**Raul António
Janeiro Costa**

Simulação e otimização do complexo de recombinação do CLIC

Simulation and optimisation of CLIC's recombination complex

**Raul António
Janeiro Costa**

Simulação e otimização do complexo de recombinação do CLIC

Simulation and optimisation of CLIC's recombination complex

o júri / the jury

presidente / president

Professor Doutor Vítor Hugo da Rosa Bonifácio
Assistant Professor at University of Aveiro's Physics Department

vogais / examiners committee

Doutor Andrea Latina
Permanent Staff Member at the European Organization for Nuclear Research (CERN)

Professora Doutora Ángeles Faus-Golfe
Senior Researcher at the Laboratoire de l'Accélérateur Linéaire (LAL)

**agradecimentos /
acknowledgements**

Palavras Chave
CLIC, *drive beam*, complexo de recombinação, anel de combinação, dinâmica de feixes, Placet2.

Resumo
Nesta dissertação apresentamos as primeiras simulações de recombinação com o software Placet2 do complexo de recombinação do feixe de aceleração (DBRC) do colisor linear compacto (CLIC). Começamos por apresentar uma revisão do projeto CLIC e do papel e design do DBRC neste projeto. Continuamos discutindo alguns princípios básicos de dinâmica de feixes e a forma como programas de seguimento como o Placet2 os implementam. De seguida, apresentamos os possíveis problemas de design levantadas pelas nossas simulações e a estratégia que propomos para os superar. A principal descoberta é uma correlação parabolica da posição logitudinal com o momento (T_{566}), que ameaça a eficiência das estruturas de extração de potência. Através de otimização iterativa do design, esta aberração foi eliminada no circuito de atraso e no anel de combinação 1. Também descobrimos que a emitância horizontal do feixe se encontra significativamente acima do orçamento ($150\,\mu\mathrm{m}$) e tentámos ir de encontro a este, reduzindo-a para $157\,\mu\mathrm{m}$. Para obter este valor de emitância, foi necessário atualizar o esquema de injeção do anel de combinação 2. No plano vertical, que tem o mesmo orçamento de emitância, esta foi mantida a $127\,\mu\mathrm{m}$.

Keywords

Abstract

In this thesis we present the first Placet2 recombination simulations of the drive beam recombination complex (DBRC) design for the compact linear collider (CLIC). We start by presenting a review of the CLIC project and the DBRC's role and design within it. We then discuss some of the core principles of beam dynamics and how tracking codes like Placet2 implement them. We follow that by presenting the design issues raised by our simulations and our proposed strategy to address them, key among which is a previously unknown parabolic dependency of the longitudinal position to the momentum (T_{566}), which threatens the efficiency of the power extraction structures. Through iterative optimisation of the design, we eliminated this aberration both in the delay loop and in combiner ring 1. We also found the beam's horizontal emittance to be significantly over the design budget ($150\,\mu\mathrm{m}$) and attempted to meet that budget, reaching $157\,\mu\mathrm{m}$. In order to obtain this emittance value, an update to the combiner ring 2's injection scheme was necessary. On the vertical plane, which has the same emittance budget, it was kept at $127\,\mu\mathrm{m}$.

Nomenclature

Abbreviations

API   Application Programming Interface

ATLAS   A Toroidal LHC Apparatus

BPM   Beam Position Monitor

CERN   European Organization for Nuclear Research

CLIC   Compact Linear Collider

CMS   Compact Muon Solenoid

CR1   Combiner Ring 1

CR2   Combiner Ring 2

CSR   Coherent Synchrotron Radiation

CTF3   CLIC Test Facility 3

DBA   Drive Beam Accelerator

DBRC   CLIC Drive Beam Recombination Complex

DB   CLIC Drive Beam

DF   Dispersion-free

DL   Delay Loop

eRHIC   Brookhaven National Laboratory Electron-Ion Collider

FCC   Future Circular Collider

FODO   Focussing-Defocussing cell

ICG cell Isochronous Chasman–Green cell

ILC   International Linear Collider

ISR   Incoherent Synchrotron Radiation

LEP   Large Electron-Positron Collider

LHC	Large Hadron Collider	
LHeC	Large Hadron Electron Collider	
linac	Linear Particle Accelerator	
PERLE	Powerful Energy Recovery Linac for Experiments	
PETS	Power Extraction Structures	
RF	Radiofrequency	
SLAC	Stanford Linear Accelerator Center	
TL1	Transfer Line 1	
TL2	Transfer Line 2	
TL3	Transfer Line 3	
TTA	Transfer Turn Around	
XFEL	X-ray Free Electron Laser	

Physics Properties

$\alpha_{x/y}$	Horizontal/Vertical Courant-Snyder alpha function	$-$
β_{rel}	Relativistic β	$-$
$\beta_{x/y}$	Horizontal/Vertical Courant-Snyder beta function	m
δ	Energy spread	%
γ_{rel}	Lorentz factor	$-$
$\gamma_{x/y}$	Horizontal/Vertical Courant-Snyder gamma function	m^{-1}
ω_{RF}	PETS frequency	GHz
ρ_0	Bending radius	m
σ_z	Bunch length	mm
$\sigma_{(\mathrm{A}\to\mathrm{B})}$	$(\mathrm{A}\to\mathrm{B})$ Event cross section	fb
$\sigma_{x/y}$	Horizontal/vertical bunch rms size	mm
$\sqrt{s}$	Collision centre of mass energy	GeV
θ	Bending angle	rad
$\tilde{\varepsilon}_{x/y}$	Pulse horizontal/vertical normalized emittance	μm
ε_0	Vacuum permittivity	$8.854\times10^{-12}\,\mathrm{Fm}^{-1}$

$\varepsilon_{g_{x/y}}$	Bunch horizontal/vertical geometric emittance	μm
$\varepsilon_{x/y}$	Bunch horizontal/vertical normalized emittance	μm
c	Speed of light	$2.998 \times 10^8 \, \text{ms}^{-1}$
E	Particle energy	GeV
e	Elementary charge	$1.602 \times 10^{-19} \, \text{C}$
E_0	Optics reference energy	GeV
f_{b}	Bunch frequency	GHz
$F_b(\lambda(z))$	PETS form factor	$-$
G	Magnetic gradient	Tm^{-1}
I	Beam current	A
K_1	Quadrupolar focal strength	$1/\text{m}^2$
L	Orbital length	m
m	Particle mass	GeV/c^2
P	Particle momentum	GeV
P_0	Reference particle momentum	GeV
Q	Beam charge	C
R_{56}	Longitudinal energy chirp	cm
t	Time	ns
x'/y'	Horizontal/vertical divergence	rad
x/y	Horizontal/vertical transverse position	mm
z	Longitudinal position	mm

Contents

Chapter 1

Introduction

This work presents the first tracking studies of the drive beam recombination complex (DBRC) to simulate multi-bunch recombination. In this chapter we briefly establish the motivation for our study. Section 1.1 offers a review of the impact of accelerator physics in modern science. Section 1.2 presents the compact linear collider (CLIC) [1] project and its scientific targets. Finally, in Section 1.3 we discuss how the novel two-beam acceleration scheme is used to power CLIC. Chapter 2 presents the DBRC's design and role within the CLIC project. Chapter 3 offers the reader a quick review of some beam dynamics concepts. Chapter 4 explores the computational tools used for our studies. In Chapter 5 we present the results of our simulations, the issues they raised, our optimisation scans and the design updates we propose. Lastly, Chapter 6 presents our conclusions and offers a brief outlook of the next development steps for this machine.

1.1   Accelerator physics

The development of modern particle accelerators started in the 1920s. At the time, there was even a public call by Ernest Rutherford in the Royal Society to increase the energy of particle beams for future experiments [2]. Earlier nuclear physics experiments, including Rutherford's gold foil experiment, relied on radioactive decay as the source. This technique is, however, very limited in beam energy and intensity. On the other hand, the evolution of particle accelerator design has been instrumental for most nuclear and particle physics breakthroughs in the last century. Particle accelerators gave us the quark in 1969 [3][1], the τ in 1975 [4], the W and Z bosons in 1983 [5, 6] and the Higgs boson in 2012 [7, 8], thus being responsible for the observation of most elementary particles.

Over the past few decades, no area of research has both stimulated and required more international cooperation than high energy physics. This is, at least partially, a consequence of the shear scale and complexity of the accelerator experiments required to push our knowledge further in this field. When SLAC [9] gave us the quark, its two-mile accelerator was nicknamed "The Monster". Nowadays, The Monster seems tamed when compared with the 27 km LHC [10]. Future accelerators like CLIC [1], the FCC [11] or the ILC [12] will push the scale of these facilities even further.

---

[1]Only the up and down quarks were observed in 1969, but most other quarks were also first observed in accelerator experiments.

Figure 1.1: Higgs production. On the left, the Feynman diagrams for the three main Higgs production channels. On the right, the cross section dependency of centre-mass energy ($\sqrt{s}$) for the various Higgs channels. Picture original from [19].

In addition to high energy physics, recent years have seen an increase in large scale accelerators used as x-ray and ultra-violet light sources (a recent example being the European XFEL [13]). As for small accelerators, they are now part of a long list of applications from package sterilization to ion implantation, medical imaging, cancer therapy and isotope production[2].

1.2 The compact linear collider

At the time of this work, the LHC Run 2 (in which the centre-mass energy was recently increased to 13 TeV) has found no new elemental particles. Additionally, the most promising recent excess, a 750 GeV di-photon signal observed both by ATLAS and CMS [15, 16], suffered a considerable reduction of its statistical significance in the data accumulated since 2016 [17, 18]. If no new particles are found through direct discovery, high energy physics must switch focus to precision measurements. For this purpose, the LHC has a significant disadvantage due to the background-rich nature of hadron collisions. This would be less of a concern in electron-positron colliders since matter-antimatter annihilation presents very little background.

CERN's next large scale electron-positron accelerator is CLIC. Unlike the LHC, CLIC will be a linear accelerator rather than a synchrotron ring. The 27 km tunnel currently housing the LHC was originally built for LEP [20] but, despite its size, LEP only achieved centre-mass energy ($\sqrt{s}$) of 209 GeV, while the LHC reaches 13 TeV in the same circumference. The huge energy disparity for two similar-sized rings is due to synchrotron radiation. This effect causes energy loss as the beam trajectory is bent, setting the maximum energy of LEP to the equilibrium between energy gained through acceleration and lost in the ring's bending

---

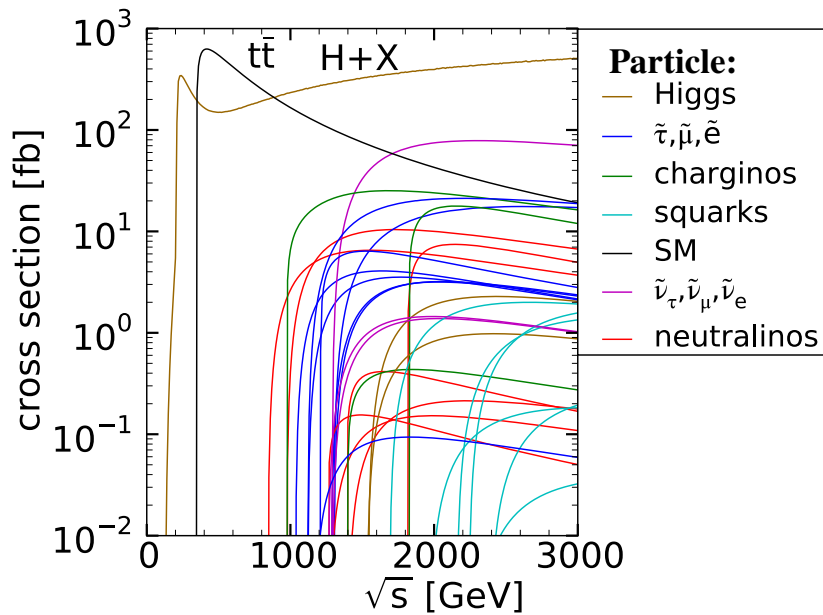[2]A good overview of industrial applications for accelerators can be found in [14].

Figure 1.2: On the horizontal axis, centre-mass energy ($\sqrt{s}$). On the vertical axis, cross sections for $t\bar{t}$, $H + X$ and supersymmetric particle production for a supersymmetry benchmark model. Picture original from [21].

magnets. This effect (that we will explore further in Section 3.5) is proportional to m^{-4} making the energy loss of proton beams thirteen orders of magnitude lower than the electron's for a similar circumference.

In addition to reducing energy loss, a linear design has the advantage of allowing a staged construction plan. CLIC is being developed has a three-stage project, in between which the accelerator and respective tunnel will be extended to increase $\sqrt{s}$. This will allow us to obtain the first physics results (at a lower energy) before extending the machine to its intermediate and full sizes. In the first stage, at 380 GeV, CLIC will be focused on Higgs and top physics. This energy will allow us to probe three Higgs production channels (see the diagrams in Figure 1.1) not available to us in the LHC. Of these processes, the one represented on the top diagram (called Higgsstrahlung) is of particular interest as it gives us a model-independent Higgs production measurement. Regarding top physics, the main focus will be in performing a precise measurement of the top's mass as the uncertainty of this quantity is currently the main source of ambiguity regarding the standard model's vacuum stability [19].

Additionally to the three channels represented by the diagrams, we can also see in Figure 1.1 three channels that will only be available at higher energies. These will be explored in CLIC's second and third stage, which will have centre-mass energies of 1.5 TeV and 3 TeV, respectively, allowing us to probe double-Higgs production channels and measure several previously unexplored coupling constants (the Higgs self-coupling for example). Additionally, these stages will explore beyond standard model theories, allowing us to either find new particles or put constrains on their production. An example of such models is shown in Figure 1.2. In the figure we can see the various supersymmetric particle cross sections at $\sqrt{s} > 1$ TeV. This model is just one among many that predict BSM particles observable at this energy range, making CLIC an ideal collider to study them.

Figure 1.3: Two beam acceleration scheme. Picture originally from [22].

1.3   Two-beam acceleration

While a linear design eliminates most energy loss through synchrotron radiation, it also implies that all acceleration must be done in a single passage through the machine. A large acceleration gradient is therefore necessary to keep the linac size compact. After extensive cost and performance studies, CLIC's accelerating gradient was set to $100\,\mathrm{MV/m}$, which allows the machine to reach $\sqrt{s} = 3\,\mathrm{TeV}$ in under $50\,\mathrm{km}$. This disqualifies superconductive cavities due to the critical magnetic field, above which superconductive material stop behaving as such.

A second concern regarding the main linac is power efficiency. There is always some power loss inherent to an accelerating structure operation. Adding to this, accelerating structures are traditionally powered by klystrons, which have their own power losses both directly in the klystron tube and due to power modulation. Using the ILC as an example, the overall power efficiency of the linac design is about 9% [12]. As for CLIC, the $12\,\mathrm{GHz}$ accelerating cavities would pose a serious challenge in the design of X-band klystrons and modulators[3].

As an alternative to the klystron approach, the CLIC project proposes a novel two-beam acceleration scheme. This approach uses a second, higher current, electron beam (the drive beam) to power the low emittance (main) beam which will be collided in the interaction point. This implies having both beams running in parallel as shown in Figure 1.3. While the main beam is being accelerated, the drive beam is being decelerated by power extraction structures (PETS). It is RF field produced in theses PETS that, through RF wave-guides, powers the main beam's X-band accelerating structures. This scheme, though more complex in nature (see Figure 1.4), was proven to be much more efficient than a klystron-based alternative.

---

[3]The choice of a $12\,\mathrm{GHz}$ frequency in the accelerating structures derives from a compromise between power efficiency and low breakdown rate. More details can be found in [1].

Figure 1.4: CLIC accelerator complex. Picture originally from [19].

Since the focus of our studies is the drive beam, it is useful to have in mind the drive beam requirements at the deceleration sections. As mentioned before, the PETS are designed to operate at a 12 GHz bunch frequency. A 100 MV/m accelerating gradient implies a 100 A beam current. Using

$$I = \frac{Q}{t} \; , \tag{1.1}$$

that translates to 8.4 nC (or 5.2×10^{10} electrons) per bunch. The length of a pulse is dictated by the main beam pulse (156 ns) and the accelerating structure loading time, amounting to 244 ns. As for bunch length (σ_z), previous studies [1, 22] showed that the field generated in a PETS is highly dependent on it. In fact, the field produced is proportional to a form factor

$$F_b(\lambda(z)) = \int_{-\infty}^{\infty} dz\, \lambda(z) \cos\left(\frac{\omega_{\mathrm{RF}}}{c} z\right) \; . \tag{1.2}$$

Assuming a Gaussian longitudinal bunch distribution, Eq. 1.2 can re-written as

$$F_b(\lambda(\sigma_z)) = \exp\left(-\frac{\sigma_z^2 \omega_{\mathrm{RF}}^2}{2c^2}\right) \; , \tag{1.3}$$

in which σ_z is the only beam parameter present (ω_{RF} is the PETS' frequency). It is therefore imperative that σ_z does not deviate from the design's requirements (1 mm).

Chapter 2

The drive beam recombination complex

In this chapter we present the design of the DBRC and its various sections. Section 2.1 offers a review of the complex's place in the CLIC project. Section 2.2 details the machine's injection parameters. In Section 2.3 we establish the notation we use in this work to classify different bunches based on their orbital path. Sections 2.4, 2.5 and 2.6 present the three recombination sectors of the beamline. Lastly, in Section 2.7 we discuss the machine's longitudinal dynamics requirements.

Considering CLIC is a project under development, the reader should not be surprised that its design is regularly updated as new studies [23–26] identify design issues and solutions. Like those before, the design updates proposed by this work are part of an iterative process, rather than the final form of the project. A full list of list of changes in the design proposed by our study can be found in Appendix A and an element by element lattice is available in [27].

2.1   Function and parameters

Let us start by presenting the need for the complex. As discussed in Section 1.3, in order to achieve an accelerator gradient of $100\,\mathrm{MV/m}$ for the main beam, we need to supply the PETS with a 100 A pulse. However, building a 100 A electron gun for the drive beam is not feasible. Instead, the drive beam gun is designed to produce 8.4 nC bunches at a frequency of 499.75 MHz which only amounts to a 4.2 A beam current.

To reach the necessary current, the drive beam gun produces trains much longer than what is required by the PETS. We then compress a $5.86\,\mu s$ train by a factor 24 into a 244 ns pulse to achieve the needed 100 A current[1]. This operation is the main function of the drive beam recombination complex. The complex is located between the drive beam accelerator (DBA) and the the long transfer lines leading to the PETS (see Figure 2.1). Figure 2.2 shows us its various recombination steps. The delay loop creates a 244 ns pulsed structure, increasing the bunch frequency from 500 MHz to 1 GHz. Interlocking multiple pulses, combiner rings one and two will further increase the frequency by a factor 3 and 4 respectively, increasing it to 3 GHz and then 12 GHz.

---

[1]The full drive beam train is actually much longer than $5.86\,\mu s$ since it needs to generate multiple pulses for multiple decelerating sections (25 on the 3 TeV stage).
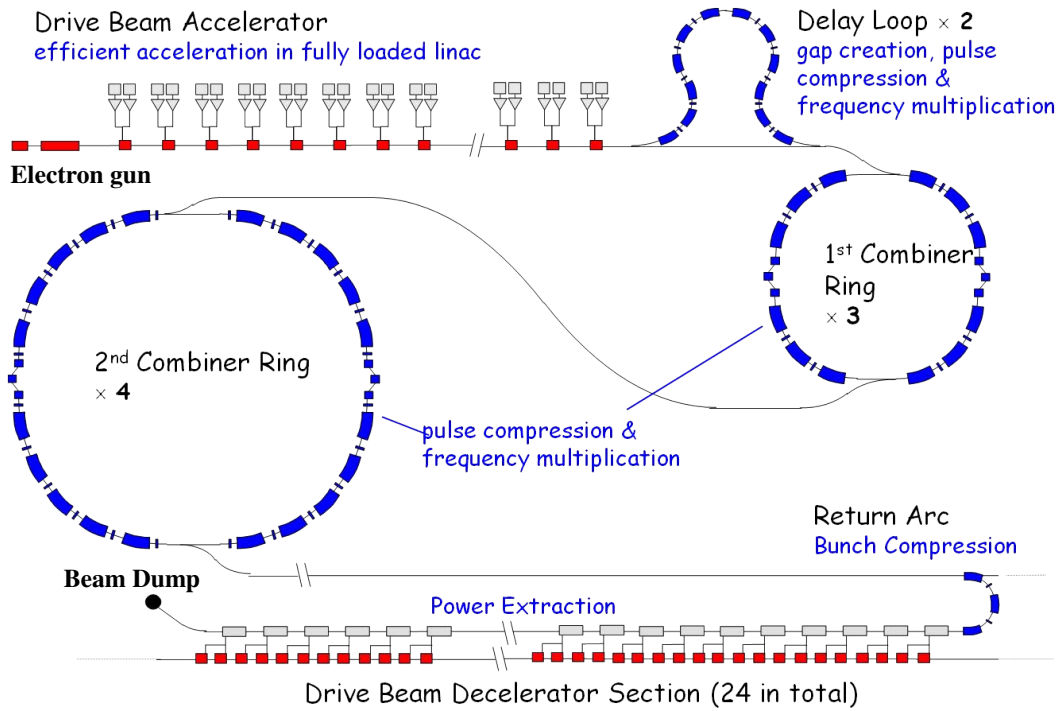
Figure 2.1: Schematic representation of the drive beam complex [1]. From the electron gun to one of the beam dumps, passing through the drive beam accelerator, the DBRC (composed by the delay loop and the two combiner rings), the long transfer line to the return arc and the PETS in the deceleration sectors.

2.2 DBRC injection

Let us now discuss the properties of the beam injected into the DBRC (see Table 2.1). The drive beam is injected into the DBRC after being accelerated to 2.38 GeV in the DBA. The most notable features of the injected beam are the existence of an longitudinal energy chirp for compression/decompression (see Section 2.7) and a π phase shift every 122 bunches.

The start of the recombination process is the formation of 244 ns subtrains at the delay loop (see Section 2.4). This operation depends on a periodic π phase shift of the drive beam train once every 122 bunches (see Figure 2.3). Since both odd and even bunches are equally accelerated, the DBA operates at twice the frequency of the drive beam train, as denoted by the sinusoidal curve in the figure.

Figure 2.2: Train structure of in the various stages of recombination.

Table 2.1: DBRC injection parameters.

Beam parameter	Unit	DBRC Injection
Bunch frequency (f_b)	MHz	499.75
Energy (E)	GeV	2.38
Energy spread (δ)	%	0.85
Charge (Q)	nC	8.4
Bunch length (σ_z)	mm	1.0
Energy chirp (R_{56})	cm	12
Horizontal emittance (ε_x)	μm	100
Vertical emittance (ε_y)	μm	100

2.3   Bunch notation

By now it should be clear to the reader that, if 24 subtrains are being recombined, there are 24 possible paths through the complex. As we will see in Chapter 5, the properties of a bunch after extraction vary considerably depending on which path it took through the complex. It is therefore useful when describing our results to classify an individual bunch by its subtrain. However, the order of bunches in the train changes dramatically throughout the machine, which makes classifying them by the order they entered or exited the complex unclear in different sections of the machine.

   With this in mind, in our results we classify bunches as $b_i^{\ j}$, where the lower index indicates the number of turns the bunch took in CR1 and the upper index the number of turns taken in CR2. Our studies do not simulate bunches going through the shortest path of the DL, so we only differentiate 12 bunch orbits with out notation. $l_i^{\ j}$ and $s_i^{\ j}$ could be used to classify the DL path in future studies. The correspondence between this notation and the injection order can be seen in Figure 2.4, where we classify 12 bunches by 244 bunches pulse, recombine them in CR1 and CR2 to show the extraction order, and then compare it with our notation. Notice that the b #i numbering at injection classifies the bunch by subtrain. There
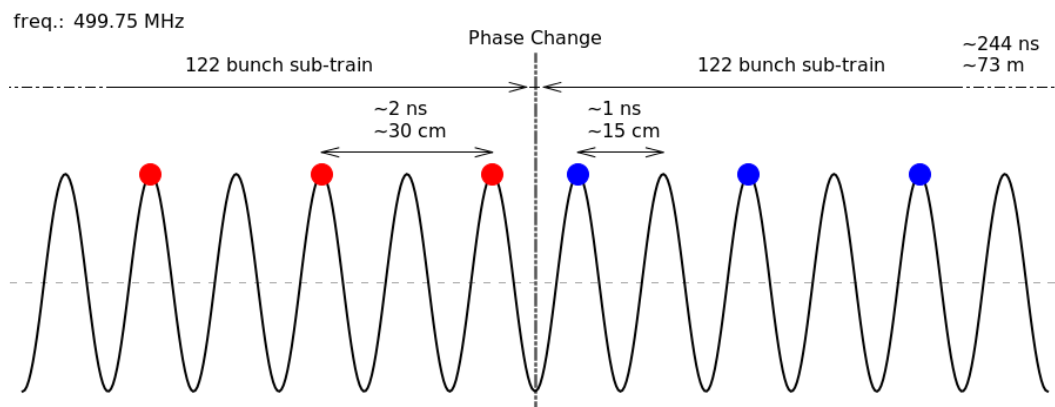

Figure 2.3: Drive beam time structure. The 73444 bunches of the 499.75 MHz DB train suffer a π phase shift once every subtrain ($\sim$ 244 ns). The DBA operates at 999.5 MHz (represented by the sinusoidal curve) in order to accelerate both odd (blue) and even (red) subtrains equally.

Figure 2.4: Bunch notation. The figure shows the change in bunch order as the trains are interlocked in both combiner rings. on the bottom of the figure we relate the bunch order to the notation introduced in Section 2.3.

are 243 bunches in between b#3 and b#4[2]. This is no longer the case at extraction, where b#1 and b#4 are only 83 ps apart.

2.4 Delay loop

The first step in recombination is to switch the DB from a continuous to a pulsed train. To achieve this, the delay loop splits the train into two beamlines, one of which is approximately twice the length of the other. In order to interlock both subtrains, the length of the DL's long path is given by

$$L_{\text{long}} = 2L_{\text{short}} \pm \frac{c}{2f_{\text{b}}} \,, \tag{2.1}$$

where the length of the short path (L_{short}) needs to be at least long enough to hold a full pulse (73 m). Since little issues are expected from the short path, its lattice is yet to be designed and our studies only track bunches that travel through the 143 m omega-shaped long path.

As shown in Figure 2.5, DL operations rely on the ability to separate odd and even subtrains with a 0.5 GHz transverse deflector. The first deflector is synchronised to send even bunches to the long path which, in turn, means odd bunches are sent through the short path. This is the reason why the phase shift presented in Section 2.2 is necessary. Approximately 488 ns later, both subtrains reach the second deflector with different angles (top right of the picture). Again, the phase difference between subtrains will make them reach the deflector in such a way that it will counteract their entry angle and bring both subtrains interlocked to a centre orbit.

---

[2]The reason why b#2 and b#13 do not appear in the scheme will become apparent in Section 2.5.

Figure 2.5: Schematic representation of DL loop operations [21]. On the top left, we can see how to the 0.5 GHz transverse deflector separates the odd and even subtrains. Both subtrains then travel through the DL beamlines, meeting on the second transverse deflector which recombines them into one pulse (top right).

Since it is a design requirement for the DL to preserve σ_z despite the beam having a strong energy chirp ($R_{56} = 12$ cm), the arc cells used in this beamline are the isochronous Chasman-Green (ICG) cells presented in [28]. A modified double-cell version of these was developed in [26] to serve as arc cells for the combiner rings.

2.5   Combiner ring 1

Once we have a pulsed train, we can recombine multiple subtrains using combiner rings. These allow us to interlock multiple pulses by having a length precisely adjusted to the pulse length and multiplication factor (we explore this further in Section 5.4).

CR1 is a six-arc ring designed to combine three 244 ns pulses with a frequency of 1 GHZ into one, increasing f_b to 3 GHz. In Figure 2.6 we see the ring's working cycle. Following the orbit of the first pulse (tr_1), we see that it takes a full turn just in time to be interlocked with tr_3. After taking a full turn, the resulting pulse ($tr_{1,3}$) is combined with tr_5 before being extracted at the opposite side of ring. Notice that, in order for the recombination to form a periodic pulse, the phase difference between tr_n and tr_{n+2} at the injection point must be $\pm\frac{2}{3}\pi$. This defines the total length of the ring (Eq. 5.5).

In order to diminish bending angles in CR1's sbends, the ring was designed to recombine two pulses simultaneously. Notice that if we had combined pulses tr_2, tr_4 and tr_6 and extracted them, the time between pulses would not be constant after extraction. We need therefore to recombine tr_4, tr_6 and tr_8. This means that the second and second to last pulse of the full cycle[3] is not used for power extraction.

---

[3]Figure 2.6 only represents the first few pulses of the DB. At 3 Tev, the DB train will be composed by 301 pulses, so the lost of two has limited significance to efficiency.
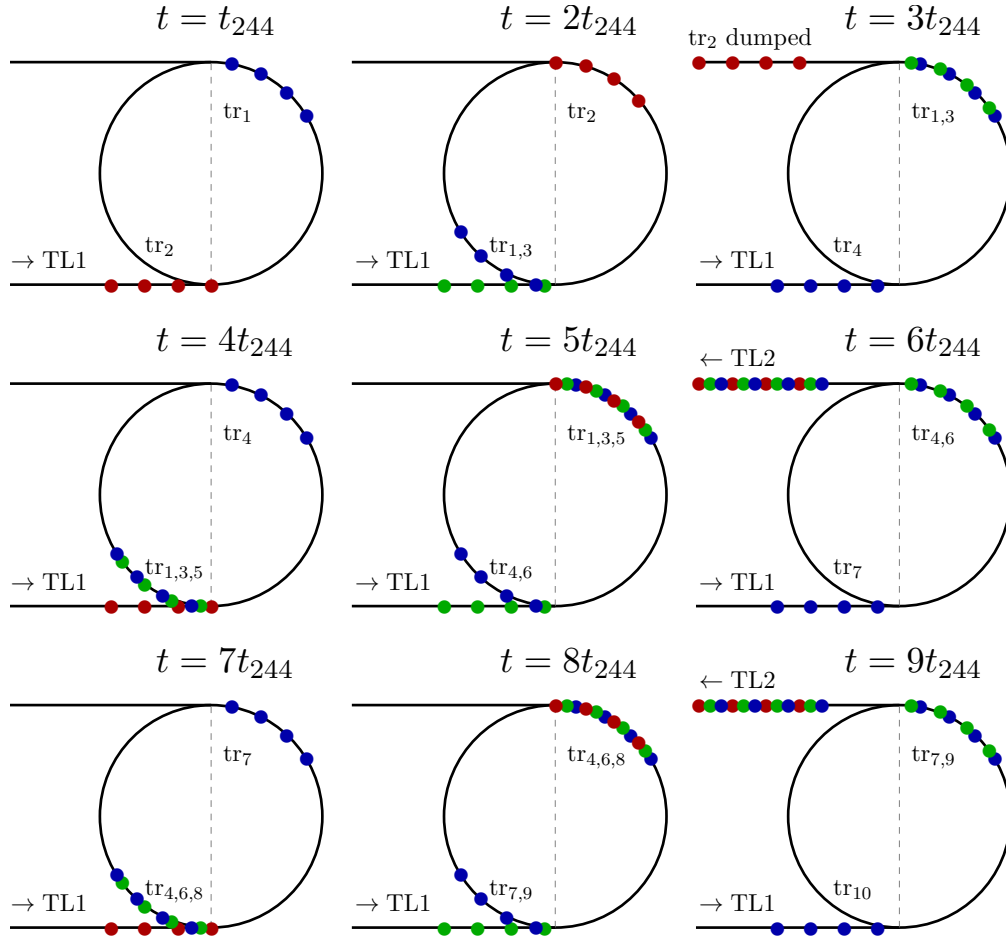
Figure 2.6: CR1 working cycle: A 244 ns pulse (tr_i) is injected into CR1 once every t_{244} (~ 244 ns). CR1's stripline extractor kicks every $3t_{244}$ sending 3 recombined pulses to TL2.

2.6 Combiner ring 2

The final recombination step is CR2. This is a 446 m combiner ring with 8 arcs in which four 3 GHz pulses are combined into one, increasing the frequency to 12 GHz and the current to 100 A. As we saw in Section 1.3, this is the frequency required at the PETS by CLIC's two-beam acceleration scheme.

One of the main issues with CR2 is its injection (Figure 2.7). When the first pulse (of the four being combined) reaches the ring, it needs to be brought to the rings central orbit. This is done by a 3 GHz transverse RF deflector (second deflector on the figure) synchronized to the bunch. In the time the first pulse takes a full turn of the ring, the second ring reaches it. The ring's length is such that, on its second turn, the bunches of the first pulse reach the deflector with a $\frac{\pi}{2}$ phase difference. This means that while the bunches of the second pulse are being kicked into orbit, the corresponding bunches of first pulse (83 ps ahead) are not moved by the deflector[4].

When the third pulse reaches the ring, the phase difference between the first and the deflector is π. If there was only one RF deflector (the one on the right), the first pulse would

---

[4]While the bunch is not deflected, the same is not true of individual particles in it. Since the bunch has non-zero length, its head and tail suffer small kicks in opposite directions.
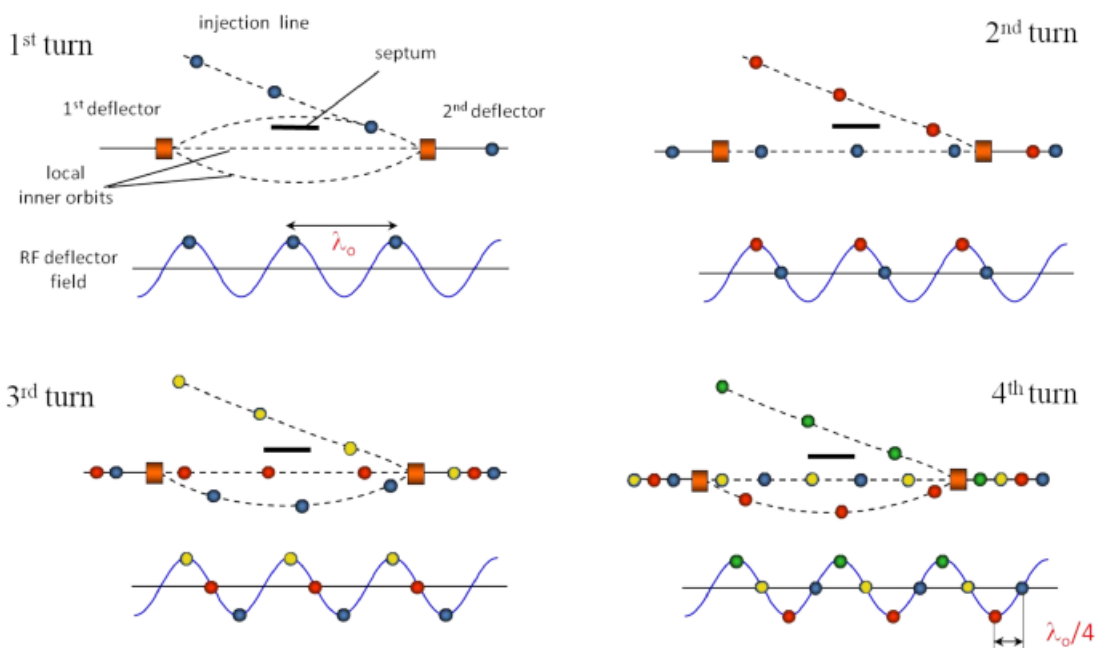
Figure 2.7: CR2 injection scheme [1]. For each passage of $b_i^{3.5}$ (blue) through CR2's injection section, we show both the orbit and the phase difference between the the bunch and the RF deflectors for $b_i^{3.5}$, $b_i^{2.5}$ (red), $b_i^{1.5}$ (yellow) and $b_i^{0.5}$ (green).

be kicked out of orbit by it. This is why we have another RF deflector with the same frequency before the injection point. On the third turn, the left deflector kicks the bunch into an off-centre-orbit. The bunch is brought back to the centre-orbit by a set of quadrupoles (not shown in the figure) between both deflectors (see Section 5.5). The design of this bump needs to ensure the bunch reaches the right deflector at a trajectory symmetric to the bunch being injected. In the fourth turn, similarly to the second, the bunch passes through the deflector without being kicked by it.

2.7   Decompression and re-compression

Earlier studies [1] of the DBRC revealed that, in the arcs of the delay loop and the combiner rings, a 1 mm long DB bunch would trigger a sufficient amount of coherent synchrotron radiation (see Section 3.5) to decrease beam quality. To minimize this effect, the design was updated in [26] by adding a 64 m decompression dogleg at the start of the complex. This dogleg was designed with $R_{56} = 5$ cm which, with an energy spread of 0.85%, nearly doubles the bunch length, turning CSR's effects on beam quality negligible. At the end of the complex, a chicane was added to the design to compress the bunch back to its original length (as we will see in Chapter 5, there are other effects that prevent this).

Chapter 3

Beam dynamics

In this chapter we offer a review of some key concepts in beam dynamics that are fundamental to describe and predict the behaviour of a particle beam travelling though a beamline. Accelerator physics is a mature science and there is extensive literature on this subject, of which we recommend [29–34]. In Section 3.1 we present the 6D coordinate system typically used in beam dynamics. Section 3.2 explores how linear optics can be adapted to track charged particles with this system. Section 3.3 presents the Courant-Snyder twiss functions as a way to describe the dynamics of a multi-particle beam. In Section 3.4 we present non-linear effects and how to address them. Section 3.5 briefly explains the damaging effects of synchrotron radiation and Section 3.6 explores the long term effects of divergence on longitudinal dynamics.

3.1 Coordinate-system

The naive approach to describe the dynamics of a particle is to define two vectors $\vec{q} = \left(\hat{q}_1, \hat{q}_2, \hat{q}_3\right)$ and $\vec{v}_q = \left(\hat{v}_1, \hat{v}_2, \hat{v}_3\right)$ using the laboratory's reference frame. This is not a particularly useful coordinate-system for accelerator physics since beamlines can range from a few meters to several km in length, might vary in topology from fairly linear to extremely complex and variations in $\vec{q}$ and $\vec{v}$ significance is highly dependent on the direction.

For these reasons, it is customary in accelerator physics and beam dynamics to use a Frenet-Serret coordinate-system. In this system, we use a reference frame following a trajectory (S) of an idealised test-particle along the beamline. Notably, while on the lab's reference frame the beamline might vary in direction (typically due to steering magnets), on our Frenet-Serret system the z direction is always aligned with the beamline's direction.

Each particle p_i can be described by a 6D vector C such that

$$C(p_i) = \begin{bmatrix} x \\ x' \\ y \\ y' \\ z \\ \delta \end{bmatrix}, \qquad (3.1)$$

in which x, y and z are, respectively, the horizontal, vertical and longitudinal offset from the

reference particle. x' and y' are the longitudinal derivatives of x and y such that

$$q' = \frac{dq}{dz} = \frac{dq}{dt}\frac{dt}{dz} = \frac{v_q}{v_z} = \frac{P_q}{P_z} \sim \frac{P_q}{P}, \ q = \{x, y\} \ . \tag{3.2}$$

In an ultrarelativistic regime most of the particle's momentum is in the z component ($P \sim P_z$), therefore $q' \ll 1$. Finally, since the longitudinal speed is $\sim c$ (is often just approximated to c), rather than parametrizing differences in speed directly we use $\delta = \frac{\Delta P}{P_0}$, in which $\Delta P = P - P_0$ and P_0 is the momentum of our reference particle. For ultrarelativistic electron beams, it is also reasonable to approximate particle energy and particle momentum ($E_0 \sim P_0$).

3.2   Linear transfer maps

One of the advantages of the Frenet-Serret coordinate-system is the ability to describe lattice elements in a matrix formalism similar to the one used in (light) optical systems. In this formalism, any linear transformation of a particle's coordinates can be described by a R-matrix such that

$$\begin{bmatrix} x \\ x' \\ y \\ y' \\ z \\ \delta \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} & R_{13} & R_{14} & R_{15} & R_{16} \\ R_{21} & R_{22} & R_{23} & R_{24} & R_{25} & R_{26} \\ R_{31} & R_{32} & R_{33} & R_{34} & R_{35} & R_{36} \\ R_{41} & R_{42} & R_{43} & R_{44} & R_{45} & R_{46} \\ R_{51} & R_{52} & R_{53} & R_{54} & R_{55} & R_{56} \\ R_{61} & R_{62} & R_{63} & R_{64} & R_{65} & R_{66} \end{bmatrix} \begin{bmatrix} x_0 \\ x'_0 \\ y_0 \\ y'_0 \\ z_0 \\ \delta_0 \end{bmatrix} . \tag{3.3}$$

As in "light optics", when computing a sequential set of transformations $\{R_1, R_2, ..., R_n\}$, one simply applies the respective matrices in reverse order ($C_{\text{final}} = R_n \cdots R_2 \cdot R_1 \cdot C_{\text{init}}$). It should be obvious that, if our reference particle is the centre of the coordinate-system, any R-matrix will not change the vector $(0, 0, 0, 0, 0, 0)$.

   Let us now establish the most common R matrices. The simplest lattice element is a drift. This is a section of the beamline with no active field and, in an ultrarelativistic regime[1], is written as

$$R_{\text{Drift}} = \begin{bmatrix} 1 & L & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & L & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} . \tag{3.4}$$

Notice that, since it isn't an active element, a drift's effect on a particle is limited to a change in its position due to divergence.

   Next, let us look into quadrupole magnets. These are perhaps the most useful optical elements in large accelerators since they can be used to focus a particle beam in the transverse directions. This allows us to counteract divergence and keep particle loss to a minimum across long distances in a linac or multiple turns in a synchrotron. Unfortunately, a quadrupole only acts as a focussing lens on one transverse direction, defocussing the beam in the direction

---

[1]In a non-ultrarelativistic regime, Matrix 3.4 has non-zero R_{56} to take into account how differences in a particle's speed affect its longitudinal position.

perpendicular to it. Consequently, we need to define $k_x = \sqrt{K_1}$ and $k_y = \sqrt{-K_1}$, where K_1 is the quadrupolar focal strength. This allows us to build the matrix

$$
R_{Quad} = \begin{bmatrix}
\cos k_x L & \frac{\sin k_x L}{k_x} & 0 & 0 & 0 & 0 \\
-k_x \sin k_x L & \cos k_x L & 0 & 0 & 0 & 0 \\
0 & 0 & \cos k_y L & \frac{\sin k_y L}{k_y} & 0 & 0 \\
0 & 0 & -k_y \sin k_y L & \cos k_y L & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}. \tag{3.5}
$$

When $K_1 > 0$, the quadrupole focusses the beam on the x axis. This is described in Matrix 3.5, for example, by having $|R_{22}| < 1$, which minimizes divergence. On the other hand, k_y is imaginary, making

$$
R_{33} = R_{44} = \cos i \sqrt{K_1} L = \cosh \sqrt{K_1} L , \tag{3.6}
$$

which is always greater than one and therefore the quadrupole increases beam divergence (defocus it). For R_{34} and R_{43}, we can use the identity $\sin i\theta = i \sinh \theta$ to show that, for $K_1 > 0$,

$$
R_{34} = \frac{\sinh \sqrt{K_1} L}{\sqrt{K_1}} \quad \text{and} \quad R_{43} = \sqrt{K_1} \sinh \sqrt{K_1} L . \tag{3.7}
$$

When $K_1 < 0$, a similar set of transformations applies to the first 2×2 block of R_{Quad} instead.

This means that the analogy between a quadrupole and a convergent lens only holds for one transverse direction. In order to focus the beam on both transverse directions we need to construct a FODO lattice. This is a modular sequence of beamline elements of the type

$$
\{... - F - O - D - O - F - O - D - O - F - ...\} , \tag{3.8}
$$

where F is a focussing quadrupole (on an arbitrary direction), D is a quadrupole focusing on the perpendicular plane and O is the drift space between them. This lattice is designed to prevent the beam from surpassing a certain size and is the basis of all beam transport over large distances.

Though it is not our intention to detail the inner-workings of individual elements, we should clarify the connection between K_i and physical devices. Beam optics tells us $K_1 L$ is the inverse of the focal length along the focussing plane (using a thin lens approximation). This is related to the quadrupole's gradient by

$$
K_1 = G \frac{e}{P} . \tag{3.9}
$$

Notice that K_1 depends on the particles momentum P. This means that a linear map (Matrix 3.5) computed for the reference particle (which has $P = P_0$) is only accurate for particles with $\delta \ll 1$. Similarly to how K_1 is used for quadrupoles, it is customary to use K_n to refer to the focal strength of higher order optics (K_2 for sextupoles, K_3 to octupoles, etc.).

The final transfer map that we will discuss is R_{Sbend}. Sector bends (sbends) are dipole magnets built in the shape of an arc sector which we use to steer the beam. Since the refer-

ence frame follows the (curved) orbit of the ideal particle, the change in direction in the lab reference frame is not obvious in

$$R_{\text{Sbend}} = \begin{bmatrix} \cos\theta & \rho_0\sin\theta & 0 & 0 & 0 & \rho_0(1-\cos\theta) \\ -\frac{\sin\theta}{\rho_0} & \cos\theta & 0 & 0 & 0 & \sin\theta \\ 0 & 0 & 1 & L & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ \sin\theta & \rho_0(1-\cos\theta) & 0 & 0 & 1 & \rho_0(\theta-\sin\theta) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{3.10}$$

where ρ_0 and θ are the sbend's radius and bending angle. In practice, it is often more cost effective to design rectangular bends but conversion between the two element's is fairly trivial and can easily be found in the literature (see [32] for example). Tracking software often only implements one of the two (sbend or rbend) and expects the user to convert the other.

There are several elements of R_{Sbend} that warrant further discussion. The first is the fact that there is a small focussing effect on a dipole (top left quadrant of the matrix). This effect is due to the curvature of the coordinate system and is present in all steering magnets with $\theta < \pi$.

The second effect to point out is called dispersion and it is present in the R_{16} and R_{26} elements of our map. Dispersion is the most notable effect of transporting particles with non-zero δ through a beamline. The origin of this effect would be more obvious in a Cartesian coordinate system, where charged particles with different momenta are bent by different angles as they travel through a fixed transverse field. On our Frenet-Serret reference frame, this is translated into particles with $\delta > 0$ diverging in one direction (outwards) and particles with $\delta < 0$ in the other. Generally, lattices are designed to minimize dispersion, and arc cells (inside which some dispersion is unavoidable) are often designed to have overall zero dispersion gain (be achromatic). In Chapter 5, we will see how this is done in the bending arcs of CR1 (Figure 5.6) and CR2 (Figure 5.8).

The last non-trivial elements of R_{Sbend} to discuss are R_{51}, R_{52} and R_{56}. These parametrize how the longitudinal position of the particle varies in an sbend and are the driving force behind how decompression doglegs and compression chicanes (see Section 2.7) work. A more detailed review of this and other longitudinal effects will be given in Section 3.6.

While the matrices above describe singular elements, an R matrix may also describe entire beamlines through the matrix multiplication each element matrix. When considering two particle distributions, $\{C\}_{\text{initial}}$ and $\{C\}_{\text{final}}$, taken before and after a section described with the matrix R, such that

$$\{C\}_{\text{final}} = R \cdot \{C\}_{\text{initial}} \ , \tag{3.11}$$

the R-matrix elements can be seen, also, as the linear correlations between the initial and final phase-space coordinates. Notice that, in these terms, a given correlation between coordinates i and j can be equally described as $R_{ij} = a$ or $R_{ji} = 1/a$.

3.3 Betatron motion and twiss parameters

The formalism presented in sections 3.1 and 3.2 was developed to compute the orbit of individual particles. We are however mostly interested in computing the dynamics of a many-

Figure 3.1: Particle orbits along a beamline with multiple quadrupoles. The envelope containing the beam on the horizontal plane is defined by the outer orbit $\sqrt{\varepsilon_{g_x} \beta_x(s)}$.

particle bunch. In order to do this, we shall use the Courant-Snyder twiss functions $\alpha(s)$, $\beta(s)$ and $\gamma(s)$ [35]. These parameters, together with the geometric emittance (ε_g), are used to create a mathematical model of a beam distribution's dynamics.

Lets start by looking at the behaviour of multiple particles with different starting coordinates along a periodic FODO lattice[2]. Figure 3.1 shows that particles that reach Q_{F_1} with high positive x suffer a stronger negative kick (this is also explicit in Eq. 3.5). This causes them to have high negative x by the time they reach Q_{F_2}, which will kick them positively to bring them back to a centre-position. Assuming $\delta = 0$, all particles will present this oscillatory orbit throughout a periodic lattice, varying only in amplitude. We call this behaviour betatron motion.

We can see that the rms beam size (σ_x) varies along the beamline. We will use the relationship between σ_x and ε_{g_x} (which we will discuss further later) to define our first twiss parameter as

$$\beta_x(s) = \frac{\sigma_x^2(s)}{\varepsilon_{g_x}} \ . \tag{3.12}$$

The other two twiss parameters can be obtained from the variation $\beta_x(s)$. This is $\alpha_x(s)$, defined as

$$\alpha_x(s) = -\frac{1}{2}\beta_x'(s) \ , \tag{3.13}$$

and $\gamma_x(s)$, defined as

$$\gamma_x(s) = \frac{1 + \alpha_x^2(s)}{\beta_x(s)} \ . \tag{3.14}$$

While β_x[3] allows us to parametrize σ_x, α_x and γ_x allow us to describe the upper and lower limits to a particles divergence (σ_x'). This three parameters together allow us to define a $\{x, x'\}$ phase-space ellipse (see Figure 3.2), in which the beam is contained.

---

[2]Notice that the lattice represented in Figure 3.1 is only periodic up to Q_{F_2}.

[3]We will suppress the the s dependency from our notation from this point forward.

Figure 3.2: Beam distribution phase-space ellipse (in arbitrary units). The ellipse is contained in the $\pm\sqrt{\varepsilon_{g_x}\beta_x}$ and $\pm\sqrt{\varepsilon_{g_x}\gamma_x}$ intervals for the x and x' planes, respectively.

The shape of this ellipse varies continuously along the beamline with the twiss functions however, on a linear system with $\delta = 0$, its area remains the same. This defines the final beam parameter, called geometric emittance (ε_{g_x}). With this final addition we can describe the transverse dynamics of any number of particles particles with just eight parameters (α_x, β_x, γ_x, ε_{g_x} and their vertical plane counterparts).

These parameters are extremely useful to accelerator design as the easiest way to find a periodic solution to a lattice is to ensure that their twiss functions at injection and extraction match. As for ε_g, since it remains constant along the beamline for an ideal beam (Gaussian, linear and monochromatic), it is a great measurement of beam quality. In our studies, matching the α and β in various beamline sections is an important starting step to the section optimisation and normalised emittance ($\varepsilon = \beta_{\text{rel}}\gamma_{\text{rel}}\varepsilon_g$, where β_{rel} and γ_{rel} are the relativistic β and γ functions, respectively) is our primary transverse beam quality target. Normalised emittance is often more useful from a optical design point of view since it remains constant as the beam is accelerated/decelerated (which is not the case for ε_g).

Now that we established how to characterize a beam in the transverse phase-space, we need to propagate it through the beamline. The beauty of the the Courant-Snyder formalism is how easily one switches between it and the single particle formalism presented in Section 3.2. The evolution of the twiss parameters along either an element or an entire beamline can be computed by computing its R-matrix and solving

$$B_{\text{final}} = R \cdot B_0 \cdot R^{\text{T}} , \tag{3.15}$$

where B_0 and B_{final} are constructed from the initial and final twiss parameters as

$$
B = \begin{bmatrix}
\beta_x & -\alpha_x & 0 & 0 & 0 & 0 \\
-\alpha_x & \gamma_x & 0 & 0 & 0 & 0 \\
0 & 0 & \beta_y & -\alpha_y & 0 & 0 \\
0 & 0 & -\alpha_y & \gamma_y & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix} .
\tag{3.16}
$$

3.4 Non-linear optics and effects

The formalism presented in Section 3.2 is invaluable to the first steps in any kind of accelerator design studies. It does not however address non-linear effects. The first effect we need to discuss is chromaticity. When designing a lattice, all optics are designed to transport particles of a reference energy[4] (E_0). In practice, however, it is difficult to produce a fully mono-chromatic beam. The orbit of particles with positive or negative δ will curve less or more than a $\delta = 0$ particle, respectively.

In sbend magnets, this is translated into dispersion and, by virtue of being a linear effect, can be represented in the R-matrix (Matrix 3.10). On the other hand, the effect on quadrupoles and higher order optics cannot be represented in a linear matrix. Instead, when tracking a particle off-momentum we need to rescale K_1 such that

$$
K_1(P) = \frac{P_0}{P} K_1(P_0) .
\tag{3.17}
$$

This shifts the focal point of the quadrupole lens in a similar manner than chromatic aberrations do on light optics (hence the name chromaticity).

Since chromaticity is not a linear effect, it cannot be corrected using linear optics. To address it and other non-linear effects, the DBRC has several sextupole magnets in its lattice. These are modelled using the thin lens approximation

$$
\begin{cases}
x' = x'_0 - K_2\left(x_0{}^2 - y_0{}^2\right) \\
y' = y'_0 + 2K_2 x_0 y_0
\end{cases} .
\tag{3.18}
$$

For higher order optics, Eq. 3.18 can be generalized as

$$
\begin{cases}
x' = x'_0 - \text{Re}\left[(x_0 + iy_0)^n\right] K_n \\
y' = y'_0 + \text{Im}\left[(x_0 + iy_0)^n\right] K_n
\end{cases} .
\tag{3.19}
$$

As we will see in Section 5.1, the presence of sextupoles in the lattice was invaluable to our design optimisation scans since they allow us to perform corrections to parabolic effects.

Describing second order effects of a lattice element of length s on the coordinate C_i can be done using a tensor T_{ijk} such that

$$
C_i(s) = C_i + \sum_j R_{ij} C_j + \sum_{jk} T_{ijk} C_j C_k .
\tag{3.20}
$$

---

[4]In most machines, E_0 is simply the bunch average energy but, as we will see in Section 3.5, the DBRC transports bunches with slightly different values of $\langle E \rangle$.

We can therefore, for instance, compute dependency of the longitudinal position to δ as

$$z(s) = z + R_{56}\delta + T_{566}\delta^2 \,, \tag{3.21}$$

which, as we will see in Section 5.1, is a significant second order effect in the DBRC design.

As in the case of linear optics, it is possible to compute the overall second order effect of multiple elements simply by knowing their T-tensors and R-matrices. It is, however, a slightly less intuitive procedure. While for two linear elements A and B, the overall R-matrix is simply $R_{[A+B]} = R_{[B]}R_{[A]}$, for the T-matrix, each of element $T_{[A+B]}$ is computed as

$$T_{ijk_{[A+B]}} = \sum_l R_{il_{[B]}} T_{ljk_{[A]}} + \sum_{lm} T_{ilm_{[B]}} R_{li_{[A]}} R_{mk_{[A]}} \,. \tag{3.22}$$

3.5   Synchrotron radiation

Electrodynamics tells us that an accelerated charged particle emits electromagnetic radiation, losing momentum in accordance with Abraham-Lorentz-Dirac's force. On a particle accelerator, this phenomenon is, for historical reasons, often refereed to as synchrotron radiation and it is studied to evaluate its (detrimental or intentional) effects on beam dynamics or the radiation produced. Thought this effect is essential for the operation of synchrotron lightsources, free-electron lasers and damping rings, for the DBRC it is purely detrimental.

Due to the higher transverse acceleration, this phenomenon is prevalent in steering magnets and we will only compute its effect on those. Depending on beam energy and bunch length, synchrotron radiation can present itself either only in the incoherent (ISR) regime or in both incoherent and coherent (CSR) regimes simultaneously.

CSR occurs when the emitted radiation has a wavelength much smaller than σ_z, allowing the head of the bunch to absorb the radiation emitted by the tail. This greatly increases the energy loss and energy spread of the bunch, which will in turn have a detrimental impact on its emittance (other effects, such as micro-bunching, might also occur). The bunch length manipulation described in Section 2.7 is done with the sole purpose of avoiding this regime.

The effect of ISR on energy spread is several orders of magnitude smaller and the most direct consequence of ISR in accelerator design is its influence when defining the optics reference energy (E_0). Typically, the optics of a beamline are set such that $E_0 = \langle E \rangle$. Since the beam is losing energy through radiation, E_0 needs to be redefined in accordance with

$$\Delta_{E_0} = \frac{e^2}{6\pi\varepsilon_0 \left(mc^2\right)^4} \frac{\alpha E^4}{\rho_0} \tag{3.23}$$

after each dipole to account for that loss. In re-circulating machines this loss is potentially more damaging since not all bunches will suffer it equally. This makes the $\langle E \rangle$ of individual bunches vary slightly so the equality $E_0 = \langle E \rangle$ can only be maintained for some of them.

A closer look at the total energy loss suffered in the DBRC showed that the energy loss in the machine varies between -0.05% for $b_{0.5}^{0.5}$ and -0.13% for $b_{2.5}^{3.5}$ (see Section 2.3 for details in the notation). Since our initial energy spread supersedes that (0.85%), we expect the difference in energy between bunches not to play a major role on the beam's emittance growth.

3.6   Longitudinal effects

It is not uncommon for particle accelerators to operate in a fully ultra-relativist regime ($v \sim c$), in which variations of the particles' momenta have negligible influence in their speed. In this regime, the relative longitudinal position of particles (z_0 and $z_0 + \delta_z$) in a bunch does not change along the beamline as a direct result of the difference in their energy/speed[5].

There are however other effects that act on the longitudinal dynamics of a particle. Rather than being dependent on the particle's speed, these originate from differences in the total orbit taken by an arbitrary particle and the reference one. Let us remember, for example, that R_{sbend} (Eq. 3.10) has non-zero R_{51} and R_{56}. This means that, while passing through this element, particles that are off-centre in the horizontal axis or off-momentum will travel in the longitudinal direction. The R_{51} effect is due to the fact that particles "taking the inside curve" take a shorter orbit than the reference particle. For particles with non-zero δ, their rigidity is different from the reference particle, meaning their trajectory is bent more ($\delta < 0$) or less ($\delta > 0$) than it (thus $R_{56} \neq 0$). This effect is often exploited to manipulate bunch length throughout a beamline (see Section 2.7).

As for the longitudinal effects of divergence ($x' \neq 0$ or $y' \neq 0$), sbends also present a linear effect (R_{52}) correlating z to x', but (for reasons that will become apparent in Section 5.1) we will chose to focus our discussion on a non-linear effect. Since our frame of reference is moving along the beamline, a variation in the longitudinal coordinate (Δz) can be better rewritten as

$$\Delta z = c \Delta t = c \left(\frac{L}{v_z} - \frac{L}{v_0} \right) \ , \tag{3.24}$$

where L is the length of an arbitrary element, v_0 is the speed of the reference particle (which has $x' = y' = 0$) and v_z is the longitudinal speed component of an arbitrary particle. Since, in an ultra-relativist regime, $v/c \sim 1$, $v_0 = v$ and

$$v^2 = v_x^2 + v_y^2 + v_z^2 \Leftrightarrow v_z = \frac{v}{\sqrt{\frac{v_x^2}{v_z^2} + \frac{v_y^2}{v_z^2} + 1}} = \frac{v}{\sqrt{x'^2 + y'^2 + 1}} \ , \tag{3.25}$$

we can re-write Eq. 3.24 as

$$\Delta z = \frac{c}{v} L \left(\sqrt{x'^2 + y'^2 + 1} - 1 \right) \sim L \left(\sqrt{x'^2 + y'^2 + 1} - 1 \right) \tag{3.26}$$

which, for low $x'^2 + y'^2$, can be approximated by a Taylor series expansion to

$$\Delta z \sim \frac{L}{2} \left(x'^2 + y'^2 \right) \ . \tag{3.27}$$

The effect presented in Eq. 3.27 is directly dependent on the geometry of our reference system, and is therefore present on all lattice elements. In non-linear optics, we can describe it by the second order matrix elements

$$T_{522} = T_{544} = \frac{L}{2} \ . \tag{3.28}$$

---

[5]While electron machines like the DBRC operate almost exclusively in the ultrarelativistic regime, that might not be the case, depending on E, for proton or ion accelerators.

Chapter 4

Simulation tools and techniques

In this chapter we discuss the computational tools and techniques used in modern accelerator design. Section 4.1 gives a brief overview of the different types of tools and their function. Section 4.2 presents `Placet2` [36, 37], the main tool for our studies. Lastly, Section 4.3 gives a brief overview of the procedure used in our studies when optimisation scans were required.

4.1   Basics of accelerator design

Large scale accelerator design has evolved up to a point that, before building one of these machines, we are required to perform exhaustive beam dynamics simulations and iterate on the design based on the performance of those (rather than trying to do it later on an expensive physical machine). After the high level design is established, the precise positions and strengths of all lattice elements are determined by first linear optics studies, then multi-particle tracking and lastly performance studies under non-ideal conditions.

As a first step, optics design software like `Mad-X` [38] or `elegant` [39] are used to design periodic cells and match the beam's twiss functions where different sections meet. This is performed using the twiss formalism explained in Section 3.3, providing an accurate approximation of the bunch behaviour under negligible energy spread and non-linear effects. For the DBRC, most design studies in [1] are of this nature.

After establishing an initial lattice solution, a large[1] multi-particle distribution is tracked through the design to take into account effects that cannot be described with linear optics. `Placet` [40–44] and `Placet2` [36, 37] were developed at CERN to conduct these studies for CLIC and other linear machines. Our studies of the DBRC fall under this category.

As a final step, element misalignments, magnet strength errors and other imperfections are randomly added to the lattice. The beam is then tracked in order to evaluate the robustness of the design and the strictness necessary in machine commissioning. This step is outside the scope of this work but studies of this nature of the DBRC are expected in the near future.

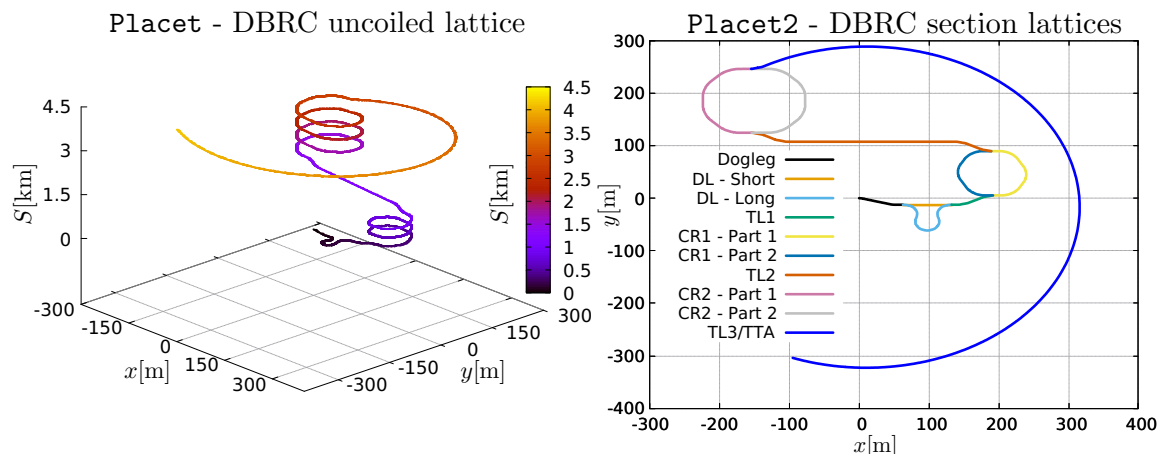[1]Our studies track 10^4 macro-particles, for example.

Figure 4.1: DBRC lattice comparison. On the left panel, we show the `Placet` longest-orbit lattice used in [26]. The vertical axis and colour scale show the total travelled orbit. On the right panel, the `Placet2` full lattice, divided by sections based on the machine's topology.

4.2 Placet2

The current state of the art software in precision beam tracking for linear accelerators is `Placet` [40–44]. Rather than only considering the beam optics transformation, it tracks the orbit of either a single particle, a multi-particle bunch or a longitudinal sliced multi-particle bunch (useful for wakefield studies). Using a multi-particle distribution allows us to take into account effects like non-linear chromatic aberrations, beam divergence's role in longitudinal dynamics, synchrotron radiation and wakefields. Evaluating how these phenomena affect the beam is necessary for an accurately determination of machine performance. Several beam dynamics studies were performed using this software for both long term projects like CLIC and the ILC, and operating accelerators like SLAC, ATF2 and FERMI@Elettra [45–49].

`Placet` was designed to study large linear accelerators and is able to efficiently track the beam through thousands of elements. However, recirculating machines like the DBRC, CTF3, LHeC, PERLE and eRHIC [50–53] are beyond `Placet`'s original scope. Simulating these machines requires the software to break the lattice according to the machine topology and send different bunches through different paths. The main goal of `Placet2`'s [36, 37] development was to extend `Placet`'s capabilities to these machines.

Using the DBRC as a specific example, some longest-orbit `Placet` studies [26] were performed by uncoiling the several turns the bunch takes in the combiner rings into a long linear lattice as shown in Figure 4.1 (left panel). However, this approach only allows for the study of one of the 24 possible DBRC orbits. In `Placet2` each orbit section is defined independently (right panel). The user is capable of defining the ten beamline sections in the figure according to the machine topology and transfer bunches from one section to the other based on timing and/or horizontal offset, allowing for a realistic description of the machine's operations.

`Placet2` is an open source software still under development. Appendix B offers a qualitative list of the updates performed to facilitate our studies. The full source code may be found in [54].

4.3   Optimisation procedure

In Chapter 5, we will present the results of several lattice optimisation scans. These were performed thanks to the API between `Placet2` and `Octave` [55], which allowed us to utilize `Octave`'s extensive optimisation libraries.

Optimisation was performed using the Nelder–Mead method [56], often called the downhill simplex method. This method is used by evaluating a loss function in $n + 1$ vertices, where n is the dimension of the optimisation parameter space, forming a $(n + 1)$-dimensional polygon. The algorithm then selects the vertex with the highest loss function result and shifts its position through the opposite face of the polygon. This process is repeated several times until the polygon converges on a minimum. This is a well known optimisation method and performs very well on parameter-spaces with unknown topology. It does, like other optimisation methods, suffer from the risk of the minimum found being local rather than global and its performance is somewhat dependent on the loss function being continuous.

In our studies, the parameter space is generally a collection of quadruple and sextupole strengths (often grouped into families to keep cell symmetry) of the beamline section we are targeting. Adding to this, if the previous section is not yet fully optimised, we sometimes have some injection parameters (twiss and dispersion) as optimisation parameters as well. As it is to be expected, the larger the parameter-space, the longer this method takes to converge, so the choice of how many optics are being optimised simultaneously need to be made with care.

Defining the loss function in our studies is also non-trivial. When there are several properties $\{l_1, l_2, ..., l_n\}$ that we wish to minimise (which is our case), the loss function needs to be defined as the weighted sum

$$L\left(x_1, x_2, ..., x_n\right) = \sum_i w_i l_i \left(x_1, x_2, ..., x_n\right) \ . \tag{4.1}$$

The weights w_i need to be selected in a way that reflects how sensitive l_i is to changes in the parameter space (which is not necessarily known information). To give a specific example, in the DBRC we needed to set w_{ε_x} at least one order of magnitude higher than w_{ε_y} (the exact ratio varied from scan to scan). If we were to use equal weights, the algorithm would quickly find a solution path that decreases ε_y faster than it increases ε_x and would pursue that path until reaching a local minimum (with unacceptable ε_x).

Chapter 5

Analysis and results

The primary objective of our studies was to develop `Placet2` simulations of the DBRC in order to assess and correct potential issues with the lattice design. In this chapter we present the results of our simulations and optimisation scans. Section 5.1 presents a previously unknown T_{566} aberration we found in the design. In Section 5.2 we give a brief overview of how we evaluate beam quality in our studies. Section 5.3 presents the results of our optimisation scans of the delay loop. Section 5.4 provides the design issues and the optimisation results for CR1. Section 5.5 presents the optimisation results of CR2 as well as a proposal to change the design of the injection lattice based on them. Lastly, Section 5.6 presents the recombination results at the end of the DBRC. We would like to point out that CLIC is still a project under development and these results are subject to significant updates in the future.

5.1   The T_{566} aberration

One of the first beam dynamics issues identified with the implementation of the DBRC's lattice in `Placet2` was a strong parabolic correlation between z and δ, to which we will refer as T_{566} in accordance with the notation established in Section 3.4. In Figure 5.1 we can see that, initially, bunch $b_{2.5}^{3.5}$ presents a longitudinal profile as required by the design. It has an energy chirp and a bunch length of 1 mm that doubles after the decompression dogleg. However, after passing the delay loop, the beam starts presenting a significant T_{566} component, which is exacerbated further after CR1 and CR2.

Since this effect was not observed in [26], we tracked $b_{2.5}^{3.5}$ in both `Placet` and `Placet2` to determine the cause of the conflicting results. We found that while `Placet2` computes the longitudinal effects of divergence (see Section 3.6) throughout the entire lattice, on the `Placet` simulations we were only doing so in the lattice sbends. While this is a common practice to diminish computational costs (beam divergence is typically higher inside bends due to dispersion) we found that, for this machine, the effect is significant enough in drift-space to impact the beam longitudinal profile.

Notice that the effect described in Section 3.6 is a parabolic correlation described by T_{522} and T_{544}, rather than T_{566}. We can, however, connect these tensor elements in a long enough beamline by using Eq. 3.22 to compute

$$T_{566_{[A+B]}} = \sum_l R_{5l_{[B]}} T_{l66_{[A]}} + \sum_{lm} T_{5lm_{[B]}} R_{l6_{[A]}} R_{m6_{[A]}} \ .$$ (5.1)

Figure 5.1: Correlation between particle energy and longitudinal position in bunch $b_{2.5}^{3.5}$ before optimisation. The first layer shows $b_{2.5}^{3.5}$ injected into the DBRC and the subsequent layers show it at the extraction point of the various sectors: decompression dogleg, DL, CR1 and CR2.

Since the only non-zero elements of T_{drift} are T_{522} and T_{544}, Eq. 5.1 can be simplified to

$$T_{566_{[A+B]}} = T_{522_{[B]}} {R_{26_{[A]}}}^2 + T_{544_{[B]}} {R_{46_{[A]}}}^2 \ . \tag{5.2}$$

We propose that this effect, though subtle in a few elements, propagates along the beamline to the point that it becomes an issue[1]. Since there are far more radical transverse transformations than longitudinal, the T_{522} and T_{544} correlations are not visible in the long term.

Having identified the issue, one needs to find a way to correct it. This is not trivial since no element in our lattice has an effect that varies over the z coordinate[2]. The strategy employed takes advantage of the areas of the lattice with high dispersion (see Section 3.2) like inside of the ICG cells. In those areas, there is a correlation between the horizontal coordinate of the particle and δ that we can take advantage of to manipulate the longitudinal profile. Since the T_{566} effect is parabolic we need to use non-linear optics to correct it.

To mitigate the chromatic effect that have a deteriorating impact on emittance, our ICG cells have multiple sextupoles which we can use to reduce T_{566}. We are however attempting to solve two issues using the same tool and, as we will show in the following sections, finding a compromise between the two is not straightforward.

---

[1]Identifying the evolution of the second-order transfer matrix along the machine is outside the scope of our studies, so this is merely an hypothesis based on our results. A deeper study into this matter would require further `Placet2` development.

[2]Strictly speaking, the lattice does have transverse RF deflectors which do have a longitudinal component. However, those are few, placed in specific locations and not of use to address this issue.

Figure 5.2: Horizontal emittance (ε_x) and dispersion-free emittance (DF$\mp\varepsilon_x$) along the decompression dogleg ($S < 64\,\mathrm{m}$) and delay loop.

5.2 Transverse beam quality

As mentioned in Section 3.2, transverse emittances (ε_x and ε_y) are often used to evaluate beam quality. It is therefore necessary to compute them from our macro-particle distribution. For a Gaussian distribution this is done by computing

$$\varepsilon_{g_q} = \sqrt{\det\left(\Sigma_q\right)}, \tag{5.3}$$

where Σ_q is a matrix defined as

$$\Sigma_q = \begin{bmatrix} \mathrm{cov}\left(q, q\right) & \mathrm{cov}\left(q, q'\right) \\ \mathrm{cov}\left(q', q\right) & \mathrm{cov}\left(q', q'\right) \end{bmatrix}. \tag{5.4}$$

For a single bunch, `Placet2` has internal routines capable of computing this along the entire beamline. The assumption that the beam is Gaussian is, however, implicit in this procedure. Deviations from that statistical distribution will lead to an over-estimation of ε_q. Since, internally, `Placet2` computes all twiss parameters using their relationship to ε_q, this will lead to a faulty estimation of those as well[3].

To address this, we implemented in `Placet2` the ability to read dispersion-free parameters from a distribution. In chromatic beams, the emittance read from a particle distribution is heavily influenced by dispersion, making it not suitable to use as a beam quality parameter in dispersive regions. Since dispersion is an expected effect in some sections of the beamline, beam quality is better gauged by dispersion-free emittance (DF$\mp\varepsilon_q$). This is obtained by creating a copy of the beam distribution where we artificially subtract the dispersion and computing the emittance of the resulting distribution. In Figure 5.2 we see how the estimation of ε_x increases by over an order of magnitude in the dispersive region of the decompression dogleg and on the twelve delay loop bending arcs. Comparatively, DF$\mp\varepsilon_x$ remains almost

---

[3]Note that `Placet2` tracks all particles individually, so the beam distribution itself is in no way affected, only the parameters estimated from it.

Figure 5.3: Final horizontal phase-spaces of bunches $b_{0.5}^{2.5}$ and $b_{1.5}^{2.5}$. The distribution tails are highlighting by denoting particles in which x, x', y or y' is above 3σ.

constant in these regions. In practice, though our results present dispersion-free emittance[4], our optimisation scans often target locations of the beamline in which we wish to suppress dispersion. For this reason, and to save computing time, our optimisation generally doesn't perform this procedure. This allows us to minimise emittance and dispersion simultaneously without adding another term to the loss function.

Additionally, since our lattice contains a lot of non-linear optics and effects, some of the bunches display very large tails when they reach extraction (Figure 5.3 shows two examples). We are tracking a limited (10^4) number of particles, theses tails are therefore composed by a very small number of particles. Since a single particle far away from the distribution has a large effect on the covariance and the bunch tails cannot be correctly modelled with so few particles, we opted to remove all particles with a transverse coordinate above 3σ from the distribution before computing the emittance at the end of the DBRC (Section 5.6). We have not used this procedure in the earlier sections of the machine.

Finally, it is important to note that, when recombining bunches with different orbits, evaluating pulse quality for our loss function becomes less trivial than it would be for a single bunch. For most machines, all bunches have the same orbit under ideal conditions so, discounting beam jitter and other dynamic effects, there is little point in even simulating more than one bunch[5]. During optimisation, one might be therefore tempted to compute the average emittance across multiple bunches $\langle \varepsilon_{x,y} \rangle$ and use it for the loss function (in fact, tracking software encourages this approach). However, since we are tracking a re-circulating machine, we need to consider the possible consequences of having different bunches exit a combiner ring at different turns. Optimising the lattice for low $\langle \varepsilon_{x,y} \rangle$ ensures only that each individual bunch has low emittance, while pulse quality can also decline due to differences from bunch to bunch. To illustrate this problem, in Figure 5.4 we present three clearly different bunches that all have the same emittance. A loss function based solely on $\langle \varepsilon_{x,y} \rangle$ will not differentiate a solution where all bunches in the pulse are similar to bunch A from one in what some bunches are similar to A and others are similar to B (having similar ε_x but different twiss parameters). The same issue can occur if part of the pulse is off-centre at the extraction point despite having low emittance (bunch C).

Using a loss function with a weighted sum of $\varepsilon_{x,y}$, $\beta_{x,y}$, $\alpha_{x,y}$, $\sigma_{x,y}$ and $\sigma'_{x,y}$ would be a solution but it would require a lot of fine tuning of its weights. As an alternative, we opted

---

[4]We will suppress the prefix from this point forward.

[5]Similarly, we simulate 12 bunches while optimising when a full pulse has 2928 (see Figure 2.2).
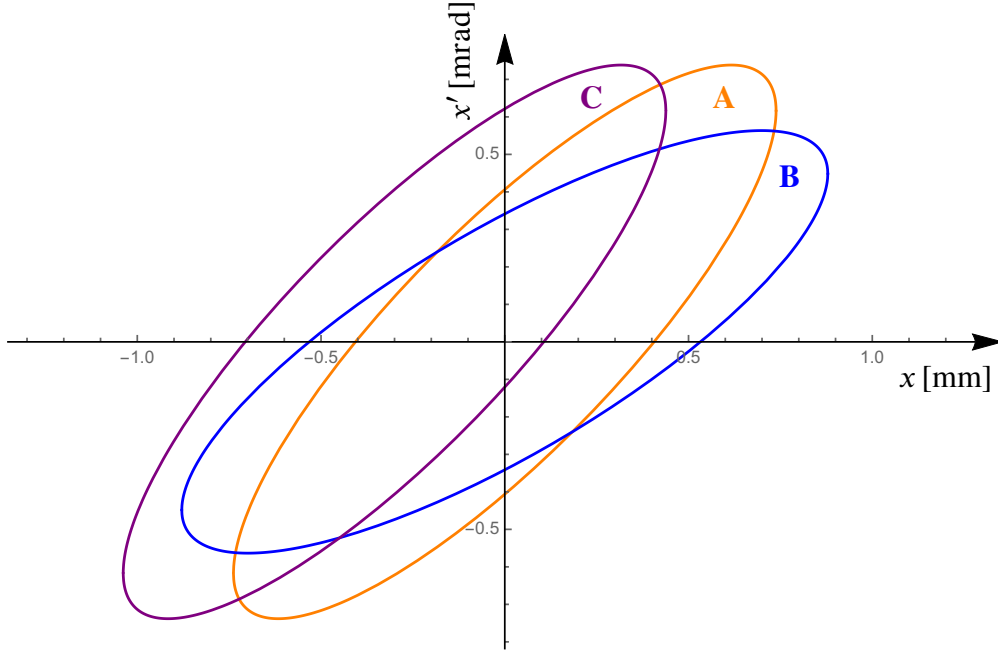
Figure 5.4: Horizontal phase-space (in arbitrary units) for three bunches with the same emittance despite having different $\langle x \rangle$ and σ_x.

to define a pulse emittance parameter ($\tilde{\varepsilon}_{x,y}$). This parameter is computed by combining all different bunches and computing the emittance of the resulting distribution. If all bunches are centred and have similar twiss parameters, $\tilde{\varepsilon}_{x,y} = \langle \varepsilon_{x,y} \rangle$ but, if that's not the case, $\tilde{\varepsilon}_{x,y} > \langle \varepsilon_{x,y} \rangle$ and optimising $\tilde{\varepsilon}_{x,y}$ will also minimise the differences between bunches.

In practice, since computing $\langle \varepsilon_{x,y} \rangle$ is simpler (and faster) with the current software, we typically start our optimisation scans targeting those parameters, and switch to $\tilde{\varepsilon}_{x,y}$ once our algorithm starts to converge. Using the final results of CR1 as example (see Table 5.1) we can see that, even after optimisation, there is a significant difference between $\tilde{\varepsilon}$ and $\langle \varepsilon \rangle$, indicating that there are still differences between different bunches in the pulse.

5.3   Delay loop optimisation

The first section targeted for optimization was the delay loop since that's where we saw the first signs of the T_{566} aberration (see Figure 5.5). The delay loop is composed of 10 ICG cells with 8 sextupoles each. Since the cell is symmetric and all 10 cells are similar we can group this sextupoles in 4 families to reduce the number of free parameters in our optimization. To this we will add the 6 sextupoles located at the injection and extraction of the loop. Since the full loop is also symmetric, we can group those by pairs, adding 3 degrees of freedom to our optimization scan.

Table 5.1: Final emittance results (individual, average and pulse) of CR1.

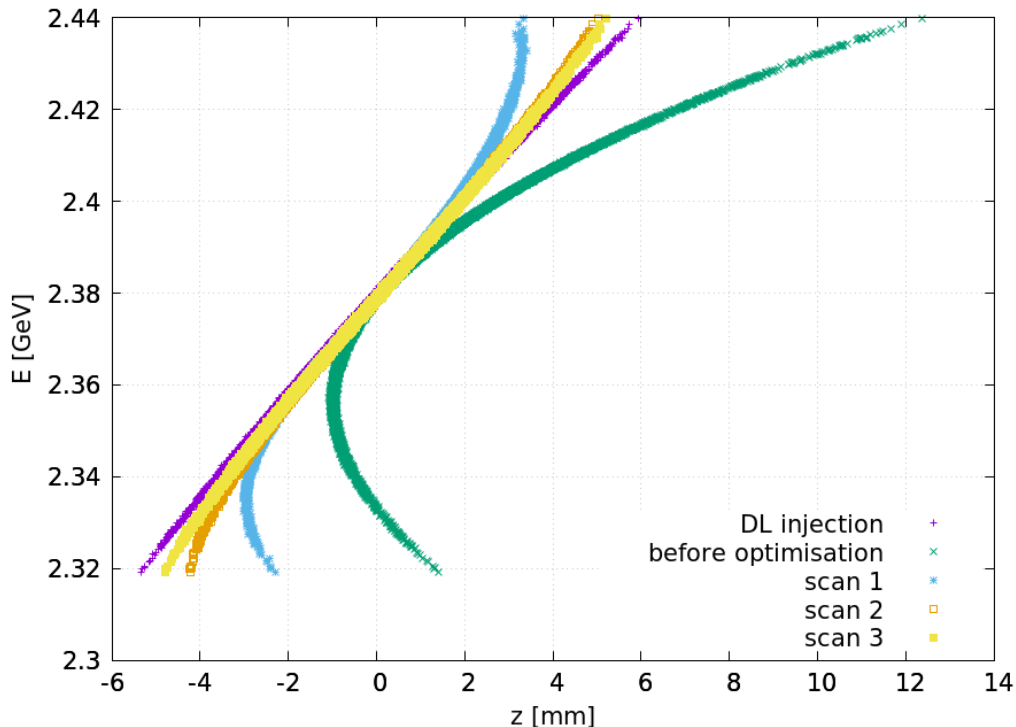| Emittance [μm]	$b_{0.5}^{j}$	$b_{1.5}^{j}$	$b_{2.5}^{j}$	$\langle \varepsilon_i \rangle$	$\tilde{\varepsilon}_i$
Horizontal	132	121	132	128	136
Vertical	116	113	130	119	122

Figure 5.5: Longitudinal energy profile at DL injection (first layer) and extraction. The second layer shows the distribution before sextupole optimisation and the subsequent layers show various optimisation results.

Optimisation scans rely heavily on the loss function used (see Section 4.3), particularly when trying to achieve diverging goals like, as we discovered, was the case when minimizing both ε_x and T_{566}. Figure 5.5 and Table 5.2 show the results of three optimisation scans performed with different weights for ε_x, ε_y and T_{566} in the loss function. After the scans, we managed to reduce ε_x and T_{566} but we suffered a small increase in ε_y as a trade-off.

Analysing Figure 5.5, the longitudinal profile of scan 1 is clearly distinguishable from the others. In this scan, the loss function contained a weighted sum of ε_x, ε_y and the quadratic coefficient of a second order polynomial interpolation fit. While that term was in fact minimized ($T_{566} \sim 0$), a third order correlation (U_{5666}) was developed instead. This revealed that, for our purposes, minimising this coefficient was not an effective strategy. For scans 2 and 3 (and, in fact, for all following scans) it was instead performed a linear polynomial fit, and the term added to the loss function was the error standard deviation value. After this study the optical strengths of the lattice sextupoles were updated to match the results of scan 3 since it offered the better trade-off between T_{566} and ε_x.

Table 5.2: Delay loop emittance before and after optimisation.

Bunch	before optimisation	scan 1	scan 2	scan 3
ε_x [μm]	123.2	115.1	120.7	119.7
ε_y [μm]	104.8	110.9	103.6	107.2

5.4 Combiner ring 1 optimisation

In order to fully recombine three 244 ns pulses, CR1's length needs to specifically match

$$L_{CR} = n_{rec} \frac{c}{f_{pulse}} \pm \frac{c}{M f_b} = 292.74 \pm 0.1 \, \text{m} \, , \tag{5.5}$$

where f_{pulse} is the pulse frequency, f_{bunch} is the bunch frequency inside a pulse, M is the ring multiplication factor (3× for CR1) and n_{rec} is the number of pulses being recombined simultaneously. As explained in Section 2.5, CR1 recombines two sets of 3 pulses simultaneously, so $n_{rec} = 2$. The first part of Eq. 5.5 ensures that the second and third pulses are injected in sync with the first, while the second part of the equation provides a small offset (depending on how many pulses we are recombining) to place the injected bunch in between bunches of the orbiting pulse.

Previous studies did not simulate recombination, focussing instead on beam dynamics. For this reason, the exact length of the lattice used (289.96 m) was not of paramount importance. `Placet2` allowed to quickly verify the discrepancy in the lattice's length and it was extended slightly to ensure CR1 recombines correctly. The choice between 292.64 m and 292.84 m is not particularly relevant to the machine design. We chose to extend the ring to 292.64 m to make the first injected bunch the head of the recombined pulse. This lengthening was performed by extending the straight sectors that connect CR1's bending arcs rather than altering the arcs themselves.

While for the DL the design of the linear optics was set in previous studies and not changed in ours, the lengthening of CR1 meant that the lattice no longer offered a periodic solution, forcing us to adjust the quadrupole optics to provide it. Since evaluating the twiss functions from a beam distribution can be problematic for chromatic beams, we started CR1's optimisation studies with a monochromatic beam and with all sextupoles turned off.

This procedure was performed by first determining the arcs' periodic solution and then optimising the straight sectors' quadrupole strengths to match this solution. Since not all straight sectors of CR1 have the same length or lattice, different sectors needed to be optimised independently.

The 52 quadrupolar optics of the straight sectors were grouped in 19 families based on lattice symmetries (see Appendix A for more details) and the lattice was optimised to match the twiss functions at the arcs' injection and extraction points to their periodic solution. In Figure 5.6 we show the evolution of β_x and β_y along the ring after optimisation. In it we can see that the six arcs all present a clearly symmetric twiss in both axis (being therefore periodic).

Once the linear optics have been set, we can utilize our chromatic beam coming from the delay loop via TL1[6]. We started by using all 74 of CR1's sextupoles grouped into 10 families to bring the emittance down and then, like for the DL, used only the sextupoles in dispersive regions of the arcs (see the blue line in Figure 5.6) to find a solution that kept both emittance and T_{566} to a minimum. Since our straight sectors are not all similar, not all arcs are attempting to correct the same T_{566}, and we found (after experimentation) that treating sextupoles in a similar position of the arc as being part of the same optimisation group (like we did for the delay loop) was insufficient for this task. We treated each arc independently

---

[6]Some TL1 optimisation was performed that we will not detail as it is similar in nature to the efforts in other sectors.

Figure 5.6: On the left axis, the horizontal (β_x) and vertical (β_y) β functions for a monochromatic beam along CR1's orbit after quadrupole optimisation. On the right axis, the horizontal dispersion (D_x) along the same orbit.

instead, which greatly increased the number of optimisation parameters and, therefore, also the required computing time.

This increased our optimisation's parameter-space from 4 to 24, which allowed us to get a solution with both low emittance and no T_{566}. The longitudinal profile of the beam before and after optimisation can be seen in Figure 5.7. Before optimisation, the beam's T_{566} gets progressively higher at each turn. After optimisation, the all bunches present negligible T_{566}, and the overall pulse horizontal and vertical emittances are 136 μm and 122 μm respectively (see Section 5.2 for more details on how pulse emittance is defined).

It is important to remember that, since we are recombining three pulses, correcting this effect must be done within the ring itself. After extraction from CR1 bunches $b_{0.5}^{j}$, $b_{1.5}^{j}$ and $b_{2.5}^{j}$ are too close together and whatever transformation we use to mitigate the effect on one of them we use in all of them. If we tried to correct this aberration after exiting the ring so that $T_{566}\left(b_{0.5}^{j}\right) = 0$, the correction would be insufficient for $b_{1.5}^{j}$ and $b_{2.5}^{j}$. On the other hand, if we targeted $T_{566}\left(b_{2.5}^{j}\right) = 0$, the other two bunches would be over-corrected, presenting a negative T_{566}. It is therefore paramount to ensure the combiner ring has no T_{566} contribution.

A final design detail is technically not an optimisation of CR1 optics, but it's an optimisation of CR1 performance. This is, of course, the injection. Once we have a periodic solution for the ring we need to ensure TL1's extraction matches the injection point's twiss. Additionally, the RF deflector that kicks the beam into centre orbit (see Section 2.5) adds some horizontal dispersion to the beam. If we were to correct this after the kicker our correction would remove dispersion from bunches being injected while adding dispersion to bunches already in orbit. To avoid this, we opt to inject a beam with some dispersion in order to nullify the amount added by the deflector.

5.5   Combiner ring 2 optimisation

Similarly to CR1, Eq. 5.5 showed that previous versions of CR2's lattice had a slight missmatch in length (~ 80 cm). We had therefore to extend the lattice to match the required length
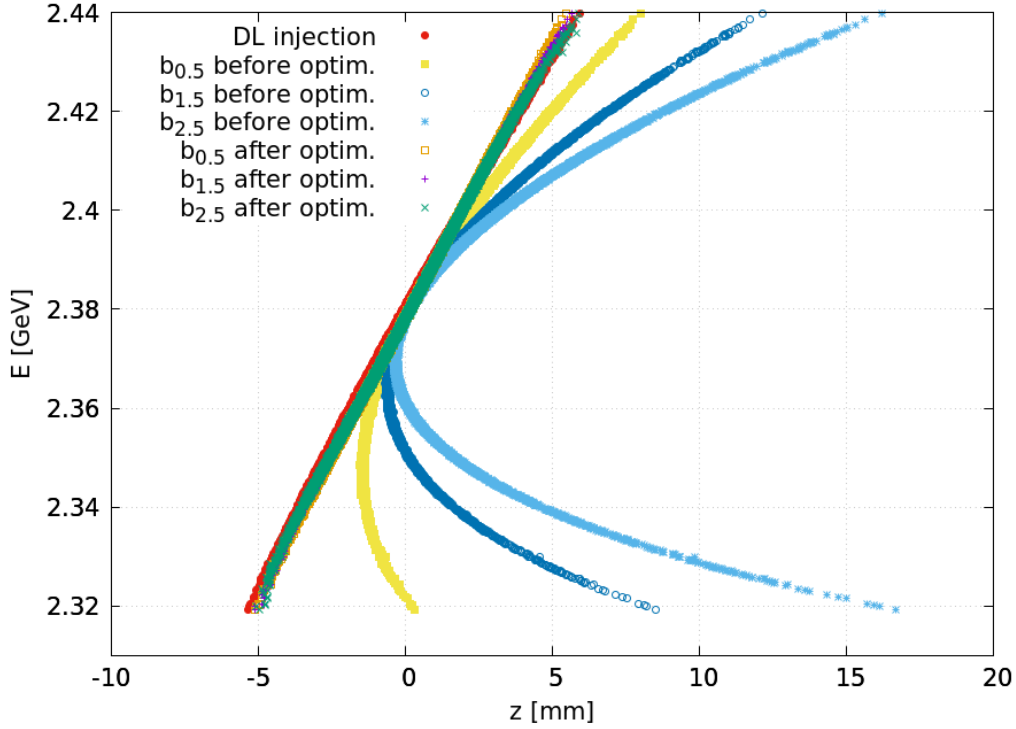
Figure 5.7: Longitudinal energy profile after decompression (first layer) and at CR1 extraction before and after sextupole optimisation.

of $439.09\,\text{m}$. This meant re-matching the quadrupolar optics to find a period orbit along the ring. This task was performed before optimising TL2 to transport the beam between the rings[7], and thus a Gaussian $\varepsilon_x = \varepsilon_y = 100\,\mu\text{m}$ beam distribution was used.

The initial optimisation procedure was similar to CR1. We used a monochromatic beam to find a periodic arc solution, and then used that solution as a target for the optimization of the straight sectors. Once we had a periodic solution for the ring (see Figure 5.8), we switched our monochromatic test beam to one with $\delta = 0.85\%$ and ran optimization scans on sextupoles to minimize emittance growth. This scan yielded a solution in which emittance growth is below 0.4% on both planes. Notice that, for this test, the beam used was fully Gaussian and didn't go through the rest of the machine. Additionally, the beam distribution used did not have a longitudinal energy chirp, and no effort was yet made to minimise the T_{566} aberration described in Section 5.1.

After finding a periodic solution for the ring, the focus of our efforts switched to optimising its injection scheme (see Figure 2.7). Since, in the third turn, the first RF deflector pushes the bunch off the centre orbit, the optics between deflectors need to bring the beam back to centre orbit at the second deflector with the correct angle and position so that the deflector kicks the bunch back to orbit. This cannot be done with more transverse deflectors since they would also kick turns 2 and 4, which are already centre-orbit. Instead, we make use of how quadrupoles and sextupoles affect the dynamics off off-centre bunches.

Starting with quadrupoles, whose thin lens approximation is

$$\begin{cases} x' = x'_0 - K_1 x_0 \\ y' = y'_0 + K_1 y_0 \end{cases}. \tag{5.6}$$

[7]TL2's lattice was also shortened by $32\,\text{m}$ to avoid beamline superpositions.

Figure 5.8: On the left axis, the horizontal (β_x) and vertical (β_y) β functions for a Gaussian beam along CR2's orbit after optimisation. On the right axis, the horizontal dispersion (D_x) along the same orbit.

If a bunch passes through the quadrupole off-centre by δ_x, we get

$$x' = x_0' - K_1 \left(x_0 + \delta_x \right) = x_0' - K_1 x_0 - K_1 \delta_x \, , \tag{5.7}$$

in which $\delta_x K_1$ is a fixed transverse kick, similar to that would be given by a static transverse deflector. This means we can bring the third turn bunch back into orbit using quadrupoles placed between the RF deflectors as shown in Figure 5.9. There are however three different requirements of this optical system. It needs to bring the bunch back to orbit, to offer a periodic twiss solution and to have zero dispersion after the second deflector. No solution was found that did all this using just the quadrupoles in the figure.

To increase the number of free parameters, sextupoles were added to the beamline (first two, then four, then six). Replacing x_0 by $x_0 + \delta_x$ in 3.18, we get

$$\begin{cases} x' = x_0' - K_2 \left(x_0^2 - y_0^2 \right) - 2\delta_x K_2 x_0 - \delta_x^2 K_2 \\ y' = y_0' + 2K_2 x_0 y_0 + 2\delta_x K_2 y_0 \end{cases} , \tag{5.8}$$
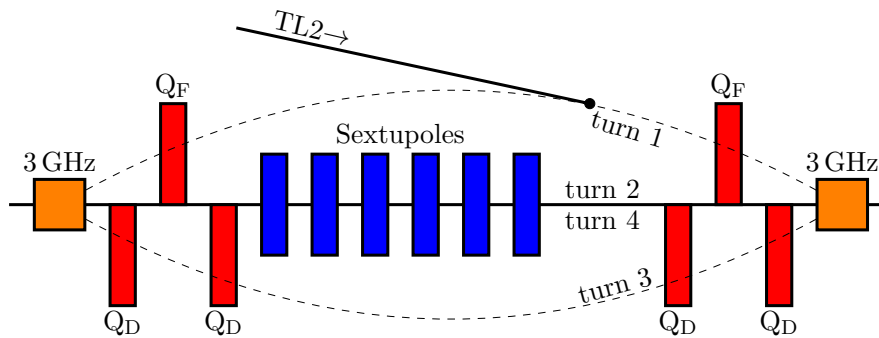

Figure 5.9: Lattice solution of CR2's injection scheme (not to scale). The horizontal line represents the ring's centre orbit with 2 RF deflectors (orange), 6 quadrupoles (red) and 6 sextupoles (blue). On the top we see the injection line and on the bottom the 3rd turn orbit bump.

Table 5.3: $b_{1.5}^{\,j}$ emittance at CR2's extraction point after optimisation. The last two columns show $\tilde{\varepsilon}_{x,y}$ for the $b_{1.5}^{\,j}$ and $b_i^{\,j}$ pulses, respectively.

Bunch	$b_{1.5}^{0.5}$	$b_{1.5}^{1.5}$	$b_{1.5}^{2.5}$	$b_{1.5}^{3.5}$	$b_{1.5}^{\,j}$	$b_i^{\,j}$
ε_x [μm]	131	129	140	150	144	189
ε_y [μm]	126	133	149	141	138	141

which describes the superposition of a sextupole of strength K_2, a quadrupole of strength $K_1 = 2\delta_x K_2$ and transverse kick of strength $\delta_x^{\,2} K_2$.

Optimising the injection optics together with TL2's extraction was, by far, the most computationally expensive of the studies being presented (over 10^5 iterations to achieve convergence). To diminish iteration time, we tracked only bunch $b_{1.5}^{\,j}$ and targeted $\tilde{\varepsilon}_{x,y}\left(b_{1.5}^{\,j}\right)$. Assuming $b_{0.5}$, $b_{1.5}$ and $b_{2.5}$ converged perfectly out of CR1, optimising for one of them should ensure all. This is however not the case. In Table 5.3, we can see that our optimisation scan managed to keep both horizontal and vertical emittance below $150\,\mu$m for $b_{1.5}^{\,j}$ but, when tracking all 12 bunches, our horizontal emittance was much higher than the budget.

It is hard to attribute this pulse emittance growth solely to CR2. If the upstream lattice provided bunches as similar as possible (as discussed in Section 3.5, there will always be a small energy difference between bunches that took different paths) the ring should be able to maintain emittance growth within the budget. The next step would then be to optimise the lattice upstream while targeting $\tilde{\varepsilon}_{x,y}$ at CR2's extraction point. This would require tracking all 12 bunch orbits simultaneously and that would be too computationally expensive to complete in this study. Similarly, we have postponed addressing the T_{566} aberration in CR2 until after the transverse emittance is within budget.

5.6   Full recombination

The final 1.26 km of the DBRC is composed of the transfer line 3 (17 m), the transfer turn around (1.22 km) and a compression chicane (23 m). The 6 quadrupoles and 6 sextupoles of TL3 were optimised to match the beam from CR2 to the TTA. Additionally, a small scan was made across the 3 sextupole families in the TTA cells in an attempt to keep emittance and particle loss low. No change was made to the chicane designed in [26].

In Table 5.4, we can see the final results of our study for each individual bunch path and the pulse results (final row). Notice that, as mentioned in Section 5.2, the emittance values presented were computed after removing the bunch particles with transverse coordinates above 3σ to reduce the influence of non-linear tails on the estimation. We can see that, in the vertical plane, both the full pulse and all individual bunches are well within the emittance budget ($\varepsilon_y \leq 150\,\mu$m). This behaviour is not unexpected considering the most relevant sources of emittance growth (bends and RF deflectors) do not operate on the vertical plane.

As for horizontal emittance, our final results yield $\tilde{\varepsilon}_x = 157\,\mu$m which, despite being above the $150\,\mu$m budget, we consider promising since our initial results (before any optimisation scans) were over an order of magnitude greater. We can see that the $b_{1.5}^{\,j}$ bunches (for which the final section of the complex was optimised) are well bellow the budget, so we expect that a full 12x recombination scan to bring $\tilde{\varepsilon}_x$ lower. Furthermore, once the DL short

Table 5.4: Total travelled orbit, individual bunch emittance (horizontal and vertical), R_{56} and T_{566} at the end of the DBRC for each of the 12 orbital paths of our study. In the final row, full pulse $\tilde{\varepsilon}_x$ and $\tilde{\varepsilon}_y$.

Bunch	S_{total} [m]	ε_x [μm]	ε_y [μm]	R_{56} [cm]	T_{566} [m]	σ_z [mm]
b_{inj}	–	100	100	12	0.0	1.00
$b_{2.5}^{3.5}$	4145	212	143	1.8	9.6	0.93
$b_{2.5}^{2.5}$	3706	220	135	1.9	7.3	0.72
$b_{2.5}^{1.5}$	3267	177	134	1.9	5.1	0.52
$b_{2.5}^{0.5}$	2828	147	128	1.9	2.8	0.32
$b_{1.5}^{3.5}$	3853	125	128	2.0	10	0.97
$b_{1.5}^{2.5}$	3414	134	123	2.0	7.8	0.76
$b_{1.5}^{1.5}$	2975	115	121	2.0	5.6	0.56
$b_{1.5}^{0.5}$	2536	116	117	1.9	3.3	0.36
$b_{0.5}^{3.5}$	3560	146	127	2.3	11	1.04
$b_{0.5}^{2.5}$	3121	147	124	2.2	8.4	0.82
$b_{0.5}^{1.5}$	2682	143	122	2.1	6.2	0.62
$b_{0.5}^{0.5}$	2243	128	116	1.9	4.0	0.42
b_i^j	–	157	127	–	–	–

path lattice is implemented, the 12 bunches going through it should present lower horizontal emittance than their long path counterparts.

Regarding longitudinal dynamics, our final results are less positive. Since neither CR2 nor TTA were optimised to minimize T_{566}, we can clearly see in the Figure 5.10 that each turn in CR2 has added to it. Furthermore, the beam seems to have lost its original R_{56} correlation in the final sections of the DBRC. Since the re-compression chicane was designed to operate in a $R_{56} = 12\,\text{cm}$ beam, it over-compresses the beam. Notice that, if all bunches had similar longitudinal properties (σ_z, R_{56} and T_{566}), the chicane design could be updated to accommodate any deviation from the original values. Our main concern is the difference in length between $b_i^{0.5}$, $b_i^{1.5}$, $b_i^{2.5}$ and $b_i^{3.5}$ at the end of the complex[8].

As mentioned in Section 1.3, the efficiency of the PETS is highly dependent on bunch length. It is therefore paramount to ensure that σ_z is stable. Furthermore, our simulations were implemented assuming the machine operates in a CSR-free regime. Should that not be the case, implementation of that effect in `Placet2` would be required for a realistic simulation of the machine.

---

[8]There is also a non-negligible difference between $b_{0.5}^j$, $b_{1.5}^j$ and $b_{2.5}^j$ that requires further study.
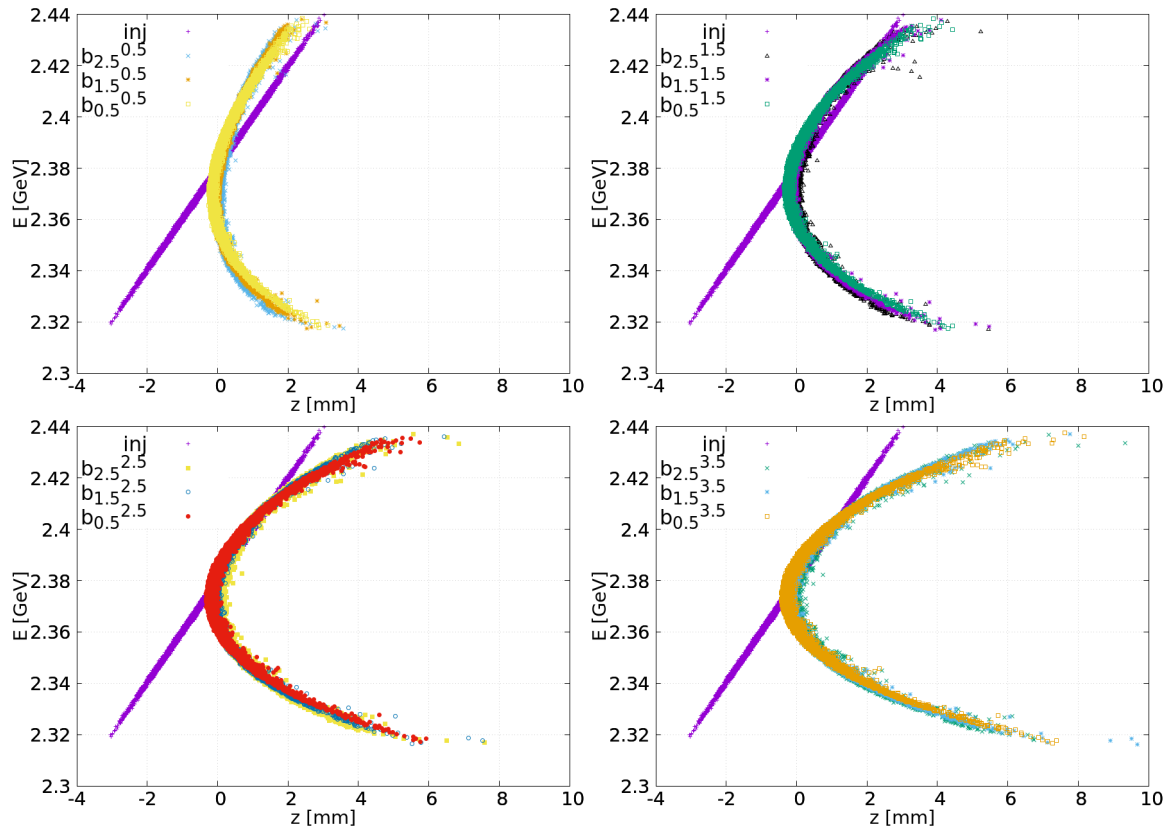
Figure 5.10: Longitudinal energy profile at injection and extraction of the DBRC for all 12 different bunch paths in comparison to the injected bunch (purple).

Chapter 6

Conclusions and Outlook

In this work, we present the first simulation and optimisation results of pulse recombination in CLIC's drive beam recombination complex. These results provided relevant inputs to the CLIC development, and were also presented in several international workshops [57–60]. To perform our studies, we implemented a DBRC lattice compatible with `Placet2` and updated the software accordingly (see Appendix B). Once recombination studies became possible, several issues were identified and the lattice has been updated to address them (see Appendix A).

The issues we found were an offset in the combiner rings' lengths, a previously unknown T_{566} aberration and a large overall horizontal emittance ($\tilde{\varepsilon}_x$) when combining bunches with different orbital paths. The rings' lengths were corrected to provide full pulse recombination and their twiss functions rematched by changing the optical strengths of the lattice's quadrupoles. After validating the origin of the T_{566} aberration with `Placet`, optimisation of sextupoles in dispersive regions of the DL and CR1 was able to mitigate it in those regions. Our final results still show T_{566} being produced in CR2 and the TTA but we expect the same procedure to be able to eliminate it. Regarding transverse emittance, our final results ($\tilde{\varepsilon}_y = 127\,\mu$m and $\tilde{\varepsilon}_x = 157\,\mu$m) are well below budget in the vertical plane and slightly above it in the horizontal. The most drastic lattice update proposed for this was the new CR2 injection lattice (see Figure 5.9), which we found necessary to keep $\tilde{\varepsilon}_x$ low after CR2.

Future studies of the DBRC should be focussed on achieving the horizontal emittance budget and eliminating the T_{566} aberration completely in a fully recombined pulse. This will require the design of the the DL's short path lattice and global optimisation while simulating 24 bunch paths[1]. In regards to longitudinal dynamics, it is imperative for PETS operations that all bunches have a similar length. It is therefore a priority to optimise CR2 to be fully isochronous. Once those goals have been achieved with an idealised lattice, the next step will be to induce static and dynamic errors in the simulations (misalignments, magnet strength errors, etc) and evaluate the robustness of the design. Beam-based alignment techniques will also need to be implemented.

---

[1]Some software development might be needed to keep these simulations fast.

Appendix A

Lattice updates

In this appendix we review the updates proposed to DBRC's lattice as a result of our work. Rather than presenting a cumbersome list of optical strengths, we give a qualitative list of changes. The full lattice is publicly available at CERN's subversion repositories [27].

Decompression dogleg

The decompression dogleg is 64 m long and contains 4 sbends, 16 quadrupoles and 2 sextupoles. In our optimisation, only the strengths of 4 quadrupoles (a matching section) were altered to meet the injection twiss of the DL.

Delay loop

The delay loop is a 143 m long lattice containing 52 sbends, 96 quadrupoles, 94 sextupoles and two 0.5 GHz RF transverse deflectors. The long path is an omega-shaped beamline composed by 12 ICG cells. The RF deflectors were added during our study and all sextupole strengths were altered as a result of our optimisation scans.

Transfer line 1

TL1 is a 63 m beamline connecting the DL to CR1. It is composed by 8 sbends, 25 quadrupoles and 22 sextupoles. Our studies altered the optical strengths of all sextupoles and the final 6 quadrupoles.

Combiner ring 1

CR1 is a 293 m long ring with 6 arcs. It contains 48 sbends, 154 quadrupoles and 74 sextupoles. Each arc is composed by a modified double ICG cell. It also contains one stripline extractor (added by our study) and 6 transverse deflectors, one of which is a 1 GHz RF kicker. Our studies added 2.9 m to CR1's length. They also altered the strengths of all sextupoles and the 52 quadrupoles in the straight sectors.

Transfer line 2

TL2 is a 350 m beamline connecting both combiner rings. It is composed by 16 sbends, 107 quadrupoles and 44 sextupoles. Our study shortened it by 32 m to avoid the superposition of

the TTA with CR1. We also altered the strengths of 45 quadrupoles at injection and extraction. The cells of the main FODO line remained unchanged.

Combiner ring 2

CR2 is a 446 m ring with with 8 arcs. It contains 64 sbends, 210 quadrupoles, 177 sextupoles, two 3 GHz RF deflectors and one stripline extractor. Each arc is composed by a modified double ICG cell. Our studies changed the injection scheme and added 80 cm to the ring's length, the stripline extractor, 6 quadrupoles and 6 of the sextupoles. They also altered the optical strength off all sextupoles and 96 quadrupoles, 24 of which formed the central matching in the arc cells.

Transfer line 3

TL3 is a 17 m long beamline connecting CR2 to the TTA. It is composed by 4 sbends, 6 quadrupoles and 6 sextupoles. All quadrupolar and sextupolar optical strengths were altered during our studies.

Transfer turn around

TTA is a 1.22 km long beamline composed by 38 bending cells which have 144 sbends, 228 quadrupoles and 228 sextupoles. Our studies only changed the optical strength of the sextupoles, grouping them into three families based on cell symmetry.

Re-compression chicane

The final beamline section of the DBRC is a compression chicane composed by 8 sbends and 10 quadrupoles. No changes were done to this section of the lattice.

Appendix B

Placet2 updates

In this appendix we give a quick summary of the changes implemented in `Placet2` for our studies. The software is open source under the GNU general public licence and the source-code is publicly available at CERN's subversion repositories [54].

BPM element

Beam position monitors (BPMs) are sensors placed long the beamline to measure the beam horizontal and/or vertical position. Thought `Placet2` offers several tracking options when defining a bunch, including its position at every element of the beamline, it is important when studying beam-based alignment techniques to have BPMs in the lattice and define their resolution and scale coefficients. As such, a BPM element was added to `Placet2` as well as a readout that collects horizontal and vertical position data from all BPMs in the lattice to an output file after adding some error noise based on the BPM resolution.

Stripline extractor element

Both combiner rings have stripline extractor kickers in their design. These elements operate as horizontal deflectors with a time dependent step function (design examples can be found in [61, 62]). They are kept off (behaving as simple drifts) while the ring is recombining and perform horizontal kick for the time of pulse (244 ns).

Losses readout

A losses output was added to track in which element of the beamline particles were lost by reaching the aperture as well as their 6D coordinates.

Dispersion-free readouts

As mentioned in Section 5.2, we implemented the ability to output dispersion-free emittance and twiss parameters along the beamline. Though our use for this procedure only involves first order dispersion, the implementation allows for the the removal of any higher order chromatic correlation.

Quadrupolar bends

Sbends may be build with non-parallel poles, which adds a quadrupolar component (K_1) to their field. Since the DBRC's design contains 164 of these bends, `Placet2`'s Sbend element was updated to allow for this option. The sbend Matrix 3.10 was therefore replaced by

$$
R_{Qb} = \begin{bmatrix}
\cos k_x L & \frac{\sin k_x L}{k_x} & 0 & 0 & 0 & \frac{1-\cos k_x L}{k_x^2 \rho_0} \\
-k_x \sin k_x L & \cos k_x L & 0 & 0 & 0 & \frac{\sin k_x L}{k_x \rho_0} \\
0 & 0 & \cos k_y L & \frac{\sin k_y L}{k_y} & 0 & 0 \\
0 & 0 & -k_y \sin k_y L & \cos k_y L & 0 & 0 \\
\frac{\sin k_x L}{k_x \rho_0} & \frac{1-\cos k_x L}{k_x^2 \rho_0} & 0 & 0 & 1 & \frac{k_x L - \sin k_x L}{k_x^3 \rho_0^2} \\
0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}, \qquad \text{(B.1)}
$$

where $k_x = \sqrt{K_1 + \frac{1}{\rho_0^2}}$ and $k_y = \sqrt{-K_1}$. Just as for quadrupoles (Matrix 3.5), the focusing transverse direction follows euclidean geometry while the defocusing follows hyperbolic.

Beamline Cartesian draw

Lattice files are just a sequence of elements, sometimes several thousand lines long, making discrepancies in the physical positioning of the beamline hard to spot. To avoid possible issues in the future, we created a function that draws the beamlines in Cartesian coordinates for verification. This was also the function used to produce Figure 4.1.

Octave and Python API

The interface between `Placet2` and generic data processing environments like `Octave` [55] and `Python` greatly increases the range of capabilities of the tracking code. For our studies in particular, interfacing with `Octave` gives us access to Nelder and Mead's simplex optimisation [56]. To allow the use of the simplex algorithm to maximize machine performance we added to the interface the capability of changing element parameters automatically (for example changing a sextupole's S_2 inside the loss function).

Arbitrary twiss and dispersion transformations

As mentioned in the previous paragraph and in Section 4.3, the ability to automate changes in the simulations is paramount for the optimisation procedure. When optimising a section of the machine (for example CR1) for which we are free to match the previous section (the only requirement of TL1's extraction is to match CR1 correctly), it is simpler to use the beam properties (twiss and dispersion) as optimisation parameters. This allows the optimisation algorithm to find the best solution for the section and, in a later study, the previous section can be independently made to match the required beam properties. With that in mind, the `Ktwiss` thin kick was implemented to make artificial changes to the beam's dispersion and twiss parameters. The technique described was used multiple times in our studies, namely to optimise the combiner rings independently from the injected beam.

Bibliography

[1] M. Aicheler, P. Burrows, M. Draper, et al. A Multi-TeV Linear Collider Based on CLIC Technology. 2012. doi: 10.5170/CERN-2012-007. URL https://cds.cern.ch/record/1500095.

[2] M. Livingstone and J. Blewett. *Particle Accelerators*. McGraw-Hill Book Company Inc., 1962.

[3] E. Bloom, D. Coward, H. DeStaebler, et al. High-energy inelastic $e - p$ scattering at 6º and 10º. *Phys. Rev. Lett.*, 23:930–934, Oct 1969. URL https://link.aps.org/doi/10.1103/PhysRevLett.23.930.

[4] M. Perl, G. Abrams, A. Boyarski, et al. Evidence for anomalous lepton production in $e^+ - e^-$ annihilation. *Phys. Rev. Lett.*, 35:1489–1492, Dec 1975. URL https://link.aps.org/doi/10.1103/PhysRevLett.35.1489.

[5] G. Arnison, A. Astbury, B. Aubert, et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s} = 540$GeV. *Phys. Lett. B*, 122(CERN-EP-83-13):103–116. 31 p, Jan 1983. URL https://cds.cern.ch/record/142059.

[6] G. Arnison, A. Astbury, B. Aubert, et al. Experimental observation of lepton pairs of invariant mass around 95GeV/c^2 at the CERN SPS collider. *Phys. Lett. B*, 126(CERN-EP-83-073):398–410. 17 p, Jun 1985. URL https://cds.cern.ch/record/163857.

[7] G. Aad, T. Abajyan, B. Abbott, et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716(CERN-PH-EP-2012-218):1–29. 39 p, Aug 2012. URL https://cds.cern.ch/record/1471031.

[8] S. Chatrchyan, V. Khachatryan, A. Sirunyan, et al. Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV. *JHEP*, 06(CMS-HIG-12-036. CMS-HIG-12-036. CERN-PH-EP-2013-035):081. 119 p, Mar 2013. URL https://cds.cern.ch/record/1529865.

[9] R. Neal. *The Stanford two-mile accelerator*. Number v. 1. W. A. Benjamin, 1968.

[10] T. Pettersson, P. Lefèvre, D. Boussard, et al. The Large Hadron Collider: conceptual design. Technical Report CERN-AC-95-05-LHC, Oct 1995. URL http://cds.cern.ch/record/291782.

[11] M. Mangano. *Physics at the FCC-hh, a 100 TeV pp collider.* CERN Yellow Reports: Monographs. CERN, Geneva, 2017. URL `https://cds.cern.ch/record/2270978`.

[12] T. Behnke, J. Brau, B. Foster, et al. The International Linear Collider Technical Design Report - Volume 1: Executive Summary. 2013. URL `https://cds.cern.ch/record/1601966`.

[13] M. Altarelli, R. Brinkmann, M. Chergui, et al., editors. *XFEL, the European X-ray free-electron laser : technical design report.* DESY XFEL Project Group, Hamburg, 2006. ISBN 978-3-935702-17-1. URL `http://repository.gsi.de/record/66728`.

[14] R. Edgecock. Industrial Application of Accelerators. Feb. 8 2017. URL `https://indico.cern.ch/event/569698/`.

[15] M. Aaboud, G. Aad, B. Abbott, et al. Search for resonances in diphoton events at $\sqrt{s}$=13 TeV with the ATLAS detector. *JHEP*, 09, Jun 2016. URL `https://cds.cern.ch/record/2160228`.

[16] V. Khachatryan, A. Sirunyan, A. Tumasyan, et al. Search for resonant production of high-mass photon pairs in proton-proton collisions at $\sqrt{s} = 8$ and 13 TeV. *Phys. Rev. Lett.*, 117, Jun 2016. URL `https://cds.cern.ch/record/2160409`.

[17] M. Aaboud, G. Aad, B. Abbott, et al. Search for new phenomena in high-mass diphoton final states using 37 fb^{-1} of proton–proton collisions collected at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Submitted to: Phys. Lett.*, 2017. URL `https://cds.cern.ch/record/2274538`.

[18] V. Khachatryan, A. Sirunyan, A. Tumasyan, et al. Search for high-mass diphoton resonances in proton-proton collisions at 13 TeV and combination with 8 TeV search. *Phys. Lett. B*, 767, Sep 2016. URL `https://cds.cern.ch/record/2214171`.

[19] M. Boland, U. Felzmann, P. Giansiracusa, et al. Updated baseline for a staged Compact Linear Collider. Technical Report CERN-2016-004. CERN-2016-004, Geneva, Aug 2016. URL `https://cds.cern.ch/record/2210892`.

[20] *LEP design report.* CERN, Geneva, 1984. URL `https://cds.cern.ch/record/102083`.

[21] B. Allanach, J. Blaising, K. Desch, et al. The physics benchmark processes for the detector performance studies used in CLIC CDR Volume 3. Jul 2012. URL `https://cds.cern.ch/record/1463554`.

[22] E. Adli, D. Schulte, and S. Stapnes. A Study of the Beam Physics in the CLIC Drive Beam Decelerator, 2009. URL `https://cds.cern.ch/record/1239173`.

[23] C. Biscari, D. Alesini, A. Ghigo, F. Marcellini, and J. Jeanneret. CLIC Drive Beam Frequency Multiplication System Design. In *Particle accelerator. Proceedings, 23rd Conference, PAC'09, Vancouver, Canada, May 4-8, 2009*, page WE6PFP076, 2010. URL `http://accelconf.web.cern.ch/AccelConf/PAC2009/papers/we6pfp076.pdf`.

[24] P. K. Skowronski, J. Barranco, F. Tecker, and C. Biscari. Design of the CLIC Drive Beam Recombination Complex. *Conf. Proc.*, C110904:1045–1047, 2011. URL `https://cds.cern.ch/record/1381582`.

[25] J. Esberg, A. Latina, R. Apsimon, and D. Schulte. Effect of CSR shielding in the compact linear collider. (CERN-ACC-2014-0152):4 p, Jun 2014. URL `https://cds.cern.ch/record/1742292`.

[26] E. Marin. Update on the CLIC Drive Beam lattice optimization studies. In *CLIC Workshop 2016*, Geneva, Switzerland, Jan. 18 2016. URL `https://indico.cern.ch/event/449801/contributions/1945359/`.

[27] Drive beam recombination complex lattices subversion repository, 2017. URL `http://svnweb.cern.ch/world/wsvn/clicsim/trunk/CLIC/Lattices/DriveBeam/DBRC_CLIC_Placet2_v08/`. CERN.

[28] W. Lau et al. Coherent Synchrotron Radiation Source Based on an Isochronous Accumulator Ring with Femtosecond Electron Bunches. *Conf. Proc.*, C110904:2085–2087, 2011. URL `http://inspirehep.net/record/1183148/`.

[29] D. Carey. *The optics of charged particle beams*. Accel. Storage Rings. Harwood, Chur, 1987. URL `https://cds.cern.ch/record/111767`.

[30] K. Wille. *The physics of particle accelerators: an introduction*. Oxford Univ. Press, Oxford, 2000. URL `https://cds.cern.ch/record/560708`.

[31] S. Lee. *Accelerator physics; 3rd ed.* World Scientific, Singapore, 2012. URL `https://cds.cern.ch/record/1425444`.

[32] A. Chao, K. Mess, M. Tigner, and F. Zimmermann. *Handbook of accelerator physics and engineering; 2nd ed.* World Scientific, Singapore, 2013. URL `https://cds.cern.ch/record/1490001`.

[33] W. Herr. Advanced Accelerator Physics. Trondheim, Norway, 2013. CERN Accelerator School, CERN. URL `https://cds.cern.ch/record/1507631`.

[34] A. Latina. Introduction to Transverse Beam Dynamics. France, 2016. Joint Universities Accelerator School. URL `https://indico.cern.ch/event/471931/contributions/1149878/`.

[35] E. Courant and H. Snyder. Theory of the Alternating-Gradient Synchrotron. *Ann. Phys.*, 3:1–48, 1958. URL `https://cds.cern.ch/record/593259`.

[36] D. Pellegrini, A. Latina, and D. Schulte. PLACET2: A Novel Code for Beam Dynamics in Recirculating Machines. (CERN-ACC-2015-322):MOPJE068. 3 p, 2015. URL `https://cds.cern.ch/record/2141792`.

[37] D. Pellegrini, L. Rivkin, and A. Latina. Beam Dynamics Studies in Recirculating Machines, 2016. URL `https://cds.cern.ch/record/2235763`. presented 8 Apr 2016.

[38] L. Deniau, R. Ghislain, I. Tecker, H. Grote, and F. Schmidt. The MAD-X Program User's Reference Manual, 2016. URL `https://mad.web.cern.ch/mad`.

[39] M. Borland. Elegant: A Flexible Sdds-compliant Code For Accelerator Simulation. Technical Report LS-287, Argonne Nat. Lab., Argonne, IL, 2000. URL `https://cds.cern.ch/record/511027`.

[40] D. Schulte, G. Guignard, N. Leros, and E. d'Amico. Simulation Package based on PLACET. In *Proc. 19th IEEE Particle Accelerator Conf., 18-22 June 2001,Chicago, Illinois, USA*, pages 3033–3035, 2001. URL `https://cds.cern.ch/record/507786`.

[41] A. Latina, E. Adli, H. Burkhardt, et al. Recent Improvements in the tracking code PLACET. In *Proc. 11th European Particle Accelerator Conf., 23-27 June 2008, Genoa, Italy*, pages 1750–1752, 2008. URL `https://cds.cern.ch/record/1122948`.

[42] A. Latina, E. Adli, B. Dalena, D. Schulte, and J. Snuverink. Improvements in the Placet tracking code. *Conf. Proc.*, C1205201(CERN-ATS-2012-166):MOPPC073. 3 p, May 2012. URL `https://cds.cern.ch/record/1463325`.

[43] A. Latina, Y. Levinsen, D. Schulte, and J. Snuverink. Evolution of the Tracking Code PLACET. In *Proceedings, 4th International Particle Accelerator Conference (IPAC 2013): Shanghai, China, May 12-17, 2013*, page MOPWO053, 2013. URL `http://cds.cern.ch/record/1574813`.

[44] D. Schulte, A. Latina, E. Adli, et al. The Tracking Code PLACET, 2015. URL `https://clicsw.web.cern.ch/clicsw/`.

[45] A. Latina, J. Pfingstner, D. Schulte, et al. Experimental demonstration of a global dispersion-free steering correction at the new linac test facility at SLAC. *Phys. Rev. ST Accel. Beams*, 17(4):042803, 2014. doi: 10.1103/PhysRevSTAB.17.042803. URL `https://cds.cern.ch/record/2135837`.

[46] A. Latina, E. Adli, N. Fuster-Martínez, et al. Tests of Wakefield-Free Steering at ATF2. In *Proceedings, IPAC 2015*, 2015. URL `http://accelconf.web.cern.ch/AccelConf/IPAC2015/papers/mopje059.pdf`.

[47] R. Bodenstein, P. Burrows, E. Marín, F. Plassard, and R. Tomás. Tuning Simulations for the CLIC Traditional Beam Delivery System. In *Proceedings, IPAC 2017*, 2017. URL `http://inspirehep.net/record/1626359/files/mopik108.pdf`.

[48] Y. Han, A. Latina, L. Ma, and D. Schulte. Static beam-based alignment for the Ring-To-Main-Linac of the Compact Linear Collider. *JINST*, 12(06):P06010. 12 p, 2017. URL `http://cds.cern.ch/record/2275925`.

[49] G. Paraskaki, A. Latina, and K. Kordas. Free Electron Lasers (FELs): Bunch Compressor Design for X FEL and Improvement of DFS algorithm for the FERMI FEL, 2016. URL `http://ikee.lib.auth.gr/record/285233/`.

[50] G. Geschonke and A. Ghigo. CTF3 Design Report. Technical Report CERN-PS-2002-008-RF. CTF3-NOTE-2002-047. LNF-2002-008-IR, CERN, Geneva, May 2002. URL `https://cds.cern.ch/record/559331`.

[51] J. Abelleira Fernandez, C. Adolphsen, A. Akay, et al. A Large Hadron Electron Collider at CERN: Report on the Physics and Design Concepts for Machine and Detector. Technical Report arXiv:1206.2913. CERN-OPEN-2012-015. CERN-LHeC-Note-2012-002-GEN, Jun 2012. URL `https://cds.cern.ch/record/1456019`.

[52] D. Angal-Kalinin, G. Arduini, B. Auchmann, et al. PERLE: Powerful Energy Recovery Linac for Experiments - Conceptual Design Report. PERLE: Powerful Energy Recovery Linac for Experiments - Conceptual Design Report. Technical Report arXiv:1705.08783, May 2017. URL `https://cds.cern.ch/record/2266048`.

[53] E. Aschenauer, M. Baker, A. Bazilevsky, et al. eRHIC Design Study: An Electron-Ion Collider at BNL. Technical Report arXiv:1409.1633, Sep 2014. URL `https://cds.cern.ch/record/1754028`.

[54] Placet2 subversion repository, 2017. URL `https://svnweb.cern.ch/world/wsvn/clicsw/trunk/placet2/`. CERN.

[55] J. Eaton, D. Bateman, S. Hauberg, and R. Wehbring. *GNU Octave version 4.0.0 manual: a high-level interactive language for numerical computations*. 2015. URL `http://www.gnu.org/software/octave/doc/interpreter`.

[56] J. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, January 1965. ISSN 1460-2067. URL `http://dx.doi.org/10.1093/comjnl/7.4.308`.

[57] R. Costa. First Start-to-End Simulations of the CLIC Drive Beam Recombination Complex. In *International Workshop on Future Linear Colliders (LCWS2016)*, Morioka, Japan, Dec. 6 2016. URL `https://agenda.linearcollider.org/event/7371/contributions/38011/`.

[58] R. Costa. Optimisation of the Drive Beam Recombination Complex. In *CLIC Workshop 2017*, Geneva, Switzerland, Mar. 6 2017. URL `https://indico.cern.ch/event/577810/contributions/2490754/`.

[59] R. Costa. Status of the CLIC Recombination Complex Lattices. In *International Workshop on Future Linear Colliders (LCWS2017)*, Strasbourg, France, Oct. 24 2017. URL `https://agenda.linearcollider.org/event/7645/sessions/4525/#20171024`.

[60] R. Costa. Drive Beam Recombination Complex Optimisation. In *CLIC Workshop 2018*, Geneva, Switzerland, Jan. 23 2018. URL `https://indico.cern.ch/event/656356/contributions/2832819/`.

[61] I. Rodriguez, A. Ghigo, L. García-Tabarés, F. Marcellini, and F. Toral. Design of a Strip-line Extraction Kicker for CTF3 Combiner Ring. page 3 p, 2006. URL `https://cds.cern.ch/record/1078557`.

[62] C. Belver-Aguilar, A. Faus-Golfe, F. Toral, and M. Barnes. Stripline design for the extraction kicker of Compact Linear Collider damping rings. *Phys. Rev. Spec. Top. Accel. Beams*, 17(7):071003. 14 p, 2014. URL `https://cds.cern.ch/record/2135841`.