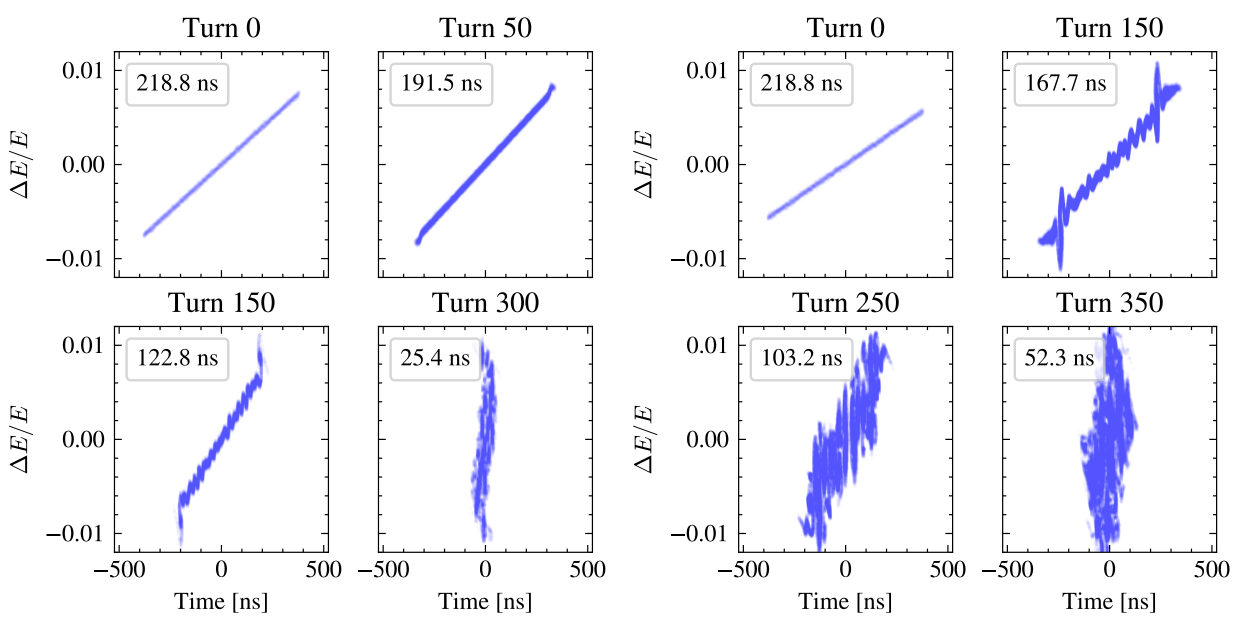

Figure 27: The compression for simulations with 1×10^{13} (left) and 5×10^{13} protons (right) with identical conditions at the nominal chirp amplitude of 30 MeV.

Although the negative mass instability is seemingly still excited, with clear microbunching occurring regardless of intensity, the energy spread increase does not lead to bunch breakup in the lower intensity cases. The final bunch length with 1×10^{13} protons is actually close to the Liouvillian limit, which is particularly impressive considering that the injection scheme was not modified to be optimised for this shorter injection. It would have been advantageous to increase the frequency shift to minimise the Liouvillian bunch length limit, while also increasing the chirp amplitude. This would not only reduce the minimum bunch length, but also significantly affect the total compression time now that the injection time is shortened to just 26 turns, reducing the detrimental impact of collective effects. Increasing the chirp amplitude to 60 MeV gives the results in Figure 28, showing the final bunch length is indeed shortened, while still avoiding particle loss with 1×10^{13} protons. The results for all intensities are given in Table 5. Here, the beam pipe shape was altered to an ellipsoidal cylinder with a horizontal half-axis twice as long as the vertical, which was kept the same as in the previous simulation, to compensate for the increased beam size due to dispersion.

Comparing the chirp amplitudes, it is clear that the bunch length is considerably shorter for low intensities at the larger amplitude, whereas for higher intensities, the difference is not as clear. This is due to the longer injection time at higher intensities, which limits the benefits of shortening the compression time for the final bunch trains. The number of lost particles is very similar between the two cases, signifying that the impact of collective effects is quite similar, even though the total compression time is shortened. The losses may occur more readily due to the larger dp/p for the large chirp amplitude, and the increased orbit excursion due to dispersion.

Table 5: The results from the intensity scan performed with the increased chirp amplitude of 60 MeV taken at the optimal number of turns for compression.

Protons/10^{13}	Turns to compression	Remaining particles [%]	RMS bunch length
1	180	100 %	13 ns
2.5	210	100 %	17 ns
5	255	94 %	27 ns
10	360	61 %	50 ns
25	700	35 %	130 ns
50	1597	23 %	N/A

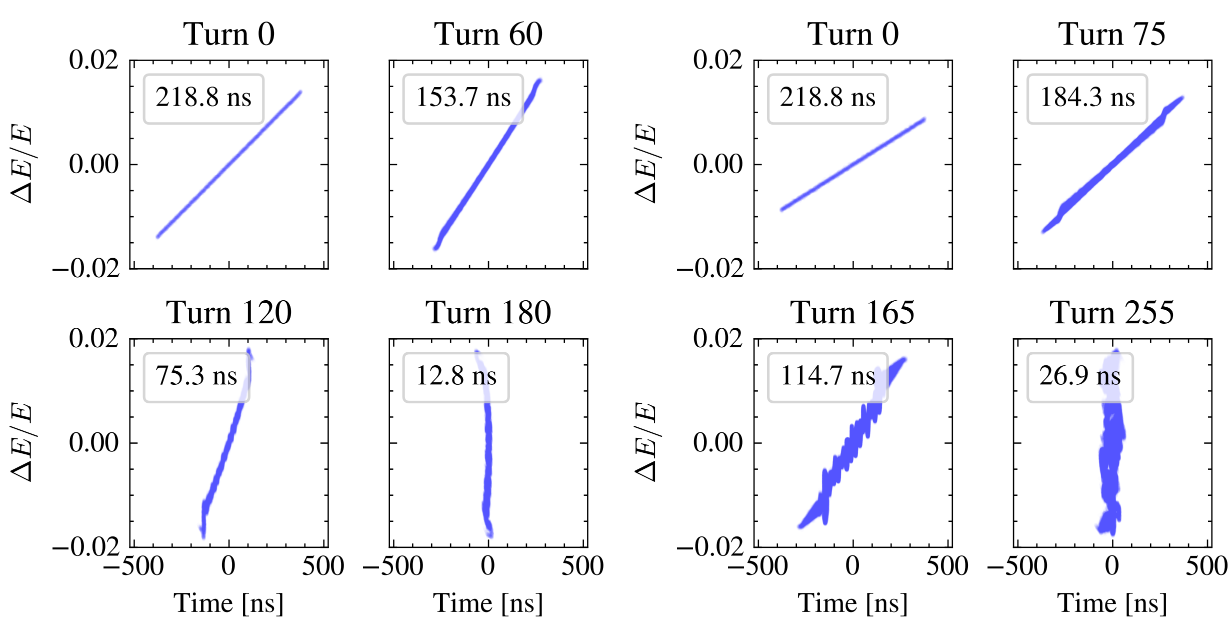

Figure 28: The compression for simulations with 1×10^{13} (left) and 5×10^{13} protons (right) with identical conditions at a chirp amplitude of 60 MeV.

6.3 Lattice Below Transition

Because of the observed issues arising from the negative mass instability, a lattice operating below transition was investigated. The lattice was adapted from one suggested by Shinji Machida [21] to circumvent some issues, including NMI, for compressors operating above transition. However, such a lattice is inherently limited in its compressive strength by Equation 36, which shows that for a real transition energy, the slippage factor cannot be greater than $\frac{1}{\gamma^2}$. In the case of this compressor lattice, operating at an energy of 5 GeV and thus $\gamma = 6.33$, this sets an upper bound for the absolute value of the slippage factor to $|\eta| < 0.025$, significantly less than the value of $\eta = 0.17$ in the lattice above transition. Therefore, the compression time increases significantly, which as it has been seen before worsens the impact of collective effects. It is also clear that now that the lattice operates below transition, the space charge force will be repulsive in the longitudinal direction, counteracting the compression from the chirp. To counteract this, the chirp amplitude was increased to 50 MeV, still allowing the bunch to fit in the beam pipe taking dispersion into account. The compression time is however nonetheless longer, estimated to 1210 turns for the final bunch to be injected, instead of 295 turns. This is seen in Figure 29, with a total time for injection and compression being 2522 turns. Compression does occur in the case without space charge, with a final bunch length of 24.7 ns. The final bunch shape is however not linear or sinusoidally shaped as expected, but displays a curvature which consequently must come from non-linearities in the lattice. This is likely a result of large non-linear terms in the slippage factor, combined with the large chirp amplitude and thus dp/p, according to Equation 37.

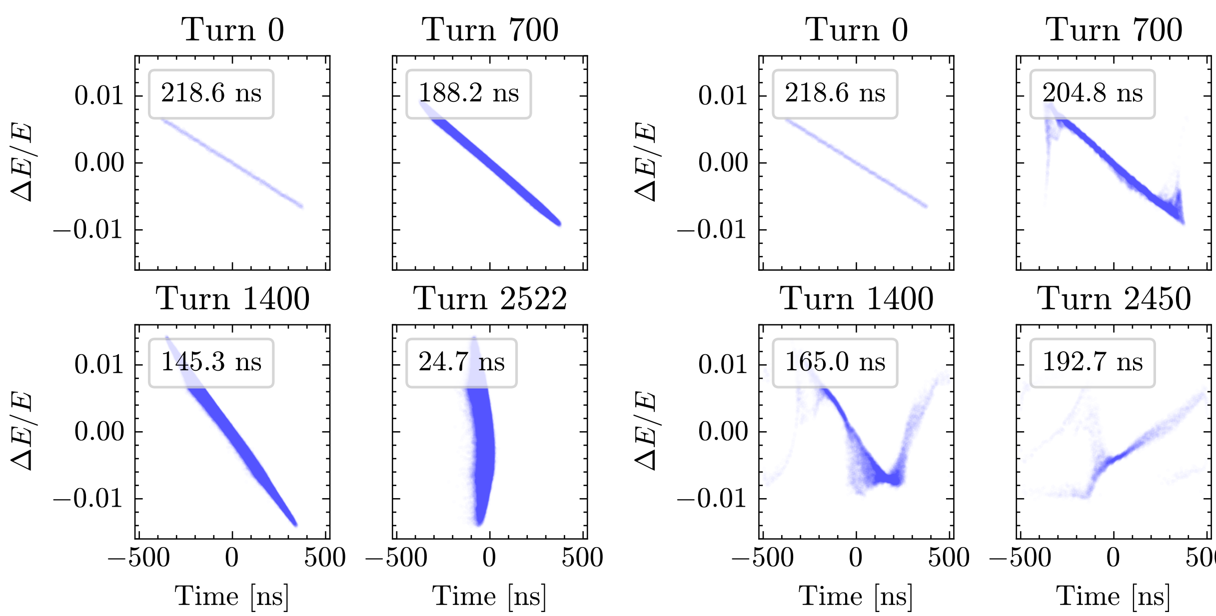

Figure 29: The compression over time for a lattice below transition without space charge (left) and with space charge (right). Compression is achieved without space charge over a prolonged compression time of 2522 turns, but a non-linear shape is seen which does not derive from the sinusoidal chirp. In the other case, compression is strongly counteracted due to the repulsive space charge force. Many particles are also lost, which can be seen from the reduced colour intensity in the plot. The RMS bunch length is shown in the top left inset.

Adding space charge, it is seen in Figure 29 that the repulsive space charge force overcomes the compressing strength from the chirp, even with the increased amplitude. This effect is strong enough that it would not be possible to find a reasonable chirp amplitude to overcome the repulsion, as this would cause a very elongated bunch in the horizontal dimension due to dispersion. Adapting the beam pipe shape to an ellipse with a longer half-axis in horizontal was attempted, but the fast particle losses mentioned in Section 5.2.1 prevented any simulation from being completed. While this may be a limitation caused by the simulation, a clear dependence is seen with bunch density, where larger bunch densities correspond to more fast losses. Thus, it is expected that this limitation also corresponds to physical limitations of the compression method, although the exact loss mechanisms are not known. Investigating this does however not fall within the scope of this thesis, and it is assumed that a compression scheme which cannot be simulated with standard simulation codes is not implementable in reality. Even in the simulations which do finish, many particles are lost, leaving only around 2–4 % of the macroparticles at the end of the simulation.

The space charge repulsion could be counteracted by adding a weak RF field, operating at constant energy and with the same slope as the compressive chirp. This RF would not drive the compression, as in the baseline case, but rather assist the chirp-based compression to overcome the repulsive forces. This was not investigated as part of this work, but could be investigated to see if a lattice below transition is feasible. A comparison of the space charge potential could give an approximate idea of the required RF voltage. Combining this idea with a lattice with an imaginary transition energy could lead to significantly faster compression without operating above transition. Such lattices have negative momentum compaction factors, and thus never cross transition since the transition energy becomes imaginary as follows from Equation 35. This circumvents the limitation on $|\eta|$ from Equation 36. Therefore, the notion of an alternative lattice is worth investigating, even though the below transition lattice investigated here does not suffice.

7 Conclusions

The results presented here for the novel compression scheme demonstrate the possibility of simultaneously injecting and compressing bunch trains from a linac, achieving a short, high-intensity proton bunch without the need for an RF cavity in the compressor ring. This is however only possible at bunch intensities much lower than those of the proton driver for the muon collider, which was the intended use case for the compression scheme. Due to the very large bunch intensity required for the creation of muons, the self-interaction of the protons through various collective effects, mainly space charge, leads to strong instabilities in the lattice operating above transition. One important such instability is the negative mass instability, which causes microbunching with increased energy spread as a result. Over the long injection time, this causes the bunch to deviate from the intended compression dynamics and eventually break up entirely. Many particles are lost in this process, largely due to particles with large energy deviations hitting the apertures due to dispersion.

Various measures were taken to reduce the impact of space charge, such as utilising a design which separates the linac pulse to three separate lines with off-frequency cavities. This was found to shorten the injection time significantly when using kicker magnets with a realistic rise and fall time of 300 ns. By also shifting the frequency of the cavities by one third the revolution frequency of the compressor ring from the nominal RF frequency, the length of the bunch train in the ring was maximised, minimising the proton density. The amplitude of the chirp was also optimised, but due to the $1/t$ dependence on the injection time, it is not possible to make any major improvements on the total compression time if the injection time is significantly longer than the compression time at maximum chirp amplitude. Although these measures were found to indeed reduce the strength of the instabilities, they were not enough to prevent the negative effects of the instabilities.

As a final attempt to avoid the negative mass instability at the required bunch intensity, a compressor lattice operating below transition was tested. As the main instability disturbing the compression was found to be the negative mass instability, which is only present above transition, such a lattice would avoid the microbunching and particle loss driven by the attractive space charge. However, below transition the repulsive longitudinal space charge force is strong enough to prevent the compression even at increased chirp amplitude. The particle loss was also found to be considerably larger, likely due to orbit excursions due to dispersion from particles which are repelled by the central space charge potential. Consequently, it is not deemed a feasible option in its current state.

Ultimately, the impact of bunch intensity on the compression was investigated in order to find a potential working point for the compression scheme. This showed that at bunch intensities reduced by a factor of 50, the compression could get close to the Liouvillian limit. Then, the instabilities did not have time to cause particle loss or beam breakup, demonstrating a case where the compression scheme would be useful.

8 Outlook

After all investigations performed, it was seen that the only way to circumvent the instabilities was to reduce the bunch intensity. Although this excludes the muon collider use case, it would in principle be possible to deliver a high average power with short bunches even at lower bunch intensities, as the repetition rate could be increased significantly. It is a requirement for the muon collider that the repetition rate is low to ensure enough muons can be captured in a single bunch, maximising the luminosity according to Equation 2. For a use case such as a neutrino factory, this constraint can be relaxed considerably. For the proposed Superconducting Proton Linac (SPL) at CERN, which was designed as a driver for a neutrino factory, and also served as the inspiration to the proposed proton driver design for the muon collider, a repetition rate of 50 Hz is used with a very similar linac [27]. The design of the full driver is also very similar to the baseline muon collider design, with an accumulator ring accumulating six bunches which are subsequently compressed one at a time in the compressor ring to a 2 ns bunch length. This delivers an average power of 4 MW at 5 GeV, very similar to the desired parameters for the muon collider.

Attempting to match these parameters with the novel compression scheme here, one can consider the higher amplitude case at a bunch intensity of 1×10^{13} protons with a chirp of 60 MeV. The injection and compression occurs in 180 turns, which is considerably less than the length of a linac pulse. In fact, ten such sets of bunch trains could be injected and compressed within a single linac pulse, which would then be almost 2 ms long, in line with the baseline design. One such pulse is then chopped for the 26 bunch trains sent to the compressor, and then the remaining bunches are chopped from the pulse for the 154 turns it takes the bunch train to compress. This bunch, which is compressed to 12.8 ns in the current simulations, is then sent to the target and the next bunch train is injected. This chopping scheme is shown in Figure 30.

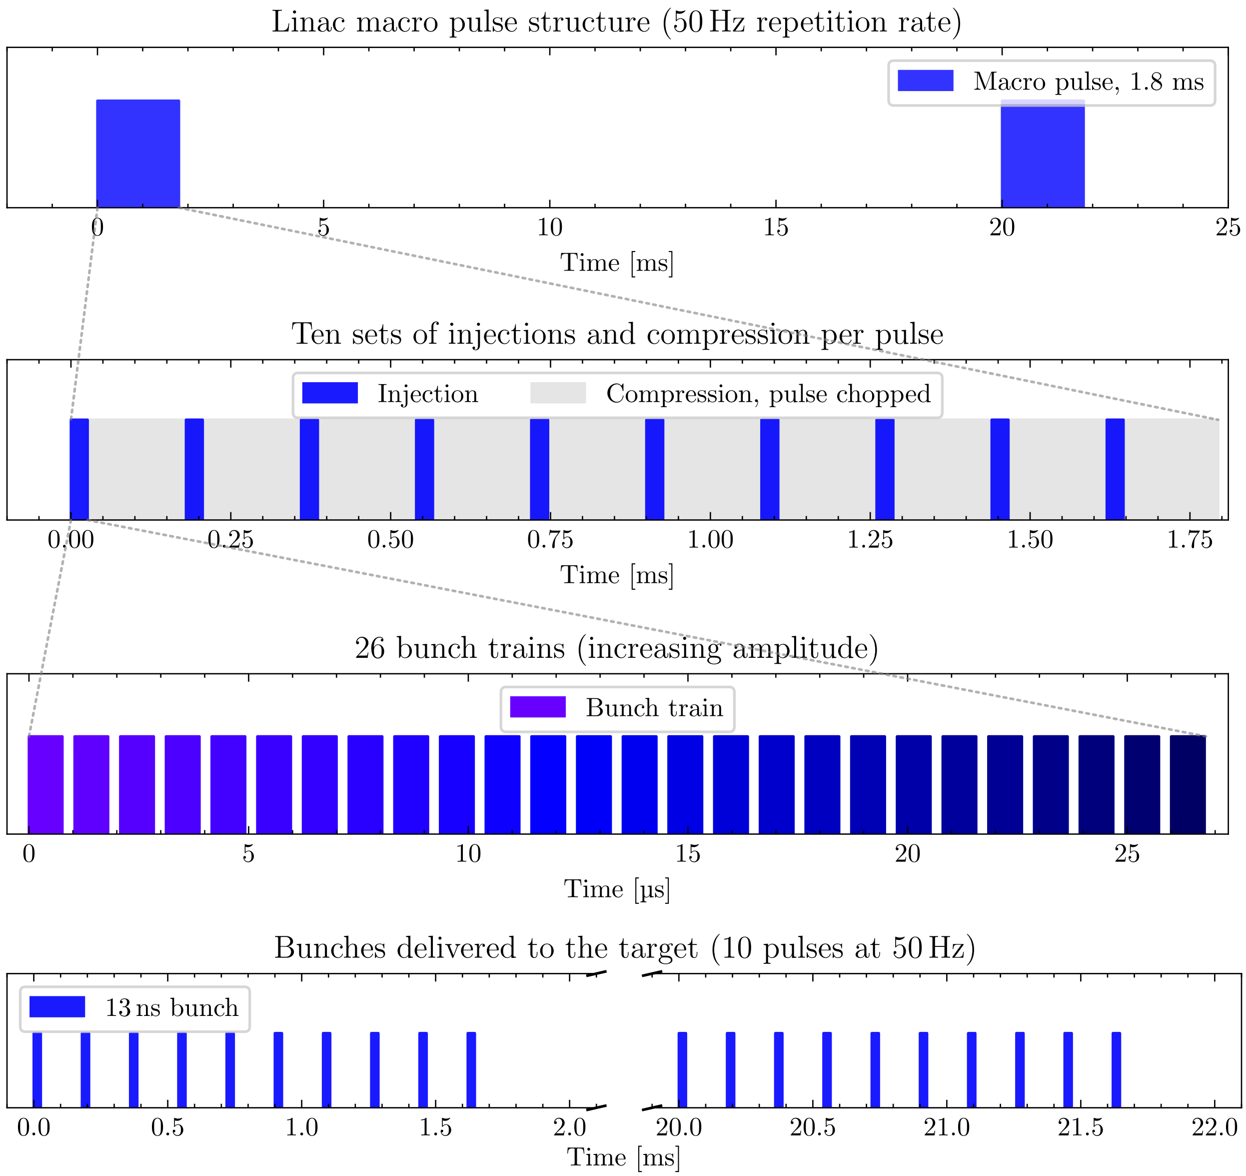

Figure 30: The chopping scheme used for maximising the repetition rate, where every linac pulse is divided into ten separate bunches which are sent to the target after compression to 13 ns. This requires an intensity of 1×10^{13} protons, at a maximum chirp amplitude of 60 MeV. The darkening colour within the bunch train represents the increasing chirp amplitude, ensuring synchronised compression.

As the bunch intensity is reduced by a factor 50 and the repetition rate is increased by a factor of 10 from the linac pulse rate, with an additional factor 10 from each pulse, the average power of such a scheme is 4 MW, as desired for the discussed neutrino factory [27]. There are also several parameters which can be optimised for such an application, for example the frequency shift, linac pulse length, energy spread, and lattice slippage factor. By tailoring these to the application, it would be possible to achieve a driver which delivers 4 MW of average power at 5 GeV with a similar pulse structure to that proposed in [27], at a bunch length of less than 10 ns. Since space charge driven instabilities are much less prevalent at this intensity, the slightly longer compression time compared to a design with an accumulator ring and compressor with RF was not found to be an issue. However, the improvement of removing the accumulator ring, while also avoiding the need for the high voltage RF in the compressor ring, which for the proposed neutrino factory is similar to the muon collider (using the same compressor lattice) at 4 MV [27], makes this compression scheme an attractive alternative. A neutrino factory is a possible first step towards constructing a muon collider, since it has very similar requirements. It can even use the same proton driver design, as seen here. Then, once the technology matures, the full muon collider complex could be constructed around the existing facility, reducing the financial risk of the construction.

Before implementing the scheme in practice however, there are several things which must be further studied, and some physical parameters of the machine which would require additional investigations. For example, the large beam pipe radius and the kicker magnet strength as well as their rise and fall times are within realistic bounds individually, but in the machine context the realisation may prove difficult. In terms of further investigations of the compression scheme itself, there are further effects which are interesting to investigate. For example quantifying the space charge potential at different stages of compression would provide useful information. Other collective effects which affect the dynamics should also be investigated further, for example wakefields from a more realistic lattice model and intrabeam scattering (IBS). It would be interesting specifically to study IBS as it typically is not relevant in proton accelerators due to the lower phase space density compared to electron accelerators which through synchrotron radiation damping can reach a considerably smaller emittance [10]. However, for the bunch compression here, the proton density is so high that even at relatively large emittances, the effect might become relevant. This holds additionally because the effect is strongly suppressed by energy, and this compression operates at a low γ. As IBS causes emittance dilution [19, p. 468], it would potentially counteract the compression significantly. The simulations themselves could also be improved, most importantly by including the off-frequency cavities at the end of the linac in the simulations and the challenging injection of pulses with strong longitudinal correlation and large energy spread into the compressor ring. It would also be relevant to perform more rigorous analysis of the numerical effects in the simulation. One could also go further in comparing the simulations between different codes to get an estimate of the error in the simulations.

There are also several ways the scheme itself could potentially be improved. One very interesting topic to investigate would be space charge compensation with impedances, which has been attempted in the past and shown success in reducing the effect of the negative mass instability at e. g. Los Alamos [28, p. 5-12]. This could increase the achievable bunch intensity if used with the compression scheme. One could also pursue the path of operating below transition, for example by introducing a lattice with imaginary transition energy and thus a larger slippage factor than the alternative lattice used here. Combining this with a weak RF cavity in the compressor could potentially overcome the space charge repulsion, at least at some intensity. Furthermore, harmonic cavities could be added to the off-frequency RF cavities in order to shape the bunch to have a more linear envelope, or potentially a non-linear one adapted to the other non-linearities experienced by the bunch.

Acknowledgements

I wish to thank everyone in the beam physics group at ESS for welcoming and helping me throughout this past year. Especially, I am very grateful to Daniel Noll for assisting me in running my simulations properly, and going through the cumbersome PyORBIT installation. I also wish to thank Mamad Eshraqi for taking me into the group and for pushing me to learn as much as possible during my time here. And last but not least I must thank my inspiring supervisor Natalia who turned me onto this very interesting topic, and also gave me the opportunity to do many additional things during my project, such as attending JUAS and going to IPAC.

Beyond this, I have received very valuable help from some people in the Muon Collider Work Package 3 on the proton driver. Particularly, I wish to thank Shinji Machida, Simon Albright, and Austin Hoover, who have all come with valuable input to my work and also given me concrete tips on how to pursue my investigations.

Furthermore, I want to acknowledge everyone I have met in Lund who have helped shape me into the person I am now: all my study comrades who braved the physics education with me, my bandmates in Törner Cryda, and of course all my dear friends from Micklagård. I consider myself lucky to have been blessed with such a rich student life during my years here, and I hope our paths cross during my future adventures.

Finally, I am incredibly grateful for the support from my family, who have supported me wholeheartedly on this journey. I am excited for you to follow me where it goes next!

References

[1] V. Papadimitriou et al. "Design of the LBNF Beamline". In: *Proc. 6th International Particle Accelerator Conference (IPAC'15), Richmond, VA, USA, May 3-8, 2015* (Richmond, VA, USA). International Particle Accelerator Conference 6. https://doi.org/10.18429/JACoW-IPAC2015-THPF120. Geneva, Switzerland: JACoW, June 2015, pp. 3992–3995. ISBN: 978-3-95450-168-7. DOI: `https://doi.org/10.18429/JACoW-IPAC2015-THPF120`. URL: `http://jacow.org/ipac2015/papers/thpf120.pdf`.

[2] Carlotta Accettura et al. *The Muon Collider*. arXiv:2504.21417 [physics]. Apr. 2025. DOI: `10.48550/arXiv.2504.21417`. URL: `http://arxiv.org/abs/2504.21417`.

[3] Mauro Migliorati. "Collective effects". en. In: *Proceedings of the Joint Universities Accelerator School (JUAS): Courses and Exercises, E. Métral (ed.)* (Nov. 2024), p. 539. DOI: `10.23730/CYRSP-2024-003.539`. URL: `https://e-publishing.cern.ch/index.php/CYRSP/article/view/1610`.

[4] Alexander Wu Chao. *Physics of collective beam instabilities in high energy accelerators*. Wiley series in beam physics and accelerator technology. New York: Wiley, 1993. ISBN: 978-0-471-55184-3.

[5] S. Cousineau et al. *The Particle Accelerator Simulation Code PyORBIT*. Proceedings. Procedia Computer Science, 51:1272 – 1281. 2015.

[6] *The European Strategy for Particle Physics: 2026 Update - Recommendations by the European Strategy Group*. Monte Verità/Ascona, Switzerland, 2025. DOI: `10.17181/CERN.423R.S20Z`. URL: `https://cds.cern.ch/record/2950671`.

[7] M. Benedikt et al. *Future Circular Collider Feasibility Study Report*. Tech. rep. 17. _eprint: 2505.00273. Geneva: CERN, 2025. DOI: `10.17181/CERN.I26X.V4VF`. URL: `https://cds.cern.ch/record/2928194` (visited on 02/24/2026).

[8] Daniel Schulte. "The Muon Collider". en. In: (2022). Ed. by Frank Zimmermann et al. Artwork Size: 6 pages, 0.372 MB ISBN: 9783954502271 Medium: application/pdf. ISSN: 2673-5490. DOI: `10.18429/JACOW-IPAC2022-TUIZSP2`. URL: `https://jacow.org/ipac2022/doi/JACoW-IPAC2022-TUIZSP2.html`.

[9] Accettura et al. *The Muon Collider*. 2025. arXiv: `2504.21417 [physics.acc-ph]`. URL: `https://arxiv.org/abs/2504.21417`.

[10] Rasmus Ischebeck. "Synchrotron radiation". In: *CERN Yellow Reports: School Proceedings* (Nov. 2024), p. 433. DOI: `10.23730/CYRSP-2024-003.433`. URL: `https://e-publishing.cern.ch/index.php/CYRSP/article/view/1608`.

[11] P A Zyla et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020), p. 083C01. ISSN: 2050-3911. DOI: `10.1093/ptep/ptaa104`. URL: `https://doi.org/10.1093/ptep/ptaa104`.

[12] M. Aicheler, ed. *A multi-TeV linear collider based on CLIC technology: CLIC conceptual design report*. en. CERN 2012,7. Geneva: CERN, 2012. ISBN: 978-92-9083-379-6. DOI: `10.5170/CERN-2012-007`.

[13] Elias Métral et al. *Proceedings of the Joint Universities Accelerator School (JUAS): Courses and exercises*. Vol. 3/2024. CERN Yellow Reports: School Proceedings. Geneva: CERN, Nov. 2024. DOI: `10.23730/CYRSP-2024-003`.

[14] Alexandre Lasheen. "Longitudinal beam dynamics". en. In: *Proceedings of the Joint Universities Accelerator School (JUAS): Courses and Exercises, E. Métral (ed.)* (Nov. 2024), p. 167. DOI: `10.23730/CYRSP-2024-003.167`. URL: `https://e-publishing.cern.ch/index.php/CYRSP/article/view/1586` (visited on 04/21/2026).

[15] Thomas P. Wangler. *RF Linear Accelerators*. en. 1st ed. Wiley, Jan. 2008. ISBN: 978-3-527-40680-7 978-3-527-62342-6. DOI: 10.1002/9783527623426.

[16] *Conceptual design report*. Lund: European Spallation Source, 2012. ISBN: 978-91-980173-0-4.

[17] David Alesini. "Linear particle accelerators". In: *CERN Yellow Reports: School Proceedings* (Nov. 2024), p. 479. DOI: 10.23730/CYRSP-2024-003.479. URL: https://e-publishing.cern.ch/index.php/CYRSP/article/view/1609.

[18] Giovanni Rumolo and I Hofmann. "A study of fast bunch rotation in the negative mass region". In: (2001). URL: https://cds.cern.ch/record/511342.

[19] Shyh-Yuan Lee. *Accelerator Physics*. eng. 4th edition. New Jersey London Singapore: World Scientific, 2019. ISBN: 978-981-327-468-6 978-981-327-469-3.

[20] King-Yuen Ng. "Emittance growth due to negative mass instability above transition". In: *Workshop on Beam Instabilities in Storage Rings*. July 1994, pp. 0251–264.

[21] Shinji Machida. *Alternative design proton complex*. Muon Collider Annual Meeting 2025. May 13$^{\text{th}}$, 2025. URL: https://indico.desy.de/event/45968/contributions/186222.

[22] Johannesson, Sofia et al. "Initial design of a proton complex for the Muon Collider". In: (July 2024). Ed. by Pilat, Fulvia et al., pp. 2528–2531. ISSN: 2673-5490. DOI: 10.18429/JACOW-IPAC2024-WEPR24. URL: https://jacow.org/ipac2024/doi/jacow-ipac2024-wepr24.

[23] W Fischer et al. "The RHIC Injection System". In: *Particle Accelerator Conference '99 Proceedings* (1999). URL: https://cds.cern.ch/record/553422.

[24] Giovanni Iadarola et al. "Xsuite: An Integrated Beam Physics Simulation Framework". In: *JACoW* HB2023 (2024), TUA2I1. DOI: 10.18429/JACoW-HB2023-TUA2I1.

[25] Natalia Milas. Personal communication. September, 2025.

[26] Shinji Machida. Personal communication. March 2$^{\text{nd}}$, 2026.

[27] R Garoby et al. *Linac-Based Proton Driver for a Neutrino Factory*. Tech. rep. Geneva: CERN, 2009. URL: https://cds.cern.ch/record/1227854.

[28] King-Yuen Ng. "Physics of collective beam instabilities". In: Fermi National Accelerator Lab., Batavia, IL (US). Aug. 2000. URL: https://www.osti.gov/biblio/760458.

[29] Sörqvist, Markus. *Proton Compression for the Muon Collider Project*. Student Paper. 2026. URL: https://lup.lub.lu.se/student-papers/record/9228185.

9 Appendix

A Chirp Creation Code

The code used to analytically create the chirp from the off-frequency cavities is found below. Combined with standard PyORBIT simulations with the parameters in Section 5.2.1, as well as choosing correctly the amplitudes for the different bunch trains according to Equation 48, it should be possible to recreate the results in this thesis. It is possible to find the full code used for the simulations on GitLab: `https://gitlab.esss.lu.se/ess-bp/mucol-lattice/-/tree/main/Novel-Compression?ref_type=heads`.

```python
def generate_bunch_train(
    amp,
    shift,
    n_particles,
    n_bunches,
    E_spread,
    emit,
    analytical_data,
    twiss_dict,
    bunch4,
    analysis,
):
    """
    Adds a new bunch train to the compressor ring, with the specified chirp given by the amplitude.
    ---
    Keyword arguments:
    amp (float): Amplitude of the off-resonance cavities, corresponding to the imprinted chirp
    n_particles (int): Number of particles per bunch. Note: the number of bunches is tripled due to
    the multi-line approach
    n_bunches (int): Number of bunches in the bunch train. If the multi-line approach is used, this
    is reflected in using three times as many bunches here
    E_spread (float): The specified energy spread of the bunches (dE/E)
    emit (float): The transverse emittance of the bunches in units of m
    analytical_data (dict): Dictionary containing bunch positions and relative energy gains based
    on analytical calculations. From the function analytical_setup
    twiss_dict (dict): Dictionary containing relevant parameters such as twiss functions and dispersion
    bunch4 (PyORBIT object): The bunch object to which the new particles are added
    analysis (PyORBIT object): The object used to perform analysis on the bunches

    Returns:
    Nothing, simply adds new bunches to the existing bunch object in pyOrbit
    """

    ### Parameter setup ###

    energy_gain = analytical_data["energy_gain"]
    t_delayed = analytical_data["t_delayed"]

    arrmuX = twiss_dict["arrmuX"]
    arrmuY = twiss_dict["arrmuY"]
```

```
    arrPosAlphaX = twiss_dict["arrPosAlphaX"]
    arrPosAlphaY = twiss_dict["arrPosAlphaY"]
    arrPosBetaX = twiss_dict["arrPosBetaX"]
    arrPosBetaY = twiss_dict["arrPosBetaY"]
    DispersionX = twiss_dict["DispersionX"]
    DispersionXP = twiss_dict["DispersionXP"]

    m_p_gev = PC.m_p * PC.c**2 / PC.e / 1e9
    gamma = (5 + m_p_gev) / m_p_gev
    beta = np.sqrt(1 - 1 / gamma**2)

    dE = energy_gain * amp
    z3 = np.zeros(int(n_particles * n_bunches))
    zp3 = np.zeros(int(n_particles * n_bunches))

    bunch_loc = t_delayed * PC.c + shift
    for i in range(int(n_bunches)):
        if i % 3 == 0:
            z3[i * n_particles : (i + 1) * n_particles] = np.random.normal(
                loc=bunch_loc[int(i // 3)], scale=20e-12 * PC.c, size=n_particles
            )  # Bunch length is 20 ps
        elif (
            i % 3 == 1
        ):  # Shifted  by third of the distance between bunches in bunch train, 950 ps
            z3[i * n_particles : (i + 1) * n_particles] = np.random.normal(
                loc=bunch_loc[int(i // 3)] + 950e-12 * PC.c,
                scale=20e-12 * PC.c,
                size=n_particles,
            )
        elif i % 3 == 2:
            z3[i * n_particles : (i + 1) * n_particles] = np.random.normal(
                loc=bunch_loc[int(i // 3)] - 950e-12 * PC.c,
                scale=20e-12 * PC.c,
                size=n_particles,
            )

        zp3[i * n_particles : (i + 1) * n_particles] = np.random.normal(
            loc=0.0, scale=E_spread, size=n_particles
        )

    # Add chirp for every particle (also within bunches)

    norm_z = (max(z3) - min(z3)) / 2
    mid_z = np.mean(z3)  # Will approximately be true
    norm_E = abs(dE[0] - dE[-1]) / 2

    def E_gain(z):
# OBS: This method for normalizing the energies and positions only works if the full bunch train is included
# Bunch train implementation:
        return (
            norm_E * (z - mid_z) / norm_z
        )  # Note that a minus sign must be added below transition!!!
```

```python

    E_gain = np.vectorize(E_gain)

    zp3 += E_gain(z3) / 1e9  # to GeV

    def z_find(idx, z):
        idx = int(idx)
        return abs(z - arrPosAlphaX[idx][0])

    print("- Matching twiss functions and dispersion\n")

    for i in range(n_bunches):
        # Twiss matching
        z_opt = scopt.minimize_scalar(
            lambda x: z_find(x, bunch_loc[int(i // 3)]),
            bounds=[0, np.size(arrPosAlphaX, axis=0)],
        )
        z_idx = int(z_opt.x)

        ### The twiss at z=0 is used for all particles in PyORBIT for the injection
        twissX = TwissContainer(
            alpha=arrPosAlphaX[0][1], beta=arrPosBetaX[0][1], emittance=emit
        )
        twissY = TwissContainer(
            alpha=arrPosAlphaY[0][1], beta=arrPosBetaY[0][1], emittance=emit
        )
        dist = GaussDist2D(twissX, twissY)

        for j in range(n_particles):
            (x3, xp3, y3, yp3) = dist.getCoordinates()

            ### Dispersion matching is only necessary when D!=0 at z=0
            x3 += (
                DispersionX[0][1] * zp3[i * n_particles + j] / 5.938 / beta**2
            )  # converting dE/E to dp/p
            xp3 += DispersionXP[0][1] * zp3[i * n_particles + j] / 5.938 / beta**2

            bunch4.addParticle(
                x3, xp3, y3, yp3, z3[i * n_particles + j], zp3[i * n_particles + j]
            )
            analysis.account(
                [x3, xp3, y3, yp3, z3[i * n_particles + j], zp3[i * n_particles + j]]
            )


def analytical_setup(A_full=10e6, n_bunches_input=89):
    """
    Note, n_bunches input gives one third of the bunches in the three-line case.
    """
    # === Set up parameters for chirp creation from analytical test ===
    E = 5e9  # eV
    gammaL = (E + 0.938e9) / 0.938e9  # kinetic plus rest energy
```

```
    betaL = np.sqrt(1 - 1 / gammaL**2)
    v = betaL * PC.c  # m/s
    pulse_length_full = 85 / 89 * 4.375e-3 / 3  # s
    # Pulse length adapted for new number of bunches (set by variable pulse_length below) and the triple line

    f0 = 2 * 352.21e6  # Hz, applies to the later ESS cavities
    T_rf = 1 / f0

    # These two are just used to define the detuning
    comp_circ = 314.1592  # m
    T_turn = comp_circ / v

    detune_multiple = 3

    df = 1 / T_turn / detune_multiple  # Hz, close to 1e6 at detune_multiple=1
    T_beat = 1 / df
    fa = f0 - df
    fb = f0 + df
    phi_a = 0
    phi_b = 0
    current = 80e-3  # A
    rep_rate = 5  # Hz
    single_bunch_length = 20e-12  # does not affect final result

    t = np.linspace(0, 2 * T_beat, 100000)

    rf_a = np.cos(2 * np.pi * fa * t + phi_a)
    rf_b = np.cos(2 * np.pi * fb * t + phi_b)

    rf_sum = rf_a + rf_b
    rf_env = 2 * np.cos(2 * np.pi * (fb - fa) * t / 2 + (phi_b - phi_a) / 2)

    pulse_length = (2 * n_bunches_input + 1) * T_rf  # s
    # Only length per cycle of detuning. Total bunch length ~2.5 ms
    bunch_f0 = 352.21e6  # Hz, applies for the earlier part of the ESS linac, i.e. this controls the bunch spacing
    bunch_T = 1 / bunch_f0

    n_bunches = (
        pulse_length // bunch_T
    )  # Note that the number of bunches is determined by the pulse_length above
    print("Number of bunches:", n_bunches)

    pulse_delay = (T_beat / 4 - pulse_length / 2) // bunch_T * bunch_T
    t_delayed = np.arange(0, pulse_length, bunch_T) + pulse_delay

    print("Time between bunch trains:", 2 / (fb - fa) * 1e6, "us")
    print("Time of one bunch train:", pulse_length * 1e6, "us")
    print(
        "Percentage of pulse utilised: {0:.1f} %".format(
            pulse_length * 3 / (2 / (fb - fa)) * 100
        )
    )
```

```
    energy_gain = np.cos(2 * np.pi * fa * t_delayed + phi_a) + np.cos(
        2 * np.pi * fb * t_delayed + phi_b
    )

    alpha = 0.190091  # =(dL/L)/(dp/p)
    eta = alpha - 1 / gammaL**2
    print("Slippage factor is:", eta)
    dE = energy_gain * A_full
    dp = (dE) / 5.938e9 / betaL**2  # Corrected dp/p calculation,
    # from dp/p = 1/beta^2 * dE/E, with E=5+0.938 GeV the total energy

    dL = eta * dp  # dT/T = dL/L, here the variable dL is dL/L

    # Convert this into phase space rotation
    dT = dL * comp_circ / v  # change in time for a bunch over one turn

    n_turns_list = (
        np.array([0, 50, 100, 200]) * 1
    )  # This list can be used to plot the compression
    #                                           after various numbers of turns
    t_updated = t_delayed + dT * n_turns_list[0]  # time of the bunches after n turns

    def bunch_length(n_turns):
        return np.std(t_delayed + n_turns * dT)

    n_turns_opt = scopt.minimize_scalar(bunch_length)
    if n_turns_opt.success:
        n_turns_max = n_turns_opt.x
    else:
        raise TimeoutError("Could not find optimal turn number analytically")

    t_delayed -= t_delayed.mean()

    data = {
        "t_delayed": t_delayed,
        "energy_gain": energy_gain,
        "n_turns_max": n_turns_max,
        "pulse_delay": pulse_delay,
        "pulse_length_full": pulse_length_full,
        "T_beat": T_beat,
        "T_turn": T_turn,
        "f0": f0,
        "fa": fa,
        "fb": fb,
        "pulse_length": pulse_length
        * 3,  # Since three lines mean that the total pulse_length is
        #                                effectively thrice as long
        "current": current,
        "rep_rate": rep_rate,
    }

    return data
```

B Transverse Results

Although the key results in this thesis regard the longitudinal dynamics, as these correlate directly to the compression, the transverse dynamics of the bunches are also very important to understanding the properties of the compression. These are studied more extensively elsewhere for the baseline case, for example in [29]. Here, the key results are shown in Figure 31 at full intensity (left) and with 1×10^{13} protons (right), demonstrating emittance growth as expected from the tune shift which occurs due to space charge. Interestingly, a four-fold instability is seen in the vertical dimension, causing further emittance growth and a characteristic resonance pattern. This has also been seen in investigations of the baseline compression scheme with the same lattice, and found to be amplified, but not driven, by space charge [29]. This instability is sextupolar, and can thus be reduced by adding additional sextupoles in the lattice. This, and other issues in the transverse behaviour of the beam can be invstigated further, but seemingly do not present unresolvable issues for the realisation of either compression scheme.

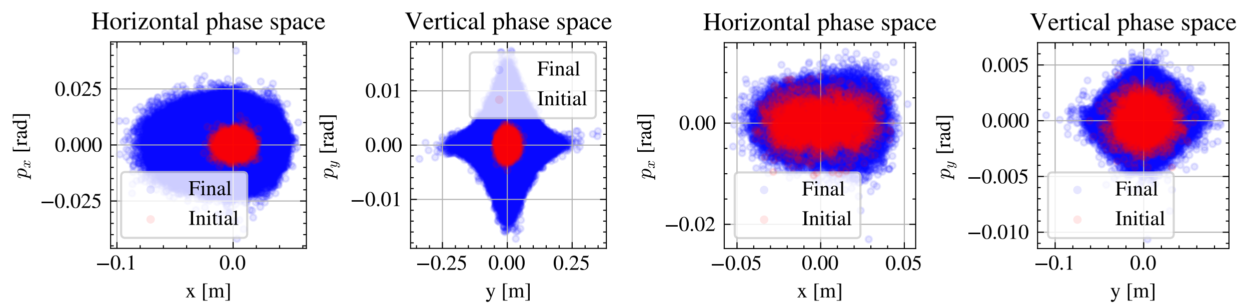

Figure 31: The transverse distribution of the simulated macroparticles in horizontal and vertical phase space at full intensity (left, 5×10^{14} protons) and with 1×10^{13} protons (right) after the number of turns required to reach maximum compression. The emittance growth is very substantial for the full intensity, which is due to the crossing of various resonances due to the space charge tune shift. Furthermore, a four-fold resonance pattern is observed in the vertical phase space, which cannot be attributed to space charge. In the lower intensity case, the emittance growth is smaller, as expected.

C Results Without Three Line Approach

Decreasing the frequency shift of the off-frequency cavities was done to reduce the impact of collective effects, as these often are proportional to the particle density or some positive power of it. Reducing it was seen qualitatively to improve the results noticeably, which can for example be seen in Figure 32 at a frequency shift of 0.94 MHz, i. e. the compressor ring revolution frequency, when compared to the result in Figure 23 with one third the frequency shift. The microbunching is much clearer in the former case, due to the larger particle density.

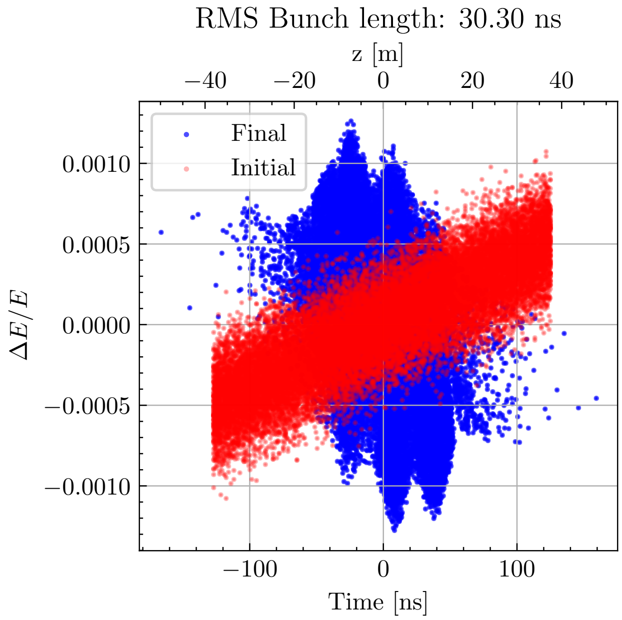

Figure 32: The compression of the first bunch train with space charge at a frequency shift of the revolution frequency, showing the difference to the three-line case with a smaller frequency shift.

D Linac Longitudinal Equations

By defining the energy deviation from the synchronous particle at phase $\phi_s = 0$ as $w = E - E_s$, and the phase slip as $\varphi = \phi - \phi_s = \phi$, the relative changes of these quantities can be calculated and a pair of differential equations emerge [14]:

$$\begin{cases} \frac{dw}{dz} = qV_{\text{rf}} \left(\cos(\phi_s + \varphi) - \cos(\phi_s)\right) \\ \frac{d\varphi}{dz} = -\frac{\omega_{\text{rf}}}{c} \frac{w}{E_0 \beta_s^3 \gamma_s^3} \end{cases} \tag{60}$$

for the rest energy E_0 of the particle, and the relativistic β and γ of the synchronous particle. These are reminiscent of the equations of a pendulum, and indeed when differentiating the second equation with respect to z, one finds a second order differential equation:

$$cE_0 \beta_s^3 \beta_s^3 \gamma_s^3 \frac{d^2\varphi}{dz^2} + cE_0 \frac{d(\beta_s^3 \gamma_s^3)}{dz} \frac{d\varphi}{dz} = -\omega_{\text{rf}} q V_{\text{rf}} \left(\cos(\phi_s + \varphi) - \cos(\phi_s)\right), \tag{61}$$

where the acceleration is typically adiabatic, i. e. slow enough that $\frac{d(\beta_s^3 \gamma_s^3)}{dz} \ll 1$ and the corresponding term can be neglected. This gives finally the second order pendulum equation:

$$cE_0 \beta_s^3 \gamma_s^3 \frac{d^2\varphi}{dz^2} = -\omega_{\text{rf}} q V_{\text{rf}} \left(\cos(\phi_s + \varphi) - \cos(\phi_s)\right). \tag{62}$$

For small deviations φ, the approximation $\cos(\phi_s + \varphi) - \cos(\phi_s) \simeq -\varphi \sin(\phi_s)$, and a linear dependence on φ has appeared in the right hand side as for a harmonic oscillator. The differential equation is now:

$$\frac{d^2\varphi}{dz} + q \frac{\omega_{\text{rf}} V_{\text{rf}} \sin(-\phi_s)}{cE_0 \beta_s^3 \gamma_s^3} \varphi = 0, \tag{63}$$

where the angular frequency of the solution is

$$\Omega_s = \sqrt{q \frac{\omega_{\text{rf}} V_{\text{rf}} \sin(-\phi_s)}{cE_0 \beta_s^3 \gamma_s^3}}. \tag{64}$$

The stable motion requires firstly that the energy gain is positive, which from Equation 12 means $\cos(\phi_s) > 0$. Secondly, the angular frequency in the harmonic oscillator must be real, giving the condition that $\sin(-\phi_s) > 0$, as all the other constants in Equation 64 are positive. This sets an allowed range of phases to

$$-\frac{\pi}{2} < \phi_s < 0. \tag{65}$$

E IPAC Paper

The following proceeding article was submitted to and accepted by the International Particle Accelerator Conference 2026 in Deauville, France.

INVESTIGATIONS OF A NOVEL BUNCH COMPRESSION TECHNIQUE FOR HADRON ACCELERATORS*

J. Holmberg† [1,2], N. Milas[2], F. Curbis[1,3], M. Eshraqi[2]
[1]Department of Physics, Lund University, Lund, Sweden
[2]European Spallation Source, Lund, Sweden
[3]MAX IV laboratory, Lund, Sweden

Abstract

Bunch compression is necessary for many accelerator applications, one of which being the creation of muons for the proposed muon collider. The design of a proton driver to deliver a short and very intense proton pulse to a target for the creation of muons is being developed. Some challenges with the proposed baseline design, notably the ambitious requirements for the radiofrequency system, motivate the investigation of alternative compression methods. A series of chirped bunch trains can be compressed in the same compressor ring as in the baseline design, removing the need for an accumulator ring and radiofrequency cavities in the compressor ring. Such a scheme has been investigated using PyORBIT simulations, where the chirp is created by off-frequency cavities at the end of the linac. The impact of space charge was investigated, and was found to be too detrimental for realising the proposed scheme with the current parameters. However, other use cases are foreseen.

INTRODUCTION

The proposed muon collider [1] has stringent requirements on the proton driver, which serve to ensure that the target produces enough muons in a short enough time to be captured and cooled to create useable bunches for the collider stage. An average power of 2 MW at a repetition rate of 5 Hz is required, which at the proton energy of 5 GeV corresponds to bunches with 5×10^{14} protons. The compression should achieve an RMS bunch length of 2 ns, which with the other parameters represents an unprecedented proton density, even though the baseline design divides the bunch in two [1, 2]. The proposed design which delivers this consists of a high intensity 5 GeV linear accelerator (linac), followed by an accumulator ring and a compressor ring, as shown in Fig. 1.

CHIRP CREATION

As the lattice for the compressor ring is designed to maximise the slippage factor for rapid compression, it can in principle fulfil its function even without radiofrequency (RF) cavities providing synchrotron motion in the bucket. A series of short bunches from the linac, which delivers 2 ps bunches, given an appropriate chirp will eventually stack on top of each other in the compressor ring. To create the required chirp, two cavities can be added at the end of the linac, operating at frequencies slightly shifted from the nominal frequency, $f = f_0 \pm \Delta f$.

The bunches will experience the superposition of the field in the two cavities, which has a beating envelope with the same frequency as the frequency shift Δf. As seen in Fig. 2, the bunches receive different energies and a chirp is created over a train of bunches. Chopping the source pulse from the linac appropriately, a series of such bunch trains with varying amplitude are injected directly into the compressor ring. A frequency shift which is an integer fraction of the revolution frequency in the ring ensures overlapping bunch trains. The bunch trains compress during the injection, but must all reach maximum compression simultaneously. Consequently, the final bunch train receives the largest chirp amplitude, and the others follow a $1/t$ dependence over the injection time t.

TOY MODEL

To identify potential parameters for the compression, a toy model was first developed which treated the compression analytically using the slippage factor from the compressor lattice. A frequency shift corresponding to the revolution frequency of the compressor ring was chosen, 0.94 MHz. Bunches are spaced at a frequency of 352.21 MHz from the linac, meaning around 100 bunches can be placed within the linear region of the envelope, and thus only around a quarter of the pulse can be utilised. The current from the linac is 80 mA, determining the total number of bunches to be 3.5×10^5 in order to reach the required intensity of 5×10^{14} protons. Consequently, the injection takes a total of 4122 turns. Based on the high-β cavities in the ESS linac, the energy gain in the off-frequency cavities was conservatively chosen to 10 MeV [3]. For the slippage factor of $\eta = 0.17$, this amplitude compresses the bunches in 248 turns according to

$$\frac{dT_{\rm rev}}{T_{\rm rev}} = \eta \frac{dp}{p}. \qquad (1)$$

The amplitude was increased from 0.6 MeV to 10 MeV across the injection, and the bunch trains are all simultaneously compressed after 4370 turns. For the baseline design, the compression in the compressor occurs in ~50 turns. In Table 1, the resulting parameters of the toy model are shown, and the compression after 4370 turns is shown in Fig. 3, where each bunch is treated as a single macroparticle.

Improving the Design

While the toy model parameters indeed achieve the required proton driver parameters, it is clear that collective

* Work supported by the European Union and endorsed by the IMCC
† johanholmberg01@gmail.com

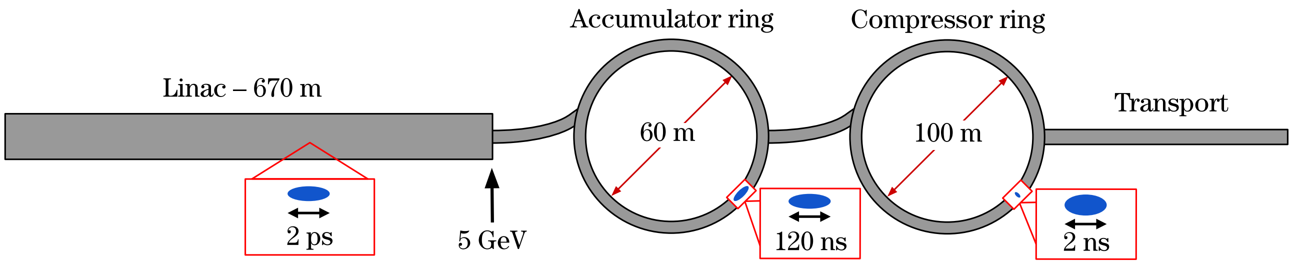

Figure 1: The baseline proton driver design [1, 2].

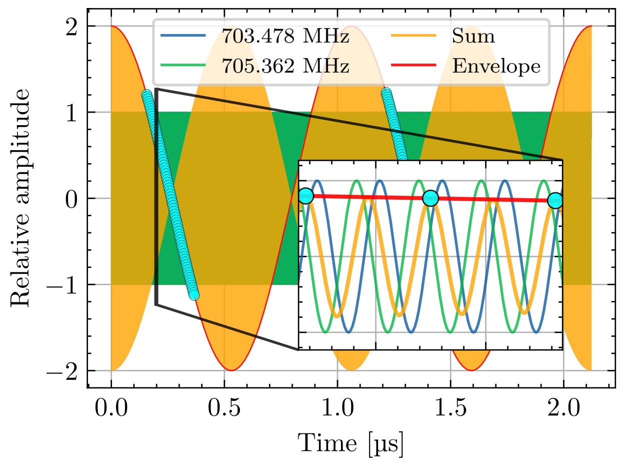

Figure 2: The superposition of the fields from the two off-frequency cavities, with the bunches in a bunch train receiving different energies.

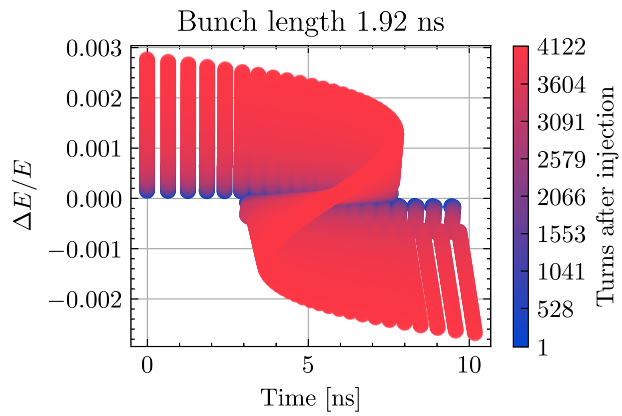

Figure 3: The compression of the 4122 bunch trains after 4370 turns as tracked by the toy model with only the slippage factor taken into account.

Table 1: Toy Model Parameters

Parameter	Value	Adjusted value
Amount of pulse kept	24 %	72 %
Pulse length	4.375 ms	1.393 ms
# turns injection	4122	1312
# turns compression	248	248
Amplitude	10 MeV	30 MeV
Final bunch length	1.92 ns	5.63 ns
Average power	2.10 MW	2.01 MW

effects will severely limit the performance of this compression scheme. As the compression time for the final bunch train can be shortened arbitrarily by adding pairs of cavities, the main limitation is the injection time. To shorten this, the chopped linac pulse can be sent to three separate lines with off-frequency cavities using fast kicker magnets. An ambitious but realistic rise and fall time of 300 ns for the magnets would not limit the number of bunches per bunch train, increasing the pulse usage by a factor three from 24 % to 72 %.The frequency shift can correspondingly be reduced by a factor of three so that one bunch train is still injected every turn, but now from each line. The total length of the bunch train is then also maximised w.r.t. the circumference of the ring, further reducing the strength of collective effects. This also has the consequence of increasing the bunch length in the toy model, as seen in Table 1, due to the increased size of the sinusoidal envelope. However, this effect is negligible compared to the impact from collective effects, as will be seen. The final design for the proposed scheme is shown schematically in Fig. 4.

SIMULATIONS

The compression was studied with particle-in-cell (PIC) simulations with PyORBIT [4]. Simulations were performed first without space charge to confirm the compressive behaviour and parameter choices, demonstrating good agreement with the analytical model. This indicates that non-linear effects from the lattice are not particularly strong. The non-zero energy spread introduced for a realistic bunch however increased the final bunch length considerably. The impact of energy spread on final bunch length can be given a lower bound using Liouville's theorem, requiring a spread of $\Delta E/E = 3 \times 10^{-5}$ for achieving the desired bunch length of 2 ns, an order of magnitude below the linac specifications.

This is however not the limiting factor for the final bunch length, as space charge simulations demonstrate. At design intensity, the sequentially injected beam starts experiencing an instability already after ~150 turns, as seen in Fig. 5. This instability then causes microbunching and violent growth in energy spread, leading to bunch deformation and particle loss. While the bunch length does decrease initially, the beam eventually breaks up before the ideal compression would be achieved after 1597 turns.

This behaviour is attributed to the *negative mass instability*, which is characterised by microbunching around local density variations. For a machine operated above transition, the space charge kicks from a region of increased parti-

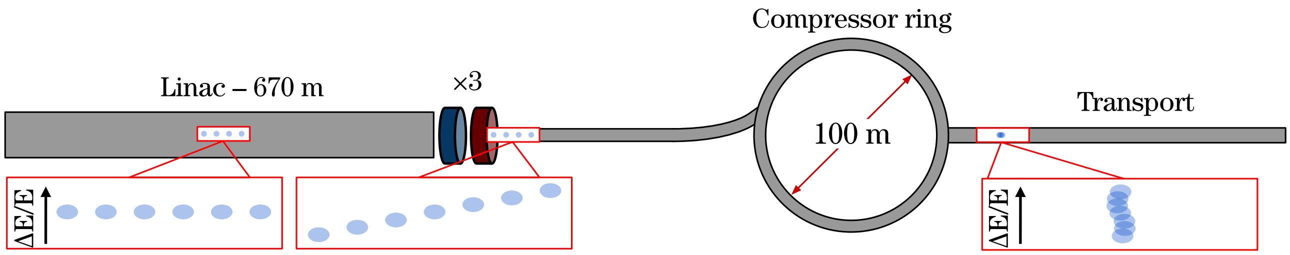

Figure 4: The proposed novel compression scheme, where a 5 GeV linac sends a chopped pulse to one of three lines with cavities shifted in frequency by $\Delta f = 314\,\mathrm{kHz}$ at different phases. The bunch trains are injected on top of each other in the compressor ring with one injection per turn.

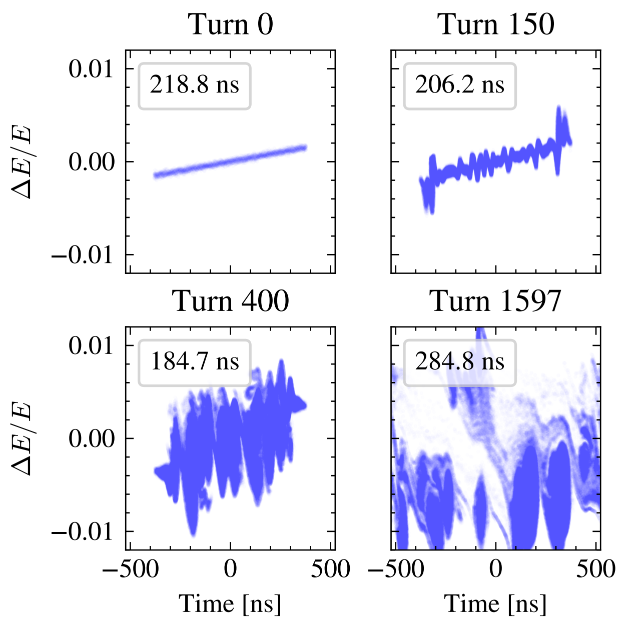

Figure 5: The longitudinal phase space distribution of the beam a certain number of turns after the first bunch train is injected, with the RMS bunch length in the top left.

cle density do not cause repulsion, but rather attraction as follows from Eq. 1 with $\eta > 0$. This increases the particle density and energy spread in these microbunches, with particles deviating enough in energy to be lost due to dispersion effects [5]. At the baseline intensity, only 23 % of the macroparticles are retained by the end of the simulation. Optimising the parameters for maximum retained intensity within realistic constraints, a transverse emittance of 20 mm mrad was chosen, with circular apertures of radius 35 cm placed around the ring. Due to the large energy spread from the chirp and microbunching, the dispersion in the ring contributes considerably to the horizontal beam size, and thus the required aperture. Using elliptical apertures to accommodate this did not help considerably.

Simulations were performed also at lower intensities, showing the same instability. The impact was smaller at reduced intensities, and it was found that with 1×10^{13} protons compression was achievable with more acceptable results. A bunch length of 21.8 ns was achieved with an injection over 26 turns and 297 turns of compression, without particle loss. Doubling the chirp reduced the bunch length to 12.8 ns after 180 turns.

Furthermore, an alternative lattice operating below compression has been cursorily investigated. In this case, η is limited by the particle energy to $\eta < 0.025$, meaning a larger amplitude chirp is required to achieve compression. However, these results, from an unoptimised lattice, show that the space charge repulsion is considerably stronger than the compression even at a 60 MeV chirp amplitude. Given the increased beam size with dispersion at larger amplitudes, such a lattice is not deemed feasible, in line with previous work [6]. A lattice with imaginary transition energy could similarly avoid the negative mass instability, while potentially having $\eta > 0.025$ and is thus an interesting object for future studies.

CONCLUSION

It has been shown that it is indeed possible to compress proton bunches with the proposed scheme based on chirp creation with off-frequency RF cavities outside the compressor ring. Without space charge effects, the simulated procedure follows expectations very well. However, a very small energy spread is required in the bunches to achieve nanosecond bunch lengths, as desired for the muon collider use case. With the introduction of space charge, the desired bunch intensities were shown to be unfeasible due to the strong space charge forces. In the baseline compressor lattice above transition these forces caused strong microbunching and deformation of the beam due to the negative mass instability. An alternative lattice below transition had a slippage factor too small to overcome the space charge repulsion at a reasonable chirp amplitude. The scheme was however demonstrated to achieve compression to 22 ns without particle losses with 1×10^{13} protons, an intensity relevant for other use cases, such as neutrino factories. An increase in repetition rate from the low 5 Hz of the muon collider is also easily achievable, limited only by the linac.

Various potential improvements to the scheme remain to investigate, for example adding weak bunching RF cavities to the compressor ring or using space charge compensation with impedance from Finemet cavities or ferrites.

REFERENCES

[1] C. Accettura *et al.*, "The Muon Collider", 2025, arXiv:2504.21417 [physics.acc-ph]. doi:10.48550/arXiv.2504.21417

[2] S. Johannesson *et al.*, "Initial design of a proton complex for the Muon Collider", in *Proc. IPAC'24*, Nashville, TN, May 2024, pp. 2528–2531. doi:10.18429/JACoW-IPAC2024-WEPR24

[3] S. Peggs *et al.*, "Conceptual design report", European Spallation Source, Lund, Sweden, Rep. ESS-2012-001, 2012, p. 135.

[4] S. Cousineau *et al.*, "The Particle Accelerator Simulation Code PyORBIT", in *Procedia Comput. Sci.*, vol. 51, pp. 1272–1281, 2015. doi:10.1016/j.procs.2015.05.312

[5] S. Y. Lee, "Synchrotron motion", in *Accelerator Physics*, Singapore: World Scientific, 4th edition, 2019, pp. 354–365. doi:10.1142/9789813274686_0003

[6] G. Rumolo and I. Hofmann, "A study of fast bunch rotation in the negative mass region", in *Proc. PAC 2001*, Chicago, IL, 2001. https://cds.cern.ch/record/511342

F IPAC Poster

The following poster was presented at the International Particle Accelerator Conference 2026 in Deauville, France.

Investigations of a Novel Bunch Compression Technique for Hadron Accelerators

J. Holmberg[1,2], N. Milas[2], F. Curbis[1,3], M. Eshraqi[2]

[1]Department of Physics, Lund University, Lund, Sweden [2]European Spallation Source, Lund, Sweden [3]MAX IV, Lund, Sweden

IPAC'26 May 17th - 22nd, 2026

ABSTRACT

Bunch compression is necessary for many accelerator applications, one of which being the creation of muons for the proposed muon collider. The design of a proton driver to deliver a short and very intense proton pulse to a target for the creation of muons is being developed. Some challenges with the proposed baseline design, notably the ambitious requirements for the radiofrequency system, motivate the investigation of alternative compression methods. A series of chirped bunch trains can be compressed in the same compressor ring as in the baseline design, removing the need for an accumulator ring and radiofrequency cavities in the compressor ring. Such a scheme has been investigated using PyORBIT simulations, where the chirp is created by off-frequency cavities at the end of the linac. The impact of space charge was investigated, and was found to be too detrimental for realising the proposed scheme with the current parameters.

INTRODUCTION

The proposed muon collider has stringent requirements on the proton driver [1]

- Average power of 2 MW
- Repetition rate of 5 Hz
- RMS bunch length of 2 ns
- Energy of 5 GeV

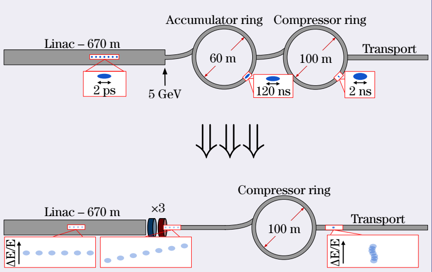

Figure 1: The baseline compression scheme for the proton driver, as compared to the novel one proposed in this thesis.

NOVEL COMPRESSION SCHEME

- Avoids the high-voltage RF in compressor
- Foregoes accumulator ring
- Uses chirp from off-frequency cavities to compress
- Same compressor ring

$$\frac{dT_{\text{rev}}}{T_{\text{rev}}} = \eta \frac{dp}{p}. \quad (1)$$

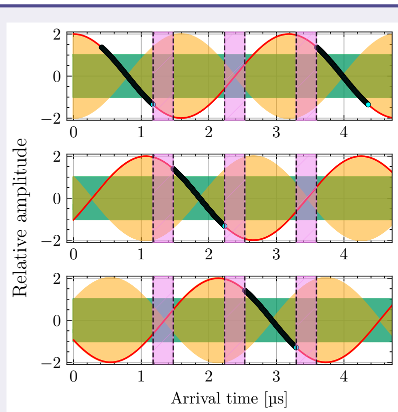

Figure 2: The relative phases of the cavity fields in the three separate lines.

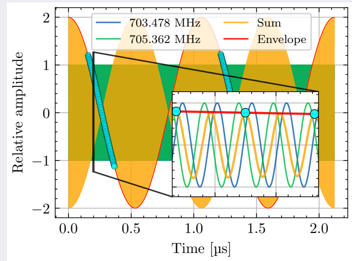

Figure 3: The superposition of the fields from the two off-frequency cavities, with the bunches in a bunch train receiving different energies.

PARAMETER IDENTIFICATION

An analytical toy model was developed to find the optimal parameters for the compression.

- Slippage factor $\eta = 0.17$
- One bunch = one macroparticle

Table 1: Toy Model Parameters

Parameter	Initial value	Adjusted value
Bunches per bunch train	89	89×3
Pulse length	4.375 ms	1.393 ms
# turns injection	4122	1312
# turns compression	248	248
Amplitude	10 MeV	30 MeV
Final bunch length	1.92 ns	5.63 ns
Average power	2.10 MW	2.01 MW

Improving Design

However, the impact of collective effects is expected to be very strong, and thus some modifications were made to the initial parameters in Table 1.

- Divide linac pulse in three lines to reduce injection time
- Frequency shift $\Delta f = 314$ kHz for longer bunch train
- Reduced injection time & long train $\implies$ weaker collective effects
- Longer bunch train means longer final bunch length due to Liouville

However, the net effect when space charge effects are taken into account is to reduce the final bunch length.

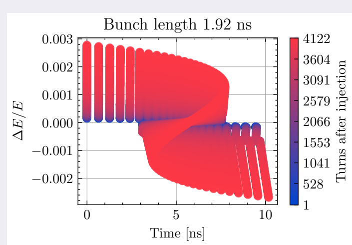

Figure 4: Analytical compression of the full series of bunch trains with the initial parameters tracked with the toy model, showing compression to the required parameters is possible.

SIMULATION RESULTS

The compression was studied with PIC simulations using PyORBIT [2].

Simulations without space charge showed:

- Turns to compress agreed with toy model
- Energy spread causes debunching

$\implies$ Liouville limits final bunch length, $\Delta E/E = 3 \times 10^{-5}$ required to reach 2 ns (factor 10 lower than linac)

With space charge, very strong instabilities were observed, leading to:

- Microbunching
- Beam break-up
- Particle loss

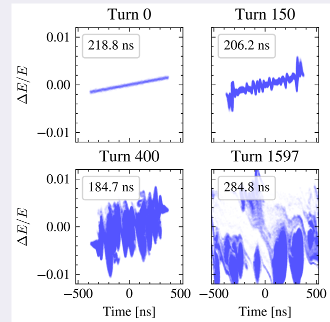

Figure 5: The longitudinal phase space distribution of the beam a certain number of turns after the first bunch train is injected, with the RMS bunch length in the top left.

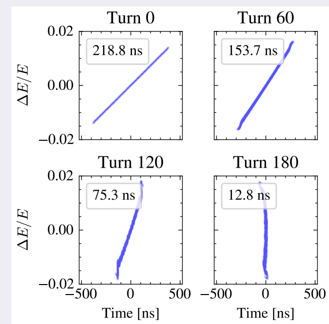

Figure 6: The compression of a beam with 1×10^{13} protons at an increased chirp amplitude of 60 MeV, with the RMS bunch length in the top left.

ANALYSIS

This behaviour is attributed to the *negative mass instability*, which is characterised by microbunching around local density variations. For a machine operated above transition, the space charge kicks from a region of increased particle density do not cause repulsion, but rather attraction. This leads to rapid compression, but unstable motion if the force is too strong [3].

If the bunch intensity is reduced, the results are considerably better. Increasing the repetition rate is a possible compensation for this. For another application, e.g. a neutrino factory, this scheme could be attractive.

Table 2: Intensity Scan Results at 60 MeV

Protons	Turns to compress	Particles left [%]	Bunch length
1×10^{13}	180	100 %	13 ns
2.5×10^{13}	210	100 %	17 ns
5×10^{13}	255	94 %	27 ns
1×10^{14}	360	61 %	50 ns
2.5×10^{14}	700	35 %	130 ns
5×10^{14}	1597	23 %	N/A

CONCLUSIONS

- Compression can be achieved using chrip creation with off-frequency cavities
- Collective effects, particularly space charge and the negative mass instability, severely degrade the performance
- Compression scheme not usable for muon collider
- However, parameters align well with neutrino factory, where repetition rate can be increased

REFERENCES

[1] C. Accettura *et al.*, "The Muon Collider", *arXiv:2504.21417*, 2025.

[2] S. Cousineau *et al.*, "The Particle Accelerator Simulation Code PyORBIT", in *Procedia Comput. Sci.*, vol. 51, 2015, pp. 1272–1281.

[3] G. Rumolo and I. Hofmann, "A study of fast bunch rotation in the negative mass region", in *Proc. PAC 2001*, Chicago, IL, 2001.

Funded by the European Union

Funded by the European Union (EU). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the EU or European Research Executive Agency (REA). Neither the EU nor the REA can be held responsible for them.