Imperial College London

IMPERIAL COLLEGE LONDON

DEPARTMENT OF PHYSICS

Design of spiral scaling magnets for LhARA and FETS-FFA

Ta-Jen Kuo

August 21, 2026

A thesis submitted to Imperial College London for the degree of Doctor of Philosophy.

Abstract

Fixed Field Alternating gradient (FFA) accelerators represent a promising technology for high intensity and medical accelerator applications, offering relativistic energy capability, freedom from magnet ramping constraints, high repetition rate potential, and variable extraction energy without an energy degrader. This thesis presents a magnet design framework for scaling FFA configurations with flat pole geometries and radially distributed trim coils, allowing the field gradient to be varied radially through adjustment of individual coil currents.

The geometric requirements of the scaling FFA magnet are derived through analysis of its focusing properties, establishing that the magnet boundaries must conform to a spiral profile, the radial sector being a special case with zero spiral angl, and that the trim coil geometry must approximate arcs centred on the machine axis. An alternative derivation of the scaling law is presented as an original contribution, emphasising the transition from the field required along the trajectory to the final field distribution in polar form. Current settings are optimised to enforce a power law radial field distribution across each magnet, with longitudinal field clamps incorporated to better satisfy the scaling condition.

The author developed a suite of Python codes to parametrically construct magnet and coil geometries, interface with OPERA-3D, optimise coil currents, and extract field maps for tracking. The framework is applied to two lattices: the FETS-FFA, a novel doublet spiral lattice for a proposed ISIS-II upgrade, achieving a field index range of k=6.20–8.75 and tunes spanning 3-4, with maximum ring tune variation of ±0.04; and LhARA, a singlet spiral FFA optimised across 57 MeV and 127 MeV extraction energies, with tune points clear of resonances up to fourth order.

Finally, a radial sector prototype magnet is presented, designed and fabricated for experimental validation and as a benchmark for OPERA models.

Acknowledgements

I am deeply grateful to the colleagues within the Intense Beam Group at Rutherford Appleton Laboratory, especially Dr. Chris Rogers, Dr. David Kelliher, and the late Allen Letchford, for their thoughtful feedback at our weekly meetings and for always being willing to lend a hand.

A special thank you goes to the group leader, Dr. Shinji Machida, who never once made me feel like any question was too small. Whether we were discussing the basic operating principles of accelerators or the focusing properties of a negative k FFA, I treasured every one of those conversations. They have left a lasting impression on me and will continue to inspire my pursuit of accelerator physics for years to come.

I also owe a great deal to the LhARA group, namely Dr. Hywel Owen, Dr. William Shields, Dr. Colin Whyte, Dr. Ajit Kurup, and Professor Ken Long, for the wealth of ideas they brought to the table and for their unwavering support throughout the project.

Outside of research, a huge part of my time at Imperial was spent with the college basketball team, and I wouldn't have it any other way. To everyone who has been part of that journey — thank you. The experiences on and off the court have been truly exceptional. My time with Francesco, Argyris and Yiyang holds a particularly special place in my heart; I always came away from those moments feeling recharged and ready for whatever came next.

To my office companions in room 1110, Rohan Kamath, Rehanah Razak, and Dr. Paul-Bogdan Jurj, thank you for making campus life so much brighter. The PhD experience would have been a very different thing without you.

To the best supervisors I could have asked for: Dr. Jaroslaw Pasternak and Dr. Jean-Baptiste Lagrange. Thank you for guiding me with such care while giving me the freedom to explore the world of accelerator physics on my own terms. The curiosity you sparked in me and the warmth you showed me along the way are things I will carry with me always.

Throughout this journey, I never once felt alone; and that is in no small part thanks to my girlfriend, Claire. You understood the demands of this work without me ever having to explain them, and you were there through all of it. The trips and adventures we shared gave me space to breathe, and your words of encouragement gave me the strength to keep going. I truly could not have done this without you.

Finally, to my parents and sister, thank you for your love, which has never once wavered. It has always brought me joy to share my passion for this field with you, and being able to show you the prototype magnet I designed in person is a memory I will always treasure. Knowing that you have my back, no matter what, has made all the difference.

This thesis exists because of the people behind it, and I feel incredibly fortunate to have had such wonderful individuals by my side. It is the lessons and attitudes I have learnt from each and every one of them that will remain with me long after these pages.

Declaration and copyright

Declaration

I hereby declare that the contents of this thesis are the product of my own original work, save where explicitly acknowledged or referenced. The development of all models, the methods of analysis employed, and the results obtained were undertaken solely by me, and I retain full ownership of this material.

Copyright

The copyright of this thesis rests with the author. Unless otherwise indicated,its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Contents

1 Introduction **4**

1.1 Assumptions and coordinate system . . . 4
1.2 Transverse motion . . . 6
1.3 Betatron oscillation and tune . . . 8
1.4 Scaling Law . . . 10
1.5 Imperfections and resonances . . . 21
1.6 Focusing properties . . . 22
1.7 Origin and motivation for FFA accelerators . . . 23
1.8 Development of FFA . . . 25
1.9 Purpose of the study . . . 26

2 Magnet design **28**

2.1 2D studies: Field gradient and gap shaping . . . 28
2.2 Radially distributed trim coils . . . 29
2.3 FEM and relevant softwares . . . 30
2.4 Geometry of pole . . . 31
2.5 Geometry of clamp . . . 32
2.6 Geometry of coils . . . 33
2.7 Coil configurations . . . 39
2.8 Boundary condition . . . 43
2.9 Methodology . . . 43
2.9.1 Target field . . . 44
2.9.2 Jacobian . . . 45
2.9.3 Gauss-Newton algorithm . . . 46
2.9.4 Fringe field optimisation . . . 46
2.9.5 Tracking . . . 48

3 FETS-FFA **50**

3.1 High intensity accelerator . . . 50
3.2 ISIS-II and FETS-FFA . . . 50
3.3 FETS-FFA lattice design . . . 51
3.4 FETS-FFA magnet design . . . 53
3.5 Optimisation procedure . . . 66

3.6 Initial optimisation . . . 70
3.6.1 Current optimisation . . . 70
3.6.2 Tracking . . . 76
3.7 Clamp optimisation . . . 78

4 LhARA **84**
4.1 Radiotherapy and Bragg peak . . . 84
4.2 LhARA . . . 85
4.3 LhARA FFA lattice design . . . 86
4.4 LhARA FFA magnet design . . . 87
4.5 Initial optimisation . . . 93
4.5.1 Current optimisation . . . 93
4.5.2 Tracking . . . 99
4.6 Clamp optimisation . . . 101
4.7 Flexibility in tune space . . . 111

5 Prototype magnet **113**
5.1 Lattice and magnet parameters . . . 113
5.2 Optimisation . . . 114
5.3 Future work . . . 116

6 Conclusion **120**

Bibliography **123**

A Numerical procedure to produce trim coil geometry **128**

B Metaopera **130**

C Perpedicular distance between spirals **140**

D Set up model with Metaopera **142**

E Current optimisation procedure **149**

Chapter 1

Introduction

The introduction chapter covers the basic transverse dynamics required to achieve a stable beam in fixed field machines. The particular field distribution required for this project, known as the scaling law, is derived using a new approach developed by the author, with the associated focusing properties explained.

1.1 Assumptions and coordinate system

In a circular accelerator, a transverse force is needed to keep the particles in orbit. The force felt by a particle of charge q and velocity $\mathbf{v}$ is described by the Lorentz equation:

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \tag{1.1}$$

where $\mathbf{E}$ is the electric field and $\mathbf{B}$ the magnetic field. Assuming $\mathbf{v} \perp \mathbf{B}$, for particles moving at relativistic speeds ($v \approx c$), producing the same force with an electric field as with a 1 T magnetic field would require an impractically large electric field of about $\mathbf{E} \approx 300$ MV/m. Consequently, in the majority of modern accelerators, magnetic fields are employed to guide and focus particles in the transverse plane, whilst electric fields are applied longitudinally, parallel to the particle trajectory, for acceleration. This separation of roles is not universal, betatron accelerators represent a notable exception, employing a time varying magnetic field to achieve acceleration. Such cases fall outside the scope of the present discussion.

In the absence of an electric field within the magnet region, the particle trajectory is determined by the equilibrium between the Lorentz force and the centripetal force:

$$\frac{\gamma m_0 \mathbf{v}^2}{\rho} = q\mathbf{v} \times \mathbf{B} \tag{1.2}$$

where m_0 is the rest mass of the particle and ρ the bending radius of the trajectory. For the case $\mathbf{v} \perp \mathbf{B}$, this expression reduces to the beam rigidity relation:

$$B\rho = \frac{p}{q} \tag{1.3}$$

where p is the momentum of the particle.

The position in the vicinity of the equilibrium orbit is described using the following right handed system of curvilinear coordinates (x, s, z):

x — the horizontal displacement from the equilibrium orbit, taken as positive in the outward radial direction,

s — the arc length along the equilibrium orbit, measured from a fixed reference point,

z — the vertical displacement from the equilibrium orbit.

This coordinate system is illustrated in Fig. 1.1. The longitudinal coordinate s serves as the independent variable, making it a natural choice for describing particle motion in the neighborhood of the equilibrium orbit.

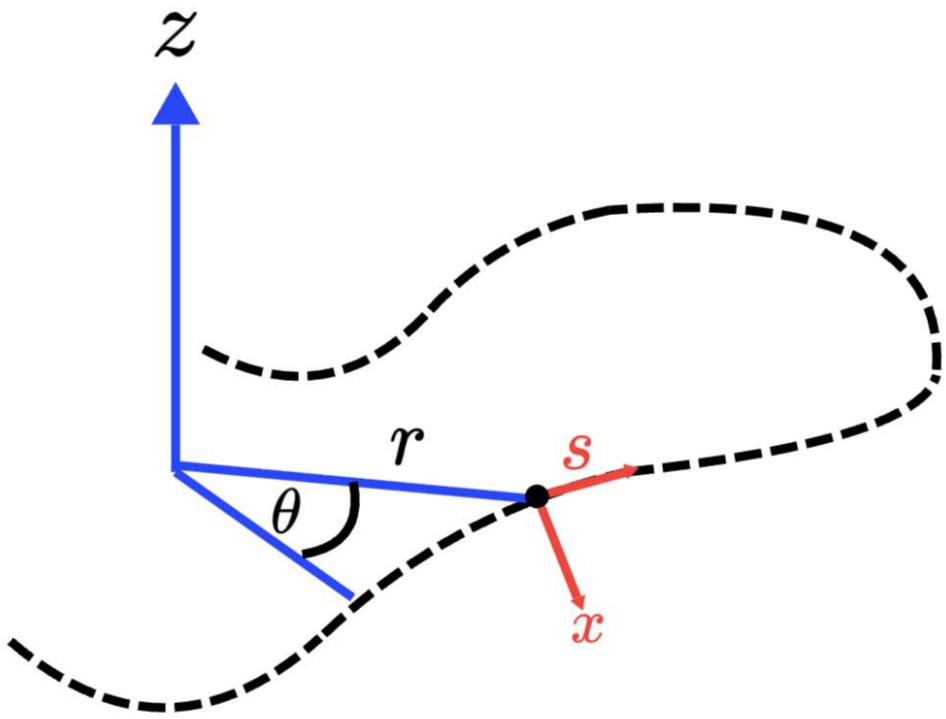


Figure 1.1: Moving coordinate system along a non circular equilibrium orbit.

A magnetic field **B** is considered such that there exists a plane in which the field vector is everywhere perpendicular to it. This plane is referred to as the median plane and is taken to be horizontal, as vertical bending is absent in most accelerator rings. In the median plane, defined by $z = 0$, the only non-vanishing component of the magnetic field is B_z

It is further assumed that, within this plane, a particle of momentum p follows a unique closed orbit, which is defined as the equilibrium orbit and is denoted by $(0, s, 0)$

1.2 Transverse motion

For a magnetic field to serve as a suitable guiding field in a particle accelerator, it must ensure the dynamical stability of particle motion in the vicinity of the equilibrium orbit. Specifically, a particle whose momentum corresponds to that of the equilibrium orbit, upon receiving a small initial displacement, must exhibit subsequent motion that remains bounded within the neighbourhood of the equilibrium orbit for all time.

By considering a small transverse displacement x from the equilibrium orbit, the vertical magnetic field component B_z can be expanded as a Taylor series:

$$B_z(x) = B_z(0) + x\left(\frac{\partial B_z}{\partial x}\right)_{\perp} + \frac{x^2}{2!}\left(\frac{\partial^2 B_z}{\partial x^2}\right)_{\perp} + \frac{x^3}{3!}\left(\frac{\partial^3 B_z}{\partial x^3}\right)_{\perp} + \mathscr{O}(x^4) \tag{1.4}$$

where $\perp$ denotes that the partial derivative is evaluated in the median plane, perpendicular to the closed orbit, at $x = 0$. Each term in the expansion corresponds to a distinct multipole component and has a different effect on the beam. The zeroth order term $B_z(0)$ is the dipole contribution, responsible for bending the beam trajectory. The first order term $\frac{dB_z}{dx}$ is the quadrupole contribution, giving rise to a linearly varying force that focuses the beam. Higher order terms introduce additional effects whose treatment falls outside the scope of the present discussion.

Synchrotrons utilise dedicated, separate magnets for each multipole component, permitting independent control of the orbit, beam focusing, and stability. As the beam momentum increases throughout acceleration, the magnetic field strengths are ramped accordingly to maintain the particles on a fixed closed orbit. Fixed field accelerators, by contrast, employ a single magnet to provide all multipole components simultaneously. In such devices, the magnetic field is static in time but varies with transverse position, such that the equilibrium orbit migrates with beam momentum. As the present work is concerned exclusively with fixed field magnets, separate function magnet systems will not be discussed further.

Applying Ampère's law along the s direction:

$$(\nabla \times \mathbf{B})_s = \left(\frac{\partial B_x}{\partial z}\right)_{\perp,v} - \left(\frac{\partial B_z}{\partial x}\right)_{\perp} = 0 \tag{1.5}$$

where $\perp, v$ signifies that the partial derivative is done perpendicularly in the vertical plane, and evaluated at $z = 0$. An analogous Taylor expansion may be performed for the B_x field component

in the z direction:

$$B_x(z) = \cancel{B_x(0)}^{0} + \underbrace{z\left(\frac{\partial B_x}{\partial z}\right)_{\perp,v}}_{=\left(\frac{\partial B_z}{\partial x}\right)_{\perp}} + \cancel{\frac{z^2}{2!}\left(\frac{\partial^2 B_x}{\partial z^2}\right)_{\perp,v}}^{0} + \frac{z^3}{3!}\left(\frac{\partial^3 B_x}{\partial z^3}\right)_{\perp,v} + \mathscr{O}(z^4) \tag{1.6}$$

By virtue of the assumed median plane symmetry, $B_x(z)$ is an antisymmetric function of z, and consequently all even order terms in the expansion vanish.

Retaining terms up to linear order and invoking Eq. 1.2, the linearised equation of motion in the vicinity of a single closed orbit may be written as:

$$\begin{aligned} F_x = \gamma m_0 \frac{d^2x}{dt^2} &= -qvB_z(x) + \frac{\gamma m_0 v^2}{\rho + x} \\ &\approx -qv\left[B_0 + \left(\frac{\partial B_z}{\partial x}\right)_{\perp} x\right] + \frac{\gamma m_0 v^2}{\rho}\left(1 - \frac{x}{\rho}\right) \\ &\Rightarrow \gamma m_0 \frac{d^2x}{dt^2} = -qvx\left(\frac{\partial B_z}{\partial x}\right)_{\perp} - \frac{\gamma m_0 v^2}{\rho^2} x \\ &= -\frac{\gamma m_0 v^2}{\rho^2}\left[\frac{\rho}{B_0}\left(\frac{\partial B_z}{\partial x}\right)_{\perp} + 1\right] x \end{aligned} \tag{1.7}$$

where $B_0 = B_z(x = 0, z = 0)$. An analogous expression may be derived for the vertical direction, in which the sole acting force is the Lorentz force:

$$\begin{aligned} F_z = \gamma m_0 \frac{d^2z}{dt^2} &= qvB_x(z) \\ &\approx qv\left(\frac{\partial B_x}{\partial z}\right)_{\perp} z = -qv\left(\frac{\partial B_z}{\partial x}\right)_{\perp} z \\ &= -\frac{\gamma m_0 v^2}{\rho^2}\left[\frac{\rho}{B_0}\left(\frac{\partial B_z}{\partial x}\right)_{\perp}\right] z \end{aligned} \tag{1.8}$$

where ρ denotes the radius of curvature of the equilibrium orbit, which varies with s according to:

$$\rho(s) = \frac{p}{qB_z(0, s, 0)} \tag{1.9}$$

Through a change of independent variable from t to s, the equation of motion may be expressed in terms of the longitudinal coordinate s:

$$
\begin{aligned}
\frac{d^2x}{dt^2} &= \frac{d}{dt}\left(\frac{dx}{dt}\right) = \frac{d}{dt}\left(\frac{dx}{ds}\frac{ds}{dt}\right) \\
&= \frac{d}{dt}\left(\frac{dx}{ds}v_s\right) = \frac{d}{ds}\left(\frac{dx}{ds}v_s\right)\frac{ds}{dt} \\
&= \frac{d}{ds}\left(\frac{dx}{ds}v_s\right)v_s = \frac{d}{ds}\left(\frac{dx}{ds}v_s^2\right) \\
&= \frac{d^2x}{ds^2}v_s^2 + \frac{dx}{\cancel{ds}}\overset{0}{\cancel{2v_s\frac{dv_s}{ds}}} \quad \text{(paraxial approximation)}
\end{aligned}
\tag{1.10}
$$

Substituting Eq. 1.10 into Eq. 1.7, the horizontal equation of motion in terms of s takes the form:

$$
\begin{aligned}
\gamma m_0 v^2 \frac{d^2x}{ds^2} &= -\frac{\gamma m_0 v^2}{\rho^2}\left[\frac{\rho}{B_0}\left(\frac{\partial B_z}{\partial x}\right)_\perp + 1\right]x \\
\Rightarrow \frac{d^2x}{ds^2} &+ \frac{1-n}{\rho^2}x = 0
\end{aligned}
\tag{1.11}
$$

A similar expression is obtained in the vertical direction:

$$
\begin{aligned}
\gamma m_0 v^2 \frac{d^2z}{ds^2} &= -\frac{\gamma m_0 v^2}{\rho^2}\left[\frac{\rho}{B_0}\left(\frac{\partial B_z}{\partial x}\right)_\perp\right]z \\
\Rightarrow \frac{d^2z}{ds^2} &+ \frac{n}{\rho^2}z = 0
\end{aligned}
\tag{1.12}
$$

A linear restoring force in each plane is characterised by the field index n, defined as

$$
n(s) = -\frac{\rho(s)}{B_0(s)}\left(\frac{\partial B_z(s)}{\partial x}\right)_\perp \tag{1.13}
$$

which appears with opposite sign in the horizontal and vertical equations of motion, reflecting the inherent coupling between transverse focusing in the two planes.

1.3 Betatron oscillation and tune

Since the equilibrium orbit is closed, the functions $n(s)$ and $p(s)$ are periodic in s. Eq. 1.11 and 1.12 therefore take the form of Hill's equation, a linear second order differential equation with periodic coefficients and no first derivative term. Both equations may be cast in the general form:

$$
\frac{d^2y}{ds^2} + K(s)y = 0, \tag{1.14}
$$

where y denotes either the horizontal or vertical displacement, and $K(s)$ satisfies the periodicity condition

$$
K(s+C) = K(s) \tag{1.15}
$$

with C denoting the circumference of the equilibrium orbit. In practice, the magnetic structure of an accelerator lattice typically comprises N identical cells, in which case $K(s)$ satisfies the stronger periodicity condition:

$$\begin{aligned} K(s+L) &= K(s) \\ L &= C/N \end{aligned} \tag{1.16}$$

At this point, it is instructive to recall certain properties of Hill's equation. The solution of any second order linear differential equation of the form of Eq. 1.14, irrespective of whether $K(s)$ is periodic, is uniquely determined by the initial values of y and its derivative y'. The general solution may be written as:

$$y = \sqrt{\varepsilon}\sqrt{\beta(s)}\cos\left(\nu\phi(s) + \phi_0\right) \tag{1.17}$$

where ε and ϕ_0 are arbitrary constants. The constant ε is commonly referred to as the emittance of the particle. The function $\beta(s)$ is periodic with period L.

The motion described by Eq. 1.17 corresponds to an oscillation of the particle about the equilibrium orbit, known as betatron oscillation. This motion is characterised by its frequency ν and by an amplitude modulated by the square root of the beta function $\beta(s)$.

The phase advance $\phi(s)$ is given by:

$$\phi = \int_s^{s+L} \frac{\mathrm{d}s}{\beta} \tag{1.18}$$

The parameter ν, referred to as the cell tune, is defined as:

$$\nu = \frac{\phi}{2\pi} \tag{1.19}$$

The ring tune Q, defined as:

$$Q = N\nu \tag{1.20}$$

represents the total number of betatron oscillations completed by a particle per revolution around the ring.

In subsequent notation, subscripts are employed to distinguish between horizontal and vertical quantities; specifically, $\nu_x$and ν_z denote the horizontal and vertical cell tunes respectively, and $\beta_x(s)$ and $\beta_z(s)$ the corresponding beta functions. Hill's equation is quasiharmonic in character, and any harmonic system is susceptible to resonance excitation. In accelerators, resonant driving terms may arise from numerous sources, including magnet imperfections and nonlinear field components. A general resonance condition relating the tunes in both transverse planes may be expressed as:

$$jQ_x + lQ_z = m \tag{1.21}$$

where (j,l,m) are all integers and the order of resonance is determined by $|j| + |l|$. In a periodic lattice without error, the condition is reduced to:

$$j\nu_x + l\nu_z = m \tag{1.22}$$

Resonances of this type are referred to as structural resonances.

When the working point (ν_x, ν_z) lies in the vicinity of a resonance condition, the amplitude of transverse oscillations may grow, leading to emittance degradation and potentially to beam loss. Accelerator designs therefore typically select working tune points well separated from resonance lines. It is further essential that the working tune point remains invariant with energy, so as to avoid the crossing of resonance lines during acceleration, unless the acceleration is sufficiently rapid that resonance crossings do not result in significant beam loss. The present discussion focuses on the former case. This motivates the concept of constant tune, which refers to the condition in which both the horizontal and vertical cell tunes remain unchanged with respect to the particle momentum.

1.4 Scaling Law

One approach to achieving constant tune in a fixed field machine is through zero chromaticity, whereby the focusing strength remains invariant with beam momentum. Since the equilibrium orbit changes with particle momentum in fixed field machines, the spatial distribution of the magnetic field must be carefully designed to meet the zero chromaticity condition. The necessary conditions may be identified by examining the linearised equations of motion, Eq. 1.11 and Eq. 1.12.

Each closed orbit may be expressed in cylindrical coordinates as $(r_c, \theta, 0)$, where:

$$r_c = r_c(p, \theta) \tag{1.23}$$

Here, r_c denotes the closed orbit radius, which may vary with the azimuthal angle θ even at a fixed momentum p. An equivalent radius R is also introduced, which for a given closed orbit (i.e. constant momentum), is defined as:

$$R = \frac{1}{2\pi} \oint ds \tag{1.24}$$

where s is the curvilinear abscissa along the closed orbit. The corresponding azimuthal coordinate Θ is defined through:

$$s = R\Theta \tag{1.25}$$

The origin of Θ is defined in Eq. 1.70. Since closed orbits do not intersect, each corresponds to a unique value of R. The uniqueness of the closed orbit for a given momentum implies that any point (r,θ) can be mapped to a unique (R,Θ), and the two coordinate systems may therefore be used interchangeably.

Field maps in fixed field magnets are typically expressed in cylindrical coordinates, it is convenient to reformulate the equations of motion in the same coordinate system. Differentiating x with respect to s along a closed orbit R, and making use of Eq.1.25:

$$\begin{aligned}\frac{dx}{ds} &= \frac{dx}{d\Theta}\left(\frac{\partial\Theta}{\partial s}\right)_R \\ &= \frac{dx}{d\Theta}\frac{1}{R}\end{aligned} \tag{1.26}$$

$$\begin{aligned}\frac{d^2x}{ds^2} &= \frac{d}{ds}\left(\frac{dx}{d\Theta}\right)\frac{1}{R} + \frac{dx}{d\Theta}\cancel{\frac{d}{ds}}\overset{0}{\cancel{\left(\frac{1}{R}\right)}} \\ &= \frac{d}{d\Theta}\left(\frac{dx}{d\Theta}\right)\left(\frac{d\Theta}{ds}\right)\frac{1}{R} \\ &= \frac{d^2x}{d\Theta^2}(\frac{1}{R})^2\end{aligned} \tag{1.27}$$

Substituting these expressions into Eq. 1.11 and 1.12, one obtains:

$$\begin{aligned}&\frac{d^2x}{d\Theta^2}\left(\frac{1}{R}\right)^2 + \frac{1-n}{\rho^2}x = 0 \\ \rightarrow\ &\frac{d^2x}{d\Theta^2} + \frac{R^2}{\rho^2}(1-n)x = 0\end{aligned} \tag{1.28}$$

$$\begin{aligned}&\frac{d^2z}{d\Theta^2}\left(\frac{1}{R}\right)^2 + \frac{n}{\rho^2}z = 0 \\ \rightarrow\ &\frac{d^2z}{d\Theta^2} + \frac{R^2}{\rho^2}nz = 0\end{aligned} \tag{1.29}$$

At a given angle Θ, a necessary and sufficient condition for the betatron oscillations to remain invariant with momentum is that:

$$\frac{\partial}{\partial p}\left(\frac{R}{\rho}\right)_\Theta = 0 \tag{1.30}$$

$$\left(\frac{\partial n}{\partial p}\right)_\Theta = 0 \tag{1.31}$$

It is important to note that these conditions represent only a sufficient, rather than necessary, criterion for achieving constant tune. Indeed, lattice configurations exhibiting constant tune over a broad energy range can be designed without strictly satisfying Eq. 1.30 and Eq. 1.31 [1]. Nev-

ertheless, magnets satisfying these conditions are generally considered more straightforward to operate, since optimisation performed at a single momentum extends naturally across the full acceleration range. This is particularly advantageous for dynamic aperture studies, where consistent beam behaviour over a range of momenta is essential.

By examining these conditions, the required magnetic field distribution can be deduced. Known as the scaling law [2, 3, 4], this was first derived in the 1950s. This work presents an alternative derivation of the scaling law, developed by the author, with particular emphasis on the transition from the magnetic field required along the particle's trajectory to the final field distribution expressed in polar form.

The first condition, Eq. 1.30, implies that:

$$\begin{aligned} \rho(p,\Theta) &= R(p)f(\Theta) \\ \Rightarrow \frac{\rho(p,\Theta)}{\rho(p_0,\Theta)} &= \frac{R(p)}{R(p_0)} = \lambda(p) \\ \Rightarrow \frac{\rho(\Theta)}{\rho_0(\Theta)} &= \frac{R}{R_0} = \lambda(p) \end{aligned} \tag{1.32}$$

where ρ_0, R_0 are evaluated at a reference momentum p_0, with dependence on momentum p implied and indicated by the corresponding subscript. That is, both the radius of curvature ρ and the equivalent radius R scale proportionally with momentum. The scalar expression for the radius of curvature of a closed orbit in polar coordinates [5] is

$$\rho = \frac{\left(r_c^2 + \left(\frac{dr_c}{d\theta}\right)^2\right)^{3/2}}{r_c^2 + 2\left(\frac{dr_c}{d\theta}\right)^2 - r_c\frac{d^2r_c}{d\theta^2}} \tag{1.33}$$

Assuming the closed orbit scales with momentum according to:

$$\begin{aligned} r_c(p,\Theta) &= r_c(p_0,\Theta)h(p,\Theta) \\ &= r_{c_0}(\Theta)h(p,\Theta) \end{aligned} \tag{1.34}$$

where r_{c_0} denotes the closed orbit at the reference momentum p_0, the derivatives of the closed orbit with respect to θ may be evaluated accordingly:

$$\begin{aligned} \frac{dr_c}{d\theta} &= \frac{d}{d\theta}(r_{c_0}h) = h\frac{dr_{c_0}}{d\theta} + r_{c_0}\frac{dh}{d\theta} \\ \frac{d^2r_c}{d\theta^2} &= \frac{d^2}{d\theta^2}(r_{c_0}h) = h\frac{d^2r_{c_0}}{d\theta^2} + 2\frac{dr_{c_0}}{d\theta}\frac{dh}{d\theta} + r_{c_0}\frac{d^2h}{d\theta^2} \end{aligned} \tag{1.35}$$

Substituting into Eq. 1.33, the resulting expression for the radius of curvature is obtained:

$$\rho = \frac{\left[(r_{c_0}h)^2 + \left(h\frac{dr_{c_0}}{d\theta} + r_{c_0}\frac{dh}{d\theta}\right)^2\right]^{3/2}}{(r_{c_0}h)^2 + 2\left(h\frac{dr_{c_0}}{d\theta} + r_{c_0}\frac{dh}{d\theta}\right)^2 - (r_{c_0}h)\left(h\frac{d^2r_{c_0}}{d\theta^2} + 2\frac{dr_{c_0}}{d\theta}\frac{dh}{d\theta} + r_{c_0}\frac{d^2h}{d\theta^2}\right)} \tag{1.36}$$

For ρ to scale correctly with momentum, it requires $h(p) = \lambda(p)$ such that:

$$\frac{dh}{d\theta} = \overset{0}{\cancel{\frac{\partial h}{\partial \Theta}}}\frac{d\Theta}{d\theta} = 0 \tag{1.37}$$

Under this condition, the expression for ρ reduces accordingly and is consistent with the scaling form of Eq. 1.32.

$$\begin{aligned}
\rho &= \frac{\left[(r_{c_0}\lambda)^2 + \left(\lambda\frac{dr_{c_0}}{d\theta}\right)^2\right]^{3/2}}{(r_{c_0}\lambda)^2 + 2\left(\lambda\frac{dr_{c_0}}{d\theta}\right)^2 - (r_{c_0}\lambda)\left(\lambda\frac{d^2r_{c_0}}{d\theta^2}\right)} \\
&= \lambda\frac{\left[r_{c_0}^2 + \left(\frac{dr_{c_0}}{d\theta}\right)^2\right]^{3/2}}{r_{c_0}^2 + 2\left(\frac{dr_{c_0}}{d\theta}\right)^2 - r_{c_0}\frac{d^2r_{c_0}}{d\theta^2}} \\
&= \lambda\rho_0
\end{aligned} \tag{1.38}$$

It therefore follows that:

$$\frac{r_c(\Theta)}{r_{c_0}(\Theta)} = \frac{R(\Theta)}{R_0(\Theta)} = \lambda(p) \tag{1.39}$$

where r_{c_0} is chosen at a reference momentum p_0. This result embodies the geometrical scaling property: closed orbits at different momenta are geometrically similar, each constituting an exact scaled enlargement of the reference orbit by the factor $\lambda(p)$.

In order to determine the magnetic field distribution required to satisfy Eq. 1.30, the magnetic field on the median plane along the closed orbit, $B_c(R,\Theta)$, is decomposed as :

$$B_c(R,\Theta) = \bar{B}_c(R)F(R,\Theta) \tag{1.40}$$

where $\bar{B}_c$ denotes the field averaged over the closed orbit defined as:

$$\begin{aligned}
\bar{B}_c(R) &= \frac{1}{2\pi}\oint B(R,\Theta)\,d\Theta \\
&= \frac{1}{2\pi R}\oint B(R,\Theta)\,ds
\end{aligned} \tag{1.41}$$

and $F(R,\Theta)$ describes its azimuthal variation.

From differential geometry, the local curvature of the closed orbit is related to the tangent angle by:

$$\frac{1}{\rho(s)} = \frac{d\psi}{ds} \tag{1.42}$$

where $d\psi(s)$ is the infinitesimal change in tangent angle along the orbit. Integrating over one complete revolution gives the accumulated bending angle:

$$\oint d\psi = \oint \frac{ds}{\rho(s)} \tag{1.43}$$

For a closed orbit, the tangent direction must return to its initial value after one full revolution, imposing the closure condition:

$$\oint d\psi = 2\pi \tag{1.44}$$

so that the total bending angle satisfies:

$$\oint \frac{ds}{\rho(s)} = 2\pi \tag{1.45}$$

The beam rigidity relation is given by:

$$B(R,\Theta)\rho(R,\Theta) = \frac{p(R)}{q} \tag{1.46}$$

Substituting Eq.1.45 and Eq.1.46 into Eq.1.41, the average bending field $\bar{B}_c$ becomes:

$$\bar{B}_c = \frac{1}{2\pi R}\oint B(R,\Theta)\,ds \tag{1.47}$$

$$= \frac{1}{2\pi R}\oint \frac{p}{q\rho}\,ds$$

$$= \frac{p}{2\pi R q}\oint \frac{1}{\rho}\,ds = \frac{p}{qR} \tag{1.48}$$

It therefore follows that:

$$B(R,\Theta)\rho(R,\Theta) = \frac{p(R)}{q} = \bar{B}_c(R)R \tag{1.49}$$

The ratio $\frac{R}{\rho}$ may then be expressed as:

$$\frac{R}{\rho(R,\Theta)} = \frac{B(R,\Theta)}{\bar{B}_c(R)} = F(R,\Theta) \tag{1.50}$$

In order to satisfy Eq. 1.30, the following must hold:

$$\begin{aligned}\frac{\partial}{\partial p}\left[F(R,\Theta)\right]_\Theta &= \frac{\partial}{\partial R}\left[F(R,\Theta)\right]_\Theta \frac{dR}{dp} \\ &= \frac{\partial}{\partial R}\left[F(R,\Theta)\right]_\Theta \left[\frac{1}{q\bar{B}_c} + \frac{p}{q}\frac{d}{dp}\left(\frac{1}{\bar{B}_c}\right)\right] = 0\end{aligned} \tag{1.51}$$

The term in the bracket is generally non zero, except in the unphysical limiting cases $\bar{B}_c = \infty$ or $\frac{\partial}{\partial p}(\bar{B}_c) = \infty$, which are not considered here. The condition for geometrical scaling of the closed orbits may therefore be stated as:

$$\frac{\partial}{\partial R}\left(F(R,\Theta)\right)_\Theta = 0 \tag{1.52}$$

That is, the function F must depend solely on Θ and be independent of R, such that:

$$B(R,\Theta) = \bar{B}_c(R)\,F(\Theta) \tag{1.53}$$

The azimuthal shape function $F(\Theta)$ must then satisfy the following normalisation condition, which follows directly from the definition of the integrated average:

$$\begin{aligned}\bar{B}_c(R) &= \frac{1}{2\pi}\oint B(R,\Theta)\,d\Theta \\ &= \frac{1}{2\pi}\oint \bar{B}_c(R)\,F(\Theta)\,d\Theta \\ &= \bar{B}_c(R)\cdot\frac{1}{2\pi}\oint F(\Theta)\,d\Theta \\ 1 &= \frac{1}{2\pi}\oint F(\Theta)\,d\Theta\end{aligned} \tag{1.54}$$

The second condition of the scaling law, given by Eq. 1.31, ensures that at corresponding points on the closed orbit, the local focusing properties remain independent of the particle momentum. Provided Eq.1.53 is satisfied, the partial derivatives of B expand as:

$$\begin{aligned}\left(\frac{\partial B}{\partial R}\right)_\Theta &= \frac{d\bar{B}_c}{dR}F + \bar{B}_c\cancelto{0}{\left(\frac{\partial F}{\partial R}\right)_\Theta} = \frac{d\bar{B}_c}{dR}F \\ \left(\frac{\partial B}{\partial \Theta}\right)_R &= \bar{B}_c\frac{dF}{d\Theta} + F\cancelto{0}{\left(\frac{\partial \bar{B}_c}{\partial \Theta}\right)_R} = \bar{B}_c\frac{dF}{d\Theta}\end{aligned} \tag{1.55}$$

Substituting into the definition of the field index yields:

$$\begin{aligned} n &= -\frac{\rho}{B}\left(\frac{\partial B}{\partial x}\right)_{\perp} \\ &= -\frac{\rho}{B}\left[\left(\frac{\partial B}{\partial R}\right)_{\Theta}\left(\frac{\partial R}{\partial x}\right)_{\perp} + \left(\frac{\partial B}{\partial \Theta}\right)_{R}\left(\frac{\partial \Theta}{\partial x}\right)_{\perp}\right] \\ &= -\frac{\rho}{B}\left[\frac{d\bar{B}_c}{dR}F\left(\frac{\partial R}{\partial x}\right)_{\perp} + \bar{B}_c\frac{dF}{d\Theta}\left(\frac{\partial \Theta}{\partial x}\right)_{\perp}\right] \\ &= -\frac{\rho}{R}\left[\frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\left(\frac{\partial R}{\partial x}\right)_{\perp} + \frac{R}{F}\frac{dF}{d\Theta}\left(\frac{\partial \Theta}{\partial x}\right)_{\perp}\right] \end{aligned} \tag{1.56}$$

It has already been established that when F depends solely on Θ, the following quantities are independent of momentum:

$$\frac{\partial}{\partial p}\left(\frac{\rho}{R}\right)_{\Theta} = \frac{\partial}{\partial p}\left(\frac{1}{F}\right)_{\Theta} = 0 \tag{1.57}$$

Differentiating the first term of Eq. 1.56 with respect to momentum:

$$\begin{aligned} \frac{\partial}{\partial p}\left[\frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\left(\frac{\partial R}{\partial x}\right)_{\perp}\right]_{\Theta} = \frac{\partial}{\partial p}\left(\frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\right)_{\Theta}\left(\frac{\partial R}{\partial x}\right)_{\perp} + \\ \frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\frac{\partial}{\partial p}\left[\left(\frac{\partial R}{\partial x}\right)_{\perp}\right]\Theta \end{aligned} \tag{1.58}$$

Differentiating the second term of Eq. 1.56 with respect to momentum:

$$\begin{aligned} \frac{\partial}{\partial p}\left[\frac{R}{F}\frac{dF}{d\Theta}\left(\frac{\partial \Theta}{\partial x}\right)_{\perp}\right]_{\Theta} = \cancel{\frac{\partial}{\partial p}\left[\frac{1}{F}\frac{dF}{d\Theta}\right]_{\Theta}}^{0} R\left(\frac{\partial \Theta}{\partial x}\right)_{\perp} + \\ + \frac{1}{F}\frac{dF}{d\Theta}\frac{\partial}{\partial p}\left[R\left(\frac{\partial \Theta}{\partial x}\right)_{\perp}\right]_{\Theta} \end{aligned} \tag{1.59}$$

Collecting all surviving terms, Eq. 1.31 reduces to:

$$\begin{aligned} \left(\frac{\partial n}{\partial p}\right)_{\Theta} = -\frac{\rho}{R}\left[\frac{\partial}{\partial p}\left(\frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\right)_{\Theta}\left(\frac{\partial R}{\partial x}\right)_{\perp}\right. \\ + \frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR}\frac{\partial}{\partial p}\left[\left(\frac{\partial R}{\partial x}\right)_{\perp}\right]_{\Theta} \\ \left. + \frac{1}{F}\frac{dF}{d\Theta}\frac{\partial}{\partial p}\left[\left(R\frac{\partial \Theta}{\partial x}\right)_{\perp}\right]_{\Theta}\right] \end{aligned} \tag{1.60}$$

The differential terms appearing in Eq. 1.60 admit a natural geometrical interpretation. We can introduce two angles, φ, defined as the angle between the normal to the closed orbit and the radial vector, and ζ, defined as the angle between the radial vector and the tangent to the curve of

constant Θ. Referring to Fig. 1.2, one may write:

$$\begin{aligned} R\left(\frac{\partial \Theta}{\partial x}\right)_{\perp} &= \tan(\varphi+\zeta) \\ \left(\frac{\partial R}{\partial x}\right)_{\perp} &= \frac{\mathrm{d}R}{\mathrm{d}l}\cos(\varphi+\zeta) \end{aligned} \tag{1.61}$$

where l denotes the longitudinal abscissa measured along curves of constant Θ.

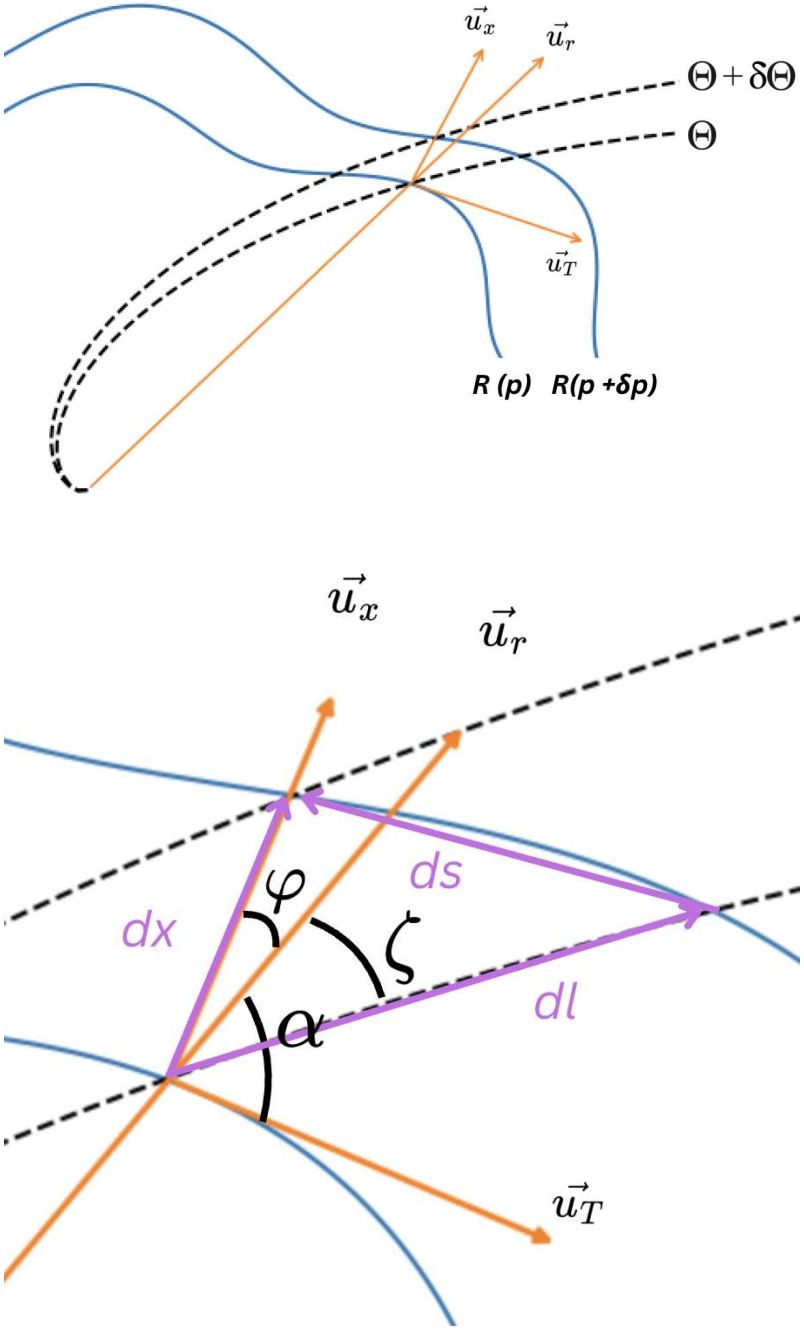


Figure 1.2: (Top) Geometrical considerations between closed orbits of different momentum. (Bottom) Zoomed in version of the top plot

Geometrical scaling implies that the angle φ is independent of momentum along curves of constant Θ, which may be demonstrated by examining the scaling of arclength. From Eq.1.25, the variation in arclength at fixed Θ across different closed orbits takes the form:

$$(ds)_\Theta = \Theta\, dR \tag{1.62}$$

which immediately implies that:

$$\begin{aligned} \left(\frac{ds}{dR}\right)_\Theta &= \Theta = \text{constant} \\ \Rightarrow \left(\frac{ds}{dr_c}\right)_\Theta \left(\frac{dr_c}{dR}\right)_\Theta &= \text{constant} \end{aligned} \tag{1.63}$$

Since Eq.1.39 establishes that $\left(\frac{dr_c}{dR}\right)_\Theta$ is constant, it follows that $\left(\frac{ds}{dr_c}\right)_\Theta$ must likewise be constant. Recalling that the arclength element in cylindrical coordinates is:

$$ds = \sqrt{dr_c^2 + r_c^2\, d\theta^2} \tag{1.64}$$

one obtains:

$$\begin{aligned} \left(\frac{\partial s}{\partial r_c}\right)_\Theta &= \sqrt{1 + \left(r_c \frac{\partial \theta}{\partial r_c}\right)_\Theta^2} = \text{constant} \\ \Rightarrow \left(r_c \frac{\partial \theta}{\partial r_c}\right)_\Theta &= \text{constant} \end{aligned} \tag{1.65}$$

Let α denote the angle between the radial vector and the tangent to the closed orbit, as illustrated in Fig.1.2. From differential geometry:

$$\left(r_c \frac{\partial \theta}{\partial r_c}\right)_\Theta = \tan(\alpha) \tag{1.66}$$

from which it follows that α must be constant. As seen in Fig.1.2, the angle between the radial vector and the normal to the closed orbit satisfies $\varphi = \frac{\pi}{2} - \alpha$, and therefore φ is likewise constant along curves of constant Θ under geometrical scaling of the orbits.

Evaluating the term $\frac{dR}{dl}_\Theta$, geometrical scaling of the orbits gives:

$$\begin{aligned} \frac{d}{dp}\left(\frac{dR}{dl}\right)_\Theta &= \frac{d}{dp}\left[\frac{dR}{dr}\frac{\partial r}{\partial l}\right] \\ &= \frac{d}{dp}\overbrace{\left(\frac{dR}{dr}\right)_\Theta}^{0}\frac{\partial r}{\partial l} + \frac{dR}{dr}\frac{d}{dp}\left(\frac{\partial r}{\partial l}\right)_\Theta \end{aligned} \tag{1.67}$$

The term $\frac{\partial r}{\partial l}_\Theta$ admits a geometrical interpretation: defining r_Θ as the radius along the curve of

constant Θ, the change in l is given by:

$$\begin{aligned} dl &= \sqrt{dr_\Theta^2 + (r_\Theta d\theta)^2} \\ \Rightarrow \frac{\partial r_\Theta}{\partial l} &= \frac{1}{\sqrt{1 + (r_\Theta \frac{\partial \theta}{\partial r_\Theta})^2}} \end{aligned} \tag{1.68}$$

where ζ denotes the angle between the radial vector and the tangent to the curve of constant Θ, such that:

$$r_\Theta \frac{\partial \theta}{\partial r_\Theta} = \tan(\zeta) \tag{1.69}$$

If ζ is chosen to remain constant along curves of constant Θ, these curves are described by the logarithmic spiral equation:

$$\Theta = \theta - \tan\zeta \cdot \ln\left(\frac{r}{r_0}\right) \tag{1.70}$$

where r_0 is a reference radius. Under this condition, combined with geometrical scaling of the orbits, both differential terms in Eq. 1.60 vanish:

$$\begin{aligned} \frac{d}{dp}\left[R\left(\frac{\partial \Theta}{\partial x}\right)_\perp\right]_\Theta &= 0 \\ \frac{d}{dp}\left[\left(\frac{\partial R}{\partial x}\right)_\perp\right]_\Theta &= 0 \end{aligned} \tag{1.71}$$

Introducing the geometrical field index k, defined as:

$$k = \frac{R}{\bar{B}_c}\frac{d\bar{B}_c}{dR} \tag{1.72}$$

the requirement of constant n reduces to:

$$\left(\frac{\partial n}{\partial p}\right)_\Theta = 0 \quad \Longleftrightarrow \quad \frac{dk}{dp} = 0. \tag{1.73}$$

The two conditions on the B field to achieve invariance of the betatron oscillations with momentum are therefore:

$$\frac{dk}{dp} = 0 \quad \left(\frac{\partial F}{\partial p}\right)_\Theta = 0 \tag{1.74}$$

If k is kept constant, Eq. 1.72 may be solved by integrating both sides between the reference radius R_0 and an arbitrary radius R:

$$\int_{\bar{B_{c_0}}}^{\bar{B_c}} \frac{d\bar{B_c}'}{\bar{B_c}'} = k \int_{R_0}^{R} \frac{dR'}{R'} \tag{1.75}$$

where $\bar{B_{c_0}} = \bar{B_c}(R_0)$ and variables with $'$ denote dummy integration variables. Evaluating the integrals yields:

$$\ln\left(\frac{\bar{B_c}}{\bar{B_{c_0}}}\right) = k \ln\left(\frac{R}{R_0}\right) \tag{1.76}$$

Exponentiating both sides and making use of Eq.1.39 gives the power law field profile:

$$\frac{\bar{B_c}}{\bar{B_{c_0}}} = \left(\frac{R}{R_0}\right)^k \tag{1.77}$$

Multiplying both sides by $F(\Theta)$ and applying Eq.1.70, a magnetic field satisfying the conditions of Eq.1.74 is given by:

$$\begin{aligned} \bar{B_c} F(\Theta) &= \bar{B_{c_0}} \left(\frac{R}{R_0}\right)^k F(\Theta) \\ B(R,\Theta) &= \bar{B_{c_0}} \left(\frac{R}{R_0}\right)^k F(\Theta) \end{aligned} \tag{1.78}$$

To transform between the coordinate systems (R,Θ) and (r,θ), consider the change in Θ at constant r using Eq.1.70:

$$\begin{aligned} \tan\zeta \ln\frac{r}{r_0} &= constant \\ \Rightarrow d\Theta &= d\theta \end{aligned} \tag{1.79}$$

It follows that $F(\Theta)$ satisfies the normalisation condition:

$$\frac{1}{2\pi} \oint F(\Theta) d\theta = 1 \tag{1.80}$$

Starting from Eq.1.39, consider the closed orbits of beams with different momenta $p(\Theta)$ and $p_0(\Theta)$ which varies with Θ such that:

$$\begin{aligned} r_c(p(\Theta),\Theta) &= r \\ r_c(p_0(\Theta),\Theta) &= r_0 \end{aligned} \tag{1.81}$$

where r is the standard polar coordinate and r_0 is the chosen reference radius. This scaling condition holds for any fixed Θ giving:

$$\begin{aligned}\frac{r_c(p(\Theta),\Theta)}{r_c(p_0(\Theta),\Theta)} &= \lambda\left(p(\Theta)\right)\\ \Rightarrow \frac{r}{r_0} &= \lambda\\ \Rightarrow \frac{r}{r_0} &= \frac{R}{R_0}\end{aligned} \tag{1.82}$$

The azimuthal average of the magnetic field at fixed r, expressed in polar coordinates, is obtained by applying Eq. 1.78 and Eq. 1.82:

$$\begin{aligned}\frac{1}{2\pi}\int_0^{2\pi} B(R(r,\theta),\Theta)\,d\theta &= \frac{\bar{B}_{c_0}}{2\pi}\int_0^{2\pi}\left(\frac{R(r,\theta)}{R_0}\right)^k F(\Theta)\,d\theta\\ &= \frac{\bar{B}_{c_0}}{2\pi}\int_0^{2\pi}\left(\frac{r}{r_0}\right)^k F(\Theta)\,d\theta\\ &= \frac{\bar{B}_{c_0}}{2\pi}\left(\frac{r}{r_0}\right)^k\int_0^{2\pi} F(\Theta)\,d\theta\end{aligned} \tag{1.83}$$

Applying the normalisation condition Eq. 1.80, it follows that at $r = r_0$ the azimuthal average magnetic field reduces to the average magnetic field on the reference closed orbit, $\bar{B}_{c_0}$:

$$\begin{aligned}B_0 &= \frac{1}{2\pi}\int_0^{2\pi} B(R(r,\theta),\Theta)\,d\theta\\ &= \bar{B}_{c_0}\end{aligned} \tag{1.84}$$

where B_0 denotes the average magnetic field evaluated at the reference radius $r = r_0$.

Eq. 1.78 may therefore be expressed in polar coordinates as:

$$B(r,\theta) = B_0\left(\frac{r}{r_0}\right)^k F(\theta - \tan\zeta\cdot\ln\left(\frac{r}{r_0}\right)) \tag{1.85}$$

A clarifying remark on notation: in an FFA lattice comprising multiple magnet families, reference is commonly made to the B_0 of individual magnets. What is more precisely meant in each case is the product $B_0F(\Theta)$, evaluated for the magnet in question.

1.5 Imperfections and resonances

In modern synchrotrons, magnets with separate functions are employed, allowing each multipole component to be independently optimised. This is not possible for magnetic fields of the form given by Eq. 1.85, where the radial dependence expands as:

$$\left(\frac{r}{r_0}\right)^k = 1 + \frac{k}{r_0}(r-r_0) + \frac{k(k-1)}{2!r_0^2}(r-r_0)^2 + \frac{k(k-1)(k-2)}{3!r_0^3}(r-r_0)^3 + \cdots \tag{1.86}$$

Except in the special cases $k = -1, 0, 1$, nonlinear field components are inherently present in the magnetic field distribution. Consequently, in addition to resonances driven by magnet imperfections, higher order multipole components intrinsic to the field distribution itself constitute further sources of resonance excitation. Careful magnet design is therefore required to mitigate the effects of these resonances; the primary objective of the present work is to satisfy Eq. 1.85 as closely as possible, thereby minimising the spread in ring tune and reducing the risk of crossing dangerous resonances.

1.6 Focusing properties

The field distribution in polar coordinates may be decomposed as:

$$B(r,\theta) = \bar{B}(r)F(\Theta(r,\theta)) \tag{1.87}$$

where $\bar{B}(r)$ denotes the average field at a given polar radius r. The field index n may be expanded in polar coordinates to identify the individual focusing contributions more clearly.

$$\begin{aligned} n = -\frac{\rho}{r}\Bigg[& \frac{r}{\bar{B}}\frac{d\bar{B}}{dr}\left(\frac{\partial r}{\partial x}\right)_{\perp} \\ & + \frac{r}{F(\Theta(r,\theta))}\left[\frac{\partial F(\Theta(r,\theta))}{\partial r}\right]_{\theta}\left(\frac{\partial r}{\partial x}\right)_{\perp} \\ & + \frac{r}{F(\Theta(r,\theta))}\left[\frac{\partial F(\Theta(r,\theta))}{\partial \theta}\right]_{r}\left(\frac{\partial \theta}{\partial x}\right)_{\perp}\Bigg] \end{aligned} \tag{1.88}$$

The first term captures the focusing contribution arising from the average field index of the magnets, which is proportional to the radius of curvature ρ. A reverse bending magnet, which deflects particles away from the machine centre, yields a focusing contribution of opposite sign relative to that of a normal bending magnet. The second and third terms are generally small except in the vicinity of magnet edges, where they give rise to edge focusing effects: spiral focusing and Thomas focusing, respectively.

The introduction of an azimuthally varying magnetic field, characterised by a non zero $\frac{\partial F}{\partial \theta}$, results in non circular particle orbits with a characteristic scalloped geometry. This configuration produces a finite angle between the normal to the particle trajectory and the magnet edge, referred to as the edge angle κ. A positive value of κ induces horizontal defocusing through edge effects whilst a negative value produces the opposite effect. Simultaneously, the coupling of the radial component of the particle momentum with the azimuthal gradient of the magnetic field generates a vertical restoring force at the magnet edges.

Shaping the magnet edges according to a logarithmic spiral (Eq. 1.70) allows the effective edge angle to be further enhanced while maintaining independence of the focusing momentum. In this configuration, the effective edge angle takes the value $\kappa + \zeta$ at the magnet entrance and $\kappa - \zeta$ at the exit, where ζ denotes the spiral angle. Characterised by a non zero $\frac{\partial F}{\partial r}$, the coupling of the

azimuthal component of the particle momentum with the radial gradient of the magnetic field produces a force whose sign is determined by ζ. In practice, ζ is chosen sufficiently large such that, for a magnet whose main body provides horizontal focusing, the entrance edge becomes horizontally defocusing whilst the exit edge is focusing; the vertical behaviour is reversed accordingly.

In Lawrence cyclotrons, stability in both the horizontal and vertical planes requires the field index n to lie in the range $-1 < n < 0$. It has been demonstrated that permitting n to alternate between positive and negative values enables strong focusing via the alternating gradient mechanism [6], substantially enhancing the stability of both horizontal and vertical betatron oscillations. When the magnetic field is additionally held constant throughout the acceleration process, as in a cyclotron, the resulting machine constitutes a Fixed Field Alternating Gradient (FFA) accelerator, refered to as FFAG in older literature. This is a concept developed in the 1950s through largely independent efforts in the United States, Japan, and the USSR [2, 3, 4]. FFA accelerators satisfying Eq. 1.30 and Eq. 1.31 are classified as scaling FFAs. Although scaling FFAs may in principle admit negative values of k [2, 4], the present work is restricted to the case of positive k.

Two principal magnet configurations are employed in scaling FFAs. In the radial sector configuration, an additional magnet with an opposite field gradient is introduced to provide sufficient alternating gradient focusing. To preserve the scaling condition, k is held constant whilst the sign of B_0 is reversed, so that the two magnets possess not only opposite field gradients but also opposite bending directions. The magnet whose main body provides horizontal focusing is designated the focusing (F) magnet and bends the beam towards the machine centre, whilst the defocusing (D) magnet bends the beam outward and produces horizontal defocusing. A notable advantage of this configuration is the ability to adjust the ratio $\frac{B_D}{B_F}$ post construction, thereby modifying κ and enabling tune adjustment during operation. However, the incorporation of reverse bends increases the total beam path length and, consequently, the overall machine circumference. Modern radial sector FFA designs typically employ three magnets per cell in either a DFD or FDF arrangement, although more elaborate multi-magnet configurations have also been proposed [7].

The second configuration employs spiral sector magnets, in which the magnet edges follow the logarithmic spiral defined by Eq. 1.70. With a sufficiently large spiral angle, adequate alternating gradient focusing can be achieved using a single normal bending magnet per cell, eliminating the need for reverse bends and thereby permitting a more compact machine circumference relative to the radial sector case. However, as each cell contains only one magnet, post construction tune adjustment is generally not possible in this configuration.

1.7 Origin and motivation for FFA accelerators

The first FFA model, Mark-II [8], was constructed in 1955 by the Midwestern Universities Research Association (MURA) in the United States. This machine operated as an electron accelerator in which each lattice cell comprised a normal and a reverse bending magnet. Two subsequent electron models followed: a ring employing spiral sector magnets, commissioned in 1956 [9], and an electron colliding ring, completed in 1961 [10]. Although cyclotrons benefited from the development of spiral shaped magnets, since the 1960s research and development activity in the field

of FFAs entered a prolonged period of dormancy spanning several decades. Interest has, however, been substantially renewed in recent years, driven by emerging applications and significant advances in computational modelling.

A primary motivation for the renewed interest in FFA accelerators is the increasing demand for high beam power. Beam power is determined not solely by the final beam energy but also by the average beam current; consequently, achieving substantial gains in beam power necessitates both high beam energies and elevated repetition rates. In FFA accelerators, the static nature of the guiding magnetic field eliminates the limitations imposed by magnet ramping in pulsed synchrotrons, thereby permitting considerably higher repetition rates. Furthermore, in contrast to cyclotrons, the field gradient in an FFA is not subject to the isochronicity constraint, enabling significantly steeper gradients to be employed. For a given energy range, a higher field gradient reduces the radial beam excursion, which in turn facilitates more compact magnet designs.

High power proton beams significantly enhance the efficiency of secondary particle production, most notably muons and neutrons, which are of considerable importance across a broad range of research disciplines. Intense neutron sources are indispensable for accelerator driven subcritical reactors (ADSR), in which sustained neutron production within subcritical assemblies underpins research into nuclear waste transmutation and advanced nuclear energy technologies. High beam power is equally critical for improving the performance of pulsed spallation neutron sources.

In addition, muon beams stored within a ring decay into neutrinos with a well defined flux and flavour composition, providing a controlled and well characterised environment for precision studies of neutrino oscillations. The large momentum acceptance and rapid acceleration capability of FFA accelerators render them particularly attractive for muon acceleration in the context of next generation lepton-antilepton colliders, where these properties are essential for mitigating beam losses arising from the short muon lifetime.

A further significant motivation lies in medical applications. The flexibility to vary the extracted beam energy without degraders makes FFA accelerators especially attractive for hadron therapy, enabling fast, depth resolved scanning of tumour volumes, improving treatment efficiency, and increasing patient throughput relative to conventional systems. Furthermore, the rapid cycling nature of FFA accelerators facilitates exploration of the ultra high dose rate (FLASH) radiotherapy regime [11], opening promising avenues for next generation treatment modalities.

Advances in computational modelling contributed substantially to the resurgence of FFA research. Contemporary FFA magnets are designed with electromagnetic simulation software using Finite Element Methods (FEM), such as COMSOL and OPERA-3D [12, 13], which facilitates accurate characterisation of highly nonlinear magnetic field distributions. In parallel, investigations of beam dynamics are conducted using stepwise particle tracking codes, including FixField, S-code, Zgoubi and Opal [14, 15, 16, 17].

1.8 Development of FFA

The first proton FFA, the Proof of Principle FFA (POP-FFA) was successfully constructed and commissioned at the High Energy Accelerator Research Organization (KEK) in Japan in 2000 [18]. Operating at repetition rates of up to 1000 Hz, this machine represented a significant milestone in FFA development, demonstrating proton acceleration from 50 keV to 500 keV. Building on this achievement, a 150 MeV radial sector proton FFA was subsequently realised at KEK in 2004 [19], establishing the feasibility of MeV scale proton acceleration within this framework. A distinguishing feature of this machine is the return yoke free triplet configuration (DFD) [20]. The triplet was connected magnetically together to improve the efficient by returning the magnetic flux between the F and D magnets. The triplet only had shunt yokes which mainly acted to control the magnetic flux rather than returning the flux like conventional yokes. This design facilitated beam injection and extraction; however, the extended fringe fields and magnetic saturation effects associated with this gave rise to a significant tune spread observed during operation.

A complete FFA proton accelerator complex was constructed at the Kyoto University Research Reactor Institute (KURRI), since renamed the Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS) [21], comprising three scaling FFA rings arranged in cascade: a spiral FFA injector followed by two radial sector rings, collectively capable of accelerating protons to energies of up to 150 MeV. This facility enabled the world's first demonstration of an accelerator driven system (ADS) utilising spallation neutrons generated by high energy proton beams, entering operation in 2009 [22].

Several additional scaling FFA designs have been proposed and investigated across a broad range of applications. In the context of boron neutron capture therapy, energy recovering internal target (ERIT) schemes [23] have been proposed and developed [24]. The production of pure and quasi monoenergetic muon beams for lepton flavour violation experiments has likewise motivated the development of the Phase Rotated Intense Slow Muon beam (PRISM) FFA ring [25] with its smaller demonstrator ring [26] constructed at RCNP, Osaka.

The Recherche en Accélérateurs et Applications Médicales (RACCAM) project [27] investigated the feasibility of FFA accelerators for proton therapy, producing a conceptual design for a variable energy accelerator based on single spiral sector magnets, with the aim of realising compact, hospital based treatment systems. A prototype spiral sector magnet was constructed and experimentally characterised [28].

All of the FFA accelerators described above are founded on the scaling principle, in which the transverse focusing properties remain invariant with beam momentum. Whilst this condition ensures stable beam dynamics throughout acceleration, it imposes significant constraints on the lattice design. To afford greater flexibility, a distinct class of FFA accelerators known as non-scaling FFAs was introduced through relaxation of the invariant focusing requirement. The first experimental realisation of this concept, the Electron Model for Many Applications (EMMA) [29], was constructed and operated at Daresbury Laboratory [30]. EMMA employed a linear non-scaling lattice with cells composed of offset quadrupole magnets, and successfully demonstrated rapid acceleration through multiple resonance crossings without significant beam loss, as well as a fixed

RF frequency acceleration scheme known as serpentine acceleration.

In a collaboration between Cornell University and Brookhaven National Laboratory (BNL), the Cornell–BNL Energy Recovery Linac Test Accelerator (CBETA) [31] demonstrated the integration of a superconducting energy recovery linac with non-scaling FFA return arcs capable of transporting beams at multiple energies simultaneously, offering a route towards more compact and efficient recirculating accelerator architectures. This concept is closely related to developments for high current recirculating linac facilities such as the Continuous Electron Beam Accelerator Facility (CEBAF) [32], which plans to extend beyond the 150 MeV maximum energy demonstrated by CBETA.

FFA concepts have been the subject of detailed design studies for a number of proposed but unrealised projects. In the medical domain, the PAMELA [33] design study and the subsequent NORMA [34] concept explored both non-scaling and scaling FFA lattice configurations as candidate solutions for hadron therapy applications. FFA arcs have also been investigated in the context of proposed future collider facilities; most notably, the eRHIC project [35] at Brookhaven National Laboratory considered non-scaling FFA arcs for deployment within an energy recovery linac-based electron-ion collider.

A further class of FFA accelerator is the vertical excursion FFA (vFFA) [36, 37, 38]. First proposed by Ohkawa in 1955, the concept attracted relatively little attention at the time of its original conception. In a vFFA, the beam undergoes vertical displacement during acceleration, in contrast to conventional horizontal FFA accelerators in which closed orbits remain confined to the horizontal plane. The concept was subsequently revisited by Brooks, who identified several potential applications [39], stimulating renewed interest in vFFA designs. Recent developments include a continuous proton beam accelerator proposed for the production of high power muon beams [40] with a demonstrator currently under construction at Kyushu University [41]. A vFFA demonstrator ring [42, 43] was also investigated as a candidate option for the upgrade of future spallation neutron sources; however, a horizontal FFA design was ultimately selected on the grounds of time constraints and engineering considerations.

1.9 Purpose of the study

The majority of FFA accelerators constructed to date have employed shaped iron poles to achieve the radial field gradient required by the scaling condition. Notable exceptions include the Ion-Beta injector at KURNS [44] and a small number of design studies [45, 46, 47], which instead utilised radially distributed coils on the magnet pole to produce the required field variation. The principal advantage of the latter is that the current settings of the individual coils may be adjusted independently, enabling variable field index k and/or variable extraction energy operation. A more detailed treatment of both technologies is presented in Sections 2.1 and 2.2.

The present work develops a general design procedure for scaling spiral sector magnets with radially distributed trim coils, as described in Chapter 2. The procedure encompasses magnet and coil geometry parametrisation in Python, finite element electromagnetic modelling and simulation,

nonlinear least-squares optimisation of coil current settings, and extraction of three dimensional field maps for particle tracking and tune evaluation.

This methodology is subsequently applied to two distinct projects with differing design requirements. The FETS-FFA, serving as a demonstrator for high-intensity FFA operation, adopts a novel doublet scaling FFA lattice to enable flexible tune space operation, as presented in Chapter 3. The Laser-hybrid Accelerator for Radiobiological Applications (LhARA), developed for hadron therapy research, employs a single scaling FFA cell to provide variable extraction energy whilst minimising the machine circumference, as discussed in Chapter 4. A prototype magnet, constructed without a spiral angle but incorporating the same trim coil principle, has additionally been fabricated to provide experimental validation of the trim coil working principle, as presented in Chapter 5.

Chapter 2

Magnet design

This chapter describes the geometry of the FFA magnet and the motivation for the technologies employed to produce the required magnetic field distribution. The author developed a series of original methods, including the identification of the integrated field, spiral angle, and fringe field extent as the quantities to be optimised, and a suite of Python codes to parametrically model and optimise the magnet in OPERA-3D.

The treatment is kept general, without reference to the specific parameters of any particular machine, with the foundational configuration being that of a singlet FFA magnet design. A doublet configuration is introduced in Chapter 3, where it is treated as an extension of the singlet design and shares the same underlying design philosophy; it is therefore not discussed further in the present chapter.

Conventional electromagnets are typically capable of generating magnetic fields of up to approximately 1.4 T before the onset of magnetic saturation. The machine circumference is accordingly selected on the basis of the final extraction energy, so as to ensure that the peak magnetic field remains within this limit. Design parameters such as the pole gap and packing factor are optimised to satisfy the specific requirements of the machine. In a scaling FFA, the field index k is chosen to position the working tune point sufficiently far from dangerous resonances whilst preserving an adequate dynamic aperture; this choice further determines the radial beam excursion and, consequently, the physical dimensions of the magnet. The remainder of this chapter proceeds on the assumption that these parameters have been established through prior lattice design studies.

2.1 2D studies: Field gradient and gap shaping

Several approaches exist for achieving the radial field distribution described by Eq. 1.85. Since the magnetic field varies inversely with the magnet pole gap height g according to:

$$B \approx \frac{\mu_0 NI}{g} \tag{2.1}$$

the conventional approach to producing the required radial gradient is to vary the gap height $g(r)$ as a function of radial position, resulting in a magnet with a profiled pole gap that is narrowest where the strongest field is desired, as illustrated in Fig. 2.1.

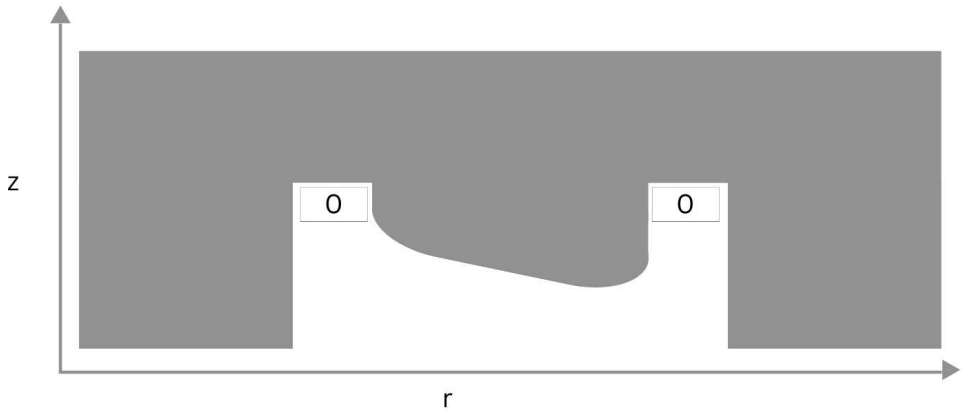


Figure 2.1: Cross sectional schematic of the upper half of an H type magnet with a radially varying pole gap. The iron yoke is shown in grey and the main coil is labelled 0. The pole gap decreases at larger radii, where a stronger magnetic field is required. Figure is illustrative and not to scale.

This approach was adopted in several previous FFA projects, including RACCAM, KURNS, and POP. The technology of mechanical iron machining is well established, permitting adjustments on a fine scale and yielding field distributions with small errors. However, once fabricated, the pole profile cannot be readily modified during operation. For machines requiring a variable geometrical field index k and/or variable extraction energy, a more flexible means of adjusting the field gradient is therefore necessary.

2.2 Radially distributed trim coils

An alternative to varying the pole gap height is to position coils at different radii on a flat pole surface. A main coil generates the dipole field, whilst trim coils either add to or subtract from the main dipole field depending on the direction of their return path. By adjusting the current settings in the individual trim coils, the radial magnetic field distribution can be controlled accordingly. This technology was ultimately selected for the present work on account of its flexibility in tailoring the radial field gradient.

Two options were initially considered. The first is the pole embedded coil configuration, in which pancake shaped coils are inserted into deep grooves and returned around the iron pole, as illustrated in Fig. 2.2. Embedding coils within the magnet pole offers several advantages: it simplifies coil fabrication by requiring only pancake shaped windings and, since the coils do not traverse the pole face, permits a reduction in the iron gap, thereby lowering power consumption. However, this configuration can give rise to localised magnetic saturation, potentially introducing significant field quality errors. It is therefore more suited to magnets operating at relatively low magnetic field strengths.

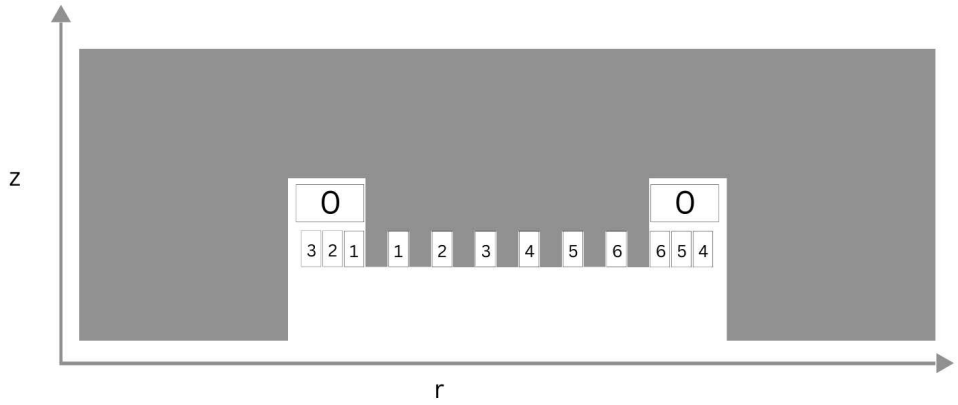


Figure 2.2: Cross sectional schematic of the upper half of an H-type magnet with pole embedded coils. The iron yoke is shown in grey and the coils are represented by numbered boxes, with the main coil labelled 0 and trim coils labelled 1 to 6. Figure is illustrative and not to scale.

The second option employs pole face crossing trim coils, in which the trim coils traverse the pole face at different radii and are routed back around the pole, as illustrated in Fig. 2.3. Each trim coil acts as an independent tuning element, enabling adjustment of the magnetic field within its respective radial range. Increasing the number of trim coils within a fixed magnet length provides additional degrees of freedom and thus improves the achievable field quality. However, this comes at the cost of a larger number of power supplies per magnet, increasing the operational complexity. The total number of trim coils must therefore be selected as a compromise between the desired field quality and the practical constraints on the number of power supplies.

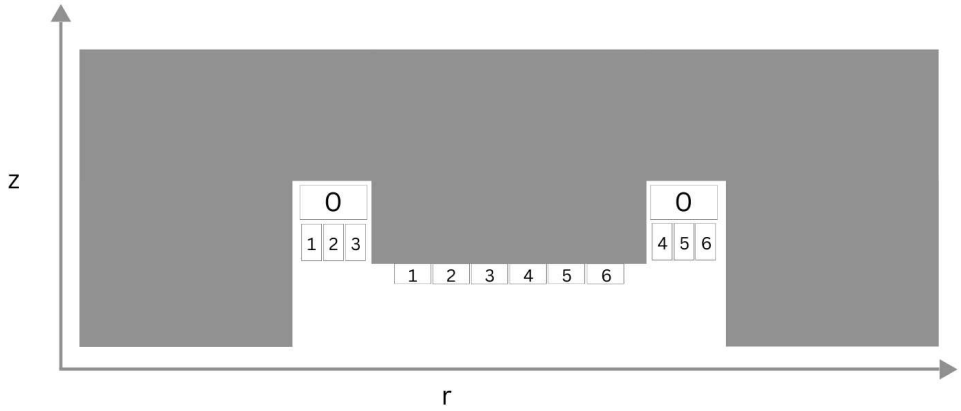


Figure 2.3: Cross sectional schematic of the upper half of an H-type magnet with pole face crossing trim coils. The iron yoke is shown in grey and the coils are represented by numbered boxes, with the main coil labelled 0 and trim coils labelled 1 to 6. Figure is illustrative and not to scale.

2.3 FEM and relevant softwares

Before presenting the magnet geometry in detail, it is necessary to introduce the modelling approach employed in this work: the Finite Element Method (FEM). Owing to the complexity of the magnet geometry, numerical techniques are required to obtain high order approximations to Maxwell's equations. In FEM, the geometry is discretised into a large number of smaller elements, referred to as mesh elements, within which the electromagnetic fields are computed numerically.

This discretisation yields a finite system of equations that can be solved numerically. The mesh resolution has a direct bearing on the accuracy of the solution: finer meshes produce more precise local field values but demand greater computational resources. The mesh resolution must therefore be selected as a compromise between solution accuracy and computational cost.

The combination of a complex geometry and a large number of pole face coils made full meshing computationally expensive. The FEM software selected for this study was OPERA [13], which provides the option to compute the magnetic vector potential of the coils in free space using the Biot–Savart law. This approach avoids the need to explicitly mesh every coil and substantially reduces the computational cost. The resulting coil fields are then incorporated as source terms in the FEM formulation to solve for the magnetic field in the complete geometry, including the iron yoke, magnetic saturation, and fringe field effects. OPERA-2D was used for the initial two dimensional investigations, while OPERA-3D was employed for the final three dimensional analysis.

The coordinates defining the coil and iron contours were generated using a custom Python script written by the author, and subsequently converted into OPERA command files to construct the full three dimensional model via Metaopera, a scripting tool originally developed by Planche and Koay. In the course of this project, Metaopera was extended by the author and Lagrange to support the generation of arbitrary three dimensional coil geometries and to embed the entire model within a parametric framework accessible through a Python interface; further details are provided in Appendix B. This parametric approach, with the workflow shown in Fig. 2.4, enables efficient variation of design parameters and greatly facilitates optimisation studies.

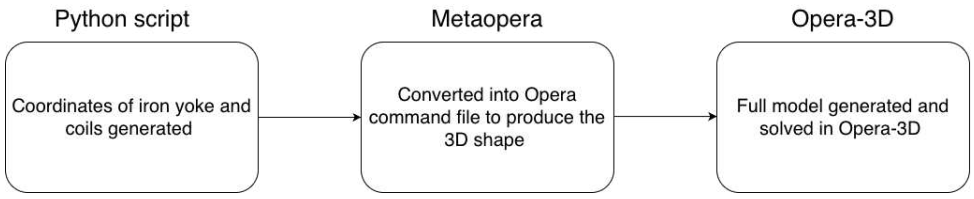


Figure 2.4: Overview of the modelling workflow, beginning with coordinates generated by a parametric Python script developed by the author. The parametric framework allows design parameters to be modified readily prior to generating the OPERA model via Metaopera.

2.4 Geometry of pole

To realise the field distribution described by Eq. 1.85, the magnet geometry must conform to a spiral profile. The two longitudinal edges of the pole are defined by curves of constant spiral angle, each characterized by a distinct value of θ_0, as given in Eq. 1.70; the difference between these values is referred to as the opening angle. These spiral edges are connected by segments of constant radius corresponding to the minimum and maximum pole radii. To avoid sharp edges, which can lead to localised magnetic saturation, the pole boundaries are rounded, resulting in the final pole profile shown in Fig. 2.5. The particle beam enters the magnet from the left hand side, defining the entrance side, while the opposite side constitutes the exit side.

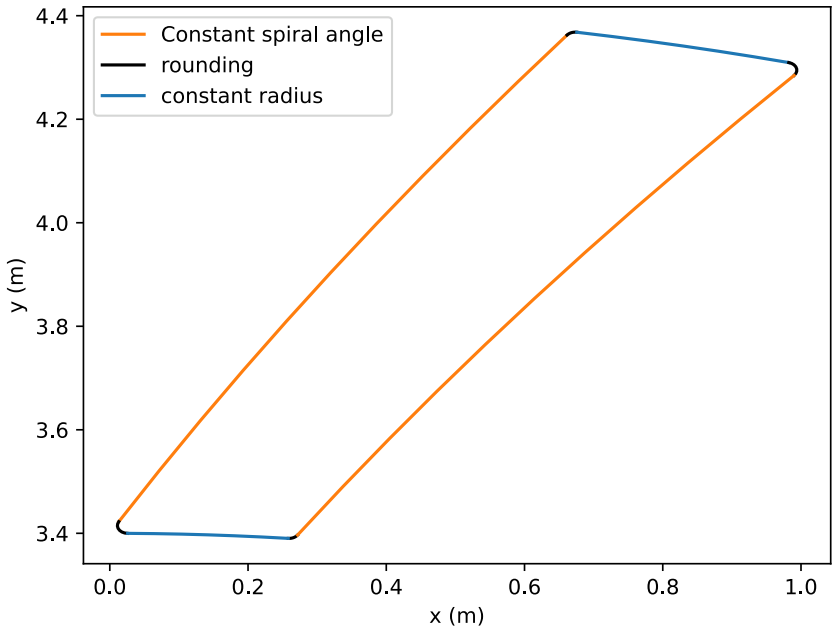


Figure 2.5: Example of the two dimensional pole profile of a spiral magnet. Coordinates are given in metres, with the machine centre located at the origin (0,0).

2.5 Geometry of clamp

The field distribution described by the scaling law in Eq. 1.85 implies a fringe field extent that increases linearlly with radius. In conventional magnet designs, the fringe field length is typically governed by the magnet gap height, which should also increase linearly with radius to adhere to the scaling law, as in the MURA electron model[48]. In the present design, the gap height remains fixed, and this inconsistency leads to deviations from the scaling law in the resulting field distribution, producing an increased vertical tune spread. To mitigate this, spiral shaped clamps with a uniform aperture are installed on both sides of the pole, as illustrated in Fig. 2.6. On the pole adjacent side, the clamps follow a constant spiral angle defined by Eq. 1.70. Since the opening angle between the pole and clamp is maintained, their physical separation increases proportionally with radius, encouraging the fringe field length to follow a similar radial dependence. However, as discussed in Section 3.6.1, this adjustment alone is insufficient, and the clamp aperture must also be modified to bring the field closer to the scaling behaviour.

On the outer side, the clamps may also follow a spiral profile. However, this leads to a clamp thickness that increases with radius, which is not space efficient. To reduce the required volume, the clamps can instead be designed with a constant thickness perpendicular to the spiral edge, chosen sufficiently large to avoid magnetic saturation.

Field clamps offer several additional advantages. Their primary effect is to accelerate the fall off of the fringe field, confining the magnetic field more tightly in the vicinity of the magnet. This confinement is further improved when the clamps and magnetic poles share a common yoke, which increases clamping efficiency and steepens the fringe field fall off, as illustrated in Section 3.4.

The practical consequence is that field-sensitive components, such as diagnostic equipment and RF cavities, can be positioned closer to the magnet without adverse effects, enabling a more compact ring lattice. In addition, the steeper fringe field gradient amplifies the Thomas focusing contribution to the field index n (Eq. 1.88), such that a reduced spiral angle is sufficient to achieve the same degree of vertical focusing.

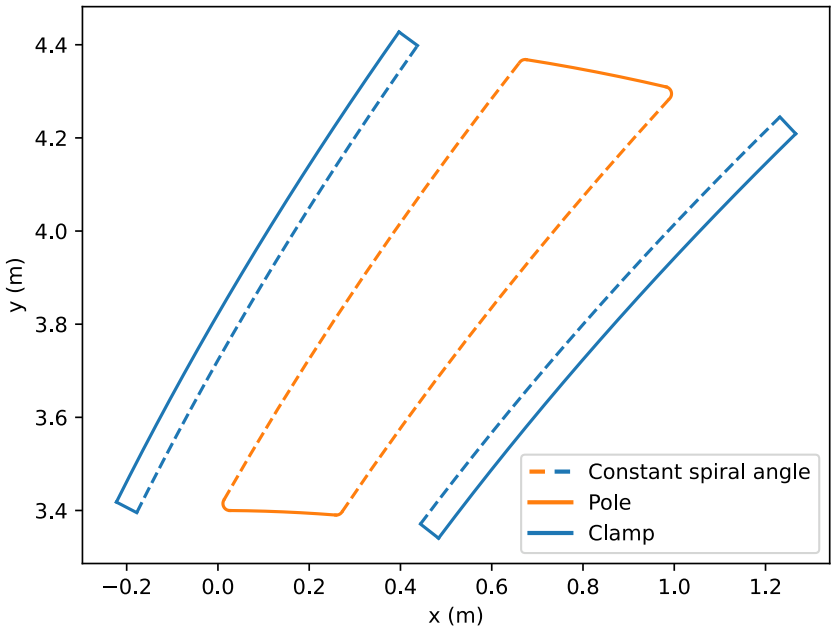


Figure 2.6: Two dimensional model of a spiral pole with field clamps installed on either side, conforming to the spiral profile. The clamps are of constant thickness, such that the outer edge of each clamp does not follow a constant spiral angle.

Based on the considerations discussed, the opening angle between the pole and the clamp should be minimised whilst avoiding any magnetic field overshoot, since such overshoot would necessitate separate modelling and thereby increase the complexity of the overall design process.

2.6 Geometry of coils

In practice, each coil consists of a conductor wound over multiple turns. To reduce modelling complexity, each coil is represented as a single equivalent conductor of larger cross section, since this simplification does not affect the resulting magnetic field.

The coils are modelled in OPERA-3D using a succession of 8-node brick elements, as shown in Fig. 2.7. Since the coil cross section remains constant throughout, only a path of coordinates is required to fully define the geometry, as illustrated in Fig. 2.8. The following description pertains to the case of a single trim coil; the full configuration in which multiple trim coils are stacked longitudinally is discussed in Section 2.7.

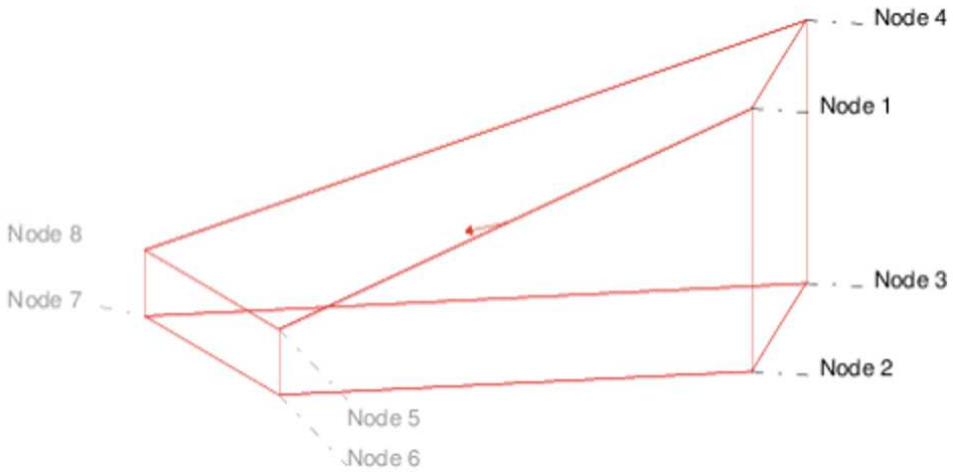


Figure 2.7: Illustration of an 8 node brick element as defined in OPERA-3D. A conductor brick is specified by providing the coordinates of all eight nodes.

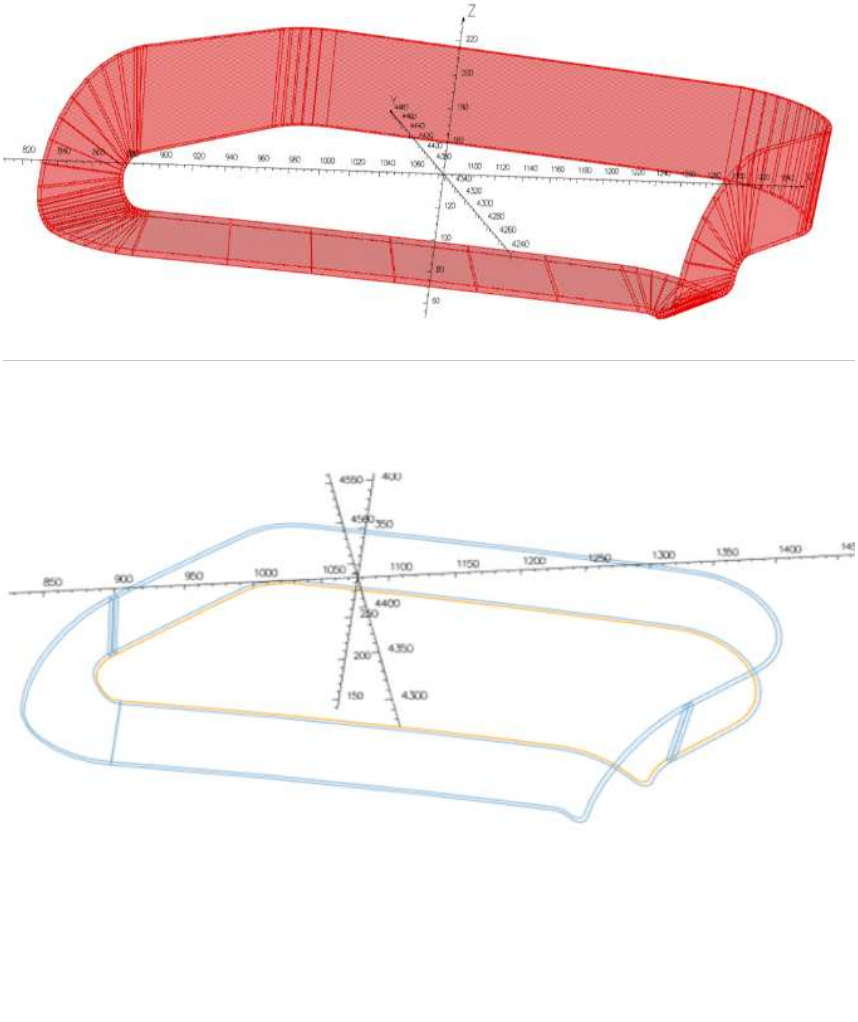


Figure 2.8: (Top) Complete model of a trim coil constructed from a succession of 8 node brick elements. (Bottom) The minimal set of coordinates required to define a trim coil, from which the remaining geometry is computed assuming a constant cross section throughout.

As established in Section 1.4, the scaling condition requires that the azimuthal variation of the magnetic field conform to Eq. 1.85. One means of achieving this is to shape the portion of the trim coil crossing the pole face as a circular arc centred on the machine axis, whilst the return path along the side of the pole follows the same logarithmic spiral. To fit the trim coils closely around the pole, however, bends are required at the magnet entrance and exit to align the coil cross section with the spiral pole profile. These bends impose a constraint on how closely the coil crossings can approximate a circular arc, and their bending radius should therefore be kept as small as possible. The resulting design, shown in Fig. 2.9, minimises the longitudinal space occupied by the coils outside the pole whilst maximising the extent of the circular arc portion of each coil crossing.

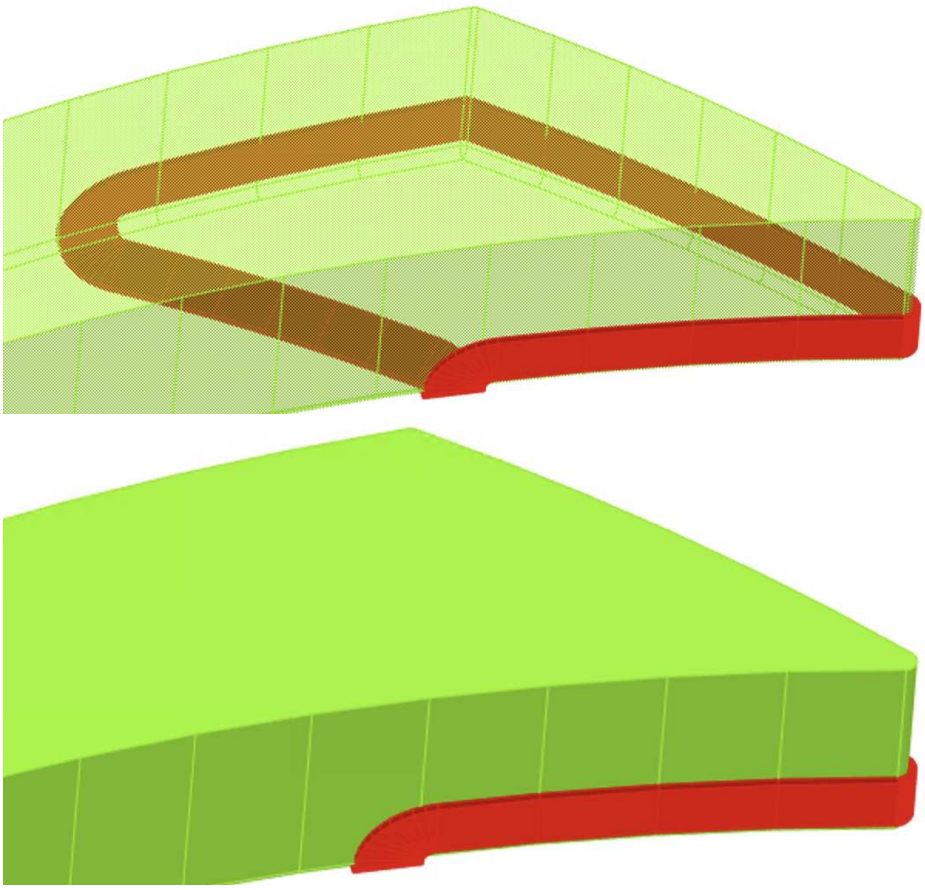


Figure 2.9: OPERA-3D model showing a segment of a trim coil (red) on a spiral pole (green). The upper portion of the spiral pole is rendered translucent and the lower portion solid, to provide complementary perspectives of the assembly.

The geometry of the trim coil on the pole face is illustrated in Fig. 2.10. Following the pole face crossing, a series of vertical bends is required before the coil can be routed back around the side of the pole. The trim coil is first bent vertically upward, as shown in Fig. 2.11. To accommodate this bend and ensure a close fit against the pole, the spiral pole edge must be chamfered to match the bending radius of the coil. A second bend is then introduced to align the trim coil with the spiral edges of the pole, as depicted in Fig. 2.12, after which the coil is wrapped around the side of the pole to complete the return path. A detailed description of the numerical procedure for computing these coordinates is provided in Appendix A.

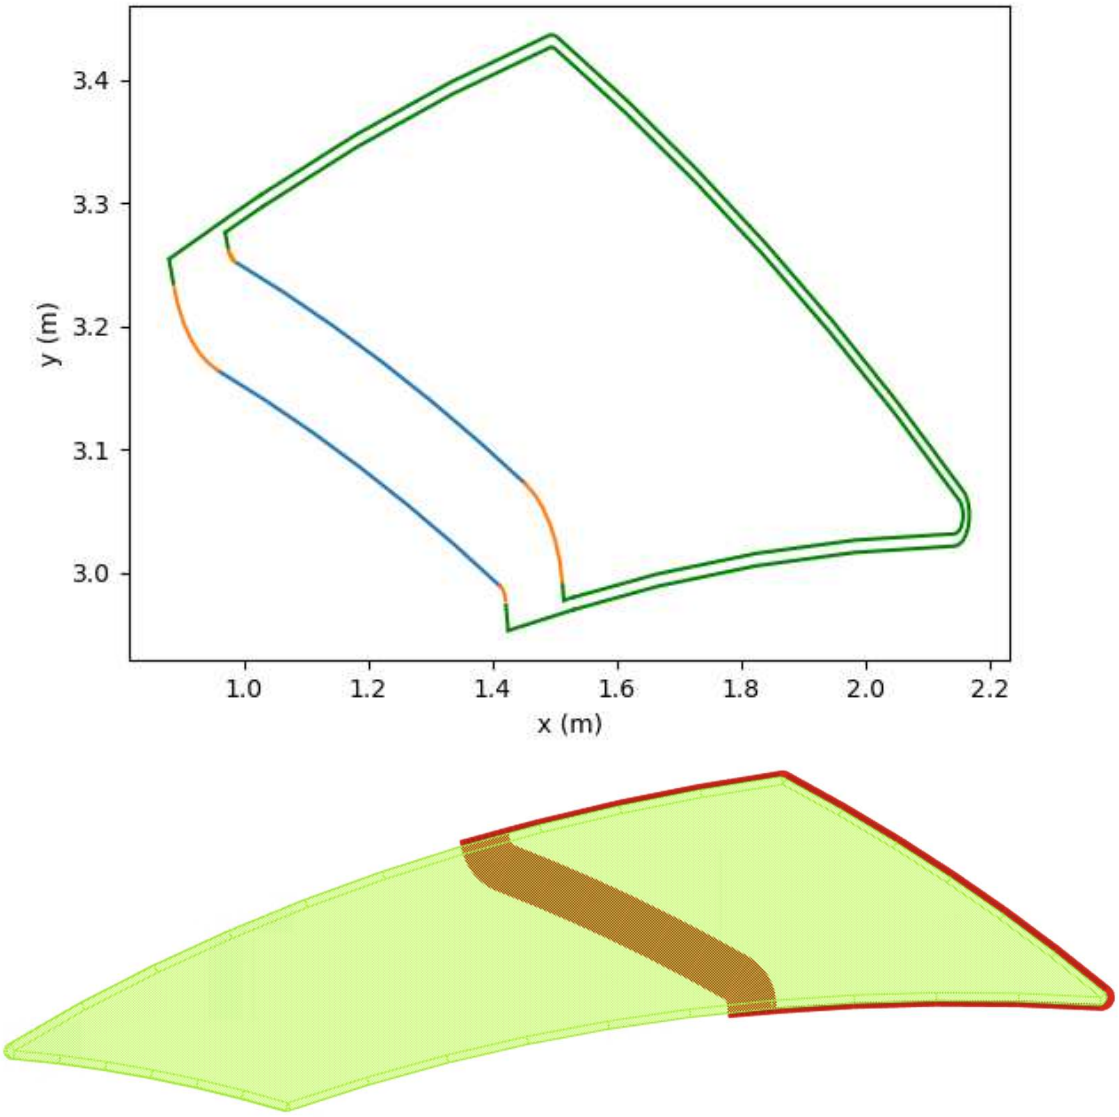


Figure 2.10: (Top) Two dimensional Python plot showing the top view of a trim coil. The circular arc of the trim coil is highlighted in blue, followed by two bends (orange) required to align the coil cross section with the spiral edges. (Bottom) Corresponding OPERA-3D model showing a segment of the trim coil (red) on a spiral pole rendered translucent (green).

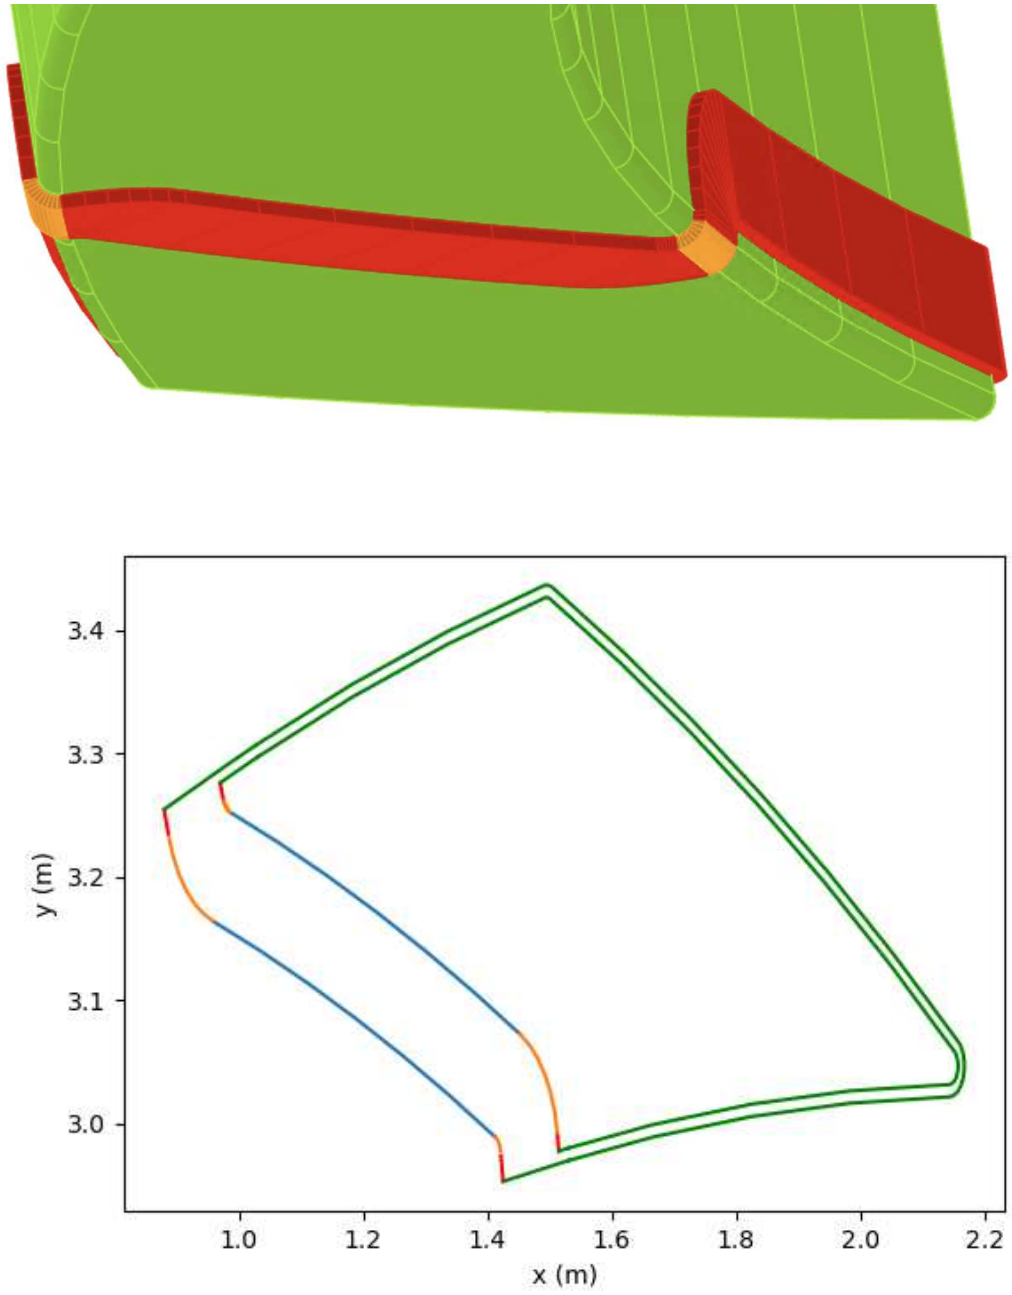


Figure 2.11: (Top) Poleface view of an OPERA-3D model illustrating a trim coil segment (red) crossing the pole face of a spiral pole (green). Two vertical bends (orange) are highlighted, which bring the trim coil vertically upward against the pole face. (Bottom) The corresponding vertical bends are indicated in red in the two dimensional Python plot.

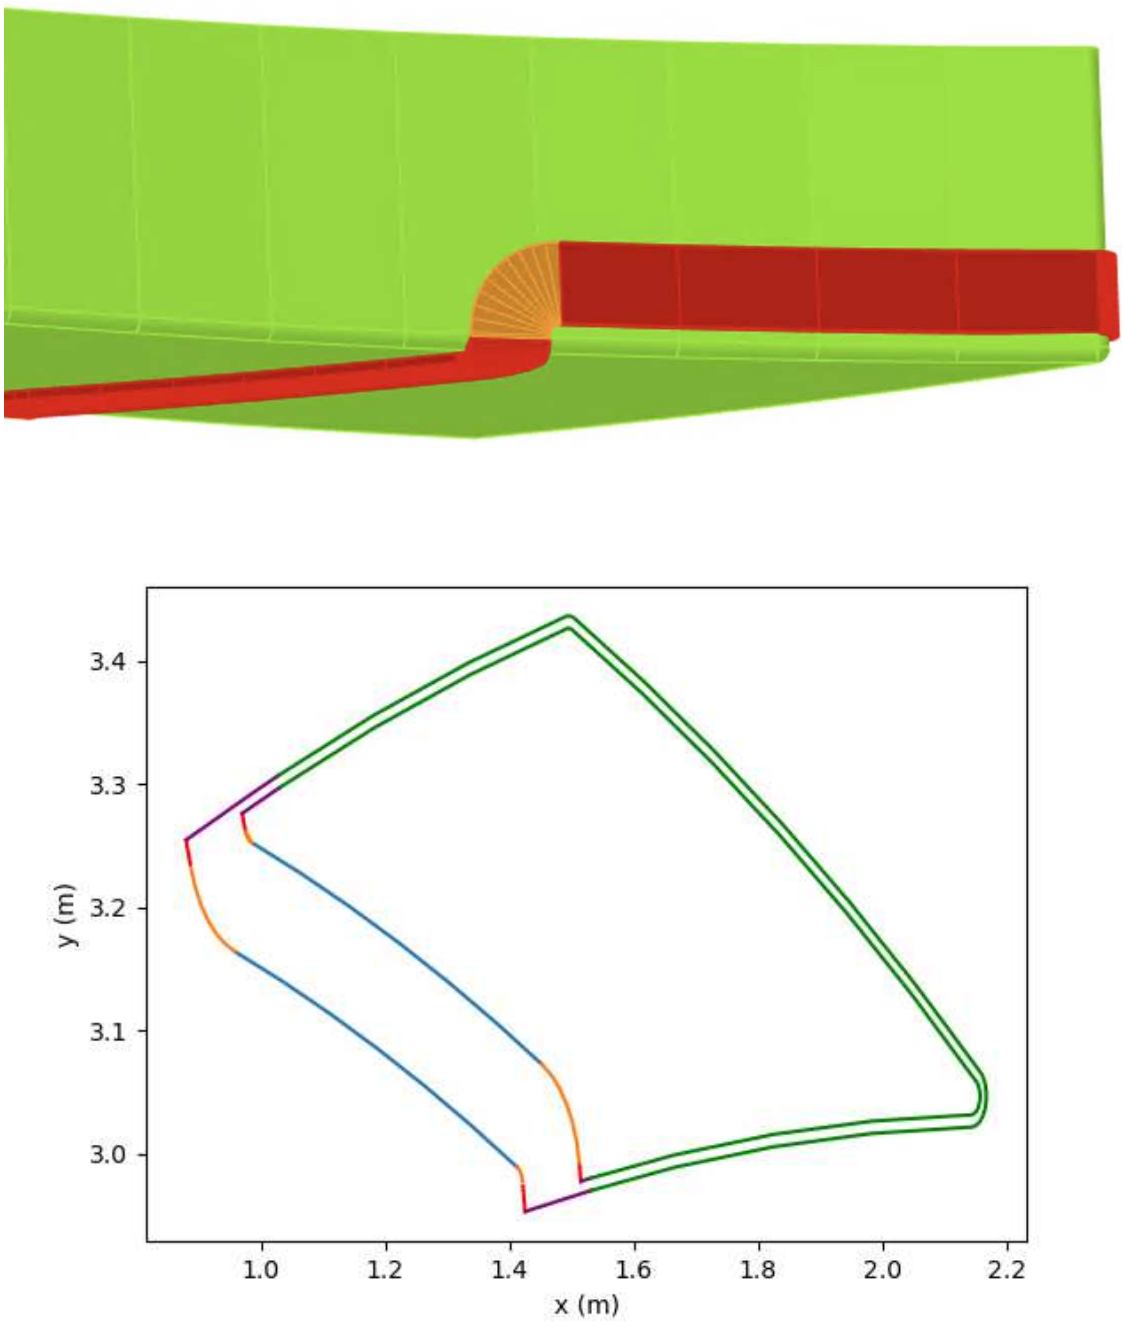


Figure 2.12: (Top) OPERA-3D model showing a trim coil segment (red) on a spiral pole (green), with the highlighted bend responsible for aligning the trim coil with the spiral edge of the pole. (Bottom) The corresponding bends are indicated in purple in the two dimensional Python plot.

2.7 Coil configurations

Trim coils intersect the pole face at different radii and are returned toward either the inner (smaller radius) or outer (larger radius) side of the pole. The magnetic field response is non-local: rather than being confined to the crossing radius on the pole face, it begins to increase at that radius and extends to the location where the trim coil is wrapped around the pole, producing a plateau like profile, as illustrated in Fig. 2.13. The polarity of the resulting magnetic field is determined by the combination of the return direction and the current orientation. To achieve the field profile corresponding to a positive k value, trim coils returning toward the inner side must reduce the main dipole field generated by the main coil, whilst those returning toward the outer side must enhance it. This uniquely determines the required current orientation in each trim coil.

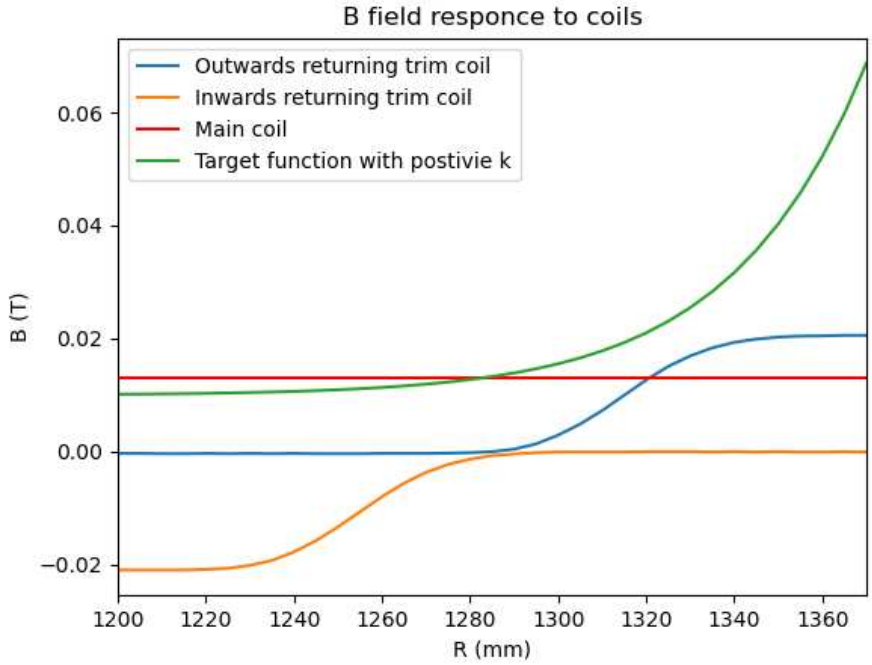


Figure 2.13: The magnetic field response of individual trim coils exhibits a plateau like profile, each either adding to or subtracting from the dipole field provided by the main coil to produce the required target field distribution. The plot shows only the central region of the pole and does not capture end effects.

Trim coils returning on the same side are stacked together, with their thickness accumulating in the direction perpendicular to the spiral pole, as shown in Fig. 2.14. This stacking constrains the available space for other components in the ring. At a point defined by r_0 and θ_0 on a spiral, the opening angle ϕ corresponding to a perpendicular separation a from another spiral with the same spiral angle may be expressed as:

$$\phi = \tan\xi \ln\left(\frac{a}{r_0}\sin\xi + \sqrt{1-\frac{a^2}{r_0^2}\cos^2\xi}\right) + \arcsin\left(\frac{a\cos\xi}{r_0}\right) \tag{2.2}$$

The full derivation of this expression is provided in Appendix C. Without loss of generality, r_0 can

be taken as the radius at which each trim coil crosses the pole, and a as the accumulated thickness of the stacked trim coils at that radius. Evaluating this expression at each crossing radius yields the minimum required opening angle ϕ.

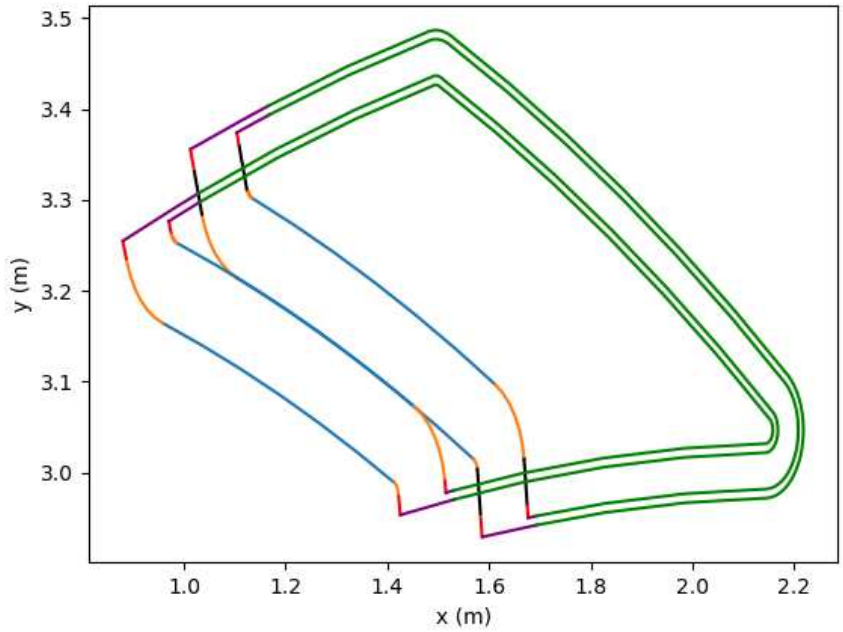


Figure 2.14: Two dimensional Python plot showing the top view of two trim coils stacked radially. Each coil crosses the pole face at a different radius; the black line represents a longitudinal extension of the trim coil that enables adjacent coils to be stacked. The length of the black line has been exaggerated for clarity; in practice, this extension should be kept as short as possible.

Returning all trim coils towards the outer side of the pole represents the most energy efficient configuration, as all trim coil contributions add to the main dipole field. However, this arrangement is space inefficient, since the opening angle ϕ must be sufficiently large to accommodate the full accumulated return thickness. Allowing some trim coils to return towards the inner side of the pole, where they subtract from the main dipole field, reduces the total accumulated return thickness and frees additional space for other ring components. An example illustrating these two configurations is shown in Fig. 2.15.

The opening angle between two spirals remains constant, implying that the physical separation between them increases with radius. Consequently, when configuring the trim coils, the coil that intersects the pole face at the smallest radius should be returned closest to the pole, with coils crossing at progressively larger radii wrapped sequentially outward, one upon the other. This approach exploits the inherent radial spacing afforded by the spiral geometry and maximises space efficiency, as illustrated in Fig. 2.16.

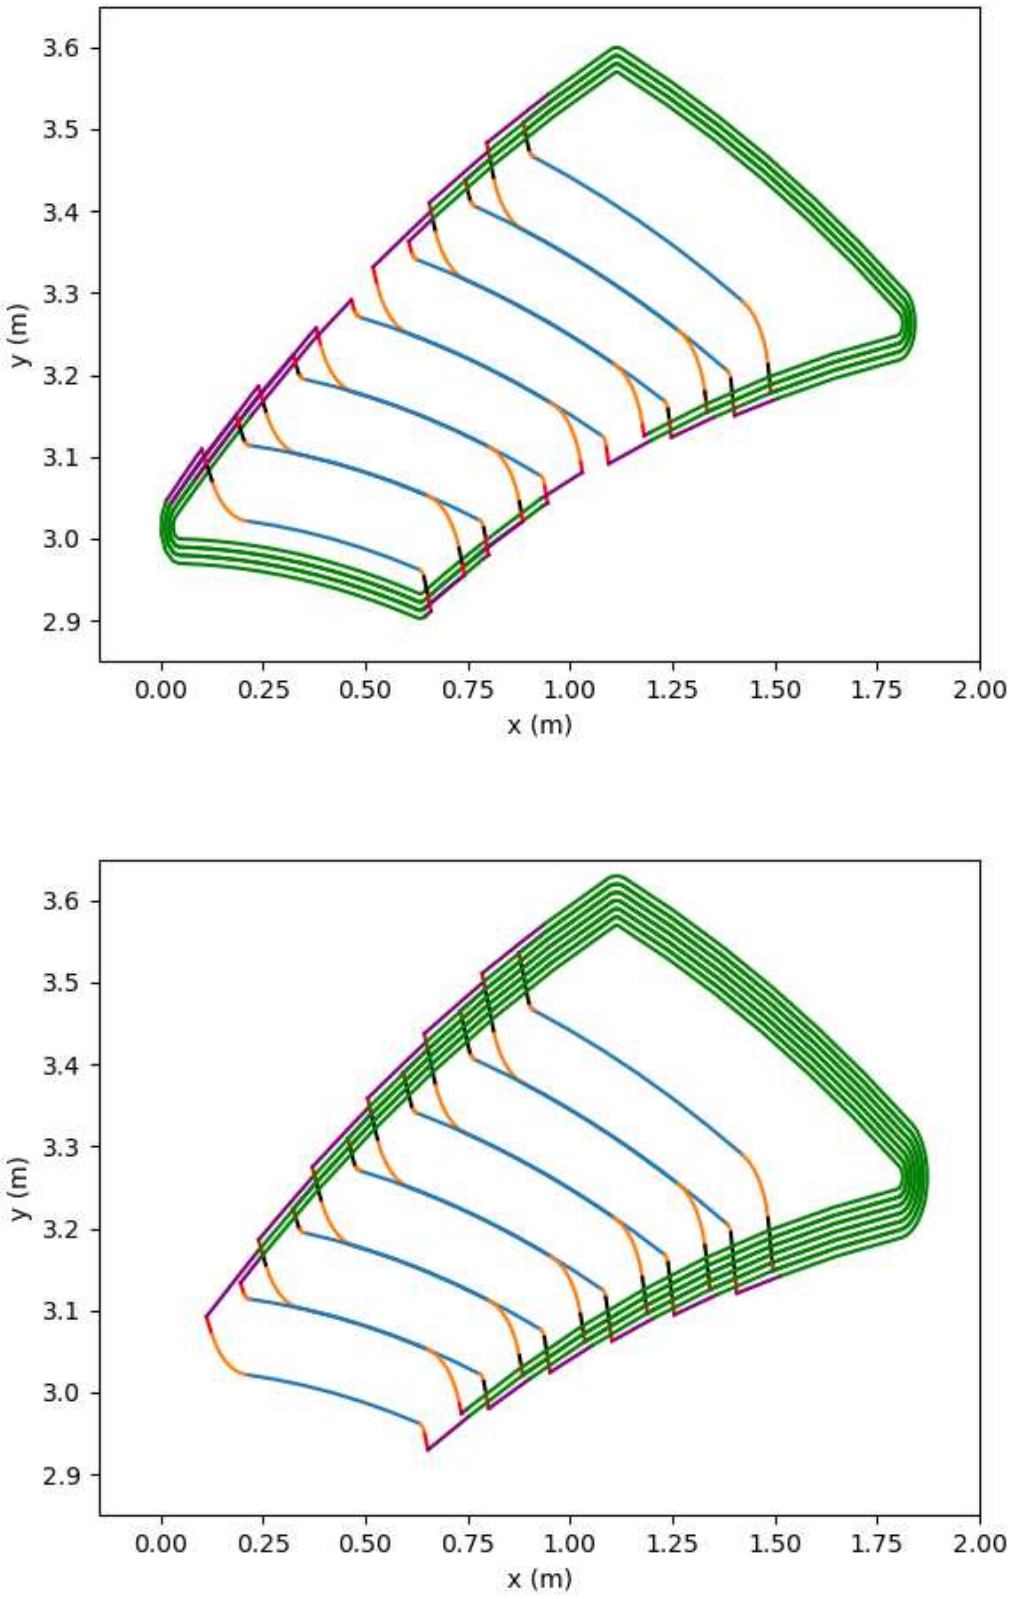


Figure 2.15: Illustration of two different trim coil return path configurations. (Top) Three coils returning towards the inner edge of the pole and three towards the outer edge. (Bottom) All six coils returning towards the outer edge. It is evident that routing all coil returns to the same side results in a greater accumulated thickness.

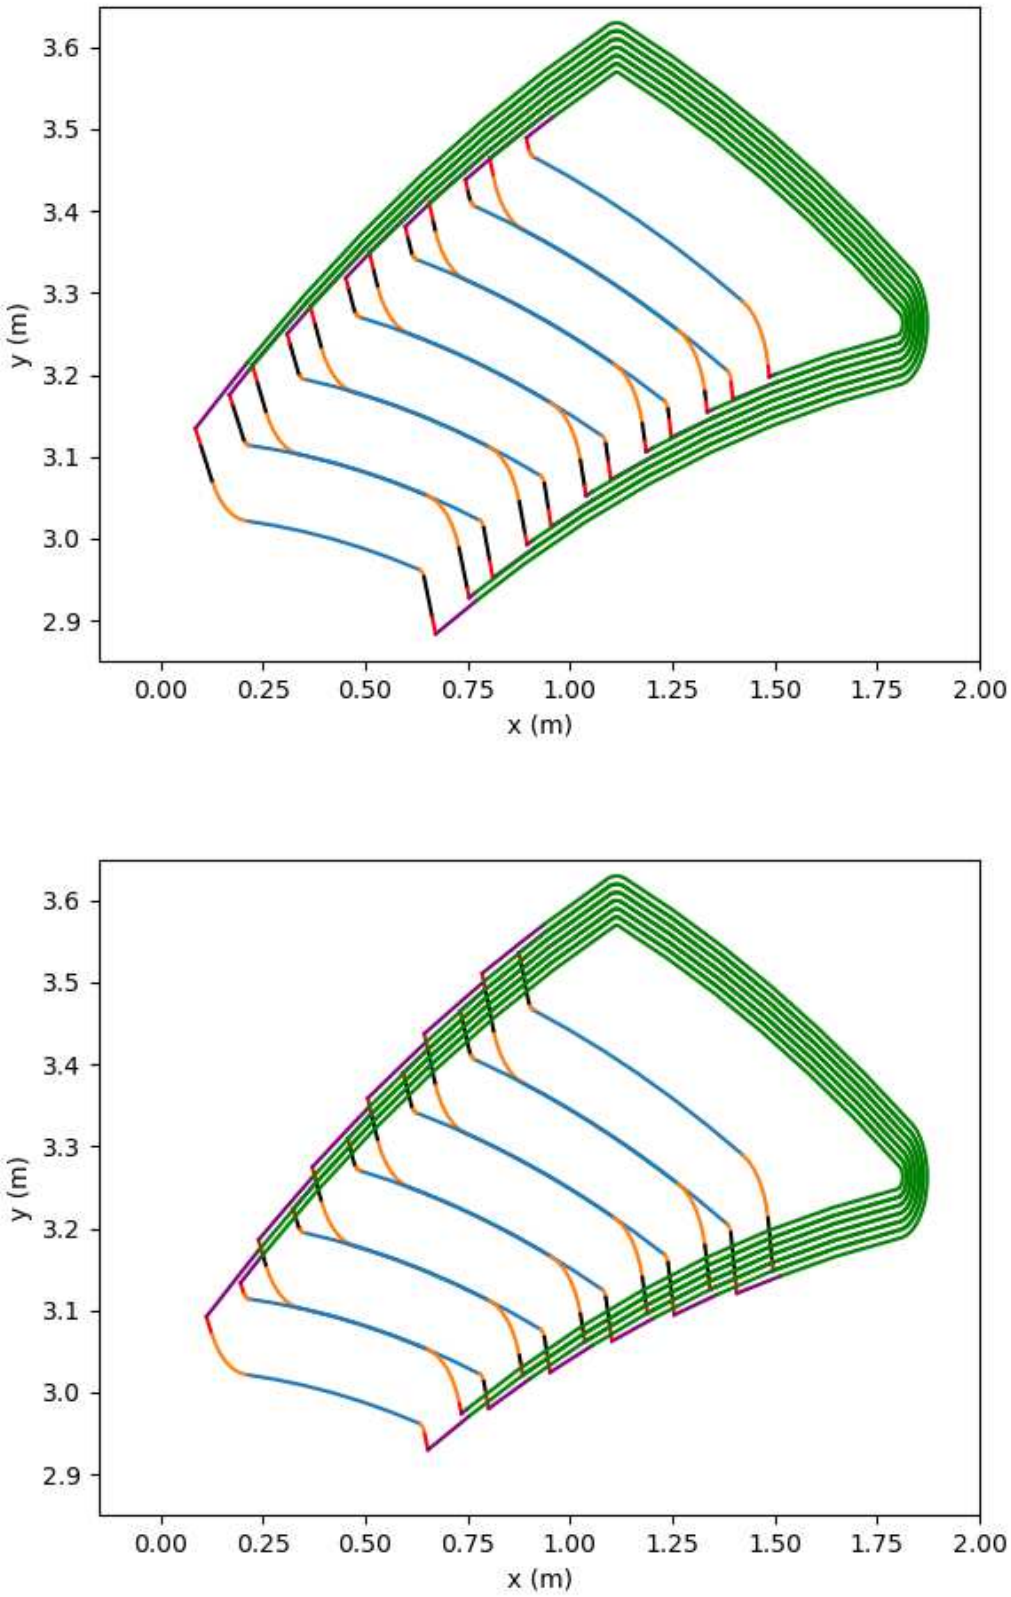


Figure 2.16: Alternative trim coil return path configurations illustrating the effect of return ordering on accumulated conductor thickness. (Top) The trim coil at the smallest radius is returned furthest from the pole, yielding a constant accumulated thickness of return conductors. (Bottom) The trim coil at the smallest radius is returned closest to the pole, such that the accumulated thickness increases with radius, taking advantage of the inherent radial spacing afforded by the spiral geometry.

Thus far, it has been assumed that all trim coils are returned on the same horizontal layer. A further means of reducing the accumulated return thickness is to employ multiple vertical return layers. An example is shown in Fig. 3.7, in which the trim coils are distributed across two separate vertical return layers, reducing the longitudinal width by half. This comes at the cost of an increased pole height, which adds to the physical dimensions and weight of the pole, and requires some coils to return at a greater height, resulting in increased leakage flux.

All of the considerations discussed above, including energy efficiency, space efficiency, and the dimensions and weight of the pole, must be evaluated in the context of the specific requirements of a given project in order to determine the optimal trim coil configuration.

2.8 Boundary condition

Alongside the magnet model, an air region must be defined within OPERA-3D to provide the domain over which the magnetic field is solved numerically. This region must be sufficiently large to capture the bulk of the magnetic flux generated by the magnet and must incorporate appropriate symmetry conditions. To improve mesh regularity, the full model is subdivided into several smaller regions, with finer meshes assigned where greater field accuracy is required and coarser meshes applied elsewhere to limit computational cost. A detailed description of the model subdivision is provided in Appendix D. Although the complete FFA ring possesses rotational symmetry, this was not exploited directly in the magnet modelling. Instead, the magnet was placed at the centre of an air region defined as a 90° sector, with translational symmetry conditions imposed at the two azimuthal boundaries of the sector (the xz and yz planes). This treatment effectively models the magnet as an isolated element, decoupled from interaction with neighbouring magnets; the global azimuthal symmetry of the ring is subsequently recovered by rotating and superimposing the individual field maps.

The vertical symmetry of the magnet is exploited by applying a normal magnetic boundary condition on the median plane ($z = 0$), such that only one half of the iron structure need be explicitly modelled in OPERA-3D, with the other half recovered through reflection. This symmetry condition does not extend to the coils, however, whose fields are computed via the Biot–Savart law prior to the finite element analysis; OPERA-3D's reflection option in the xy plane is employed to conveniently generate the complete coil geometry.

No boundary conditions were imposed on the remaining faces of the model, namely the inner and outer radial boundaries of the sector and the top surface, as these were positioned sufficiently far from the magnet to ensure that the residual field at those boundaries is negligible ($B < 0.0001$).

2.9 Methodology

The objective of the magnet design is to produce a magnetic field that conforms to the scaling law of Eq. 1.85 as closely as possible. Through application of the procedures described in Sections 2.4 and 2.6, the physical geometry of the magnet is configured such that the azimuthal variation of the magnetic field adheres to Eq. 1.85 as closely as possible. Equivalently, the second and third

terms in Eq. 1.88 are rendered as nearly constant as possible through the geometric design of the magnet. This leaves the k index, the first term in Eq. 1.88, as the primary quantity to be optimised. Owing to magnetic saturation and geometric constraints, the effective spiral angles and fringe field extents of the magnetic field may deviate from their nominal design values and must therefore be refined iteratively following optimisation of the integrated field.

To carry out this optimisation, the author developed a series of original methods, including the identification of the integrated field, the spiral angle of the magnetic field, and the extent of the fringe field as the parameters to be optimised. A suite of Python codes was written by the author to extract the field from OPERA-3D, calculate these parameters, and iteratively update the current values in OPERA-3D accordingly.

2.9.1 Target field

Once the OPERA-3D model has been constructed in Metaopera, the current settings obtained from two dimensional studies are used as an initial estimate, and the model is solved using the magnetostatic solver. A three dimensional magnetic field map is then extracted from OPERA-3D for this current configuration. Since, by Maxwell's equations, the magnetic field is fully determined by its distribution on the median plane, the optimisation may be restricted to this plane. Starting from the scaling law in Eq. 1.85, both sides of the equation are multiplied by r and subsequently integrated:

$$\begin{aligned}
\int_0^\theta rB\,d\theta' &= \int_0^\theta rB_0\left(\frac{r}{r_0}\right)^k F\left(\theta' - \tan\zeta \ln\left(\frac{r}{r_0}\right)\right) d\theta' \\
\Rightarrow BL(r) &= r\left(\frac{r}{r_0}\right)^k B_0 \int_0^\theta F\left(\theta' - \tan\zeta\right) d\theta' \\
&= \left(\frac{r}{r_0}\right)^{k+1} r_0 B_0 \int_0^\theta F\left(\theta' - \tan\zeta\right) d\theta' \\
&= BL_0\left(\frac{r}{r_0}\right)^{k+1}
\end{aligned} \tag{2.3}$$

where primed variables denote dummy integration variables. The integrated field BL is defined as:

$$BL(r) = \int_0^\theta rB\,d\theta' \tag{2.4}$$

The integration limits assume that the magnetic field within the opening angle of the magnet cell (0 to θ) is of uniform sign, that is, no field overshoot occurs due to the longitudinal field clamps. The quantity BL_0 denotes the integrated magnetic field at the reference radius r_0.

In anticipation of slight deviations of the magnetic field from the ideal scaling law, the geometrical field index of the magnet k_F is defined as:

$$k_F(r) = \frac{r}{BL(r)}\frac{\partial BL(r)}{\partial r} - 1 \tag{2.5}$$

where the subcript F denotes the horizontal focusing of the magnet. The goal is to make k_F as close as possible to the target field index k.

2.9.2 Jacobian

The next step is to evaluate the Jacobian characterising the sensitivity of the integrated magnetic field to changes in the coil currents. To this end, the current in each coil is perturbed by an amount sufficiently large to exceed the numerical noise associated with the finite element mesh, whilst all other coil currents are held fixed. The total number of simulations therefore equals the number of coils. For each perturbed configuration, a new field map is generated and the corresponding change in the integrated magnetic field is computed. Comparison with the reference configuration yields the sensitivity of the field to each coil current, from which the Jacobian is assembled.

Since the Jacobian depends on the degree of magnetic saturation, it may require recomputation during the optimisation procedure, particularly when the initial field configuration deviates substantially from the target. In such cases, it may be more practical to adjust the currents manually before reapplying the full Jacobian based optimisation algorithm.

The Jacobian is formulated as follows. Suppose the integrated magnetic field is evaluated at n discrete radii, yielding n residual functions collected in the vector $\boldsymbol{\Delta} = (\Delta_0, \Delta_1, ...\Delta_{n-1})$, where Δ_i denotes the residual function at radius r_i. The current settings of the m coils are represented by the vector $\boldsymbol{c} = (c_0, c_1, ...c_{m-1})$, where c_j is the coil setting of coil number j.

The target function is the integrated scaling law:

$$\begin{aligned} y_i = BL_0 \left(\frac{r_i}{r_0} \right)^{k+1} \text{ and} \\ i = 0, 1, 2, ...n-1 \end{aligned} \tag{2.6}$$

The residual is defined as a measure of the deviation from the scaling law for a given set of current settings:

$$\Delta_i(r_i, \boldsymbol{c}) = f(r_i, \boldsymbol{c}) - y_i \tag{2.7}$$

where $f(r_i, \boldsymbol{c})$ denotes the integrated magnetic field at radius r_i for current settings $\boldsymbol{c}$.

The Jacobian of the residuals, $\mathbf{J_{re}}$, is defined as the sensitivity of the residuals to changes in the coil currents:

$$\begin{aligned} (\mathbf{J_{re}})_{ij} = \frac{\partial \Delta_i(r_i, \boldsymbol{c})}{\partial c_j} \\ i = 0, 1, 2, ...n-1 \\ j = 0, 1, 2, ...m-1 \end{aligned} \tag{2.8}$$

This is equivalent to the Jacobian of the integrated magnetic field, $\mathbf{J_{BL}}$, which describes how

the integrated magnetic field at radius r_i varies with the coil currents. This formulation is more straightforward to compute and offers clearer physical insight into the effect of individual coil current adjustments; it is therefore adopted throughout the remainder of this work:

$$\begin{aligned}(\mathbf{J_{BL}})_{ij} &= \frac{\partial f_i(r_i, \boldsymbol{c})}{\partial c_j} \\ i &= 0, 1, 2, ..., n-1 \\ j &= 0, 1, 2, ..., m-1\end{aligned} \tag{2.9}$$

Using this Jacobian, the coil current settings are iteratively optimised until the desired scaling law is achieved.

2.9.3 Gauss-Newton algorithm

The Gauss–Newton algorithm is employed in conjunction with the Jacobian to iteratively update the parameter vector $\boldsymbol{c}$ so as to minimise the sum of squares of the residuals:

$$S(\boldsymbol{c}) = \sum_{i=0}^{m-1} \Delta_i(\boldsymbol{c})^2 \tag{2.10}$$

The procedure is terminated once the sum of squares falls below a prescribed accuracy threshold. Throughout the optimisation, the Jacobian may require recomputation at multiple iterations to accurately account for variations in magnetic saturation.

Beginning from an initial current setting $\boldsymbol{c}^{(0)}$, the iterative update rule is given by:

$$\boldsymbol{c}^{(s+1)} = \boldsymbol{c}^{(s)} - \left(\mathbf{J_{BL}}^{\mathrm{T}}\mathbf{J_{BL}}\right)^{-1}\mathbf{J_{BL}}^{\mathrm{T}}\boldsymbol{\Delta}\left(\boldsymbol{c}^{(s)}\right) \tag{2.11}$$

In practice, the algorithm is implemented within the SciPy module of Python and integrated with OPERA-3D via MetaOpera; further details of the implementation are provided in Appendix E.

2.9.4 Fringe field optimisation

Once the integrated magnetic field has been optimised, the spiral angle must be evaluated, as it may deviate from the design value due to magnetic saturation or other effects. Since a unique spiral angle cannot be straightforwardly defined at every point in space, a combined spiral angle is introduced, whose interpretation is determined by the sign of the focusing contribution. The focusing effects due to the spiral angle are of opposite sign at the two edges: the entrance (En) side $\kappa + \zeta_{En}$, which is horizontal defocusing, whilst the exit (Ex) side $\kappa - \zeta_{Ex}$, is horizontal focusing. The spiral angles $\zeta_{En,Ex}$ are determined from the angle of the center of moment θ_{COM}, of the longitudinal field gradient evaluated along a constant radius.

$$\theta(r)_{COM_{En,Ex}} = \frac{\int \frac{\partial B(r,\theta)}{\partial \theta}\theta d\theta}{\int \frac{\partial B(r,\theta)}{\partial \theta} d\theta}\Big|_{\frac{\partial B(r,\theta)}{\partial \theta}<0, \frac{\partial B(r,\theta)}{\partial \theta}>0} \tag{2.12}$$

The quantity θ_{COM} is evaluated separately for each edge using the boundaries illustrated in Fig. 2.17. It should be noted that $\frac{\partial B(r,\theta)}{\partial \theta}$ is employed solely to identify an effective magnet boundary; its sign does not represent the sign of focusing.

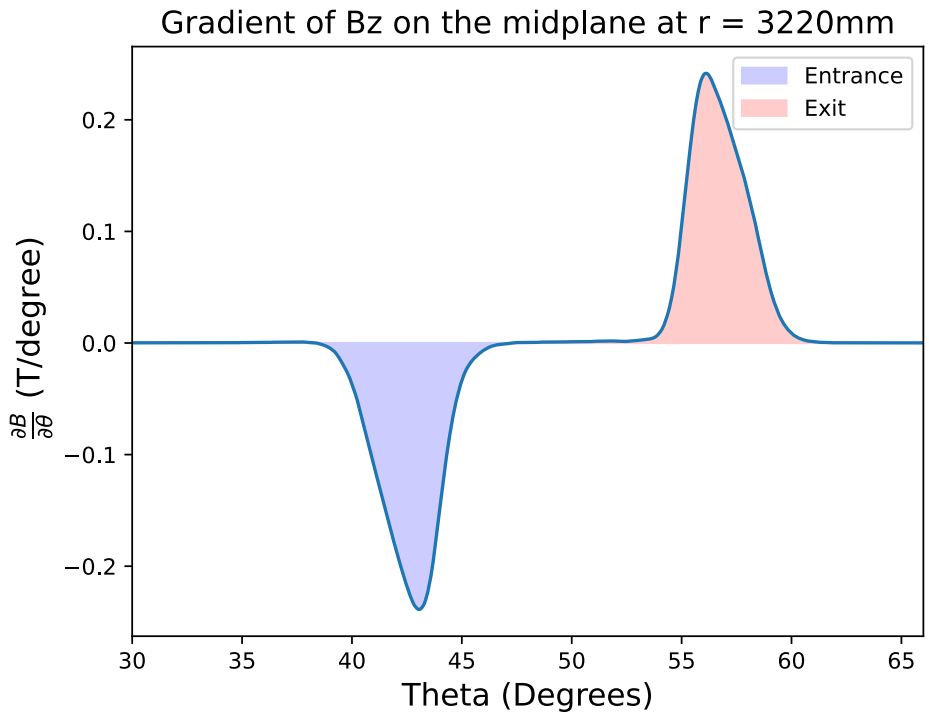


Figure 2.17: Azimuthal field gradient at a fixed radius in the median plane. Using the zero-crossing of the gradient as a boundary, the integrated field may be decomposed into two contributions corresponding to the entrance and exit regions of the magnet.

The spiral angles are subsequently computed as:

$$\zeta(\mathrm{r})_{\mathrm{En,Ex}} = \arctan\left(r\frac{\partial\theta(r)_{COM_{\mathrm{En,Ex}}}}{\partial r}\right) \tag{2.13}$$

The fringe field is a parameter that significantly influences the Thomas focusing properties of the magnet. An effective approach to modelling this behaviour is through Enge's field form factor [49]:

$$\begin{aligned} B(r,\theta) &= \frac{B_{max}(r)}{(1+\exp(P_{En}))\,(1+\exp(P_{Ex}))} \\ P_{En}(\theta) &= C1_{En}\frac{r_0}{\lambda_0}\,(\theta_{COM_{En}} - \theta) \\ P_{Ex}(\theta) &= C1_{Ex}\frac{r_0}{\lambda_0}\,(\theta + \theta_{COM_{Ex}}) \end{aligned} \tag{2.14}$$

where r_0 is the reference radius, λ_0 is the full iron gap height of the magnet. The $C1$ coeffients and B_{max} are obtained by fitting to interpolated magnetic field data. For a perfectly scaling field, the $C1_{En/Ex}$ parameters are expected to remain constant, corresponding to a fringe field extent that increases linearly with radius. When the azimuthal extent of a magnet is small compared to its gap height, B_{max} may not coincide with the true maximum field achieved; it is therefore treated as a free fitted parameter.

The fringe field extent is proportional to the magnet gap height, which in the present design remains fixed. Spiral clamps with a uniform gap height were introduced to promote a linear increase of the fringe field extent with radius; however, as shown in Section 3.6.1 and Section 4.5.1, this measure alone is insufficient. In addition, magnetic saturation can cause the effective spiral angle to deviate from its design value. These fringe field parameters must therefore be corrected following optimisation of the integrated field.

Two corrective strategies may be considered. The first involves modifying the physical spiral angle of the pole. Saturation typically reduces the spiral angle at the exit side, which can be compensated by increasing the pole edge spiral angle at the exit. However, this approach requires additional space within the ring and necessitates a complete redesign of the coils, thereby increasing modelling complexity.

Alternatively, the clamp gap heights may be adjusted to modify the effective magnetic spiral angle, whilst simultaneously affecting the fringe field extent characterised by the $C1$ parameters. At the entrance or exit, increasing the clamp gap height decreases or increases the local magnetic spiral angle, respectively, whilst also extending the fringe field and reducing the corresponding $C1$ values. In either case, the integrated magnetic field is affected by the modification and must be reoptimised accordingly.

In a lattice designed to operate across multiple working tune points or extraction energies, a single magnet/clamp configuration may be insufficient to maintain the desired spiral angle and fringe field characteristics under all conditions. This can give rise to a small increase in the vertical tune spread, since the spiral angle and fringe field predominantly influence the vertical tune whilst having a comparatively minor effect on the horizontal tune. In the present work, the field clamps were optimised to bring the magnetic field closer to the scaling law, and particle tracking was employed to verify that the vertical tune spread across all operating points was minimised.

The strategy adopted in this work was to adjust only the entrance clamp. As discussed in Section 1.6, the edge angle at the entrance, $\kappa + \zeta$, is of greater magnitude than that at the exit, $\kappa - \zeta$, and therefore exerts a stronger influence on the vertical tune. Restricting the adjustment to the entrance clamp reduces the optimisation to a single free parameter, thereby simplifying the procedure.

2.9.5 Tracking

Single particle tracking was performed on field maps following optimisation of the integrated magnetic field. Since a three dimensional magnetic field map can be obtained directly from the optimised OPERA-3D simulation, it was used directly for tracking in preference to a two dimen-

sional median plane field map. This approach avoids the need to extrapolate the field off the median plane and simplifies the overall procedure.

With translational symmetry conditions applied at the two longitudinal boundaries of the magnet model, the influence of neighbouring magnets must be reconstructed by rotating and superimposing the computed field maps. A 90° three dimensional field map of the magnet was extracted from OPERA-3D, which was then rotated in both the positive and negative azimuthal directions by the cell opening angle and superimposed. The central region of the resulting combined map represents a single magnet with its neighbouring magnets on either side. The final field map used for tracking is extracted by selecting the central sector corresponding to one cell opening angle. A number of tracking codes are available for single particle tracking in three dimensional field maps, including Scode [15] and Zgoubi [16]. The code selected for this work was Fixfield [14], a step wise tracking code that computes transfer matrices to evaluate the betatron tune of the particle.

Chapter 3

FETS-FFA

The methods described in Chap. 2 are applied by the author to the design of the FETS-FFA magnets, with the specific parameters for this application introduced below. As a doublet magnet design is required, extensions to the methods are also introduced.

3.1 High intensity accelerator

One of the principal motivations for developing FFA accelerators is the production of high power beams. Unconstrained by the isochronous condition, FFA accelerators can sustain higher field gradients than cyclotrons, which, for a given energy range, reduces the radial beam excursion and enables more compact magnet designs. Furthermore, the static guiding magnetic field eliminates the limitations imposed by magnet ramping in synchrotrons, permitting repetition rates to be increased up to the limits set by the RF system. This characteristic is particularly advantageous for applications demanding high beam power. However, an FFA accelerator specifically designed and optimised for high intensity operation has not yet been realised, and dedicated demonstrator rings are required to validate this capability.

3.2 ISIS-II and FETS-FFA

ISIS is a neutron and muon source located at the Rutherford Appleton Laboratory [50]. It employs a 70 MeV linac followed by an 800 MeV synchrotron to accelerate protons at a repetition rate of up to 50 Hz. The protons are accelerated to approximately 84% of the speed of light before being directed into one of two target stations. In the first target station, the proton beam passes through a thin carbon target to produce muons for experimental programmes before striking the primary tungsten target, where neutrons are generated via spallation. The second target station is optimised for the production of low energy neutrons, providing enhanced flexibility for the study of atomic structure in materials. ISIS-II, the planned upgrade of ISIS, aims to operate as a standalone spallation neutron source delivering proton beams at energies up to 1.2 GeV. The upgraded lattice must support high intensity beam operation at a high repetition rate, with the capability to deliver beams to multiple target stations simultaneously in order to meet capacity

requirements. It must also accommodate a variety of user demands by supporting the generation of different pulse structures [51]. Several upgrade options are currently under consideration, one of which is an FFA ring [52].

To demonstrate the feasibility of an FFA accelerator for this application, a compact proof of principle ring is currently under design. This machine will accept a 3 MeV proton beam from the Front End Test Stand (FETS) linac [53] and accelerate it to 12 MeV, corresponding to a momentum ratio of 2, consistent with that envisaged for ISIS-II. This demonstrator, referred to as the FETS-FFA [54], is intended to validate high intensity FFA operation with minimal beam loss, whilst also enabling controlled variation of the working point within the tune space.

3.3 FETS-FFA lattice design

At an injection energy of 3 MeV, the FETS-FFA is designed to achieve a space charge tune shift of approximately -0.3 [55], comparable to that observed in the SNS and J-PARC accelerators [56, 57]. Experience from these machines indicates that minimising beam loss requires the nominal ring tune to lie close to the diagonal in transverse tune space [58, 59, 60]. Accordingly, the nominal working point was selected as $[3.41, 3.39]$. Although the underlying reason why this operating point yields optimal performance is not yet fully understood, a key lesson from these facilities is the importance of precise control over the working point during beam commissioning to mitigate beam loss [61, 62]. To this end, the FETS-FFA is designed with flexible optics — the range of achievable k values is provided in Table 3.1 — enabling the bare tune to be scanned across the integer range from 3-4 in both the horizontal and vertical planes.

To achieve this flexibility, the novel FD spiral lattice [63] was adapted to incorporate radially distributed trim coils [64]. The transverse tune may be controlled via two parameters: the field index k, which is varied through the trim coil current settings, and the ratio between the magnetic fields of the F magnet (B_F) and the D magnet (B_D). The use of spiral magnets reduces the reverse bending strength B_D required to achieve a given tune, thereby permitting a more compact machine circumference. A spiral angle of 30° was selected, as this was demonstrated to maximise the dynamic aperture for this lattice [65].

For high power accelerator design, sufficient space in the straight sections is essential to accommodate injection and extraction systems. Starting from a 16-fold symmetric lattice, selected straight sections are lengthened whilst others are shortened to yield a 4-fold symmetric lattice with extended straight sections [54], as shown in Fig. 3.1. This configuration remains consistent with the field requirements of Eq. 1.85, but now comprises four doublets within each superperiod. The total circumference is preserved, whilst the longest straight section is increased by 50% to accommodate the injection and extraction systems.

Table 3.1: Specifications of FETS-FFA from lattice design

Parameter	Value
Cell type	FD spiral
Symmetry	4
Number of doublets in the ring	16
Spiral angle (deg)	30.0
k index range	6.401–8.748
Nominal k index	7.717
Injection, extraction proton energy (MeV)	3, 12
Fixed injection radius (m)	3.6

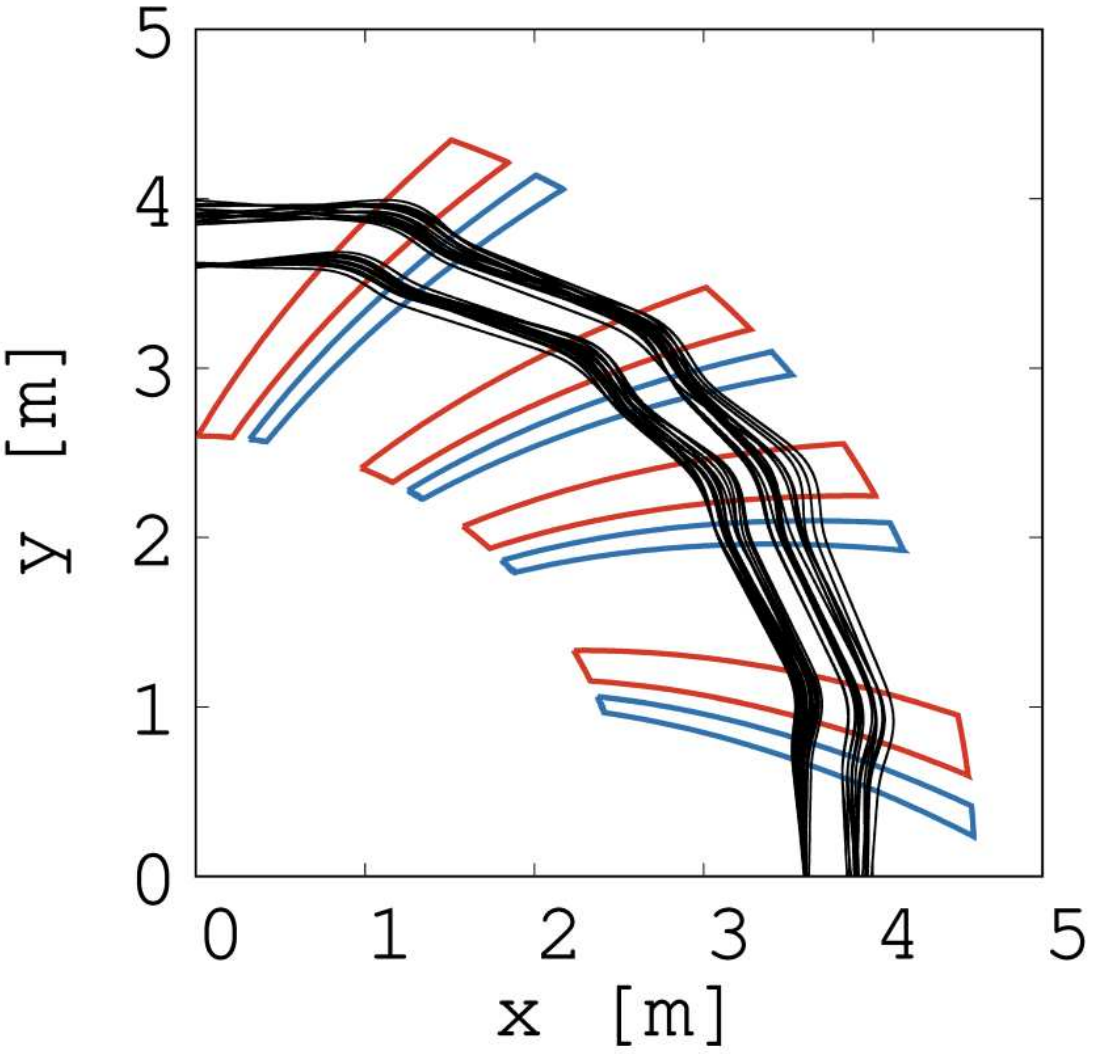


Figure 3.1: Closed orbits at 3 MeV and 12 MeV (shown as black trajectories) are plotted over one superperiod, comprising of four lattice doublets, for different tune points. The radial beam excursion varies with the selected working point and must be accounted for in magnet design.

3.4 FETS-FFA magnet design

Whilst varying the field index k enables access to different tune points, it also modifies the horizontal beam excursion. A radial excursion of 650 mm was therefore specified to cover the full range of working tunes. In addition, a horizontal physical beam acceptance of 40 mm on both the inner and outer sides of the orbit is included, together with an additional 8 mm margin to account for closed orbit distortion. This yields a required good field region of 738 mm [64]. This is the range of radius occupied by the equilibrium orbits, within which the coil currents must be optimised such that the integrated magnetic field follows the desired profile.

A two-dimensional study was carried out by Lagrange [65] to determine the field accuracy as a function of the trim coil configuration. The root mean square (rms) error decreases with increasing number of trim coils, as shown in Fig. 3.2. The final number of trim coils per layer was set to 10, representing a compromise between the number of required power supplies and the achievable field quality. To fully optimise the magnetic field within the good field region, trim coils must be present on either side beyond the boundaries of the good field region. Each trim coil is modelled as a single continuous conductor with a width of 92.5 mm ($\approx \frac{738\ mm}{8}$) and a thickness of 10 mm.

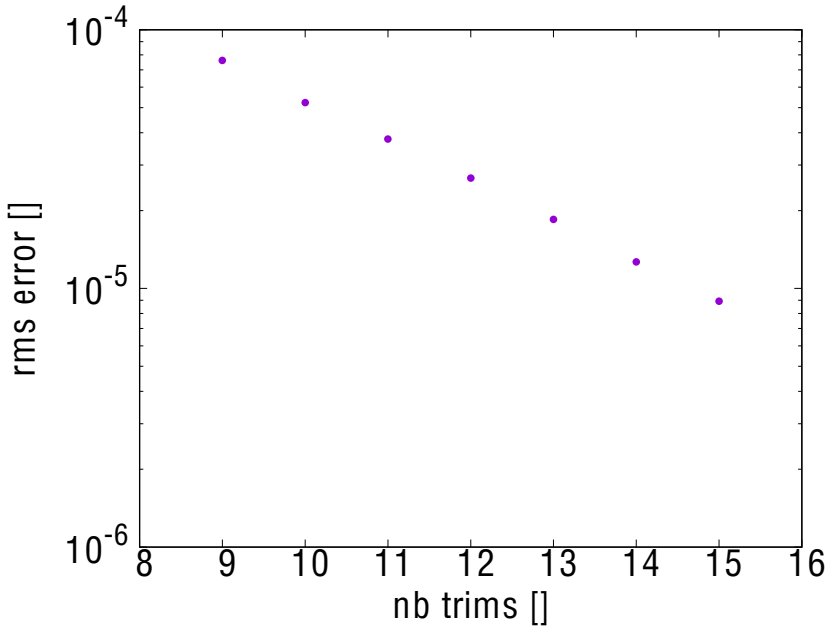


Figure 3.2: Rms field error as a function of the number of trim coils with optimised current settings, obtained from two dimensional studies.

It was observed that introducing an overlap between two layers of trim coils on the pole face, as illustrated in Fig. 3.3, reduces the field oscillations arising from the finite number of coils by approximately a factor of five. Both pole layers operate with identical current settings, avoiding the need for additional power supplies. As shown in Fig. 3.4, the optimal field accuracy is achieved with an overlap equal to half the trim coil width (0.5).

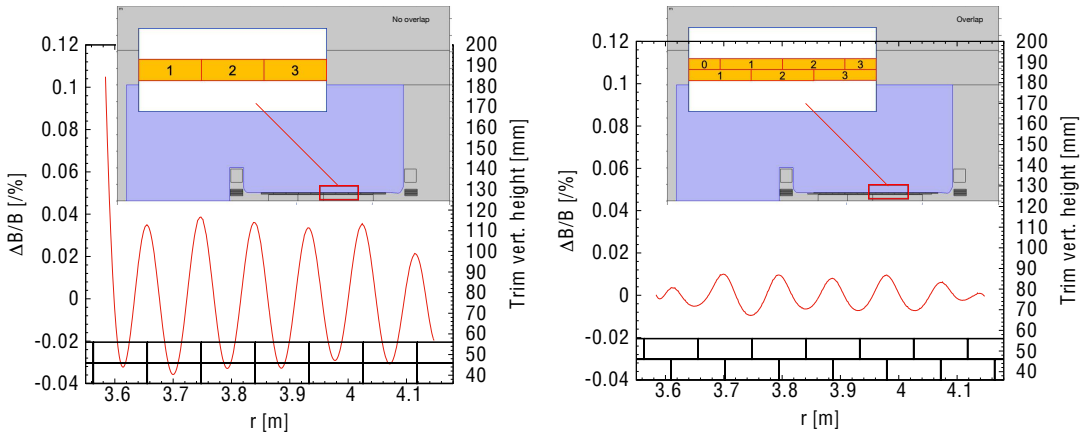


Figure 3.3: Field error with optimised current settings in two dimensions, comparing non-overlapping and half-width overlapping trim coil configurations. The black boxes at the bottom indicate the geometry of the two layers of trim coils on the pole face.

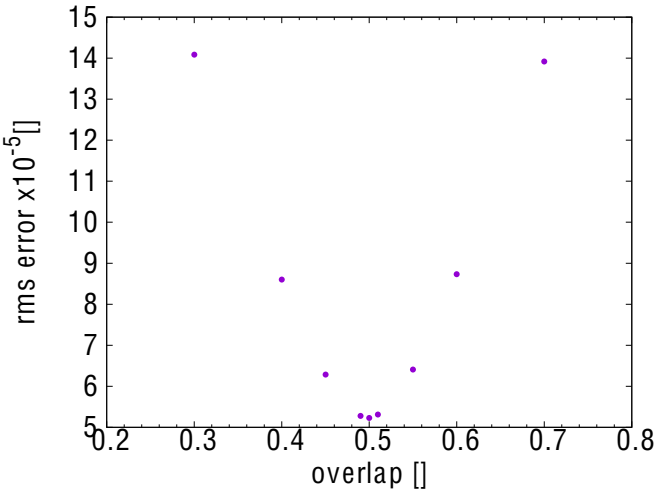


Figure 3.4: Rms field error as a function of the overlap fraction between the two pole layers of trim coils. An overlap of 0.5 corresponds to a displacement of half the trim coil width.

The main parameters of the fourth doublet of the FETS-FFA are summarised in Table 3.2. With 10×92.5 mm wide trim coils on each pole layer and an overlap of 0.5, the trim coils cover a radial range from 3420 mm to 4390 mm. The vertical physical beam acceptance was determined to be ± 23 mm, with an additional ± 3 mm allowance for closed orbit distortion, the vertical beam stay clear is ± 26 mm [65]. Accounting for a 10 mm vacuum chamber, two pole layers of 10 mm thick trim coils, and a 0.5 mm clearance between the coils and the iron pole, the resulting pole half gap is set to $\pm$ 56.5 mm.

Table 3.2: Specifications of FETS-FFA from magnet design

Parameter	Value
F pole opening angle (deg)	4.5
D pole opening angle (deg)	2.5
F/D drift opening angle (deg)	3.55
Opening angle between pole and clamp (deg)	1.92
Gap height of iron pole (mm)	±56.5
Min/Max radius of the trim coils (m)	3.42 / 4.39
Min/Max radius of the poles (m)	3.40 / 4.42
Number of trim coils per magnet	20
Number of power supplies per magnet	10

The opening angle of the D magnet was chosen to be 2.5 °, ensuring that the circular sections of the trim coils on the pole face span at least half of this angle. The opening angle of the F magnet was set to 4.5 ° from lattice design parameters.

Due to space constraints, it was decided not to include a field clamp between the F and D magnets. The opening angle between the F and D poles is therefore entirely determined by the trim coil configuration and is described by Eq. 2.2. The minimum required opening angle can be computed analytically, as shown in Fig. 3.5, and is achieved with 4 trim coils on each pole layer returning towards the inner side, as illustrated in Fig. 3.6.

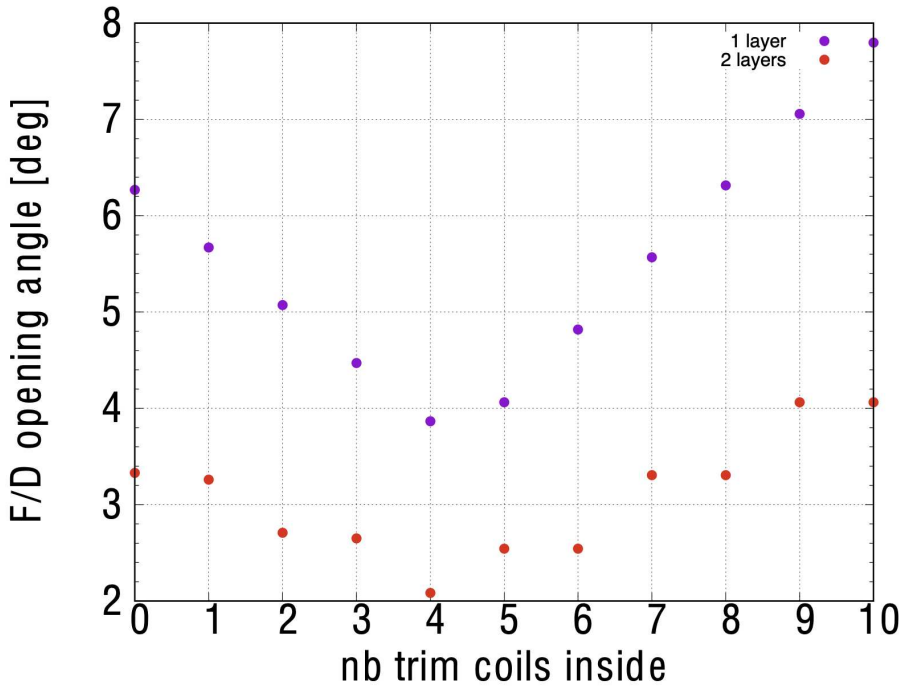


Figure 3.5: Minimum opening angle between the F and D magnets as a function of the number of trim coils per pole layer returning towards the inner side, shown for two configurations: one and two vertical return layers.

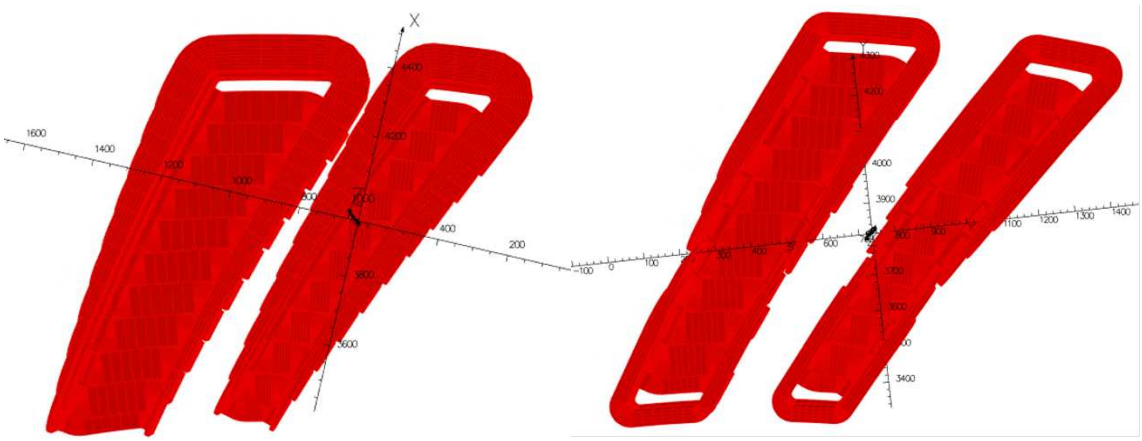


Figure 3.6: Comparison of two trim coil return configurations with identical pole geometry. In the left configuration, all coils return towards the outer side, whilst in the right configuration, 4 coils return towards the inner side and 6 towards the outer side on each pole layer. The right hand configuration permits a reduction in the pole opening angle, enabling a more compact overall design.

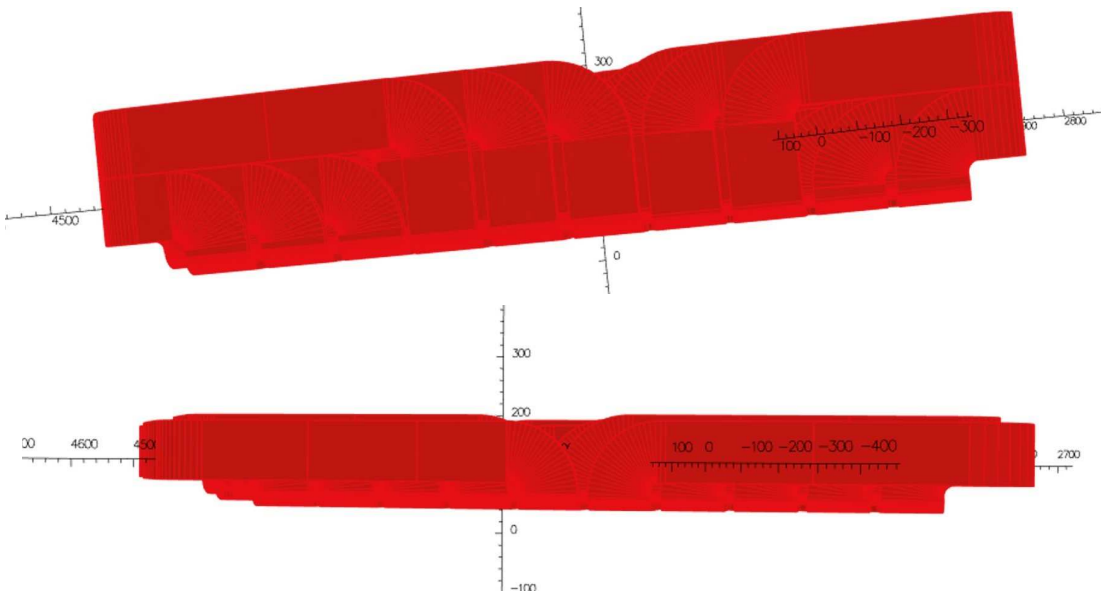


Figure 3.7: Side view of the trim coil return path arrangements for (top) two vertical return layers and (bottom) one vertical return layer.

For this project, all trim coils are configured to return towards the outer side (larger radius) to maximise power efficiency. To minimise the longitudinal thickness of the return conductors, the coils are arranged in two vertical return layers, as shown in Fig. 3.7. It should be noted that this arrangement concerns only the configuration of the coil returns and is independent of the pole layer concept introduced earlier. The use of two vertical return layers comes at the cost of an increased pole height, which renders the design more susceptible to field leakage and magnetic saturation. This configuration was nonetheless adopted as the final solution for power efficiency, with the complete coil arrangement shown in Fig. 3.8.

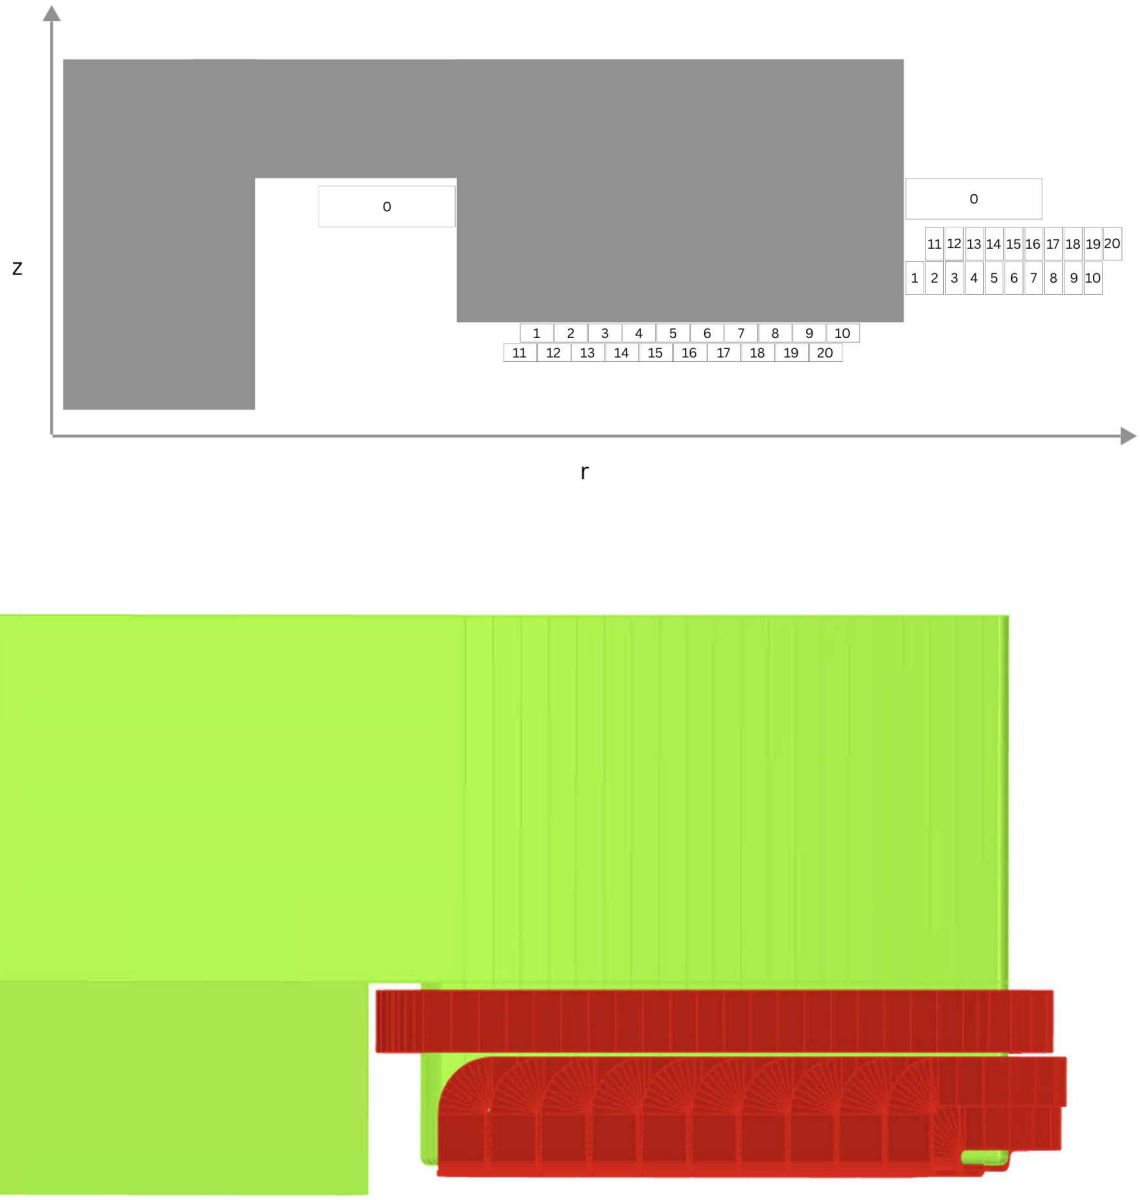


Figure 3.8: Cross-sectional view of the trim coils wrapped around the magnet pole in the FETS-FFA. (Top) Schematic diagram with the main coil labelled 0 and trim coils labelled 1–20. (Bottom) Corresponding OPERA-3D model showing the same cross sectional view.

Accounting for the accumulated thickness of the trim coils between the two magnets, together with a 10 mm margin, the opening angle between the two poles is determined to be 3.55 ° using Eq. 2.2. As shown in Fig. 3.9, the closest point of contact between adjacent coils occurs at the larger radii.

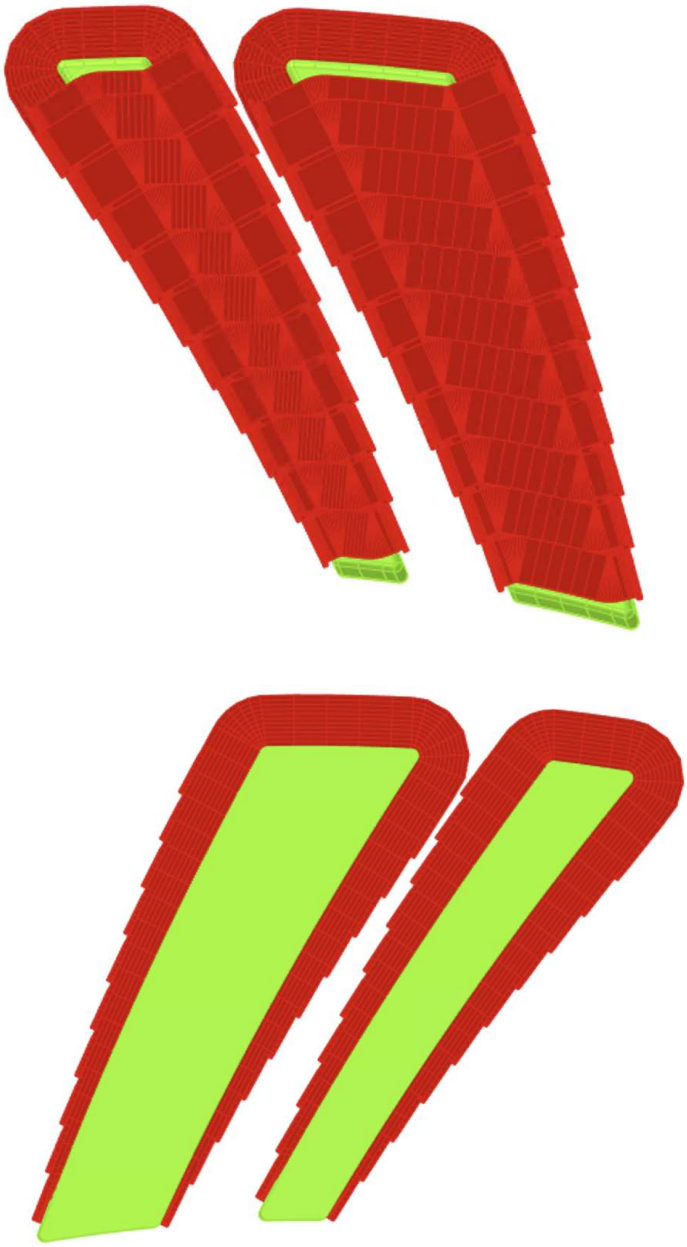


Figure 3.9: Different views of the trim coils (red) wrapped around the iron poles (green), with the centre of the machine towards the bottom of the image.

(Top) Pole face view showing the D pole on the left and the F pole on the right. The trim coils cross the pole face at different radii and are wrapped tightly adjacent to one another.

(Bottom) Top view showing the F pole on the left and the D pole on the right. Since the trim coils stack longitudinally, the available space between the poles decreases with increasing radius, necessitating an opening angle of 3.55° between the two poles.

Once the opening angle between the two poles is fixed, the main coil width is maximised within the available space. The coil height is selected to ensure that the current density remains within acceptable limits. For the FETS-FFA magnet design, with an opening angle of 3.55°, the main coil dimensions were set to a width of 85 mm to fit within this opening angle, as shown in Fig. 3.10, and a height of 117 mm.

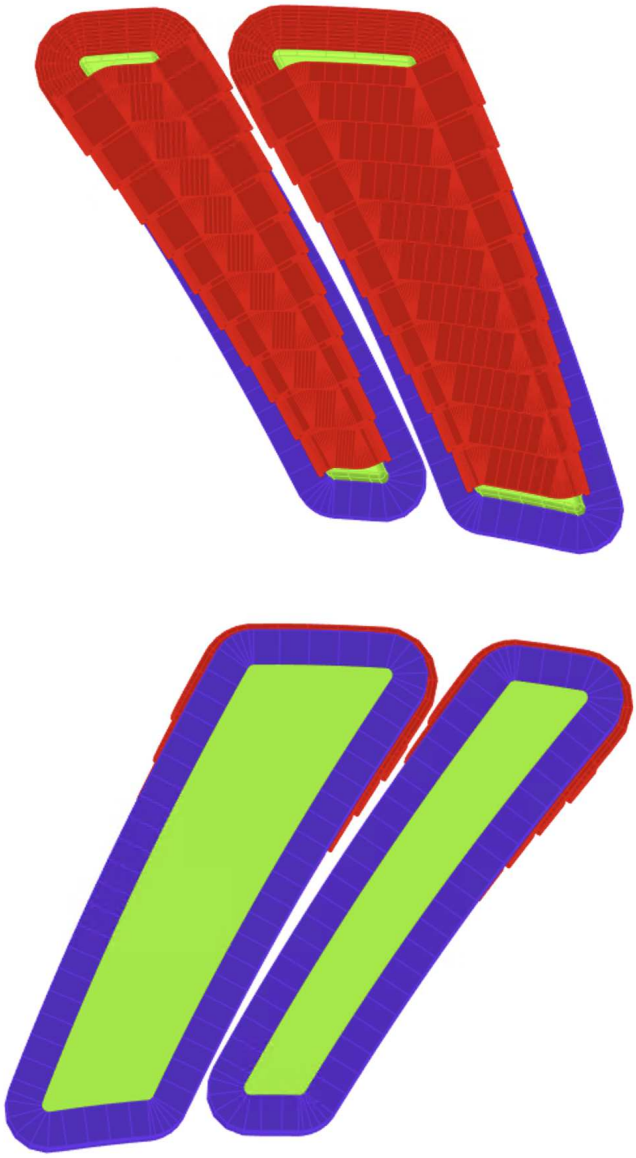


Figure 3.10: Different views of the trim coils (red) and main coils (purple) wrapped around the iron poles (green) are shown, with the machine centre located towards the bottom of the image.

(Top) Pole face view of all coils, with those surrounding the D pole shown on the left and those surrounding the F pole on the right.

(Bottom) Top view of all coils, with those surrounding the F pole on the left and the D pole on the right. Since the main coil thickness is constant, the minimum separation between adjacent main coils occurs at the smallest radius, in contrast to the behaviour observed for the trim coils.

As shown in Fig. 3.11, clamps with a thickness of 50 mm are added on either side of the doublet to control the fringe field. These clamps have a constant aperture of 36 mm, consistent with the height at which the trim coils are positioned. Using Eq. 2.2, the opening angle between the clamps and the poles was determined to be 1.92°. To accommodate the ring within the planned building at the Rutherford Appleton Laboratory, a C-type magnet configuration was adopted, with the return yoke located at the inner radius. A radial clamp was additionally introduced to suppress field leakage in the radial direction.

An investigation of four alternative configurations was conducted using identical coil current settings and yoke height (700 mm) to identify the optimal design in terms of magnetic efficiency. The configurations considered were: (i) all components fully separated, with independent return yokes for the magnets and clamps ("All separated"); (ii) each magnet sharing a common return path with its corresponding clamp ("Clamp magnet connected"); (iii) the F and D magnets sharing a common return yoke, with the clamps remaining independent ("FD connected"); and (iv) a fully connected configuration in which both the F and D magnets and the clamps share a common magnetic return path ("All connected"). The models were solved in OPERA-3D and the resulting field contours for each configuration are shown in Fig. 3.12 and Fig. 3.13.

The magnetic field profile on the median plane, evaluated at the centre of the good field region, is shown in Fig. 3.14. The highest power efficiency is achieved in the "All connected" configuration, in which the clamps and magnet yokes are magnetically linked. In this configuration, the magnetic field attains the largest amplitude and exhibits the sharpest fringe field fall off for the same coil current settings. The "All connected" configuration was therefore selected as the final design for the FETS-FFA magnet.

As shown in Fig. 3.15, connecting the poles magnetically allows the flux from the F pole to return through the D pole, improving the overall magnetic efficiency. As a result, the yoke height was reduced from 700 mm to 300 mm, decreasing the total iron weight of the doublet from 19 tonnes in the "All separated" configuration to 15 tonnes. Accounting for this reduction in yoke height, the final doublet model is presented in Fig. 3.16.

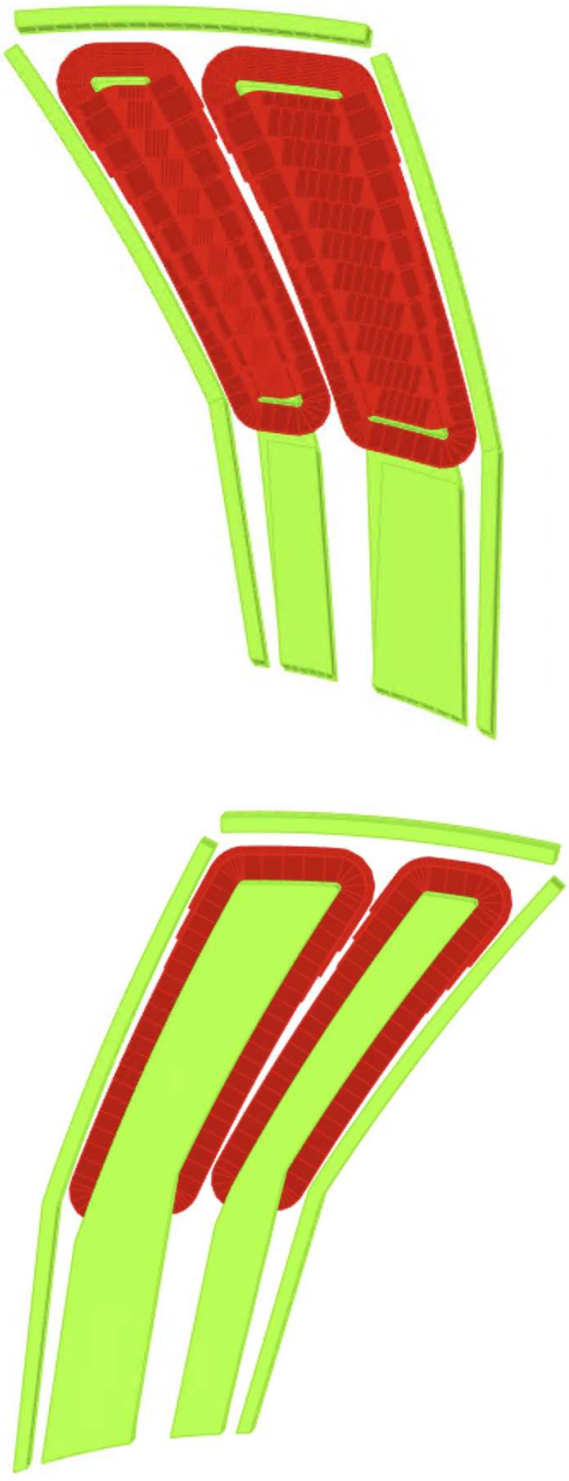


Figure 3.11: Different views of the complete doublet are presented, with the coils (red) wound around the F and D poles, each connected to its respective return yoke; the machine centre is located towards the bottom of the image.

(Top) Pole face view of the doublet. Two longitudinal clamps are positioned on either side of the magnet assembly, with the D clamp on the left and the F clamp on the right, each connected to its corresponding return yoke. A transverse clamp is positioned at the larger radius to suppress field leakage in the radial direction, as required by the C-type magnet configuration.

(Bottom) Top view of the doublet, showing the F magnet with its associated clamp on the left and the D magnet with its associated clamp on the right.

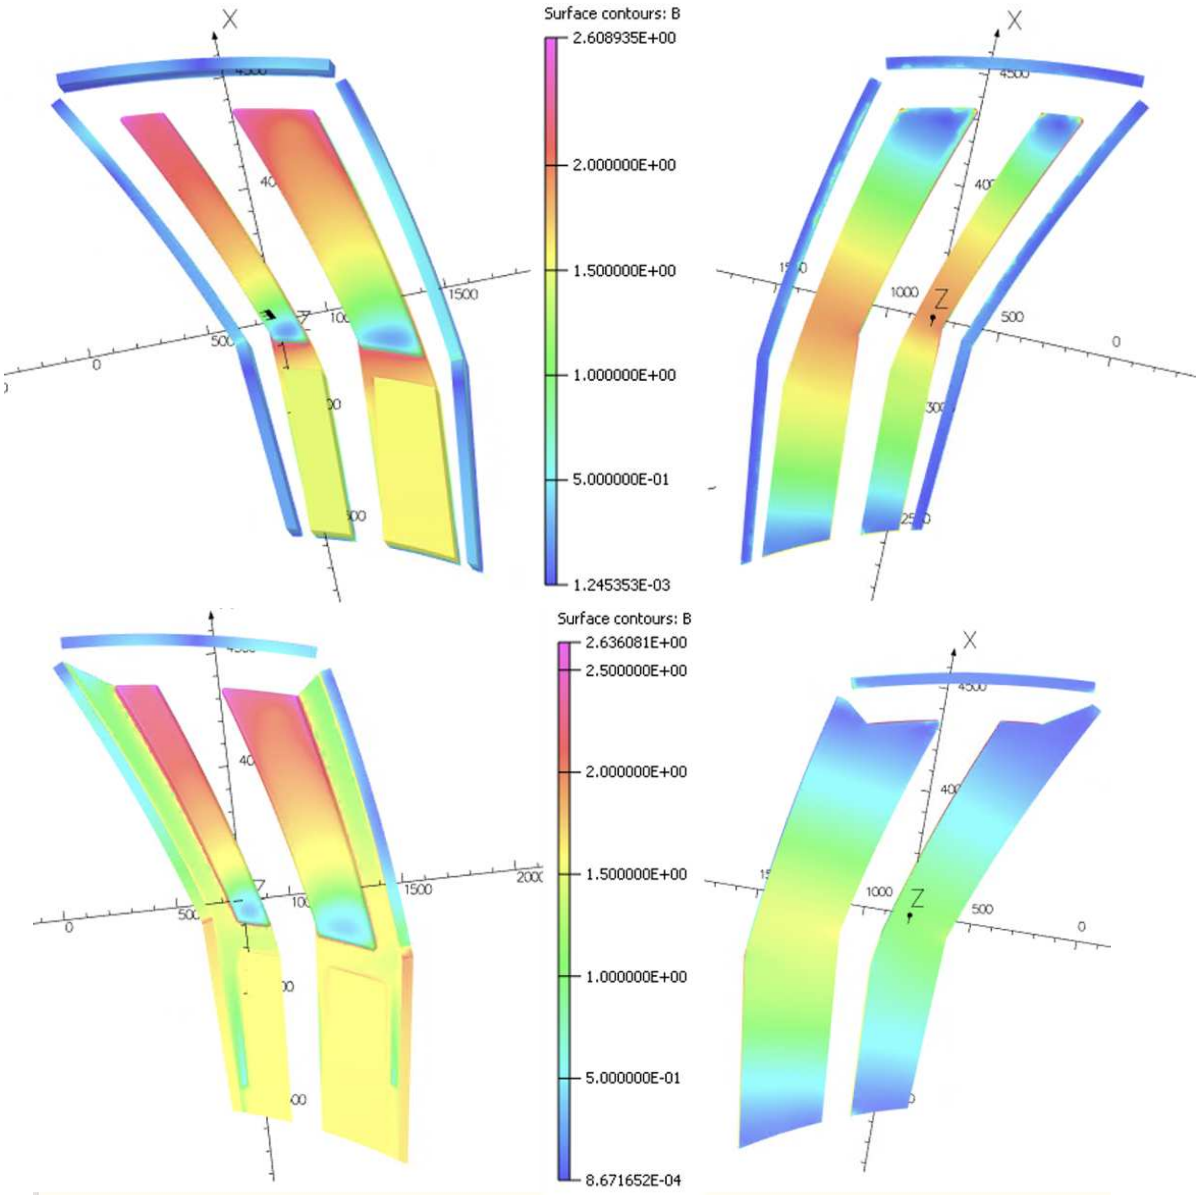


Figure 3.12: Field contours of the "All separated" configuration (top) and the "Clamp magnet connected" configuration (bottom), shown without coils for visual clarity. The left panels present the pole face view and the right panels show the top view; surface contours are given in units of Tesla. The field distributions on the pole face are similar between the two configurations; however, the flux density in the return yoke differs, with reduced magnetic saturation observed in the "clamp magnet connected" case.

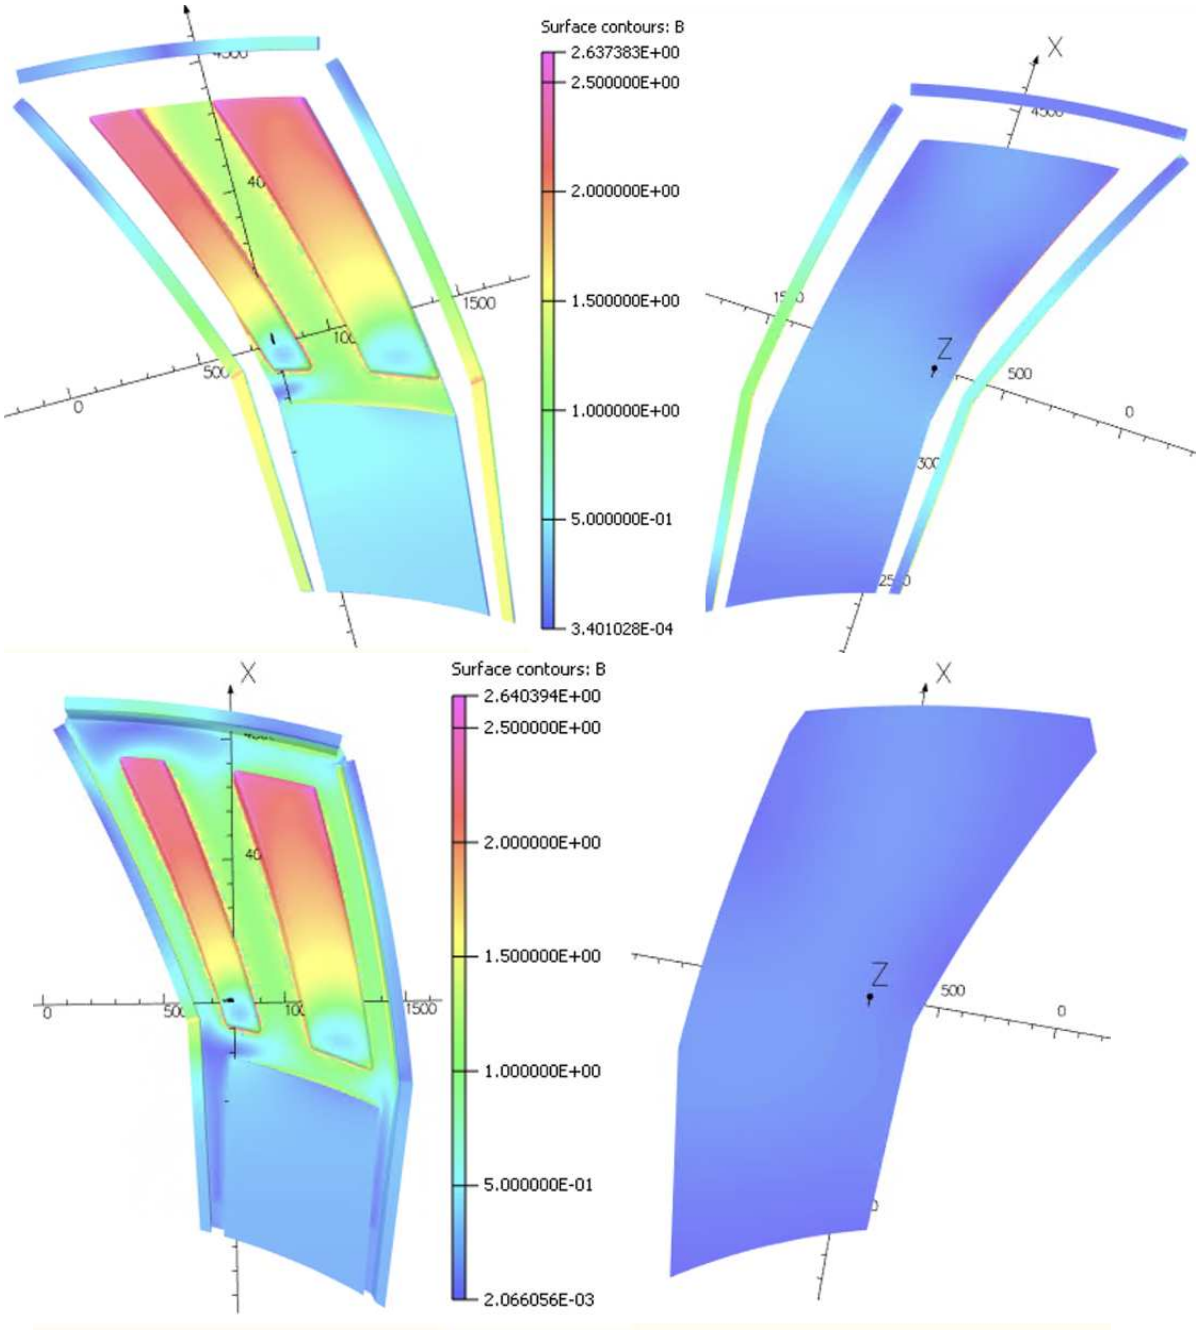


Figure 3.13: Field contours of the "FD connected" configuration (top) and the "All connected" configuration (bottom), shown without coils for visual clarity. The left panels present the pole face view and the right panels show the top view; surface contours are given in units of Tesla. The field distributions on the pole face remain similar to those of the preceding configurations; however, the flux density in the return yoke differs considerably and, in both cases, remains well below the saturation regime.

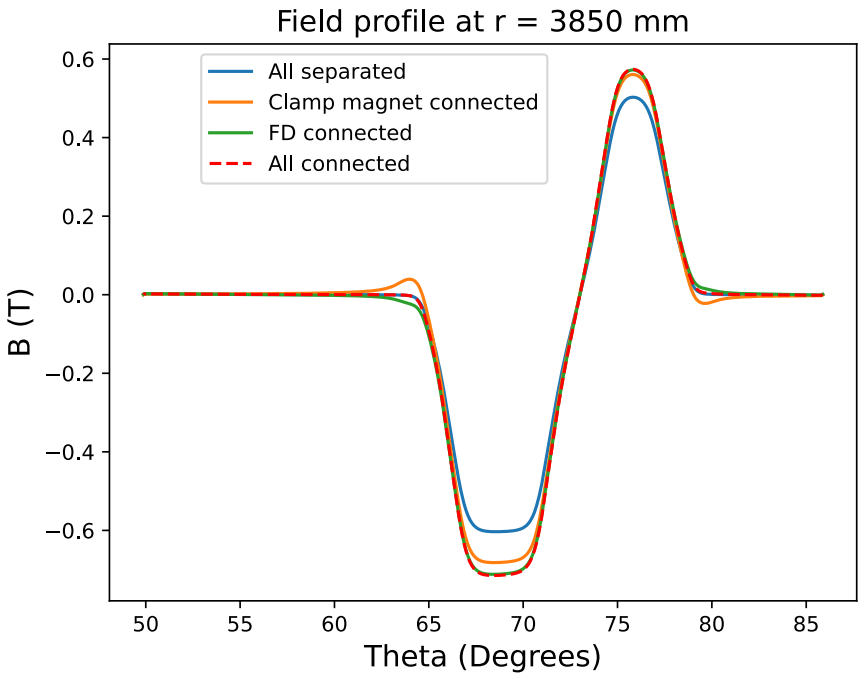


Figure 3.14: Azimuthal field profiles at a fixed radius on the median plane for different iron connection configurations are shown. Despite identical current settings, the magnetic field exhibits distinct behaviors depending on the configuration.

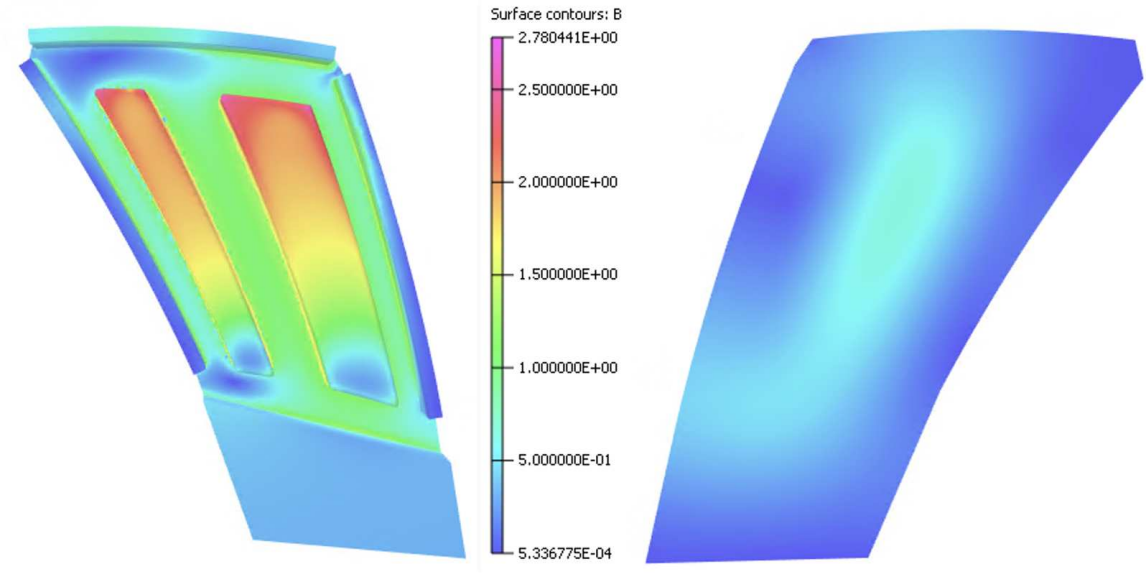
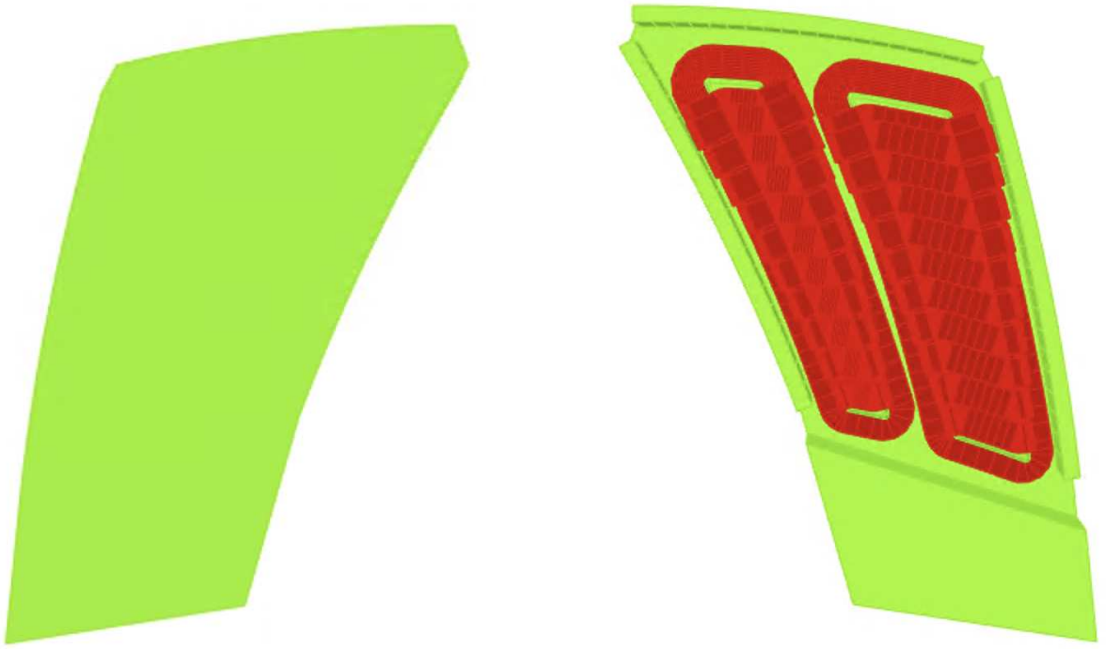
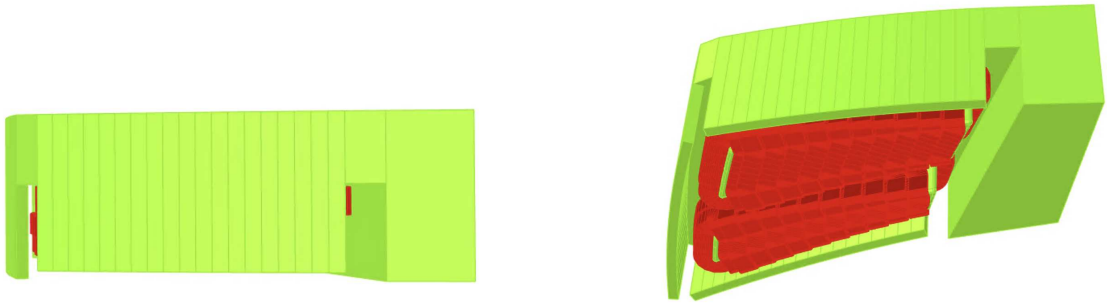


Figure 3.15: Field contours for the "All connected" case after yoke height reduction. The left panel present the pole face view, while the right panel show the top view which illustrates the flux flowing between the F and D pole.



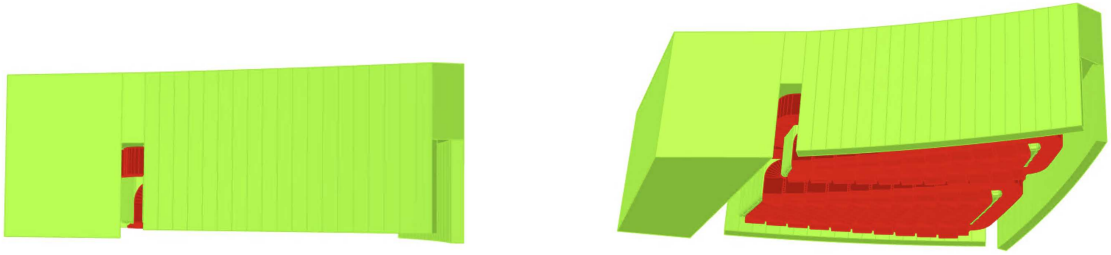
(a) Top view

(b) Pole face view



(c) View of the F_{En} clamp

(d) Second view of the F_{En} clamp



(e) View of the D_{Ex} clamp

(f) Second view of the D_{Ex} clamp

Figure 3.16: OPERA-3D model of the final design of the fourth doublet of FETS-FFA. Coils (red) are wrapped around the F and D pole which shares a common magnetic return yoke with each other and all the clamps.

3.5 Optimisation procedure

The optimisation procedure follows that introduced in Section 2.9, with additional considerations arising from the doublet spiral design. In order to achieve a field distribution consistent with Eq. 1.85, the integrated field $BL(r)$ for both the F and D magnets must each be parametrised by a constant field index k:

$$BL(r)_{F,D} = BL_{0F,D}\left(\frac{r}{r_0}\right)^{k+1} \tag{3.1}$$

where r_0 is the chosen reference radius and $BL_{0F,D}$ is the integrated magnetic field at that radius for the respective magnet.

The integrated B field BL is defined as:

$$BL(r)_{F,D} = r\int B(r,\theta)\,d\theta \tag{3.2}$$

where r and θ are cylindrical coordinates. Since the F and D poles produce vertical bending fields of opposite sign, the integrals are evaluated separately for each magnet. As described in Section 2.8, the doublet magnet is placed at the centre of an air region defined as a 90° sector. The integration limits extend from one side of the sector to a boundary located between the F and D poles, as illustrated in Fig. 3.17.

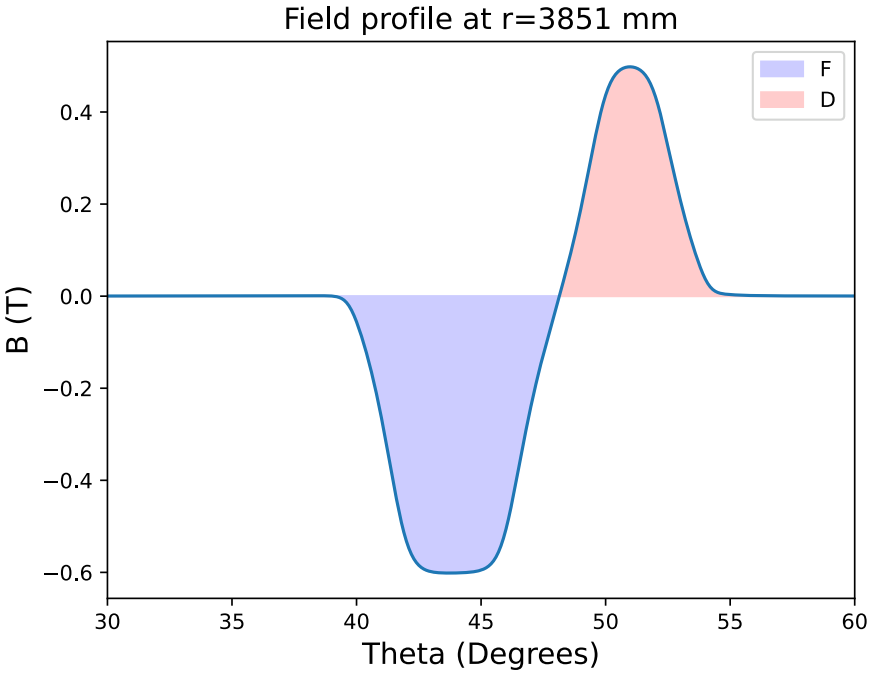


Figure 3.17: Example of the azimuthal magnetic field profile at a constant radius on the median plane of the FETS-FFA magnet. The zero crossing of the field is used as a boundary to decompose the integrated field into separate F and D contributions.

In anticipation of deviations of the modelled magnetic field from the ideal scaling law, a field index $k(r)_{F,D}$ is assigned to each magnet individually:

$$k(r)_{F,D} = \frac{r}{BL(r)_{F,D}} \frac{dBL(r)_{F,D}}{dr} - 1 \tag{3.3}$$

which may differ from the ideal value. The objective of the optimisation is to adjust the coil current settings such that $k(r)_{F,D}$ conforms as closely as possible to the target field index k.

Due to strong magnetic cross talk within the doublet configuration, variations in the coil currents of the F magnet influence the magnetic field in the D magnet, and vice versa. Consequently, the current settings of both magnets must be optimised simultaneously. To account for this coupling, a Jacobian matrix analogous to that defined in Eq. 2.9 is constructed:

$$\begin{aligned} (\mathbf{J_{BL}})^h_{ij\ell} &= \frac{\partial f_i^h(r_i, \boldsymbol{c})}{\partial c_j^\ell}, \\ i &= 1, 2, \ldots, n \\ j &= 0, 1, \ldots, 10 \\ h, \ell &\in \{F, D\}. \end{aligned} \tag{3.4}$$

Here, the index h denotes whether the integrated field of the F or D magnet is being evaluated, whilst ℓ specifies whether the perturbation is applied to the coil currents of the F or D magnet. The index j labels the individual coil within each magnet, according to the numbering scheme defined in Fig. 3.8 and i corresponds to the radial position at which the integrated field is evaluated. This formulation explicitly captures the magnetic cross talk coupling between the two magnets. An example of the Jacobian for the FETS-FFA is shown in Fig. 3.18.

Although each magnet is equipped with 20 trim coils, the two pole face layers of trim coils are operated with identical current settings, such that coils 1 and 11 share the same current, coils 2 and 12 share the same current, and so on. When computing the Jacobian, the current in each corresponding pair of trim coils is perturbed by the same amount. In the Jacobian indexing, each pair is represented by the lower coil number.

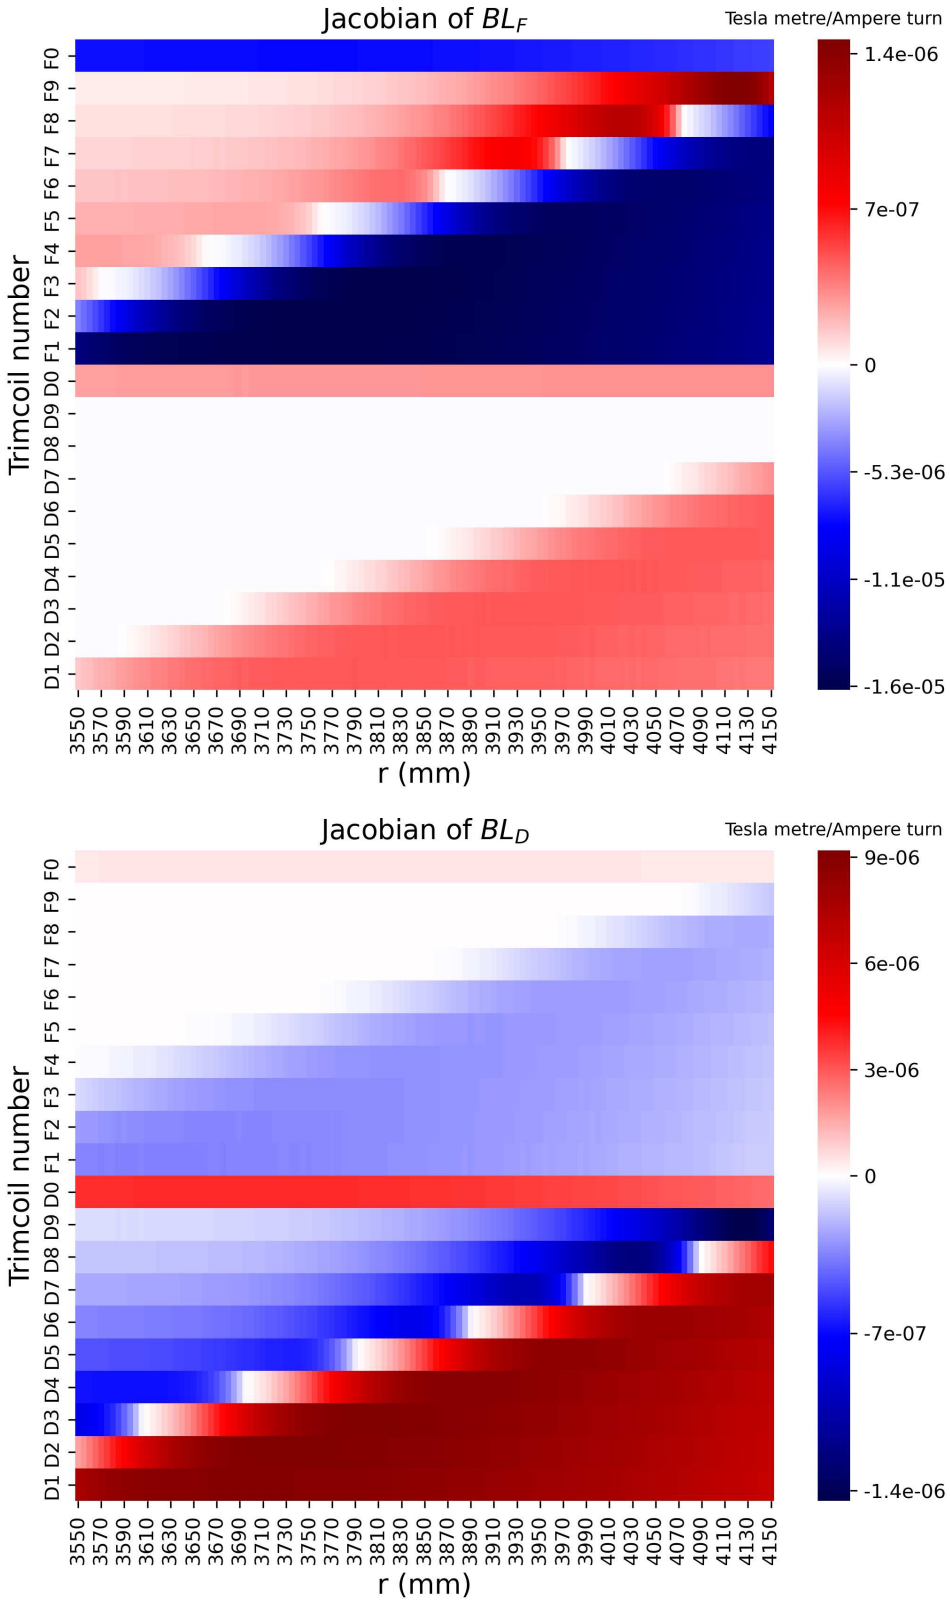


Figure 3.18: The response matrix illustrating how the integrated magnetic field of the F magnet (top) and D magnet (bottom) varies with changes in the respective coil currents. Trim coil 10 is switched off in both magnets as it does not contribute within the given radial range. Although the colour map is centred at zero, it is intentionally asymmetric to emphasise the separation between positive and negative variations: negative deviations are shown with greater magnitude in the top plot, whilst positive deviations are emphasised in the bottom plot. Magnetic cross talk between the two magnets is evident, as variations in the current of a coil in one magnet influence the integrated field in both magnets.

The spiral angles of the field should likewise be maintained as constant as possible. They are determined from the centre of moment angle θ_{COM} of the longitudinal field gradient evaluated along a constant radius:

$$\theta(r)_{COM_{F_{En},FD,D_{Ex}}} = \frac{\int \frac{\partial B(r,\theta)}{\partial \theta} \theta d\theta}{\int \frac{\partial B(r,\theta)}{\partial \theta} d\theta} \tag{3.5}$$

In a doublet spiral configuration, four spiral angles must in principle be considered, as the sign of the edge focusing differs at each magnet boundary: vertical focusing occurs at the entrance of the F magnet and the exit of the D magnet, whilst vertical defocusing occurs at the exit of the F magnet and the entrance of the D magnet. Although all four contributions should ideally be evaluated independently, the integrals corresponding to the F exit and D entrance share the same sign and are difficult to separate due to numerical noise. Since these two contributions have the same focusing sign, they are grouped into a single region. The integration is therefore performed over three regions: the entrance of the F magnet (F_{En}), the region between the F and D magnets (FD), and the exit of the D magnet (D_{Ex}), with the corresponding boundaries illustrated in Fig. 3.19.

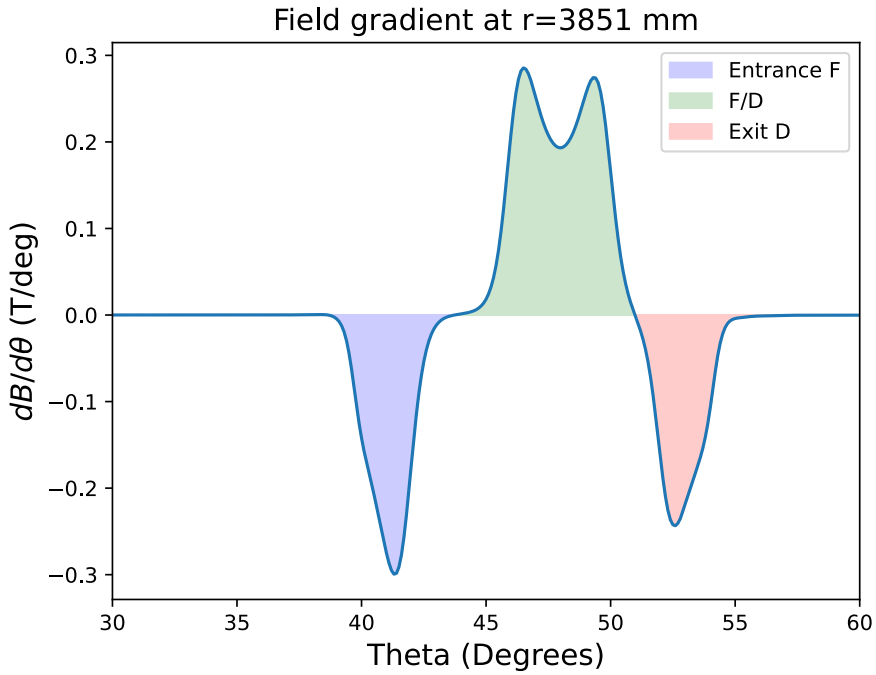


Figure 3.19: Example of the azimuthal magnetic field gradient at a constant radius on the median plane of the FETS-FFA magnet. The zero crossings of the gradient are used as boundaries to decompose the integral into three distinct sections.

The spiral angles $\zeta(r)$ can then be calculated as:

$$\zeta_{F_{En},FD,D_{Ex}}(r) = \arctan\left(r\frac{\partial \theta_{COM_{F_{En},FD,D_{Ex}}}(r)}{\partial r}\right) \tag{3.6}$$

For a doublet magnet, the fringe field can be modelled by addition of the contributions of each spiral magnet taken separately [66], namely

$$
\begin{aligned}
B(r,\theta) = \frac{B_{F_{\max}}(r)}{(1+\exp(P_{F_{En}}))\,(1+\exp(P_{F_{Ex}}))} + \frac{B_{D_{\max}}(r)}{(1+\exp(P_{D_{En}}))\,(1+\exp(P_{D_{Ex}}))} \\
P_{F_{En}}(\theta) = C1_{F_{En}} \frac{r_0}{\lambda_0} \left(\theta_{COM_{F_{En}}} - \theta\right) \\
P_{F_{Ex}}(\theta) = C1_{F_{Ex}} \frac{r_0}{\lambda_0} \left(\theta + \theta_{COM_{F_{Ex}}}\right) \\
P_{D_{En}}(\theta) = C1_{D_{En}} \frac{r_0}{\lambda_0} \left(\theta_{COM_{D_{En}}} - \theta\right) \\
P_{D_{Ex}}(\theta) = C1_{D_{Ex}} \frac{r_0}{\lambda_0} \left(\theta + \theta_{COM_{D_{Ex}}}\right)
\end{aligned}
\tag{3.7}
$$

where r_0 is the reference radius, λ_0 is the full iron gap height of the magnet at the reference radius. The $C1$ coeffients and $B_{F/D_{max}}$ are obtained by fitting to interpolated magnetic field data. Owing to the difficulty in separating the integrals corresponding to the F exit and D entrance, $\theta_{COM_{F_{Ex}}}$ and $\theta_{COM_{D_{En}}}$ were also treated as free parameters in the fit, whilst $\theta_{COM_{F_{En}}}$ and $\theta_{COM_{D_{Ex}}}$ were determined from the calculations in Eq.3.6.

The aperture of the clamp at the entrance of the F magnet (F clamp) was optimised to bring both the effective magnetic spiral angle and the $C1$ parameters closer to their ideal scaling law values — that is, towards constant values — whilst minimising the vertical tune spread across all working points.

3.6 Initial optimisation

3.6.1 Current optimisation

Five working tune points were selected within the tune space targeted by the magnet design: a nominal operating point at $[Q_x = 3.41,\ Q_z = 3.39]$, a notation adopted throughout the following discussion without repeated specification of Q_x and Q_z, together with four corner points spanning the range $3.06 \leq Q_x, Q_z \leq 3.76$. The target parameters corresponding to each working point are summarised in Table 3.3, and the required integrated field values are presented in Fig. 3.20.

Table 3.3: Parameters for different tune points of FETS-FFA

	Parameter	Value
$Q_x = 3.06,\ Q_z = 3.06$	k	6.2013
	GFR $r_{\text{min,max}}$ [m]	(3.57, 4.16)
	$(BL_0)_{F,D}$ [T·m]	(-0.29126, 0.11261)
$Q_x = 3.06,\ Q_z = 3.76$	k	6.3515
	GFR $r_{\text{min,max}}$ [m]	(3.62, 4.26)
	$(BL_0)_{F,D}$ [T·m]	(-0.25922, 0.11516)
$Q_x = 3.41,\ Q_z = 3.39$	k	7.4561
	GFR $r_{\text{min,max}}$ [m]	(3.55, 4.15)
	$(BL_0)_{F,D}$ [T·m]	(-0.31674, 0.16811)
$Q_x = 3.76,\ Q_z = 3.06$	k	8.7484
	GFR $r_{\text{min,max}}$ [m]	(3.54, 4.02)
	$(BL_0)_{F,D}$ [T·m]	(-0.35179, 0.16488)
$Q_x = 3.76,\ Q_z = 3.76$	k	8.5597
	$r_{\text{min,max}}$ [m]	(3.57, 4.10)
	$(BL_0)_{F,D}$ [T·m]	(-0.32911, 0.19514)

Since the radial beam excursion depends on the selected working tune point, trim coils that do not contribute to the integrated field within the relevant radial range are switched off to improve power efficiency. Consequently, a separate Jacobian is required for each working point, reflecting the different radial ranges and sets of active trim coils. These Jacobians are used in conjunction with the Gauss–Newton algorithm (Eq. 2.11) to optimise the integrated magnetic field, with the resulting optimised $k_{F,D}$ values shown in Fig. 3.21.

The oscillations observed in the error arise from the discretisation of the trim coils into elements of finite size. This behaviour is consistently present across all evaluated parameters, including the $k_{F,D}$ values, spiral angles (Fig. 3.22), $C1$ coefficients (Fig. 3.23), and the resulting tunes.

The spiral angle behaviour across the five working tune points is broadly similar. In addition to the oscillatory error, $\zeta_{D_{Ex}}$ exhibits an increasing trend with radius, whereas $\zeta_{F_{En}}$ and ζ_{FD} remain approximately constant, albeit at slightly different values, with ζ_{FD} decreasing at the start and end of the radius range. An exception is the working point [3.06, 3.76], where the radial beam excursion is largest and the outer radial region is more susceptible to magnetic saturation, leading to a pronounced decrease in ζ_{FD} of more than 1.5°, as shown in Fig. 3.22.

The $C1$ coefficients display a similar trend across all five working tune points. Although their absolute values differ slightly, all exhibit a weak increasing trend with radius, whereas a perfectly scaling field would require them to remain constant. This indicates that achieving a scaling fringe field cannot be accomplished solely through the use of longitudinal field clamps with a constant aperture following a constant spiral angle.

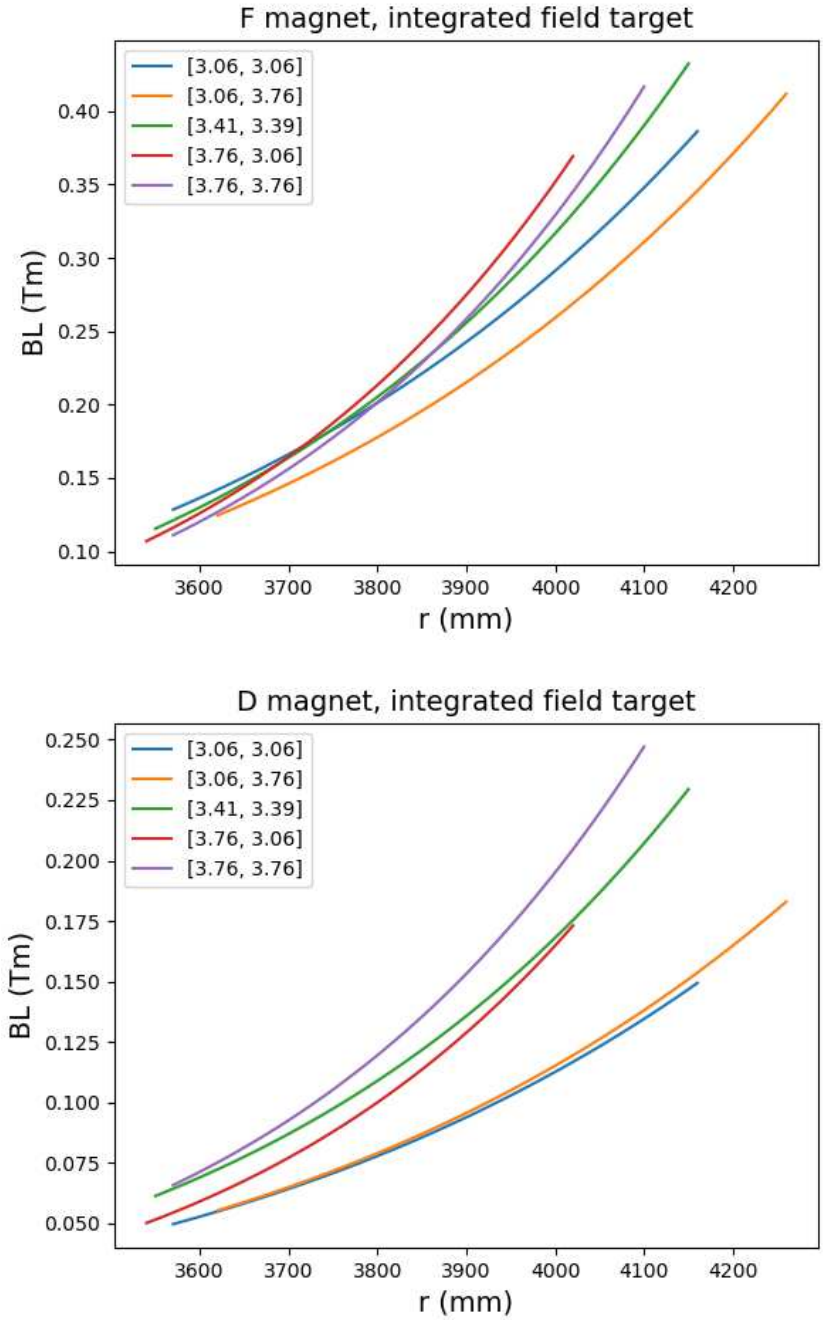


Figure 3.20: The integrated field required for each working tune point differs accordingly for the F magnet (top) and D magnet (bottom). The magnetic field in the F magnet is negative in sign; its absolute value is plotted here for clarity. In addition to differences in field strength and gradient, the extent of the good field region varies between working points, reflecting the change in radial beam excursion associated with each tune.

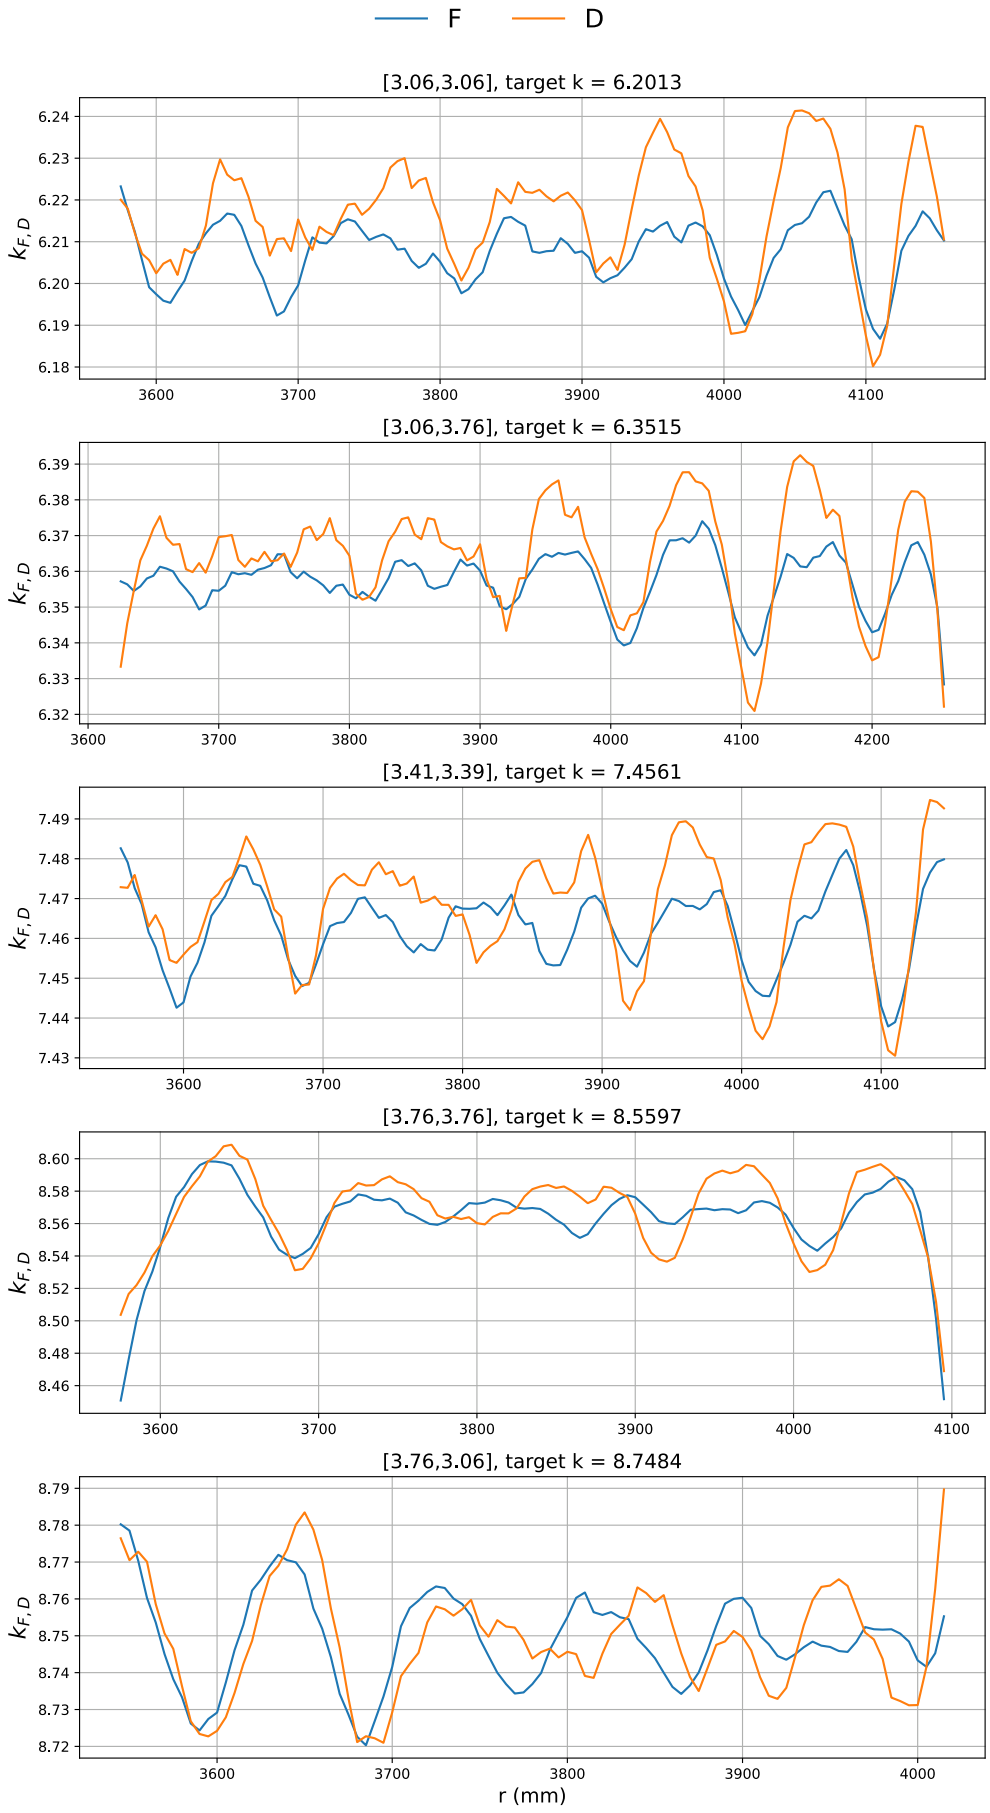


Figure 3.21: The field index k for both the F and D magnets evaluated across the five working tune points.

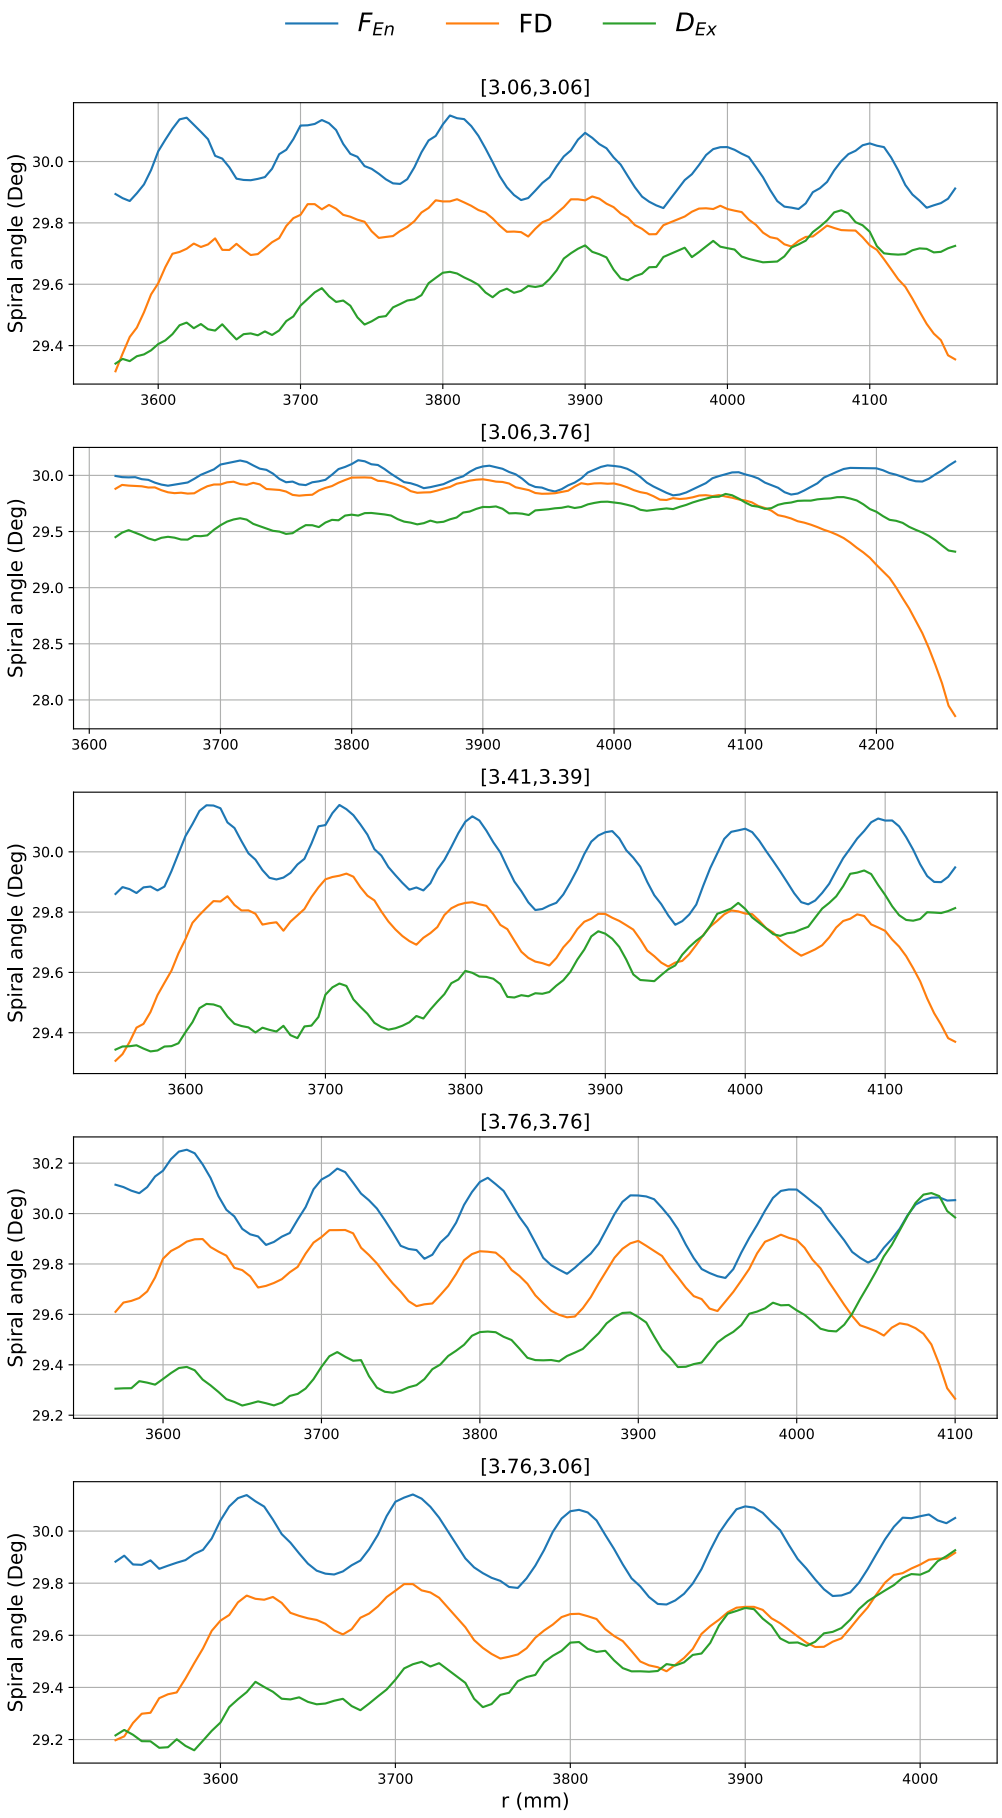


Figure 3.22: The effective spiral angles evaluated across all five working tune points.

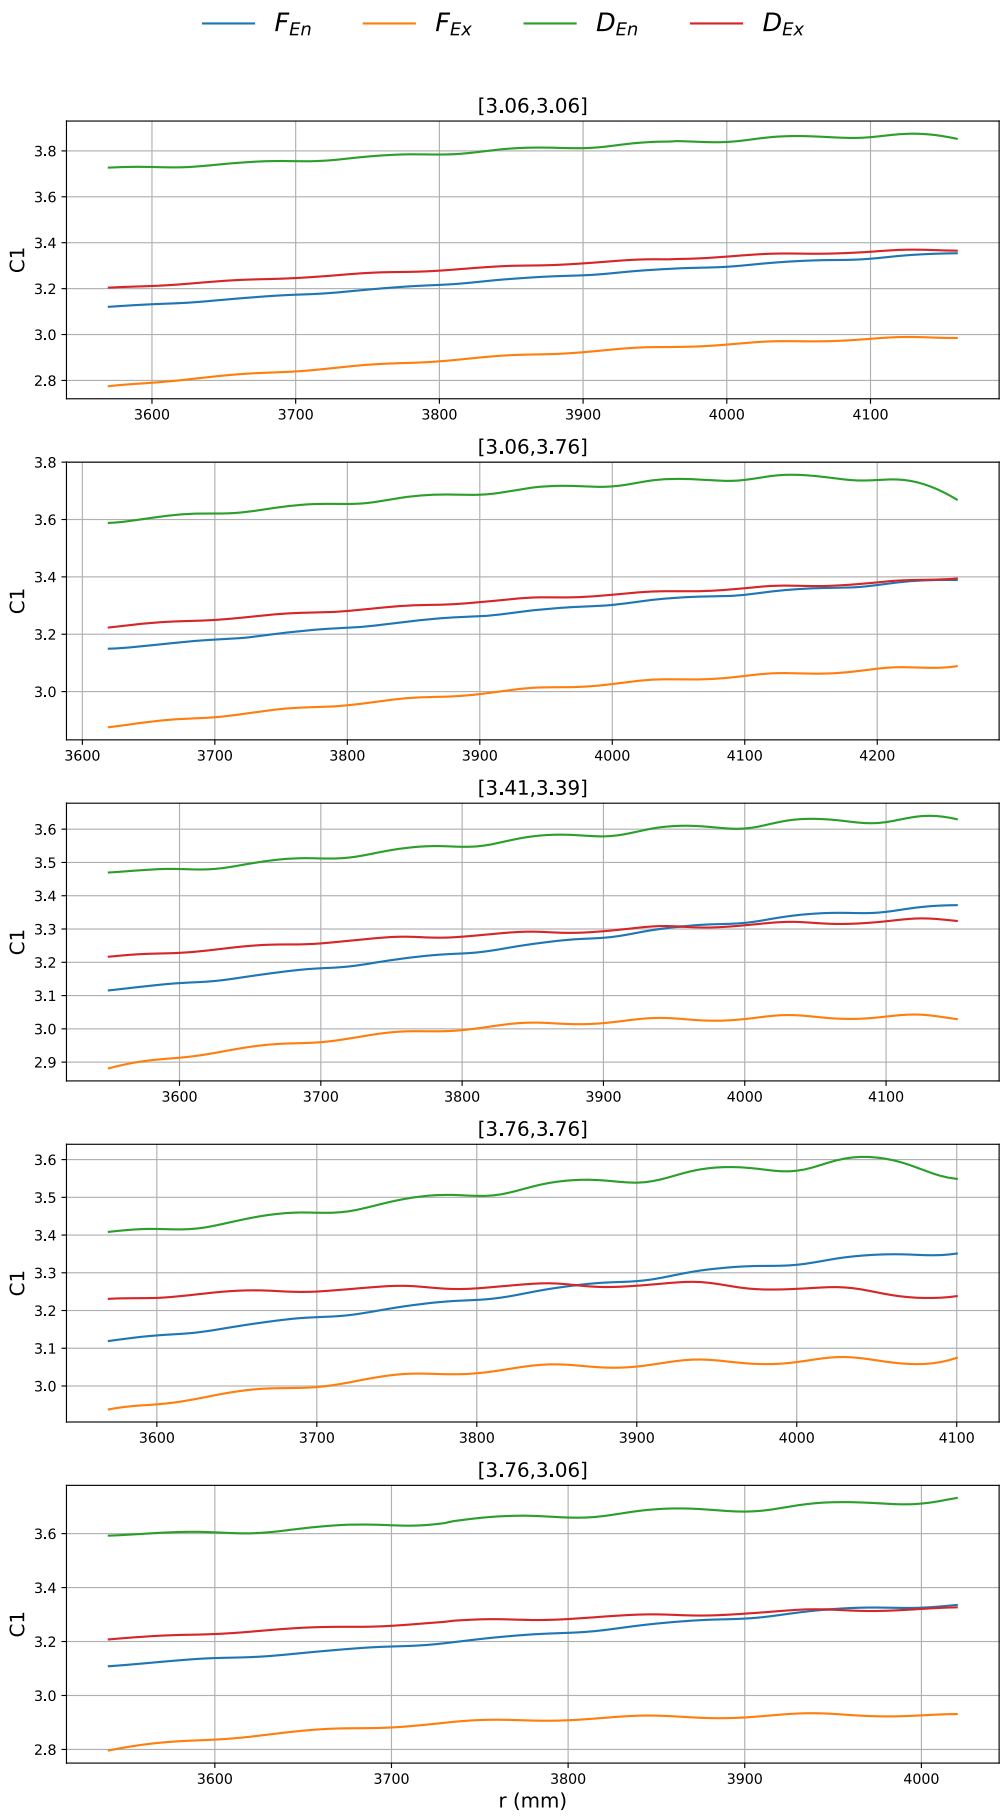


Figure 3.23: The $C1$ coefficients evaluated across all five working tune points.

3.6.2 Tracking

The full FETS-FFA ring exhibits fourfold symmetry, and both the good field region and the required integrated field are defined accordingly. However, only the fourth doublet has been modelled in the three dimensional simulations. Tracking studies were therefore performed assuming sixteenfold symmetry, with this doublet treated as the sole magnet family. This approach introduces two discrepancies relative to tracking in the true fourfold symmetric lattice. First, the ring tunes obtained under sixteenfold symmetry differ from those specified by the lattice design (Table 3.6.1). Whilst these working points are retained for labelling purposes, the tracking analysis focuses on tune spread rather than absolute tune values. Second, the energy range associated with the good field region differs from the design range of 3–12 MeV. In tracking, particles are propagated across the full specified good field region for each working point, irrespective of energy; consequently, the effective energy range varies between working points in the final results.

The aim of the magnet design was to achieve an overall ring tune spread of ± 0.01, one order of magnitude smaller than the space charge tune shift , in both the horizontal and vertical planes. The ring tune variation for the different working points is shown in Fig. 3.24. The horizontal ring tune variation exhibits oscillatory behaviour about zero, which is expected since the horizontal tune is predominantly governed by the k index, which has been optimised. The vertical ring tune variation, however, increases with radius, arising from the failure of the spiral angles and $C1$ coefficients to scale correctly. It was therefore determined that further optimisation of the field clamps was required to reduce the vertical tune spread.

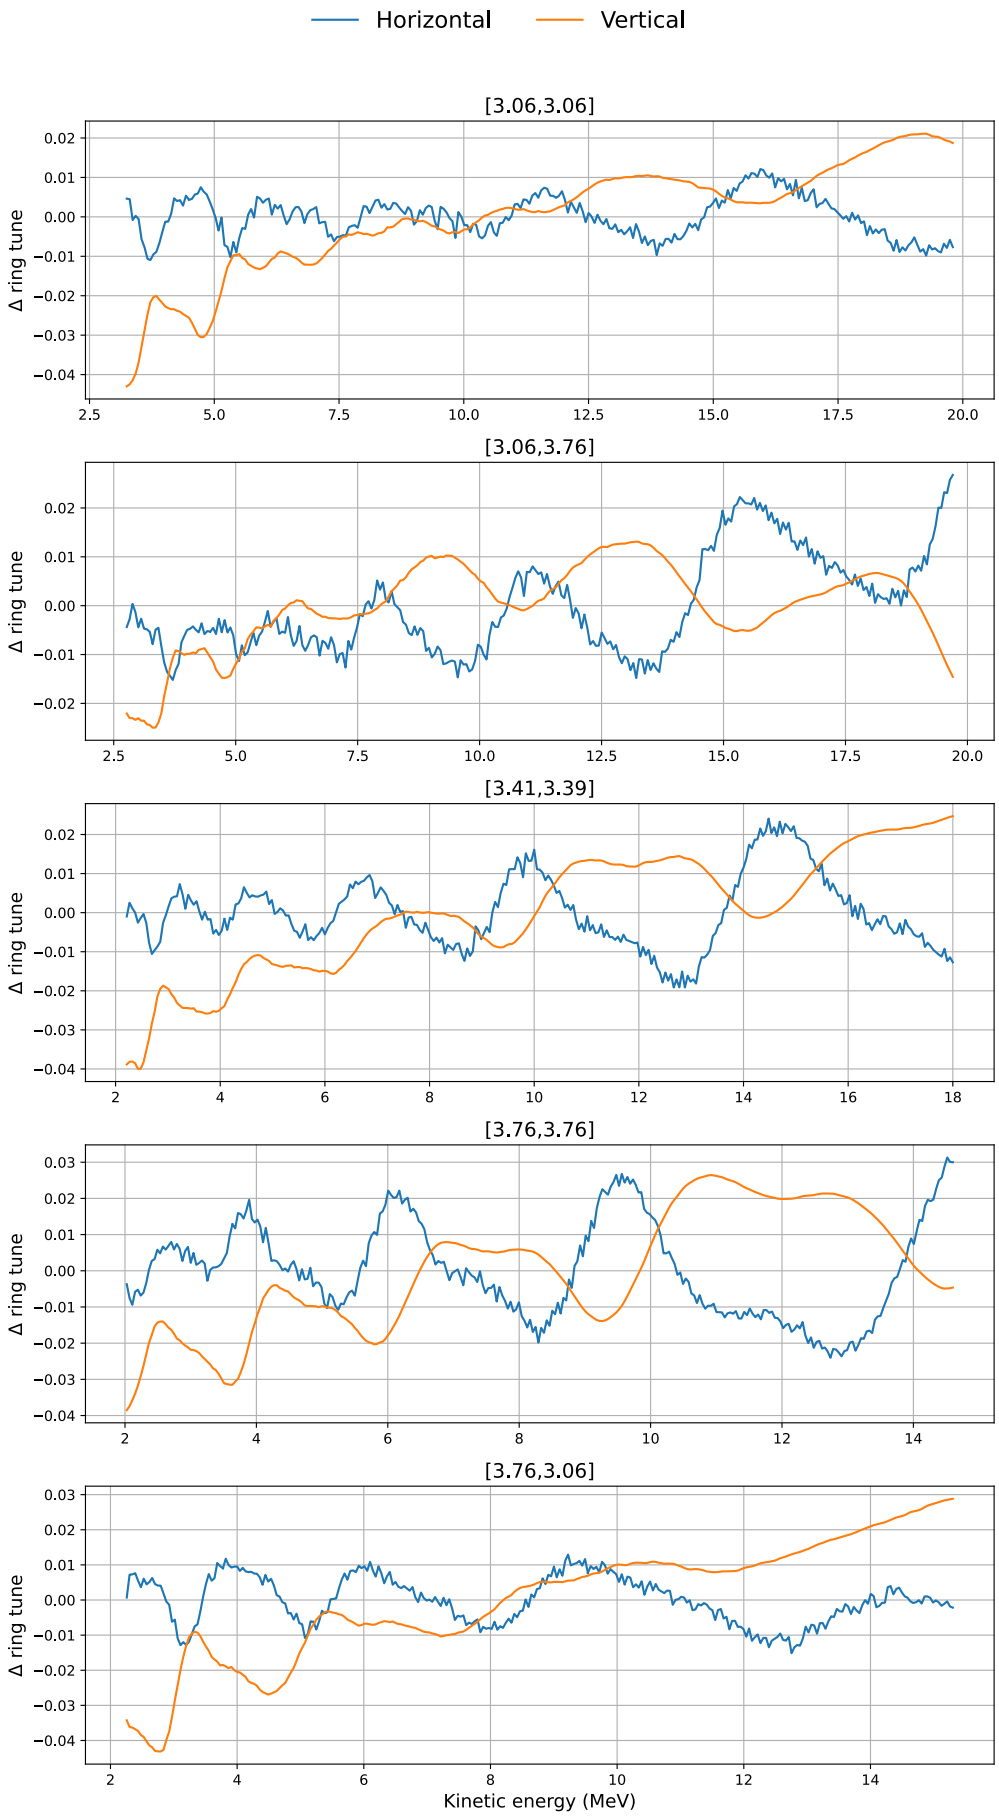


Figure 3.24: Variation of ring tune across all five working tune points.

3.7 Clamp optimisation

Given the similar fringe field behaviour across all working points, the clamp aperture can be optimised to reduce the overall tune spread simultaneously across all working points. As discussed in Section 2.9.4, only the F_{En} clamp aperture was shaped in the present work. The final aperture profile of the F_{En} clamp is shown in Fig. 3.25. The aperture at r=3540 mm was kept at 36 mm from the median plane, increasing linearly to 96 mm from the median plane at r=4300 mm.

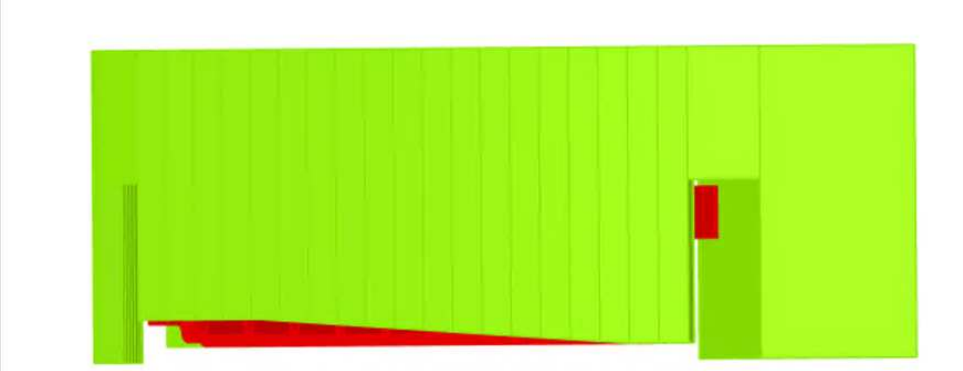


Figure 3.25: Side view of the doublet model in OPERA-3D, the centre of the machine is towards the right of the image. The aperture of the F_{En} clamp increases linearly with radius across the entire good field region.

Modifying the clamp aperture alters the integrated magnetic field and necessitates reoptimisation of the coil currents. The optimised current values can be found in the CDR report for FETS-FFA [65]. The reoptimised $k_{F,D}$ values are shown in Fig. 3.26. Throughout the remainder of this section, references to the optimised clamp aperture implicitly assume that the current settings have also been reoptimised accordingly.

The impact on the spiral angle is presented in Fig. 3.27. Since only the F_{En} clamp was adjusted, the primary effect is observed in $\zeta_{F_{En}}$. Whilst the overall structure remains unchanged, the values are slightly reduced relative to the original case and are now more consistent with ζ_{FD} and $\zeta_{D_{Ex}}$, indicating improved scaling behaviour. A corresponding improvement is also observed in $C1_{F_{En}}$, which becomes more uniform across all five working tune points, as shown in Fig. 3.28.

The resulting ring tunes after clamp aperture optimization are shown in Fig. 3.29. Whilst the horizontal tune variation remains largely unchanged, the vertical tune spread is reduced across all five working points, with the majority of the oscillatory error now centred about zero. Two minor exceptions are observed: at high energy, the [3.76,3.06] working point exhibits a slight increase in vertical tune, whilst the [3.06,3.76] working point shows a corresponding decrease. The overall tune spread across all five working points is reduced; however, variations in radial beam excursion, k index, and magnetic saturation between working points prevent a single clamp geometry from fully optimising each tune point independently.

The tune spread in both planes across all five working points remains slightly above the target of ± 0.01 in ring tune. Following clamp aperture optimisation, the errors are oscillatory about a constant value and are primarily determined by their amplitude, which is governed by the number

of trim coils within the respective good field region. Reducing this amplitude would require either smaller trim coils, allowing a greater number to be accommodated within the same radial range, or independent powering of the two pole face trim coil layers, increasing the number of adjustable current settings from 10 to 20 per pole.

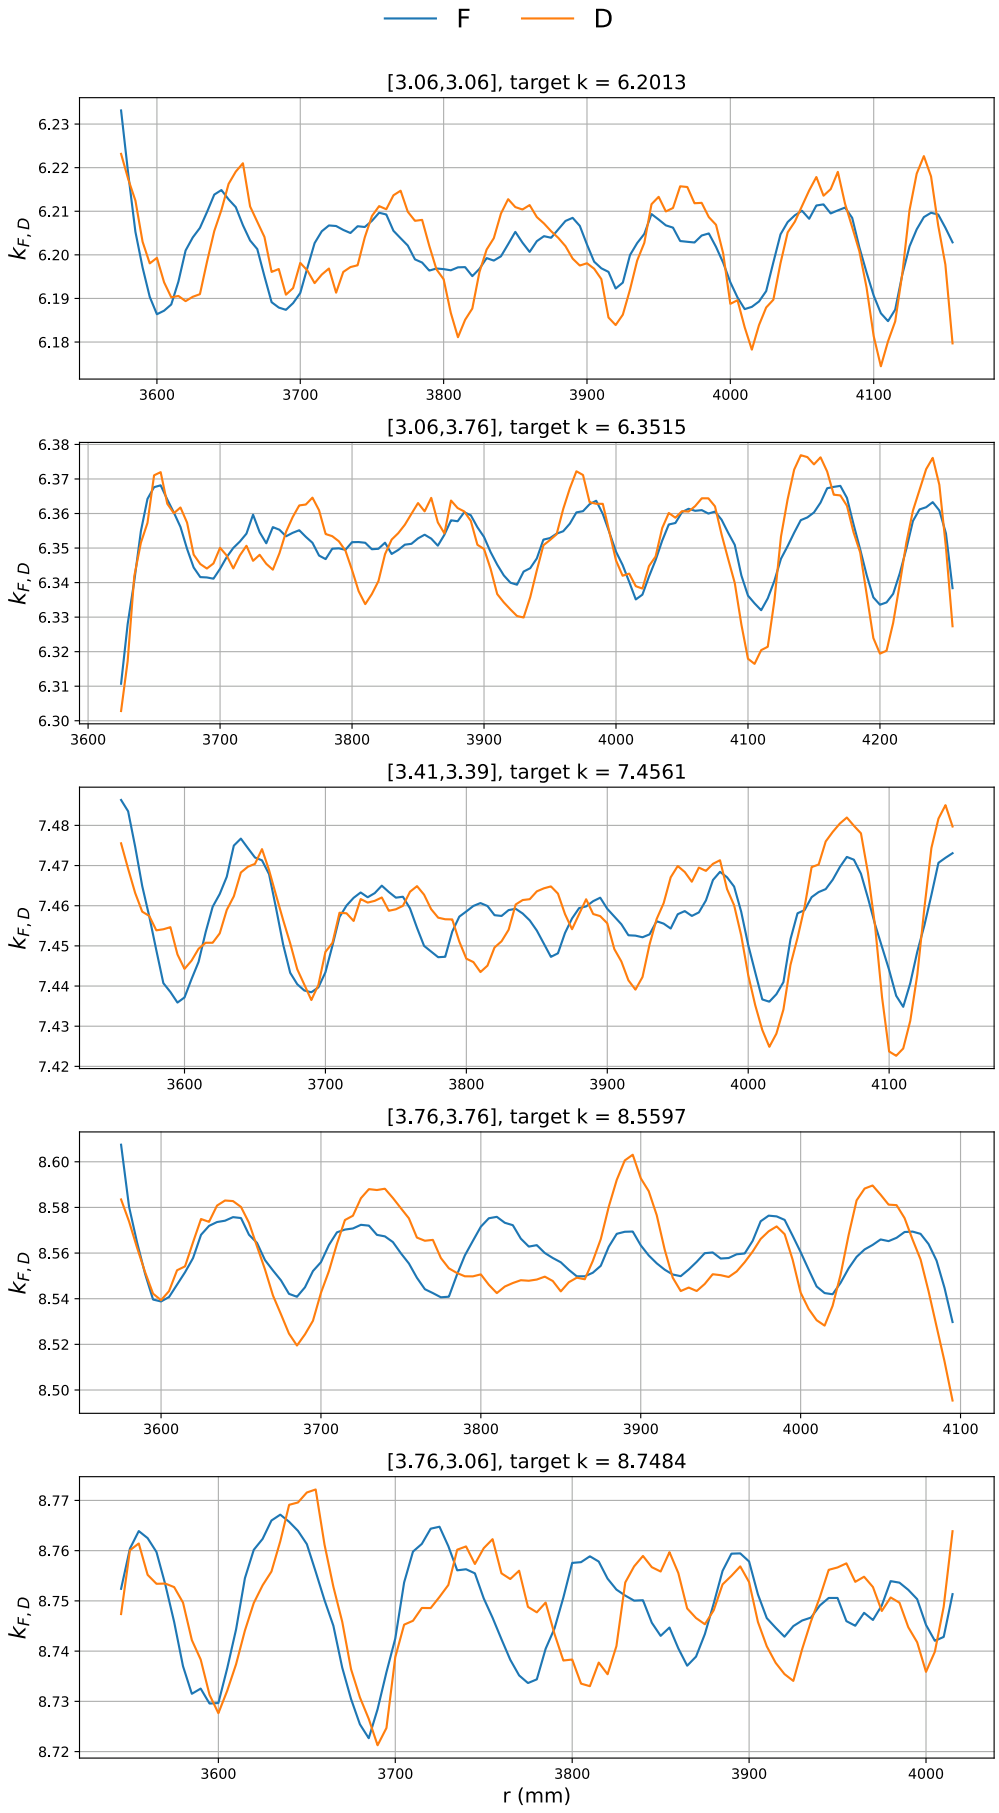


Figure 3.26: The field index k for both F and D magnets evaluated across the five working tune points following clamp aperture optimisation.

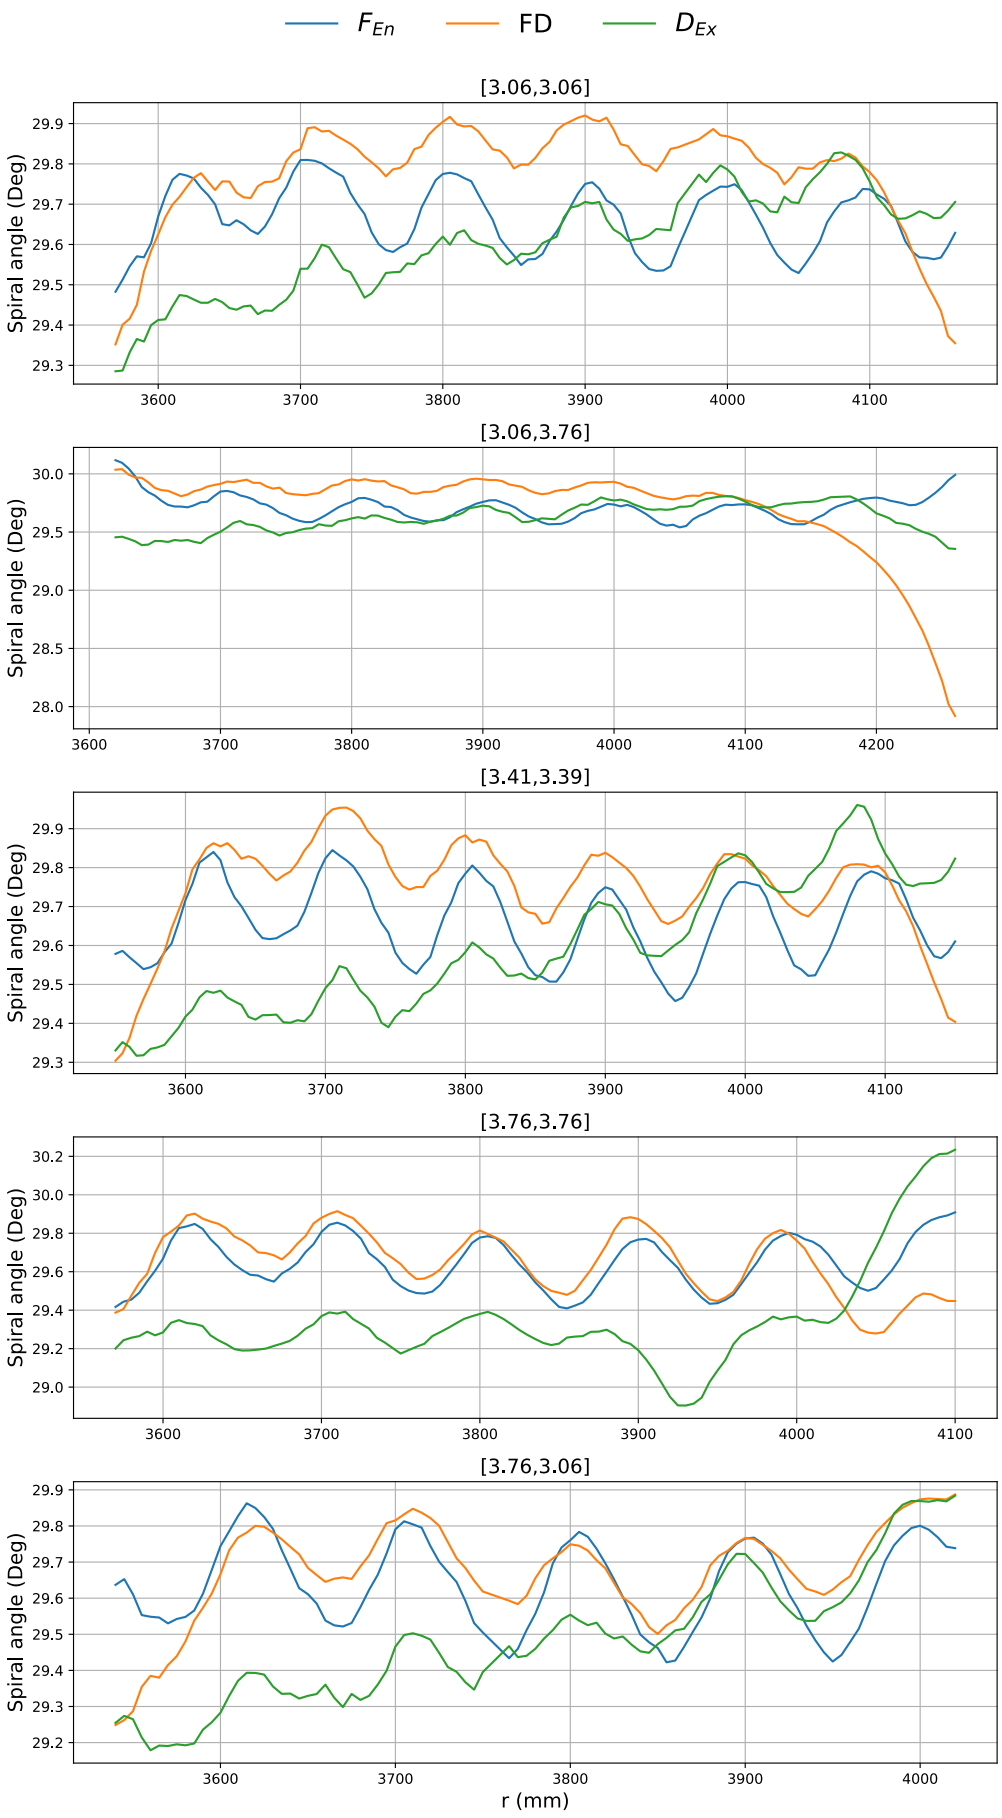


Figure 3.27: The effective spiral angles evaluated across all five working tune points following clamp aperture optimisation.

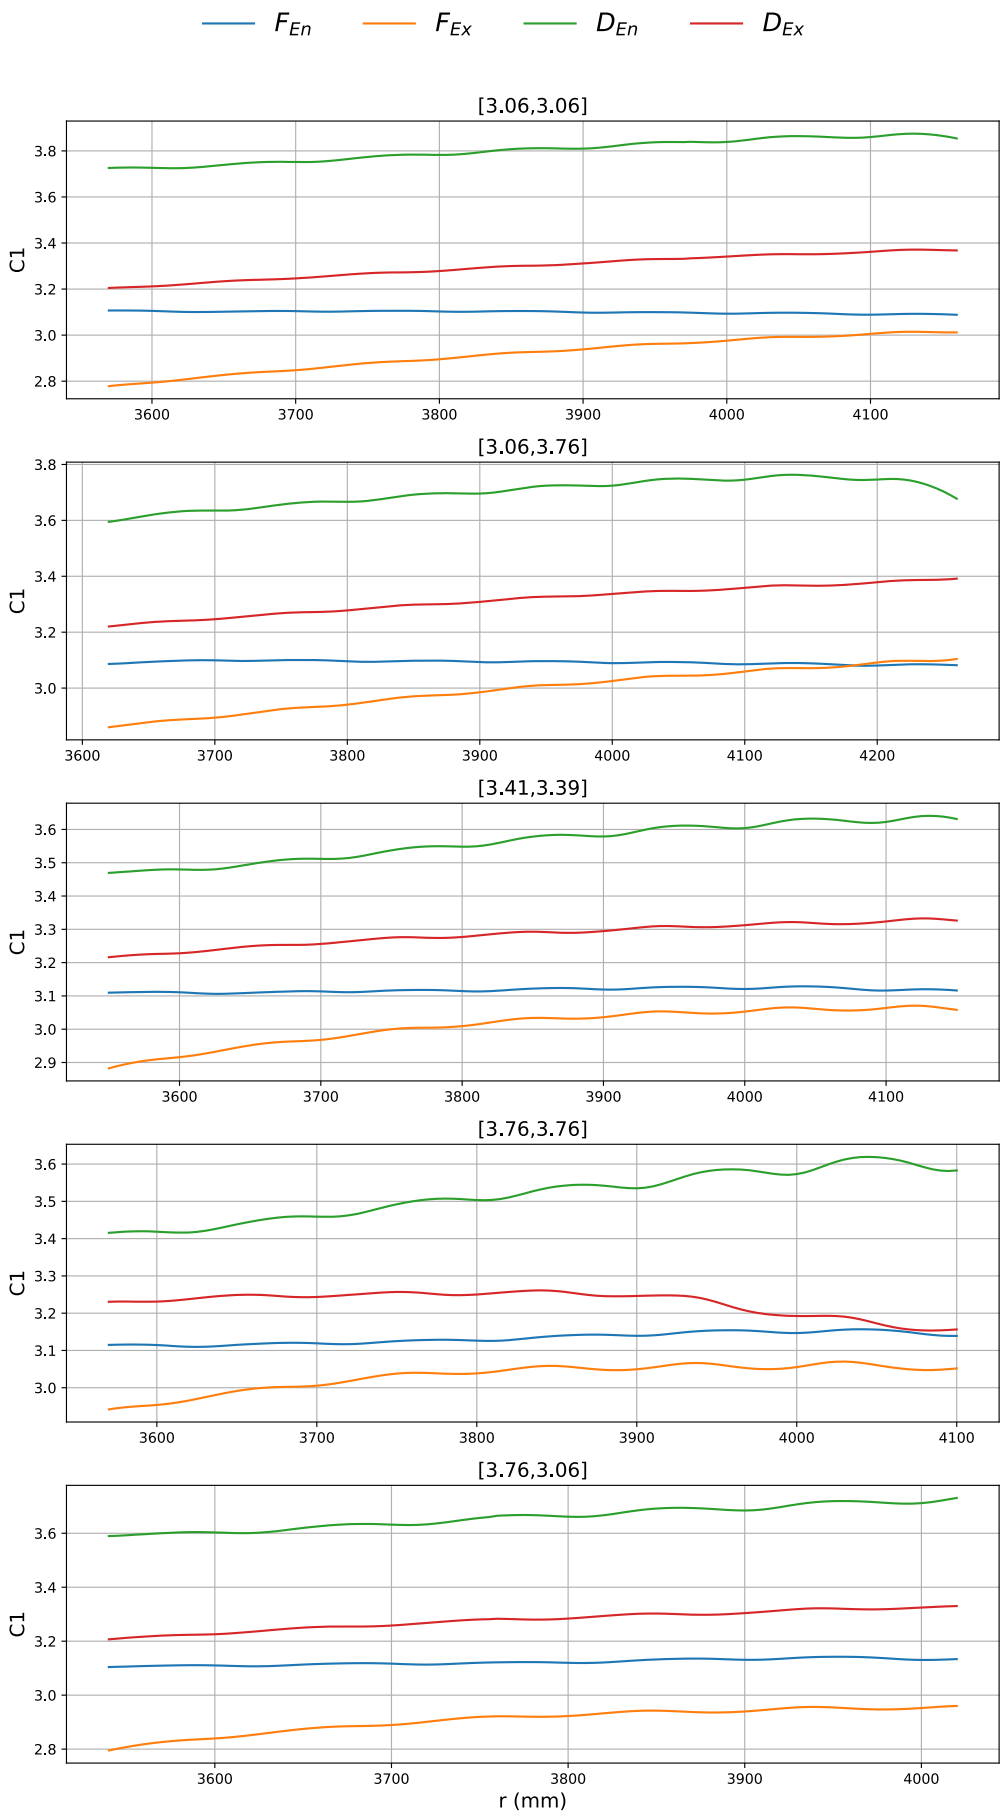


Figure 3.28: The $C1$ coefficients evaluated across all five working tune points following clamp aperture optimisation.

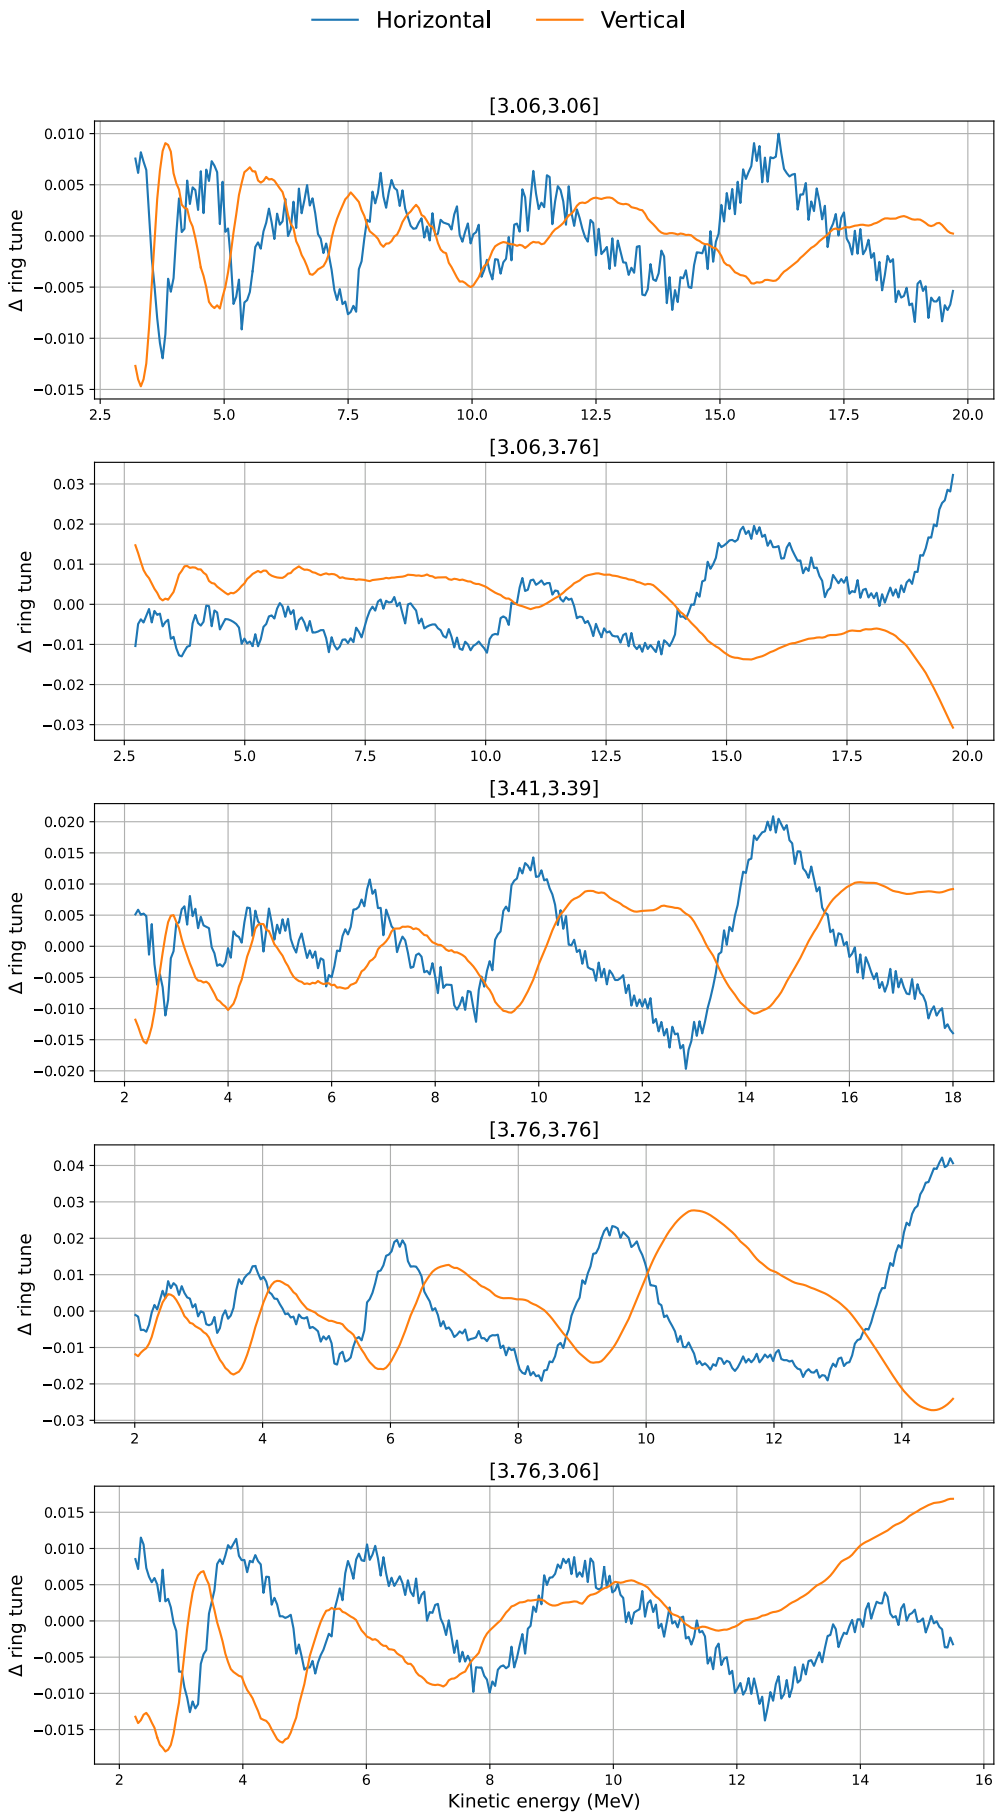


Figure 3.29: Variation of ring tune across all five working tune points following clamp aperture optimisation.

Chapter 4

LhARA

The methods desribed in Chap. 2 are applied by the author to the design of the LhARA FFA magnets, with the specific parameters for this application introduced below.

4.1 Radiotherapy and Bragg peak

Radiation therapy is a cornerstone of cancer treatment, employed in over half of all patients [67]. Whilst conventional radiotherapy typically employs MeV scale photon or electron beams, proton and ion beams offer significant advantages owing to the Bragg peak, which enables highly localised dose deposition within the tumour whilst minimising damage to surrounding healthy tissue.

The characteristics of the Bragg peak depend strongly on the particle species, mass, and energy. For a given initial energy, heavier ions exhibit a sharper and more pronounced Bragg peak, attributable to their reduced lateral scattering and higher linear energy transfer near the end of their range. As the beam energy increases, the penetration depth of the Bragg peak increases, enabling treatment of deeper seated tumours, whilst the peak itself broadens slightly due to range straggling. In addition, heavier ions produce a fragmentation tail beyond the Bragg peak arising from nuclear interactions, which modifies the distal dose distribution. These dependencies provide flexibility in tailoring dose profiles, but also introduce complexity in treatment planning and radiobiological response.

A pristine Bragg peak is intrinsically narrow, meaning that a monoenergetic proton or ion beam deposits the majority of its dose over a very limited depth. Whilst this is advantageous for targeting a specific point, most tumours extend over a finite volume in depth. Irradiation with a single Bragg peak would therefore result in underdosing of tumour regions lying outside the peak, whilst failing to deliver a sufficiently uniform dose across the target volume. To address this, a spread out Bragg peak is created by superimposing multiple Bragg peaks at different energies. Each energy corresponds to a different penetration depth, and their combination produces a relatively flat, uniform dose distribution across the full tumour thickness. This ensures that the entire target volume receives the prescribed therapeutic dose whilst preserving the key advantage of a sharp

distal fall off, which spares healthy tissue beyond the tumour.

Current proton beam therapy is generally delivered in daily fractions of approximately 2 Gy at dose rates below 5 Gy/min. However, emerging approaches such as ultra-high dose rate "FLASH" radiotherapy (>40 Gy/s) [68, 11] and microbeam irradiation have demonstrated promising therapeutic benefits [69], although the underlying radiobiological mechanisms remain incompletely understood. Achieving the dose rates required for FLASH therapy demands beam intensities significantly beyond those of conventional clinical accelerators, typically requiring beam currents on the order of several hundred nanoamperes or more at the treatment site, depending on the field size and delivery method.

A deeper understanding of how radiation damages cancer cells at the fundamental level is required, necessitating close collaboration between physicists and biologists. To advance this field, the next generation of radiotherapy research facilities must be capable of delivering diverse ion species with customisable beam timing, energy levels, and spatial patterns, including dose rates at or above the FLASH regime. Such capability will enable researchers to develop more effective cancer treatment protocols.

FFA accelerators are particularly well suited for this purpose. They preserve the structural integrity of the incoming beam whilst permitting energy adjustment without the need for physical degraders. FFAs can accelerate multiple ion species simultaneously whilst producing both high and variable doses at rapid cycling rates, potentially reaching 100 Hz or faster.

4.2 LhARA

The Laser-hybrid Accelerator for Radiobiological Applications (LhARA) [70, 71] is a proposed flexible accelerator facility dedicated to radiobiological research, designed to accommodate multiple ion species, with the present work focused on proton beam production.

As illustrated in Fig. 4.1, LhARA is a multi-stage facility employing a laser-driven source. The laser pulse initiating proton production may be triggered at repetition rates of up to 10 Hz, producing proton bunches with durations in the range of 10 ns to 40 ns, thereby enabling the time structure of the beam to be varied in order to probe the chemical and biological pathways governing the biological response to ionising radiation. The facility proposes to employ a 100 TW laser source, enabling the direct production of short bunches containing up to 10^9 protons at a nominal energy of up to 15 MeV $\pm 2\%$, thereby circumventing the space charge limitations inherent to lower energy injection sources. The resulting highly divergent beams are captured using Gabor plasma lenses [72], which provide focusing strengths comparable to high field solenoids at considerably reduced cost, yielding bunches of 1–3 cm in diameter.

Captured bunches may be delivered directly to a low energy *in vitro* end station, or alternatively transported through the injection line and injected into an FFA ring for further acceleration. The FFA provides both variable extraction energy and fast extraction, both of which are identified requirements of the radiobiology community. Protons are accelerated within the FFA to a maximum energy of 127 MeV, after which beams may be directed to one of two high energy end stations.

The repetition rate of the FFA is currently limited to 10 Hz by the laser source frequency; however, the lattice design is fully compatible with future operation at higher repetition rates of 100 Hz or above.

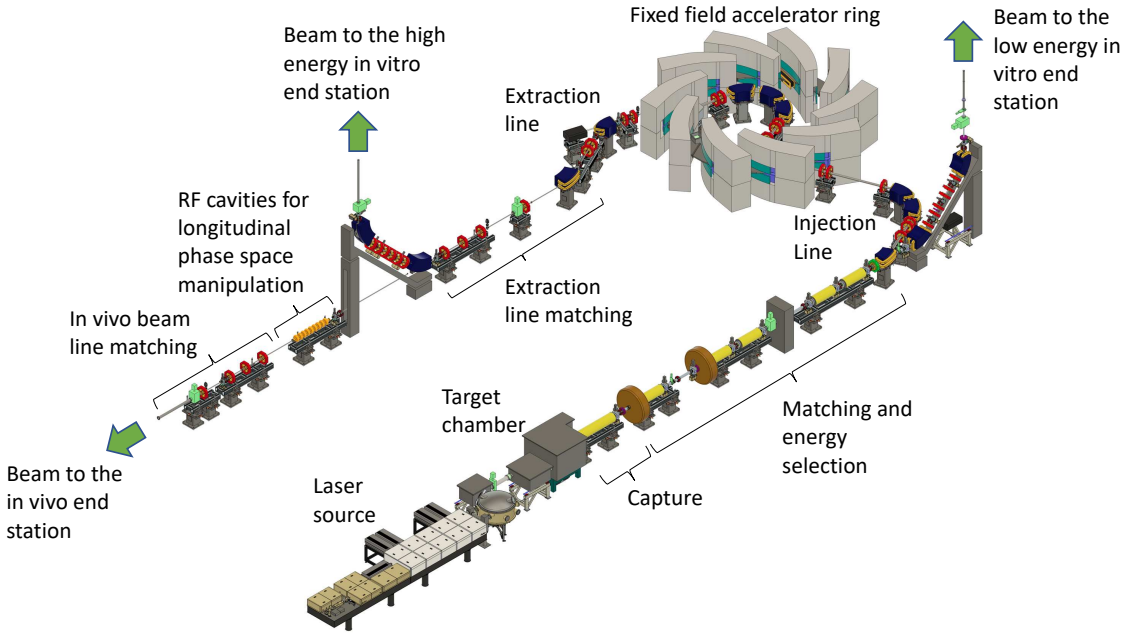


Figure 4.1: Schematic design of the LhARA facility. Image taken from [71].

4.3 LhARA FFA lattice design

The beam energy at extraction from the FFA is determined by the injection energy, which is varied by adjusting the focusing strengths along the Stage 1 beam line, from the capture section through to the FFA ring. The FFA subsequently accelerates the beam, acting effectively as a threefold momentum multiplier. This approach simplifies the injection and extraction systems, as their geometry and physical location remain fixed throughout operation. The purpose of the LhARA FFA is therefore to provide variable extraction energy at a fixed tune point, avoiding the crossing of potentially deleterious resonances that could otherwise lead to beam loss.

A singlet spiral magnet configuration was selected to minimise the size of the ring; since the tune point is held fixed, a doublet spiral configuration is not required. The principal lattice parameters of the LhARA FFA are summarised in Table 4.1. The ring comprises ten identical cells, a choice representing a compromise between limiting the orbit excursion, which governs the radial extent of the magnets, and providing sufficient straight section length to accommodate the injection and extraction systems. The working tune point was selected at approximately [2.79,1.17], to achieve this the required magnet parameters comprise of a field index of k = 5.23 and a spiral angle of 53.9°. The dynamic acceptance has been evaluated in both the horizontal and vertical planes and is found to be substantially larger than the beam emittance [70]; this remains the case even under the conservative assumption that the emittance is ten times its nominal value, confirming that a robust and stable operating point has been identified.

Table 4.1: Specifications of the LhARA FFA from lattice design

Parameter	Value
Cell type	Singlet spiral
Number of cells in the ring	10
Spiral angle (deg)	53.9
k index	5.23
Max extraction energy (MeV)	127
Target horizontal ring tune	2.79
Target vertical ring tune	1.17

4.4 LhARA FFA magnet design

The main paramters of the LhARA FFA magnet are summarised in Table 4.2. The pole opening angle was set to 12.24° to ensure that the peak magnetic field remains within approximately 1.4 T. The radial excursion of the beam was determined to be 0.56 m, accounting for the horizontal beam acceptance and closed orbit distortion, a good field region spanning 2880 mm to 3560 mm was established; within this region, the integrated magnetic field must be optimised to follow the relevant scaling law.

Table 4.2: Parameters of the magnet

Parameter	Value
Pole opening angle (deg)	12.24
Opening angle between pole and clamp (deg)	2.89
Gap height of iron pole (mm)	$\pm$48
Min/Max radius of the trim coils (m)	2.797 / 3.676
Min/Max radius of the poles (m)	2.77 / 3.74
Min/Max radius of the GFR (m)	2.88 / 3.56
Number of trim coils per magnet	18
Number of power supplies per magnet	9

The required field gradient is produced through a combination of a main coil and radially distributed trim coils. The use of trim coils is motivated by the expectation that the magnet will operate at varying levels of magnetic saturation across different extraction energies; by adjusting the individual coil current settings, the same field index k, and hence the same working tune point, can be maintained across the full operational energy range. Informed by the FETS-FFA magnet design described in Section 3.4, trim coils of the same dimensions, 92.5 mm $\times$ 10 mm, were adopted.

A beam stay clear region of 17 mm was specified. Accounting for a 10 mm vacuum chamber wall, two layers of 10 mm trim coils, and a clearance of 0.5 mm between the vacuum chamber and the coil assembly and between the coil assembly and the iron pole, a pole height of $\pm$48 mm was determined.

The number of trim coils was chosen to be 9, with 1 outside either radial end of the good field region and 7 within ($\approx 680\,mm/92.5\,mm$). An overlap of 0.5 was introduced between the two pole

face trim coil layers, which are operated with identical current settings to minimise the number of required power supplies. This yields a total of 18 trim coils per magnet, covering a radial range from 2797 mm to 3676 mm.

As shown in Fig. 4.2, all trim coils are routed to return towards the outer radius of the pole in order to maximise power efficiency. To reduce the longitudinal thickness, the coils are arranged in two vertical return layers. Using Eq. 2.2, the opening angle between the pole and the clamp was calculated to be 2.89°. The resulting trim coil configuration is illustrated in Fig. 4.3. The main coil was then sized to be as large as possible within this constraint, with a width of 75 mm (Fig. 4.4) and a height of 117 mm.

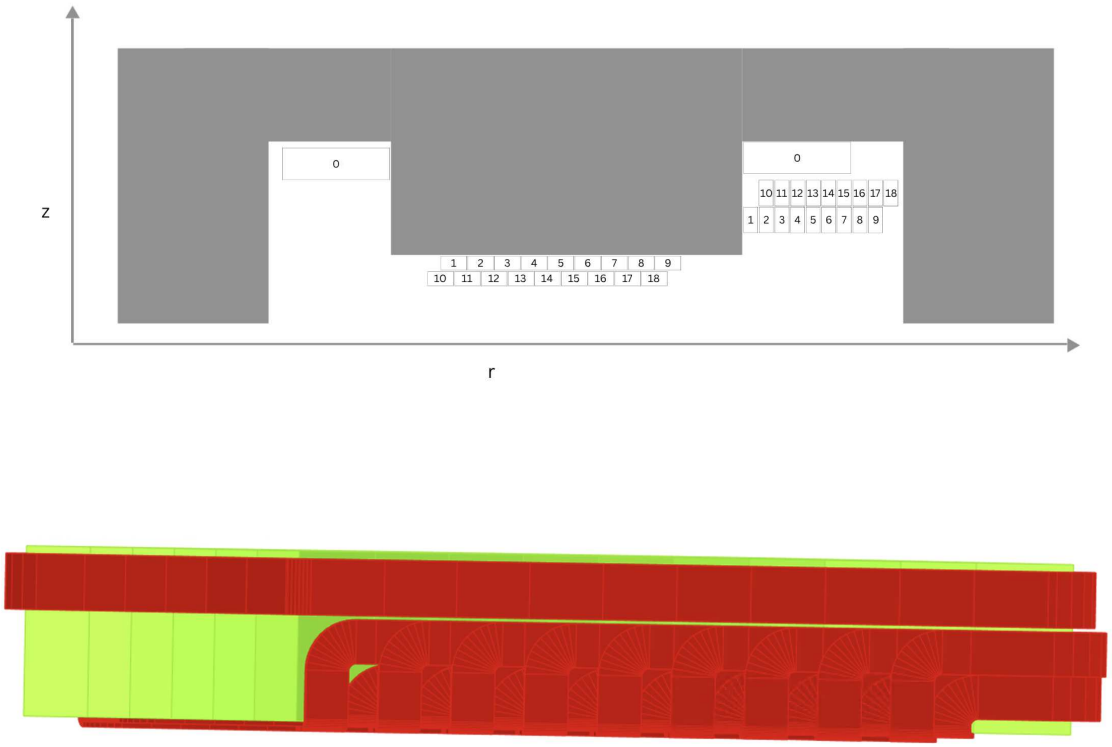


Figure 4.2: Cross sectional view of the coils wrapped around the magnet pole in the LhARA FFA. (Top) Schematic diagram of the cross section of the LhARA FFA magnet, with the main coil labelled 0 and trim coils labelled 1–18. (Bottom) Corresponding OPERA-3D model showing the pole and coils without the return yokes.

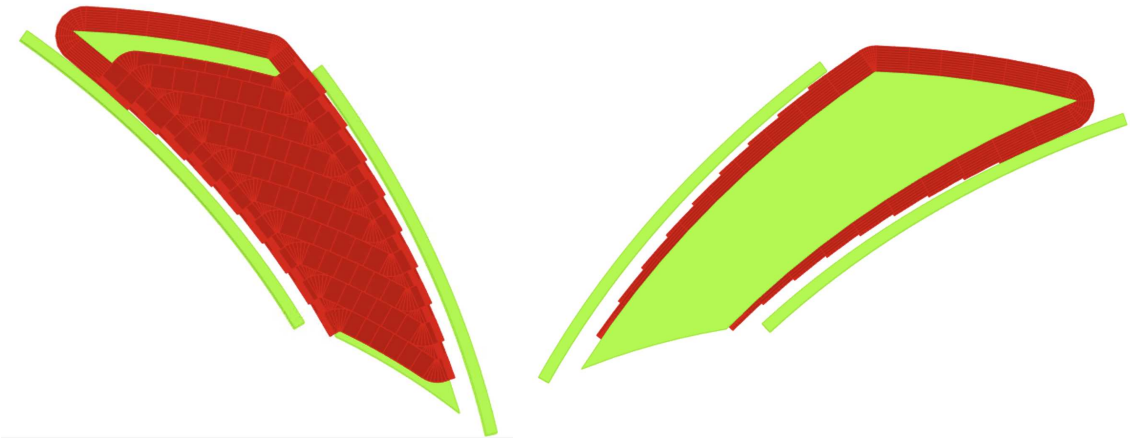


Figure 4.3: (Left) Pole face view, showing the trim coils crossing the pole face at different radii, wrapped tightly adjacent to one another. (Right) Top view, showing that as the trim coils stack longitudinally, the available space between pole and clamps decreases with radius,necessitating an opening angle of 2.89 °.

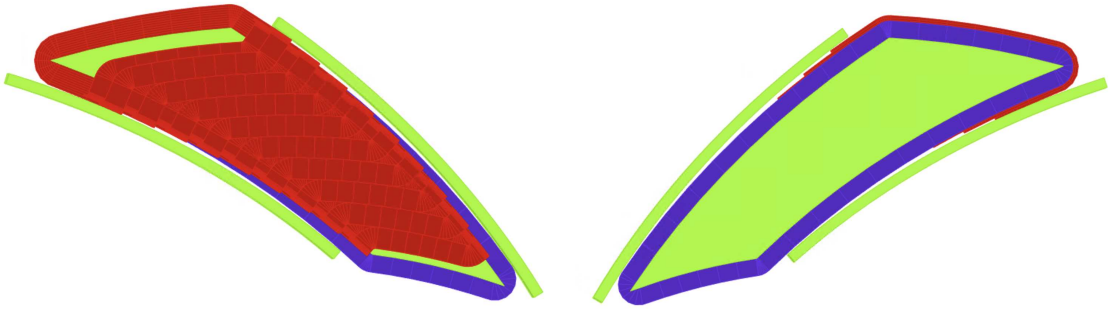


Figure 4.4: Since the main coil (purple) thickness is constant, the minimum separation between adjacent main coils occurs at the smallest radius, in contrast to the behaviour observed for the trim coils. (Left) Pole face view of all coils. (Right) Top view of all coils.

Two iron configurations were investigated using identical current settings: (1) a configuration in which the clamps and pole are connected to a separate yoke (Fig. 4.5), and (2) a configuration in which the clamps and pole share a common yoke (Fig. 4.6). The results were compared against a reference configuration without field clamps (Fig. 4.7). The clamps have a thickness of 35 mm and an aperture of 52 mm. The resulting field distributions at the centre of the good field region for all three configurations are shown in Fig. 4.8. The inclusion of clamps produces a more rapid fringe field fall off, reducing field leakage into the straight sections whilst simultaneously increasing the magnetic field amplitude within the magnet body. The configuration with a shared common yoke was therefore selected for the LhARA FFA magnet design. The geometry of the return yoke was further modified to better align it with the direction of the pole, freeing additional space for other elements in the straight sections. The final LhARA FFA magnet configuration is shown in Fig. 4.9.

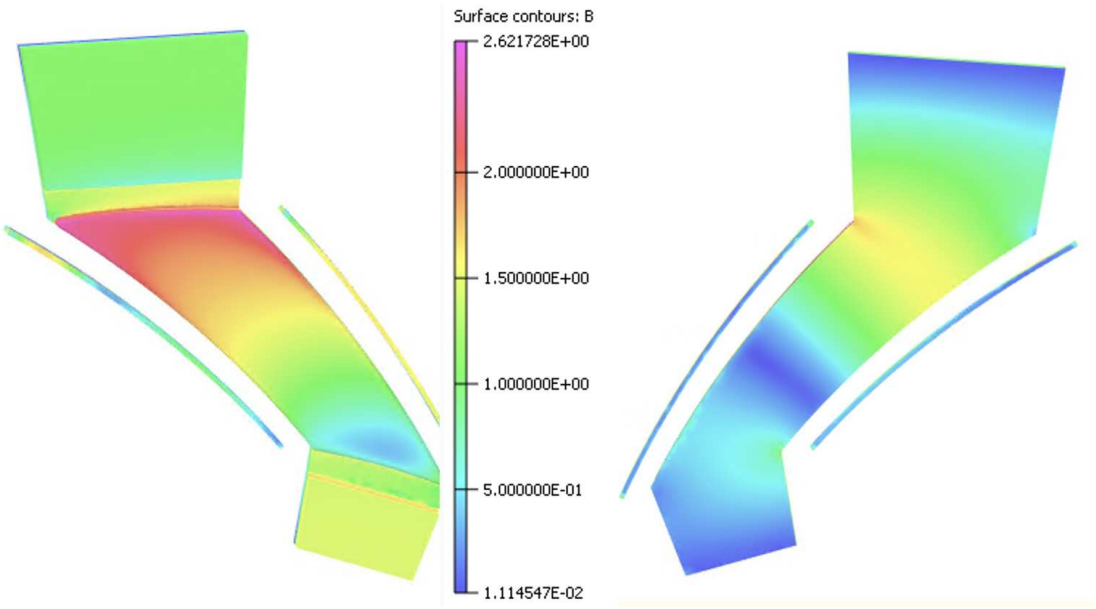


Figure 4.5: Field contours of the LhARA magnet in the configuration where the clamps and pole are connected to separate yokes, shown without coils for visual clarity. The left panel presents the pole face view and the right panel shows the top view; surface contours are given in units of Tesla.

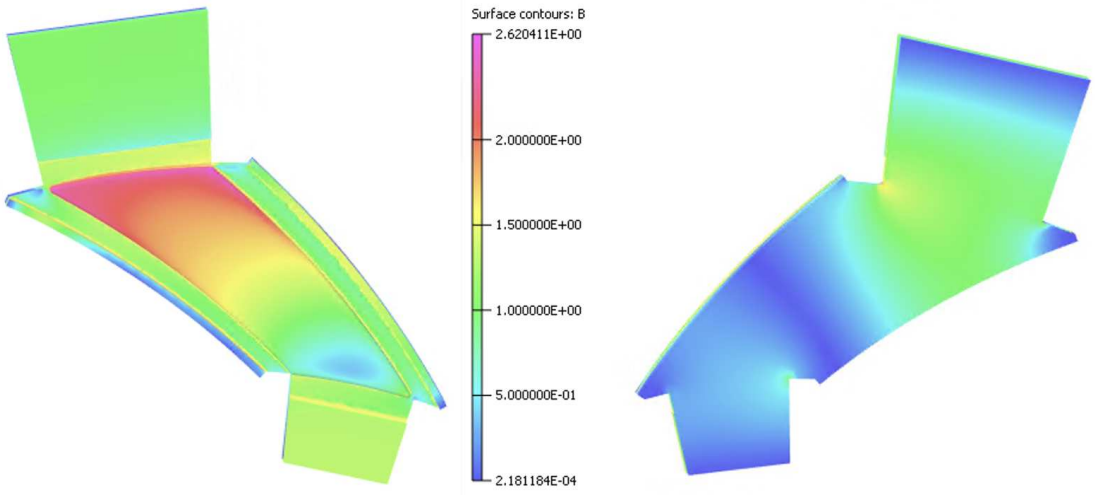


Figure 4.6: Field contours of the LhARA magnet in the configuration where the pole and clamps share a common yoke, shown without coils for visual clarity. The left panel presents the pole face view and the right panel shows the top view; surface contours are given in units of Tesla.

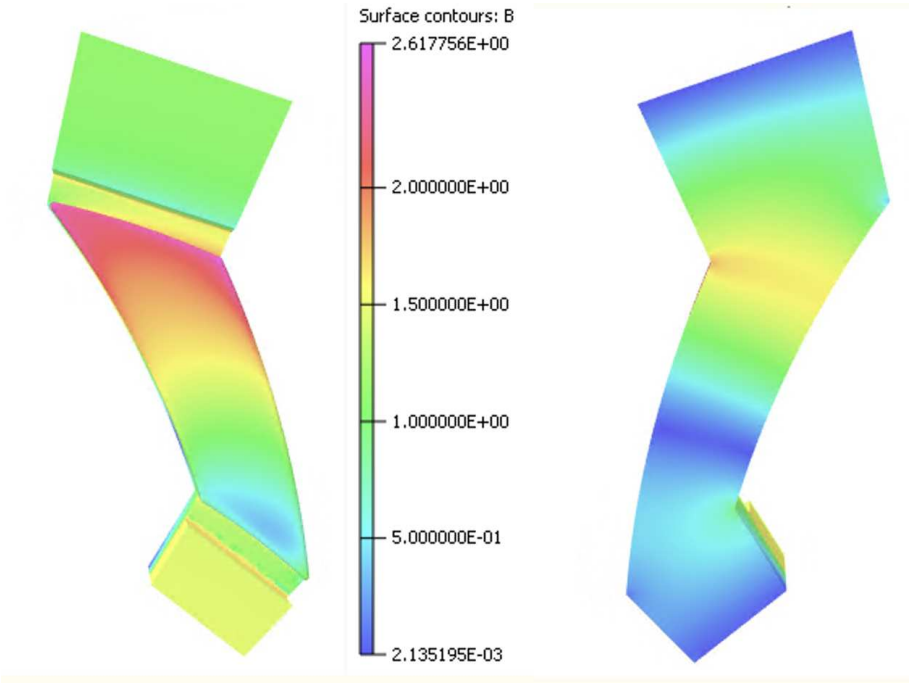


Figure 4.7: Field contours of the LhARA magnet in the configuration without field clamps, shown without coils for visual clarity. The left panel presents the pole face view and the right panel shows the top view; surface contours are given in units of Tesla.

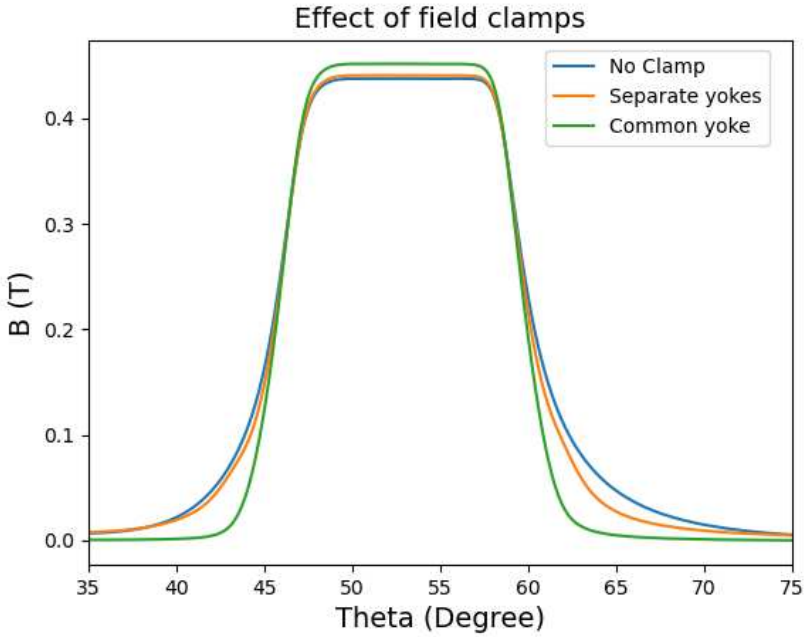
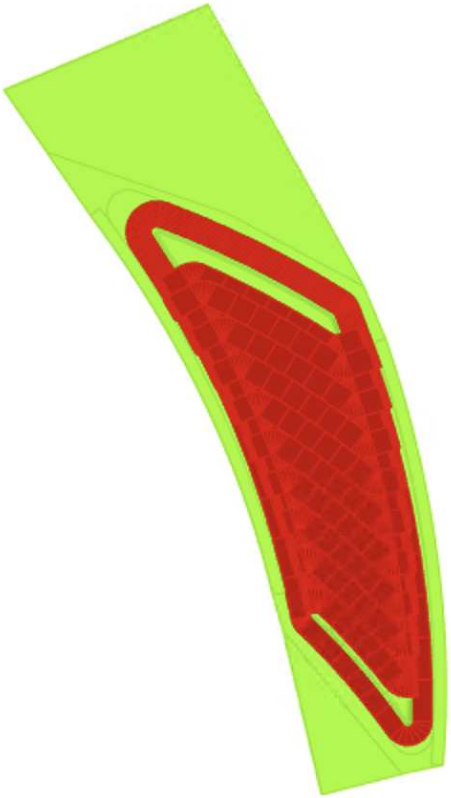


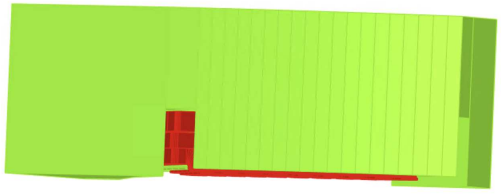
Figure 4.8: Azimuthal field profile at a fixed radius within the good field region for the three configurations. Despite identical current settings, the magnetic field exhibits distinct behaviour depending on the configuration.



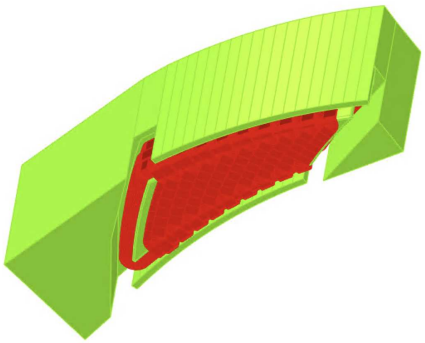
(a) Top view



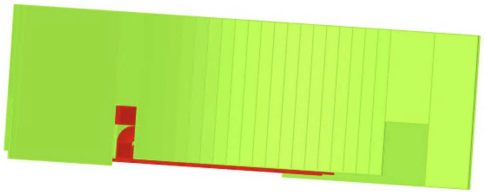
(b) Pole face view



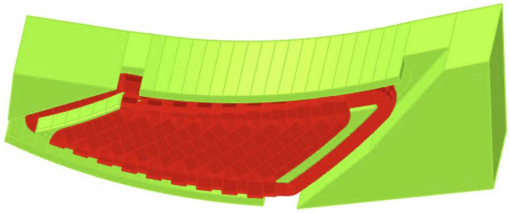
(c) View of the entrance clamp



(d) Second view of the entrance clamp



(e) View of the exit clamp



(f) Second view of the exit clamp

Figure 4.9: OPERA-3D model of the final LhARA FFA magnet design. The upper half of the magnet is shown, with the coils (red) wound around the pole, which shares a common magnetic return yoke with all field clamps.

4.5 Initial optimisation

4.5.1 Current optimisation

The optimisation procedure follows that introduced in Section 2.9, with the integrated magnetic field of Eq. 2.3 serving as the target function. For the maximum extraction energy of 127 MeV, an injection energy of 15 MeV is required, with corresponding parameters $r_0 = 3.477$ m $BL_0 = 0.96$ Tm. A lower energy range was also investigated, in which the integrated field was reduced to $\frac{2}{3}$ of the maximum value, corresponding to BL_0= 0.64 Tm.

The Jacobians evaluated at these two energies are shown in Fig. 4.10. The two pole face layers of trim coils are operated with identical current settings, such that coils 1 and 10 share the same current, and so forth. When computing the Jacobian, the current in each corresponding pair of trim coils is varied by the same amount. In the Jacobian indexing, each pair is represented by the lower coil number.

The field quality following current optimisation is shown in Fig. 4.11; the error of the integrated magnetic field relative to the scaling law is within $\pm 0.015\%$, whilst the error in k lies within ± 0.04. The observed oscillations in the error arise from the discretisation of the trim coils with finite size, with the error amplitude similar between the maximum energy case (127 MeV) and the lower energy case (57 MeV).

The fringe field parameters are shown in Fig.4.12 and Fig.4.13. The spiral angles oscillate about a broadly constant value, with the exception of ζ_{Ex} at the highest energy, which decreases to below 50°. For a perfectly scaling field, the $C1$ coefficients should remain constant; however, most exhibit an increasing trend with radius. The exception is $C1_{Ex}$ at the highest energy, which initially increases before subsequently decreasing.

The differing behaviour of the fringe field parameters across extraction energies is attributable to the varying levels of magnetic saturation within the magnet, as illustrated in Fig. 4.14. At the highest extraction energy, the magnet enters the saturation regime, causing ζ_{Ex} and $C1_{Ex}$ to decrease in value.

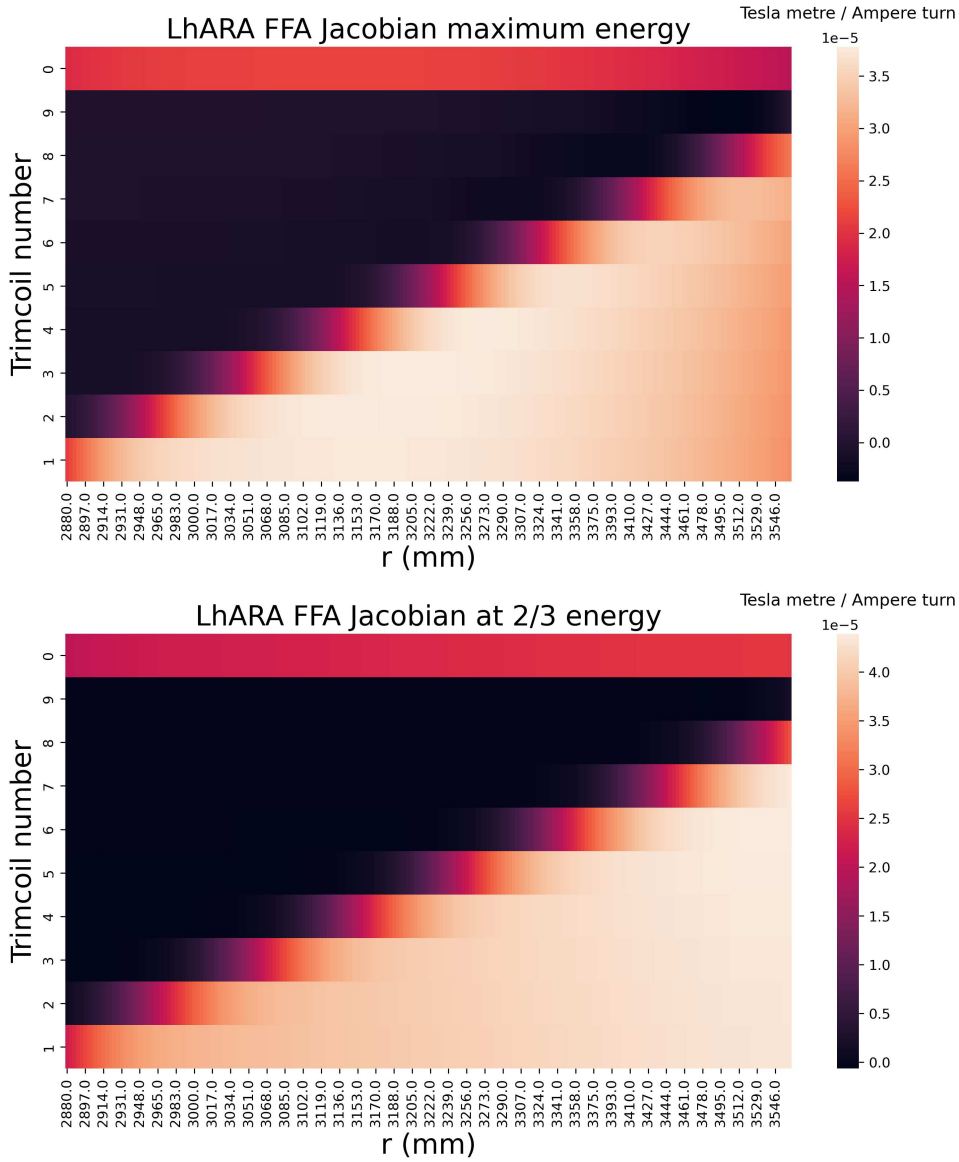


Figure 4.10: The response matrix illustrating how the integrated magnetic field varies with changes in the coil currents at (Top) maximum energy and (Bottom) two thirds of the maximum energy. The two Jacobian matrices are very similar in structure; however, the maximum energy Jacobian exhibits a smaller amplitude, indicating that the magnet is beginning to enter the magnetic saturation regime.

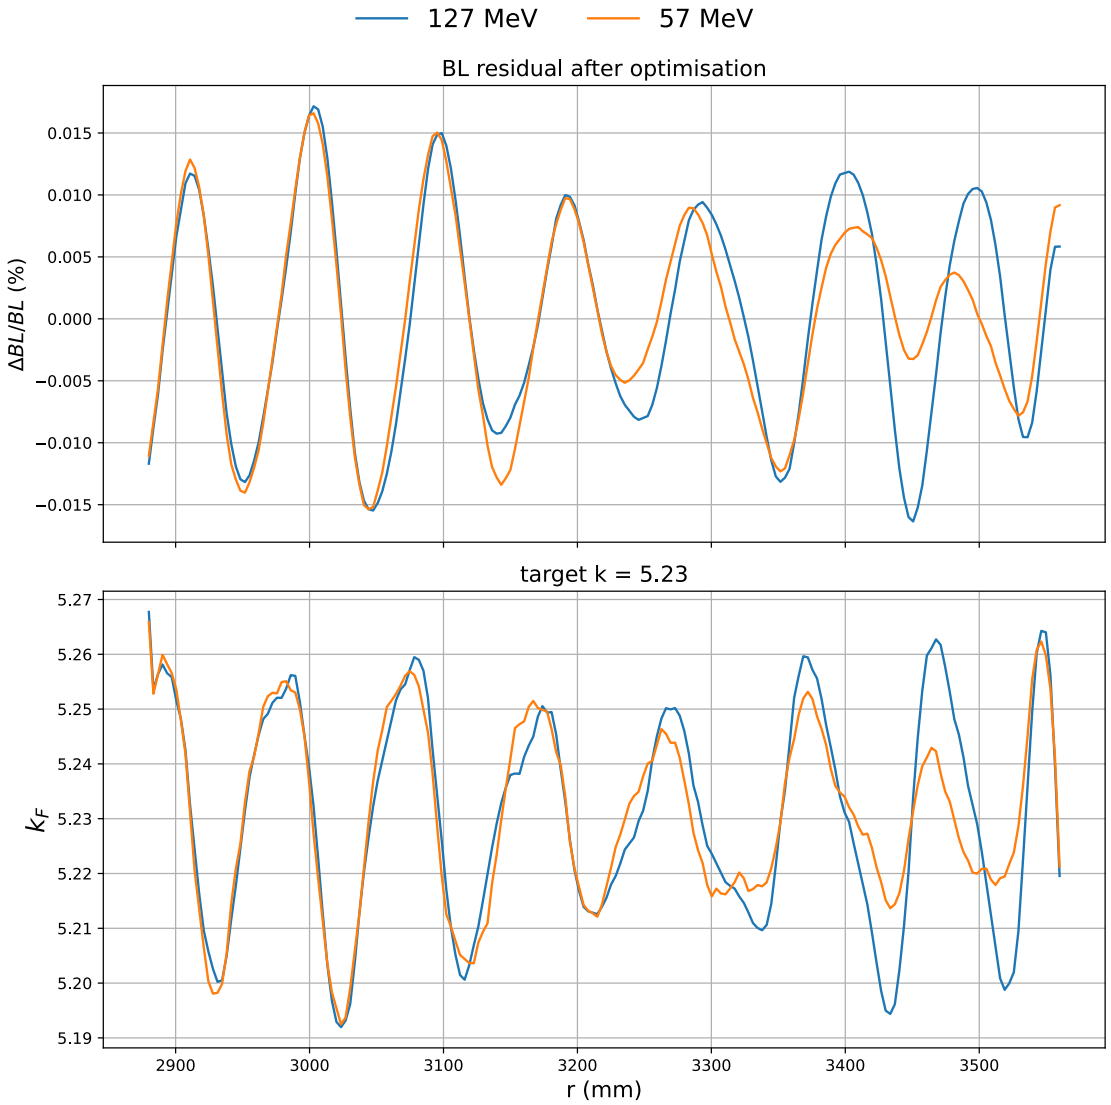


Figure 4.11: Field quality for different extraction energies. (Top) Deviation of the integrated field BL in the model from the target scaling law. (Bottom) Field index k_F calculated from the magnet model.

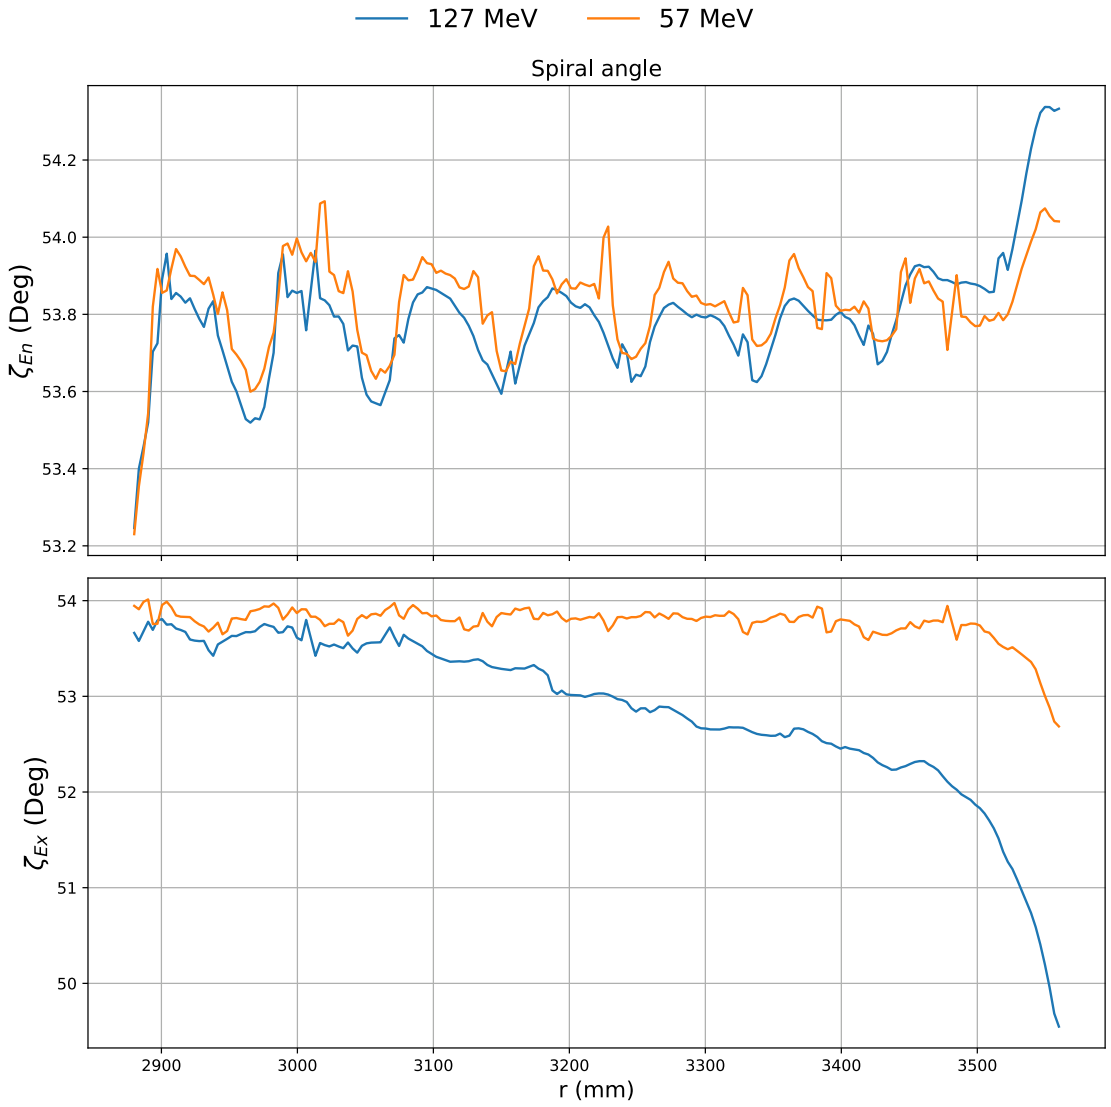


Figure 4.12: Spiral angles of the magnetic field at (Top) the entrance and (Bottom) the exit for different energy ranges.

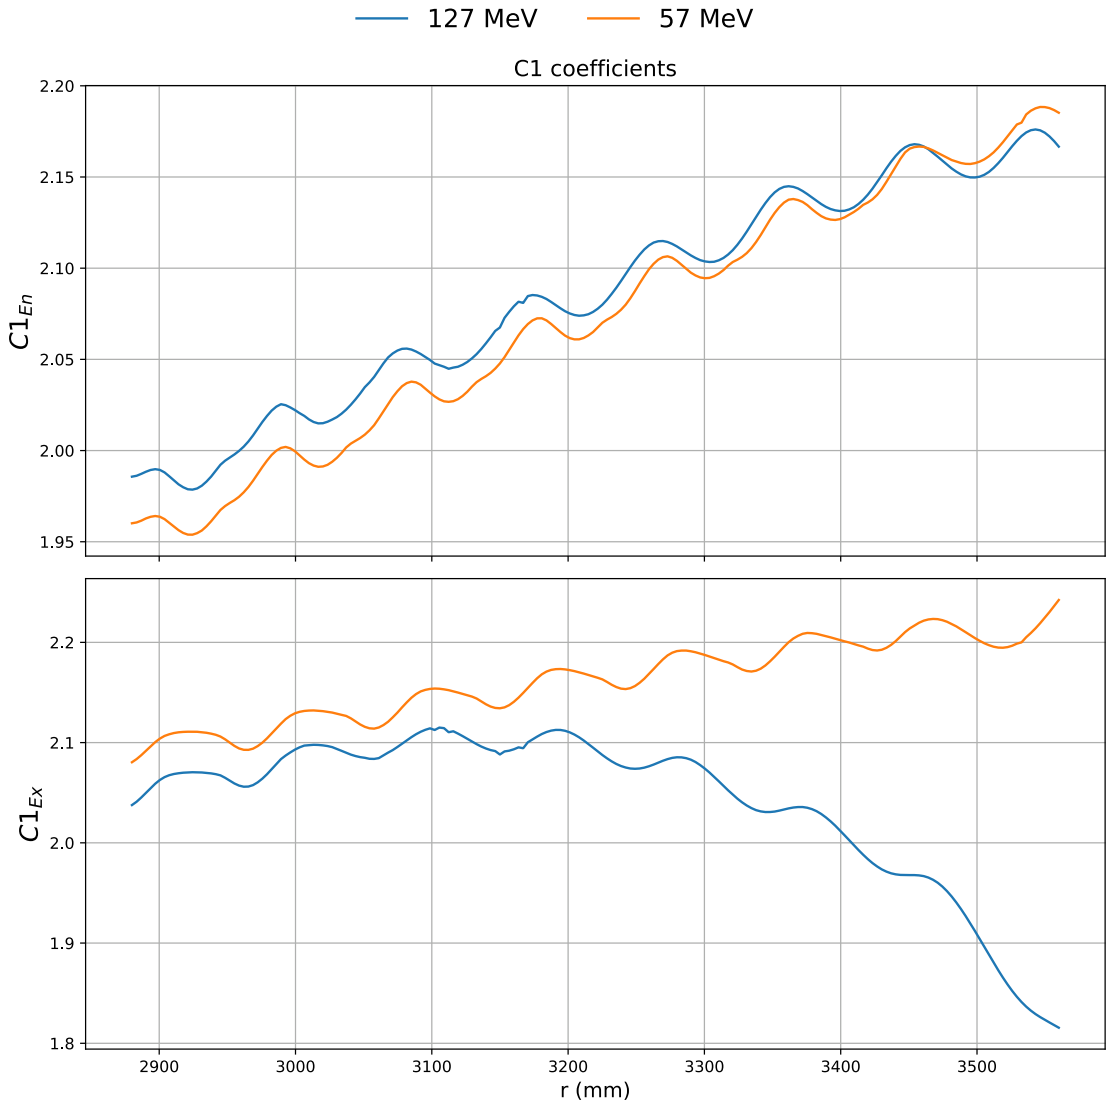


Figure 4.13: $C1$ coefficients of the magnetic field at (Top) the entrance and (Bottom) the exit for different energy ranges.

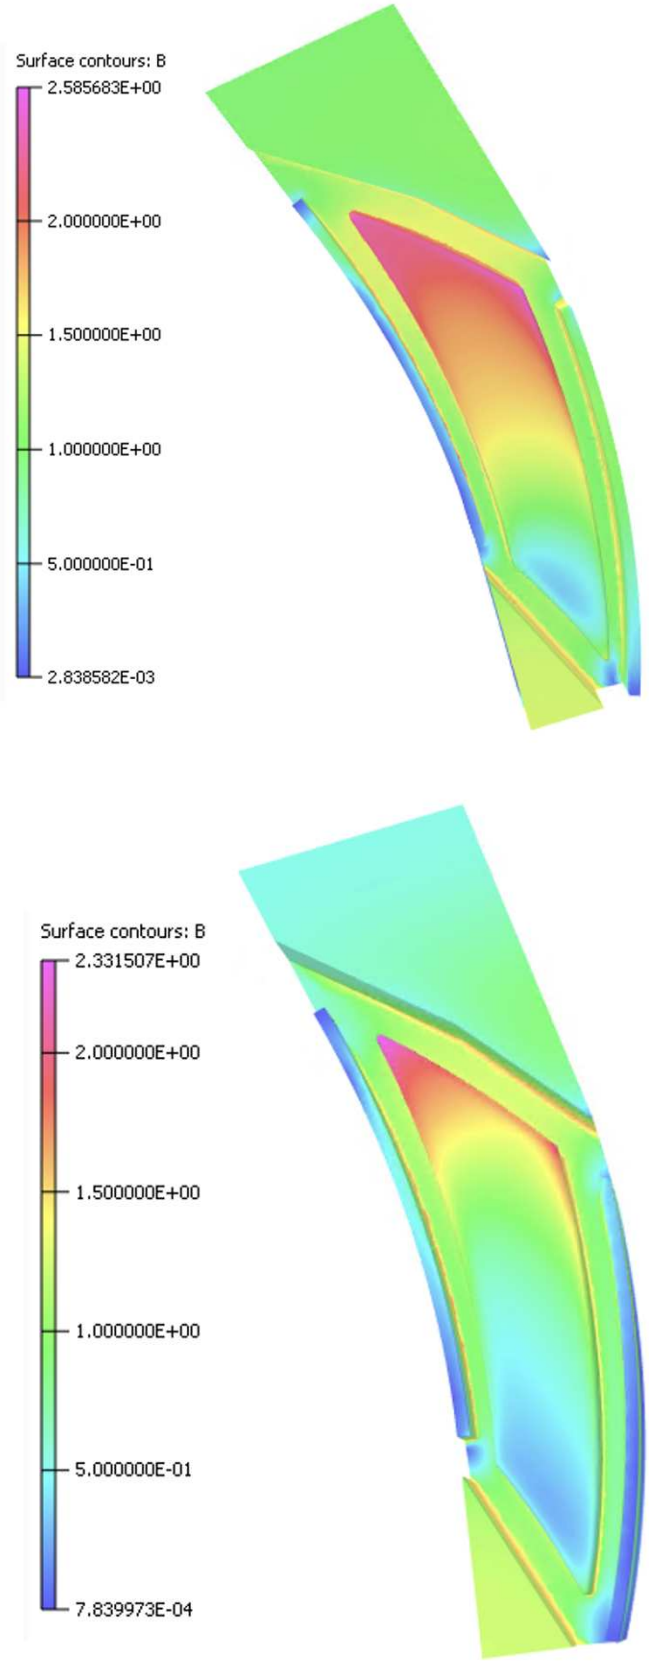


Figure 4.14: Field contours of the LhARA FFA magnet for (Top)127 MeV extraction energy and (Bottom) 57 MeV extraction energy

4.5.2 Tracking

Three dimensional field maps were extracted from the OPERA-3D model and used as input for single particle tracking in FixField, with the results presented in Fig. 4.15. Whilst the horizontal ring tune remains approximately constant throughout acceleration, a gradual upward trend is observed in the vertical ring tune. This behaviour is attributed to the non constant $C1$ field coefficientsl; the clamp aperture is therefore further optimised in the subsequent section. Nevertheless, the tune variation across the two extraction energy ranges is plotted on a resonance diagram in Fig. 4.16, where it can be seen that the tune footprints overlap and remain clear of all resonances up to fourth order. The tune spread in the horizontal plane is greater at the highest extraction energy, whilst in the vertical plane it is larger for the lower extraction energy case.

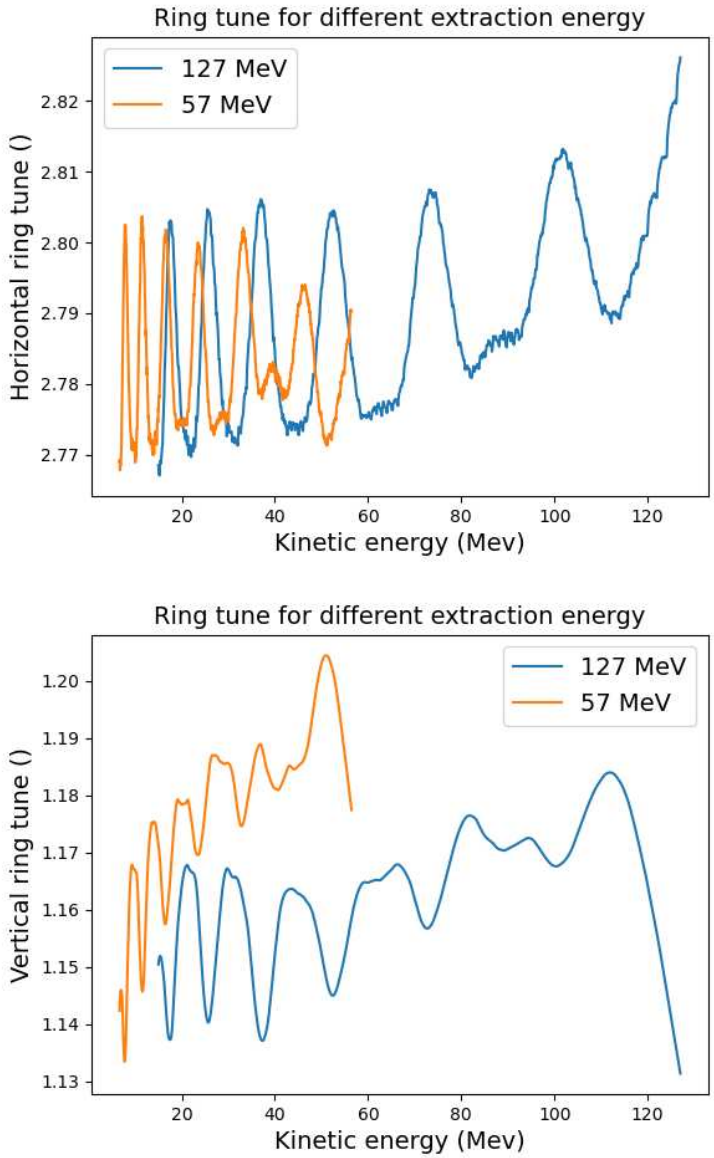


Figure 4.15: The variation of (Top) horizontal and (Bottom) vertical ring tune for different extraction energies.

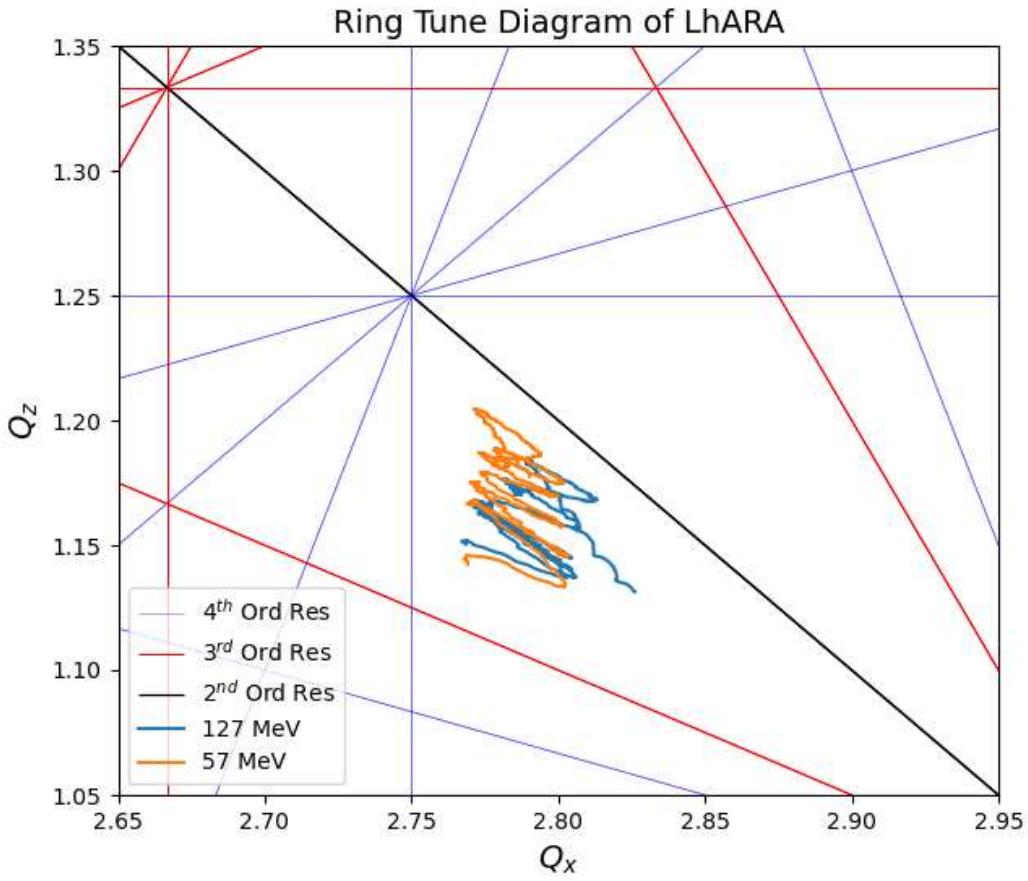


Figure 4.16: Ring tune footprints for two different extraction energy ranges plotted on a resonance diagram including resonance lines up to fourth order.

4.6 Clamp optimisation

As discussed in Section 2.9.4, only the entrance clamp aperture was shaped in the present work; the final aperture profile is shown in Fig. 4.17. The clamp aperture was set to 50 mm at the smallest radius and 52 mm at the largest radius, with additional cuts introduced to control the vertical tune.

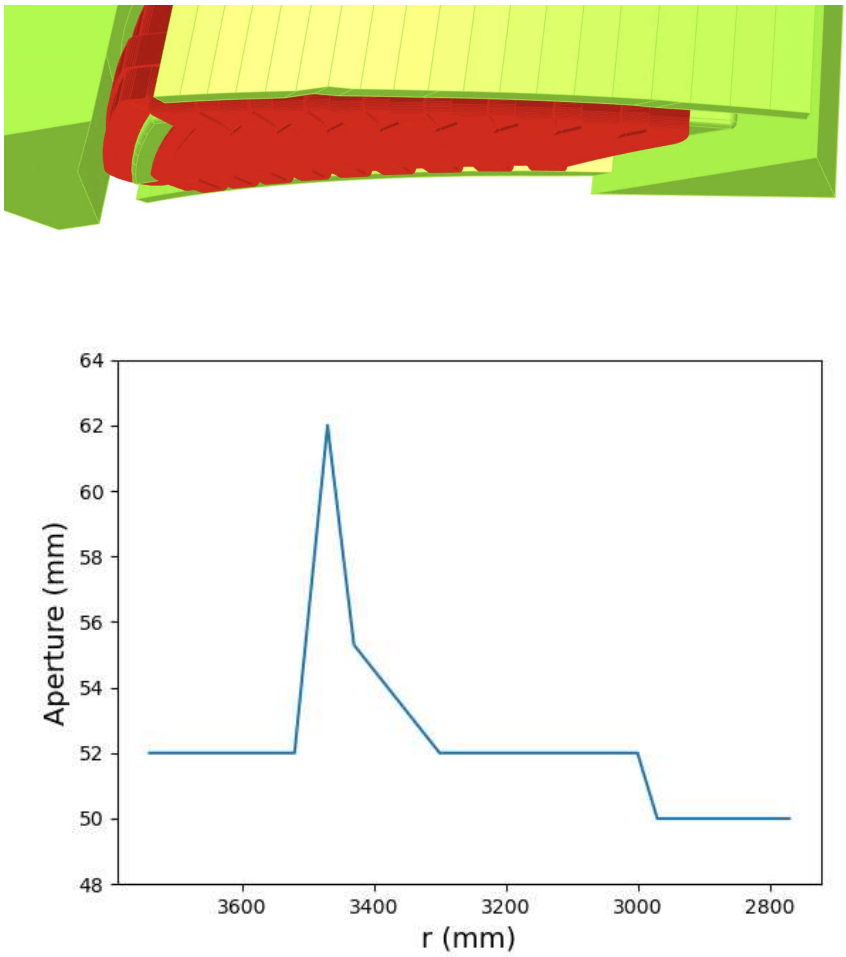


Figure 4.17: (Top) Side view of the OPERA-3D model showing the aperture profile of the entrance clamp; the machine centre is located towards the right of the image. (Bottom) Plot showing the variation of the entrance clamp aperture with radius. The plot is not to scale in order to better highlight the aperture changes.

Modifying the clamp aperture affects the integrated magnetic field, necessitating reoptimisation of the coil currents. The final current settings can be found in the CDR report of LhARA [73]. The reoptimised field qualities are shown in Fig. 4.18 and are comparable to those obtained prior to clamp aperture shaping. Throughout the remainder of this section, references to the optimised clamp aperture implicitly assume that the current settings have also been reoptimised accordingly.

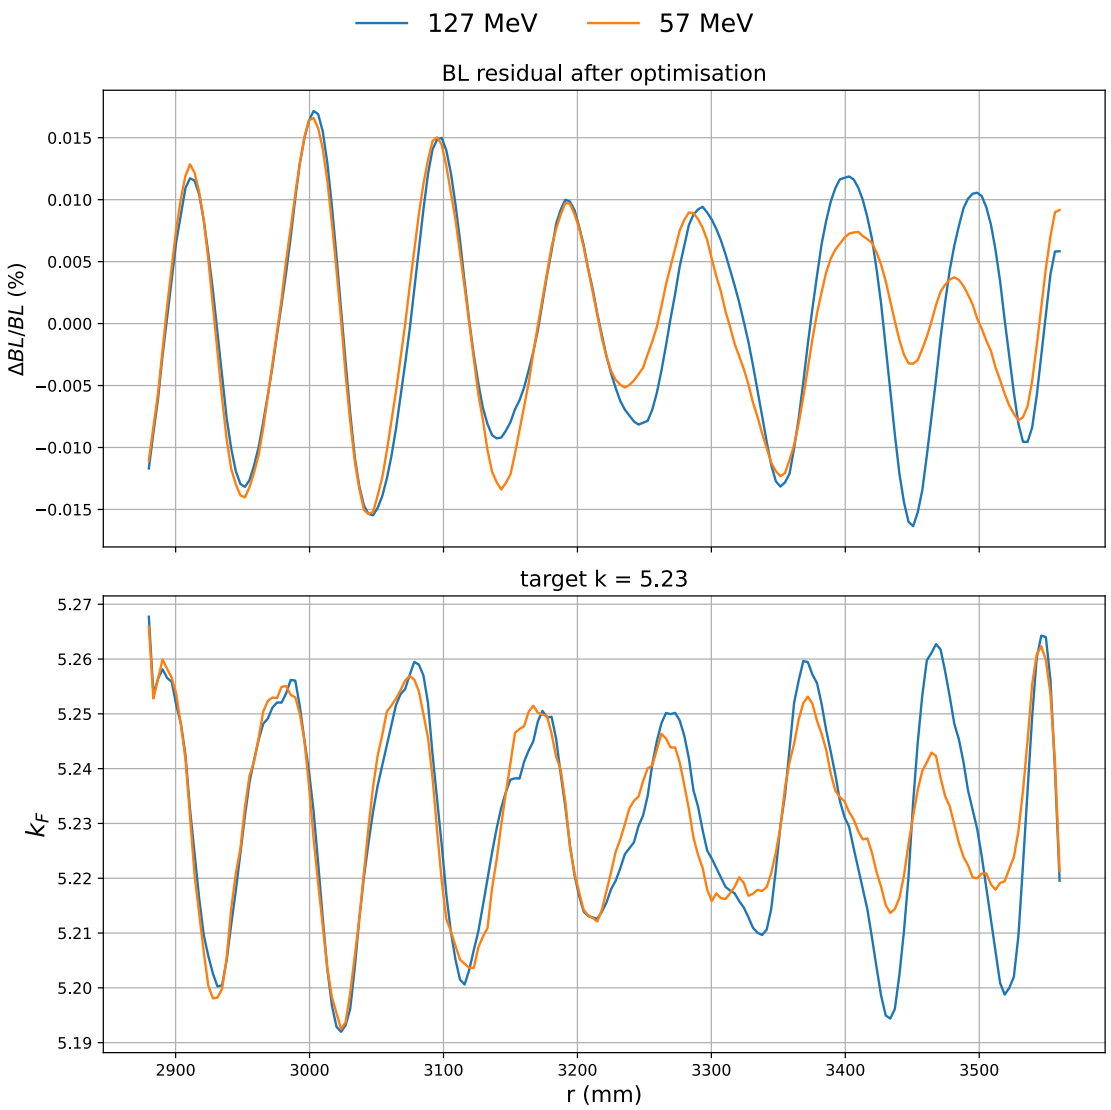


Figure 4.18: Field quality for different extraction energies following clamp aperture optimisation. (Top) Deviation of the integrated field BL in the model from the target scaling law. (Bottom) Field index k_F calculated from the magnet model.

The fringe field parameters before and after clamp optimisation are compared for the highest extraction energy in Fig. 4.19 and Fig. 4.21, and for the lower extraction energy in Fig. 4.20 and Fig. 4.22.

The change in fringe field parameters is modest compared to that observed in the FETS-FFA project (Section 3.7); nevertheless, it has a significant impact on the vertical tune. This indicates that the fringe field parameters exert a considerably stronger influence on the vertical tune in the singlet spiral design than in the doublet spiral configuration.

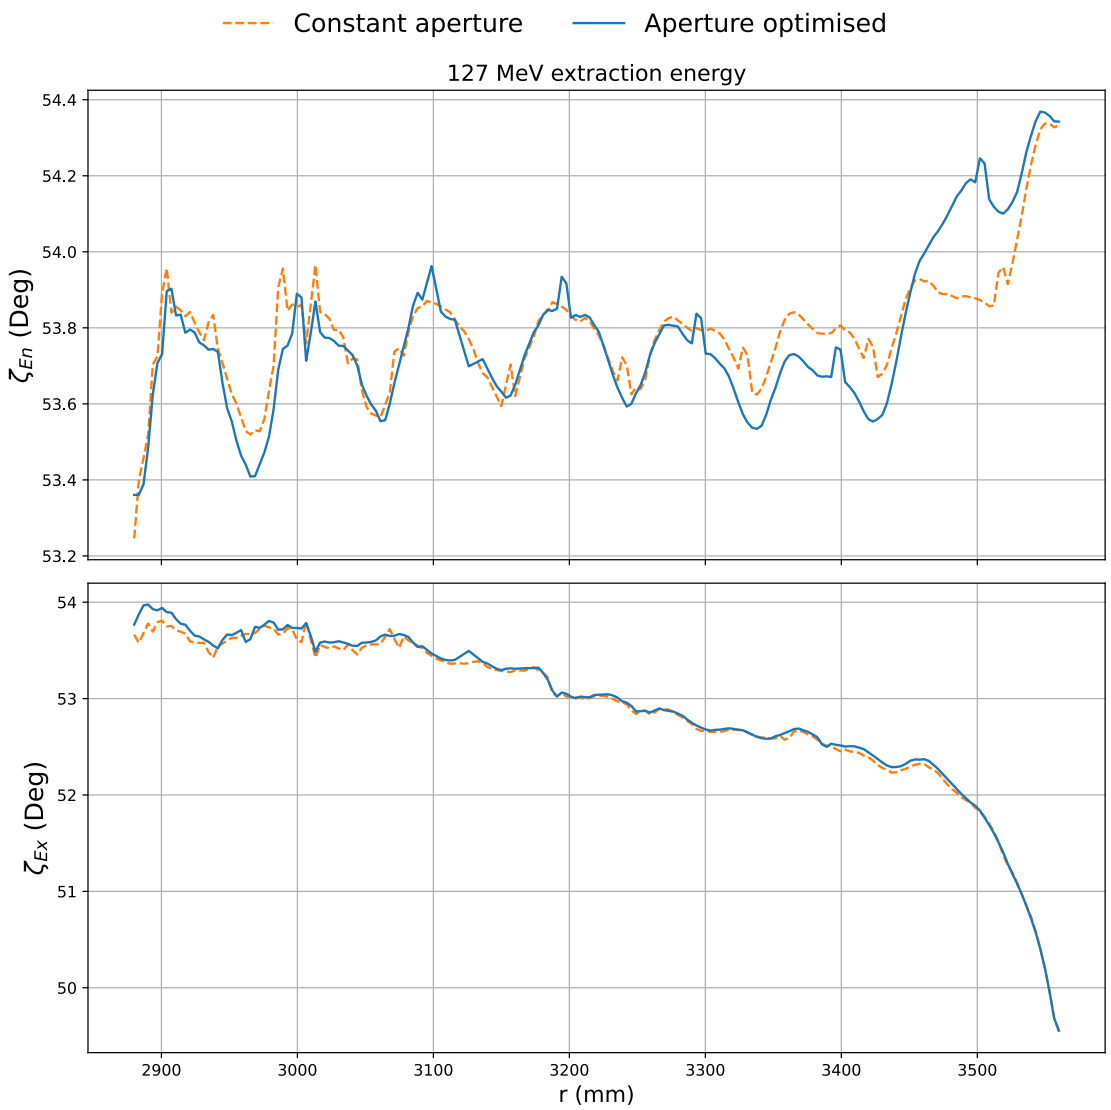


Figure 4.19: Spiral angle of the magnetic field before and after clamp aperture optimisation at the (Top) entrance and (Bottom) exit for the highest extraction energy.

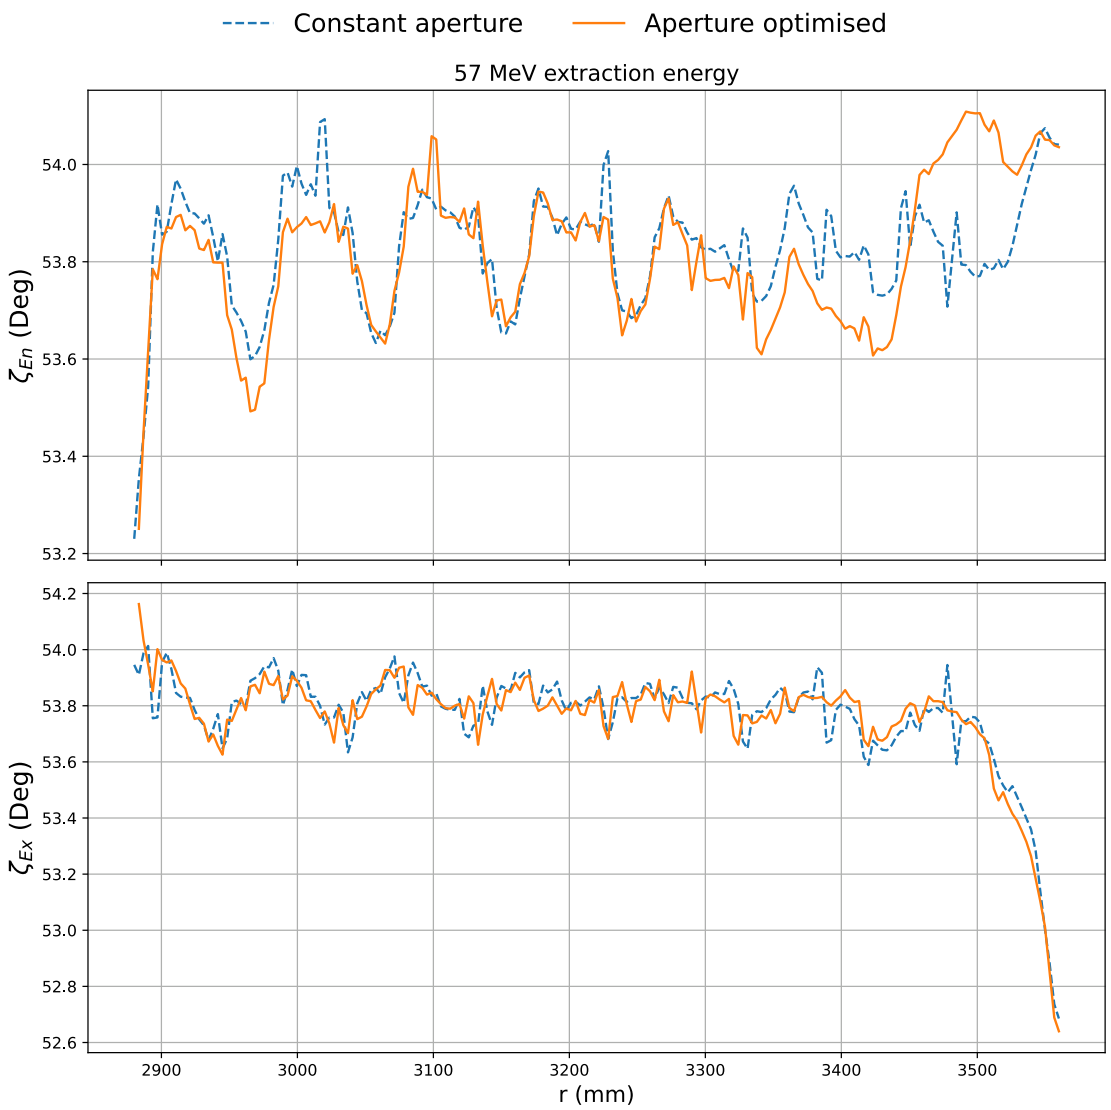


Figure 4.20: Spiral angle of the magnetic field before and after clamp aperture optimisation at the (Top) entrance and (Bottom) exit for the lower extraction energy.

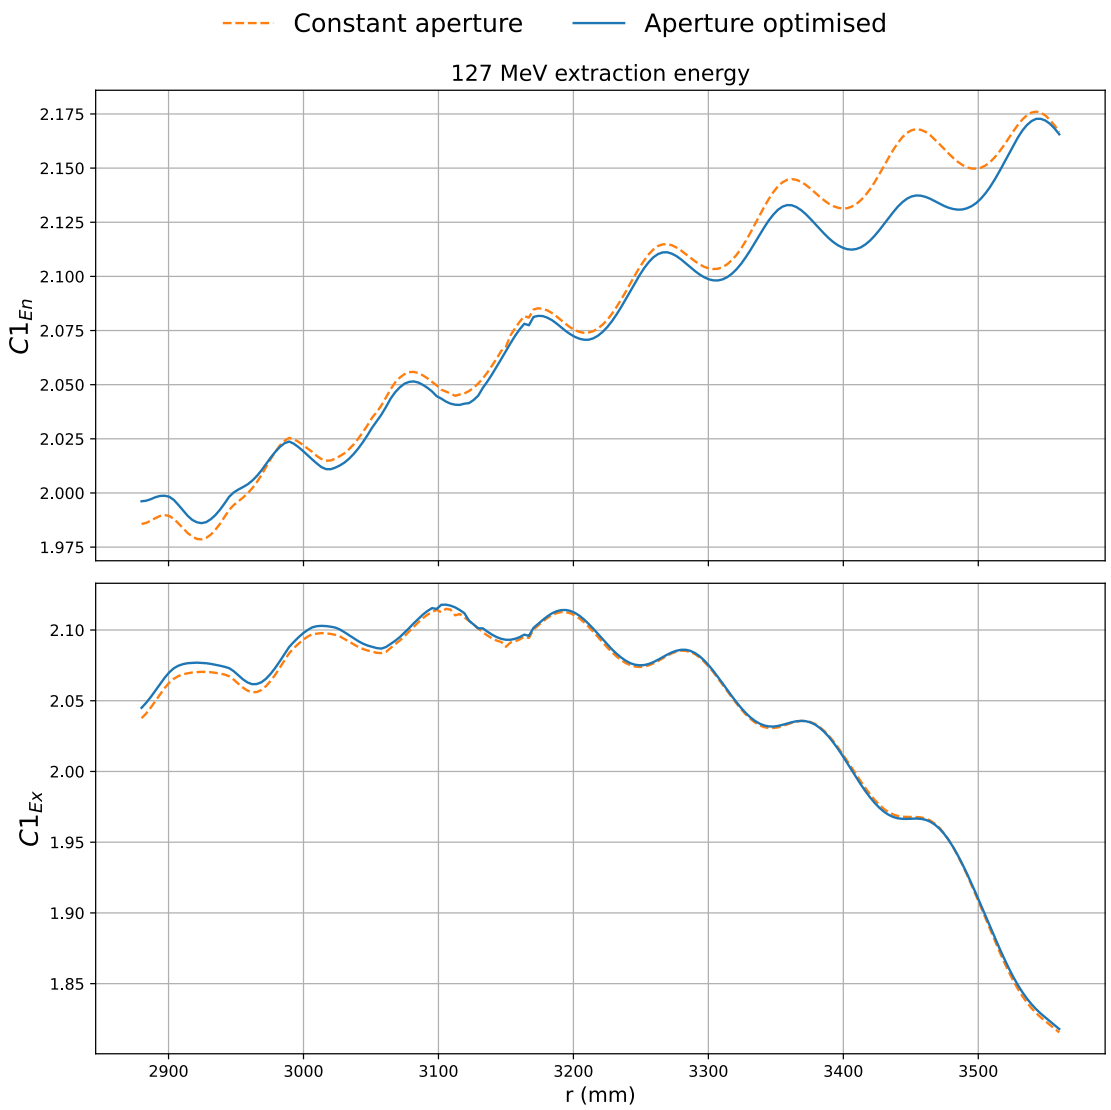


Figure 4.21: $C1$ coefficients of the magnetic field before and after clamp aperture optimisation at (Top) the entrance and (Bottom) the exit for the highest extraction energy.

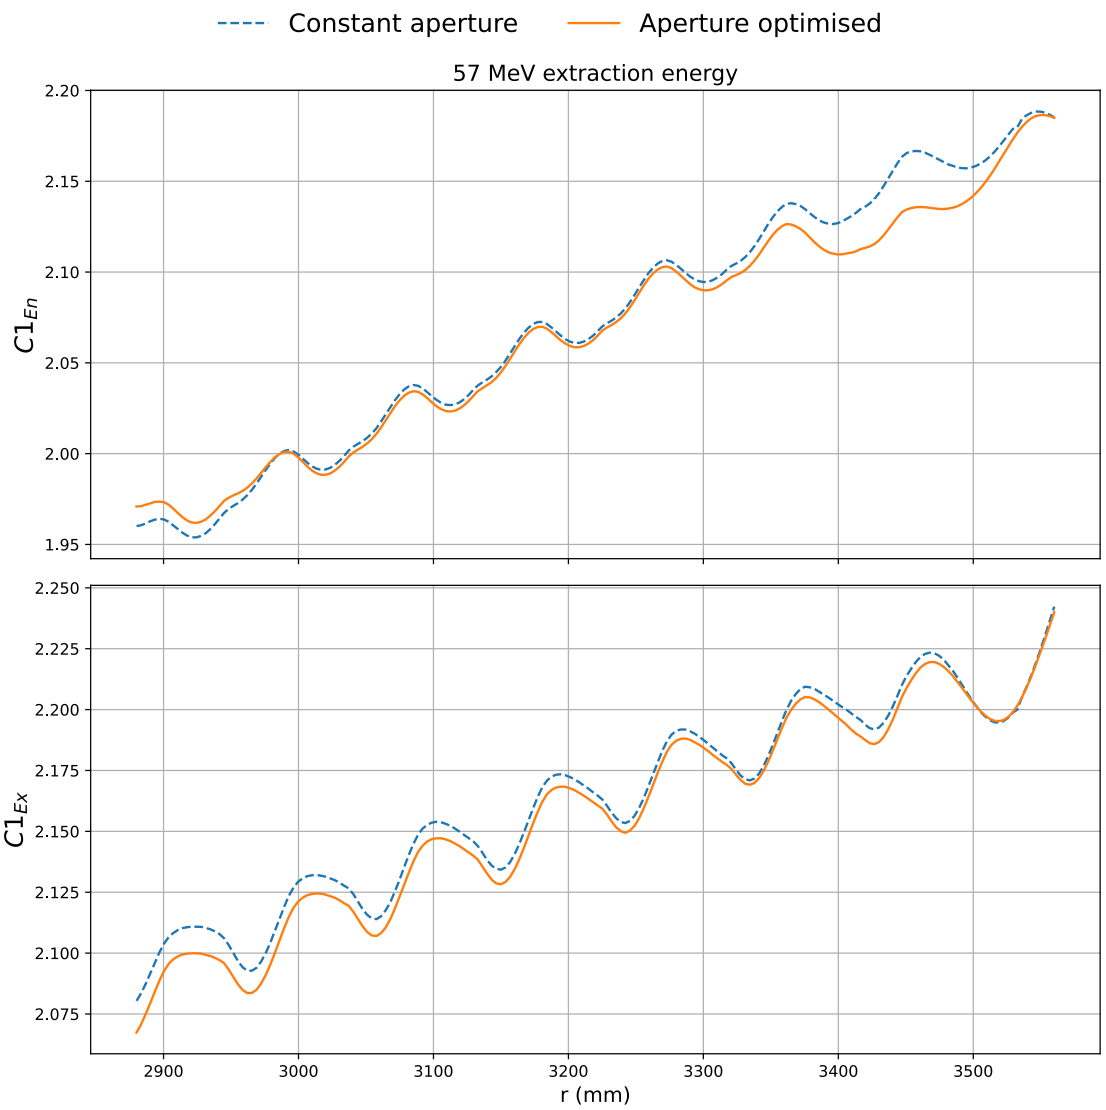


Figure 4.22: $C1$ coefficients of the magnetic field before and after clamp aperture optimisation at (Top) the entrance and (Bottom) the exit for the lower extraction energy.

The ring tune variation following clamp aperture optimisation is presented in Fig. 4.23 and Fig. 4.24. Since a single clamp aperture cannot simultaneously optimise the tune behaviour across both extraction energy ranges, priority was given to minimising the vertical tune variation at the highest extraction energy, as this case exhibits a larger horizontal tune variation, thereby achieving the greatest overall reduction in tune excursion across both energy ranges.

The vertical ring tune at the highest extraction energy was deliberately optimised to exhibit a slight decreasing trend towards the end of acceleration. The rationale for this choice is illustrated in Fig. 4.25: by aligning the tune footprint along the direction of the second-order resonance line, a greater clearance is maintained between the operating tune and the resonance compared to the case of a constant vertical tune. For the lower extraction energy case, the reduced horizontal tune variation results in a more compact tune footprint, and the vertical tune spread has likewise been reduced; in both cases, the overall tune excursion is improved relative to the constant clamp aperture design. To improve the design further, the apertures of both the entrance and exit clamps should be optimised such that the resulting fringe fields conform to the scaling law.

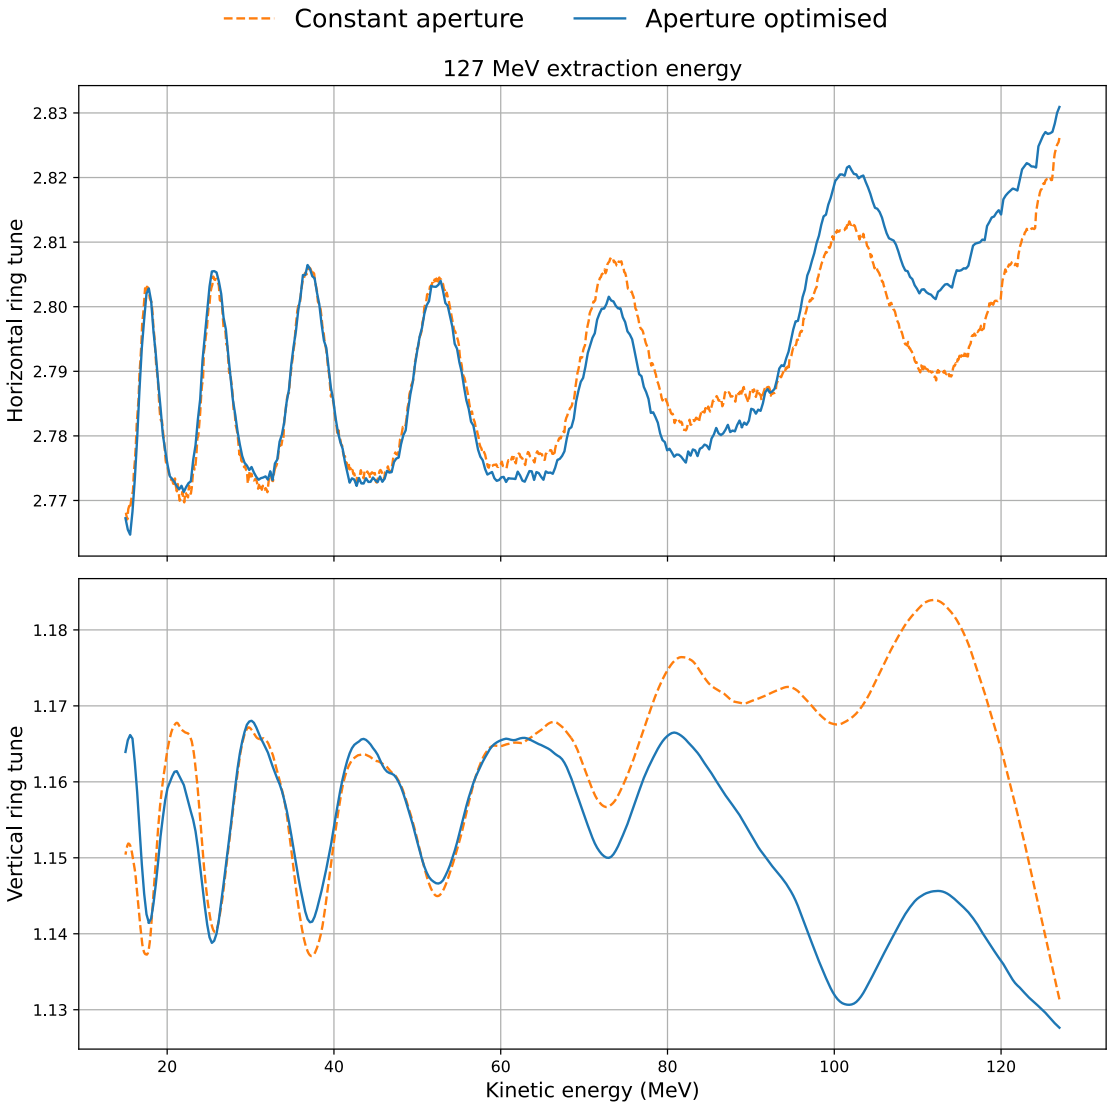


Figure 4.23: Variation of the (Top) horizontal and (Bottom) vertical ring tune before and after clamp aperture optimisation for the highest extraction energy case.

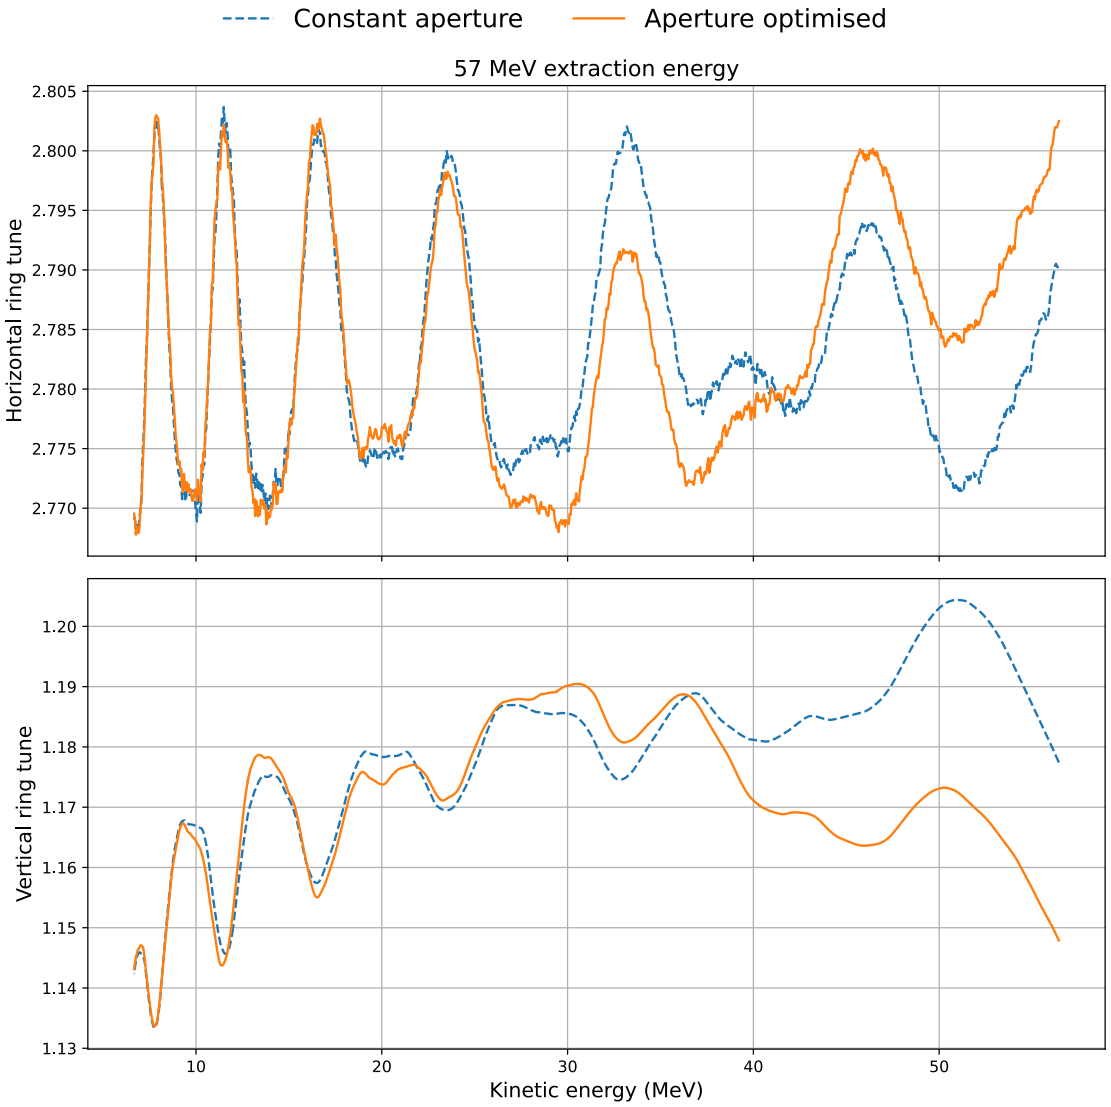


Figure 4.24: Variation of the (Top) horizontal and (Bottom) vertical ring tune before and after clamp aperture optimisation for the lower extraction energy case.

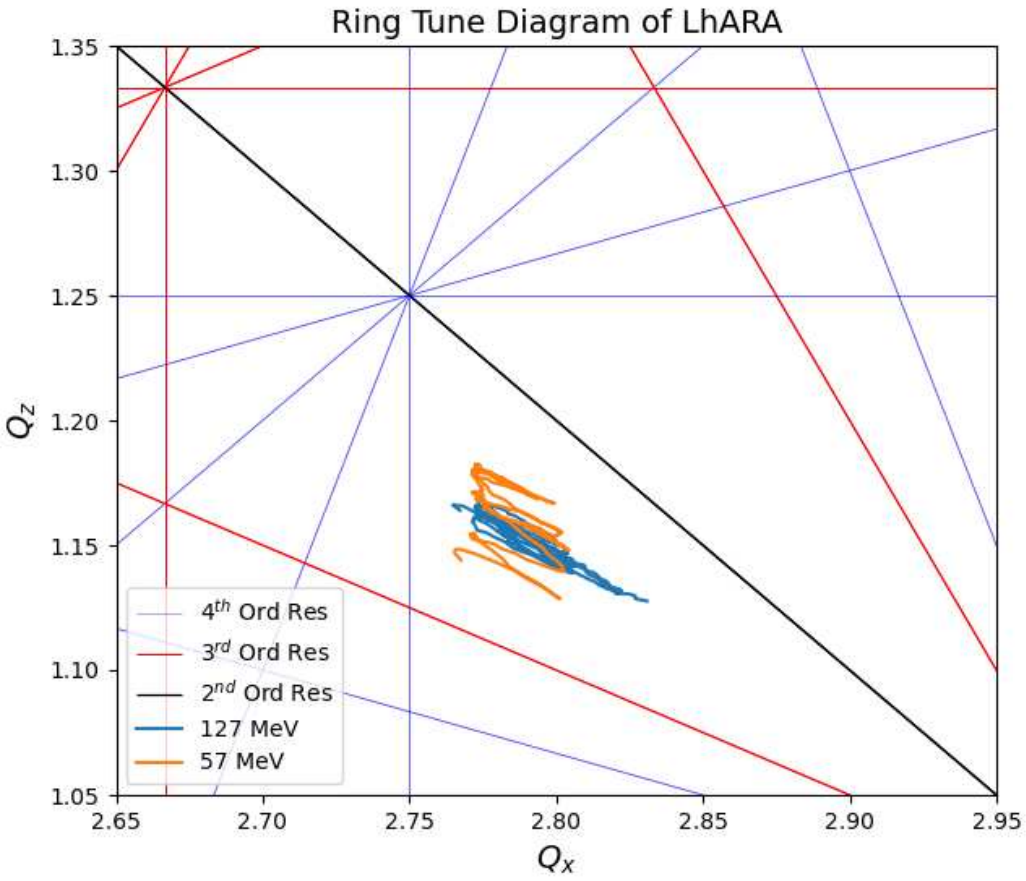


Figure 4.25: Ring tune footprints for two different extraction energy ranges after clamp aperture optimisation, plotted on a resonance diagram including resonance lines up to fourth order.

4.7 Flexibility in tune space

The LhARA FFA employs a singlet spiral magnet design, which inherently offers more limited flexibility in adjusting the working tune point post construction compared to a doublet configuration. Nevertheless, several adjustments remain available following manufacture to recover the target working tune point in the event of deviations arising from discrepancies between the physical magnet and the model, space charge tune shifts, or other operational effects.

The first adjustment involves modifying the field index k by varying the trim coil current settings, whilst retaining the clamp aperture geometry shown in Fig. 4.17. As illustrated in Fig. 4.26, increasing k results in an increase in the horizontal tune and a corresponding decrease in the vertical tune. The tune spread remains broadly similar across different values of k; however, the clamp aperture may be independently optimised for each k setting to further reduce the tune excursion.

The second adjustment consists of translating the entrance and exit field clamps vertically, whilst preserving the clamp aperture profile of Fig. 4.17. The range of this adjustment is constrained by the vacuum chamber height and the extent to which the clamp may be displaced upward before introducing significant field leakage into the straight section. Fig. 4.27 presents the resulting tune variation for vertical clamp displacements of $+20$ mm and -20 mm. Lowering the clamp aperture is found to increase the vertical tune, whilst raising it produces the opposite effect; however, the tune spread changes appreciably in both cases, necessitating a reoptimisation of the clamp aperture profile for each displacement setting.

These studies demonstrate the flexibility afforded by the trim coil design in adjusting the operating tune of the LhARA FFA, and indicate that the field clamps should be mechanically designed to permit post-construction adjustment.

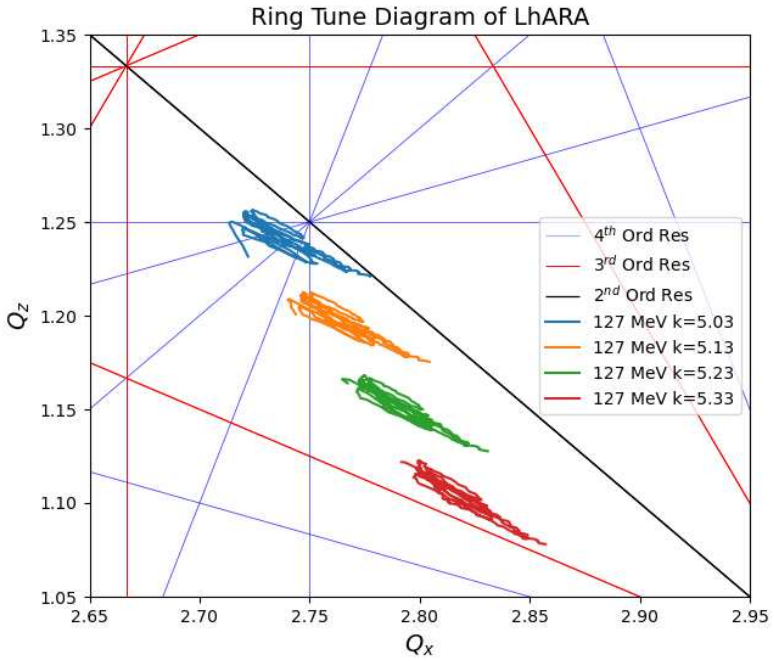


Figure 4.26: Ring tune variation for different values of the field index k.

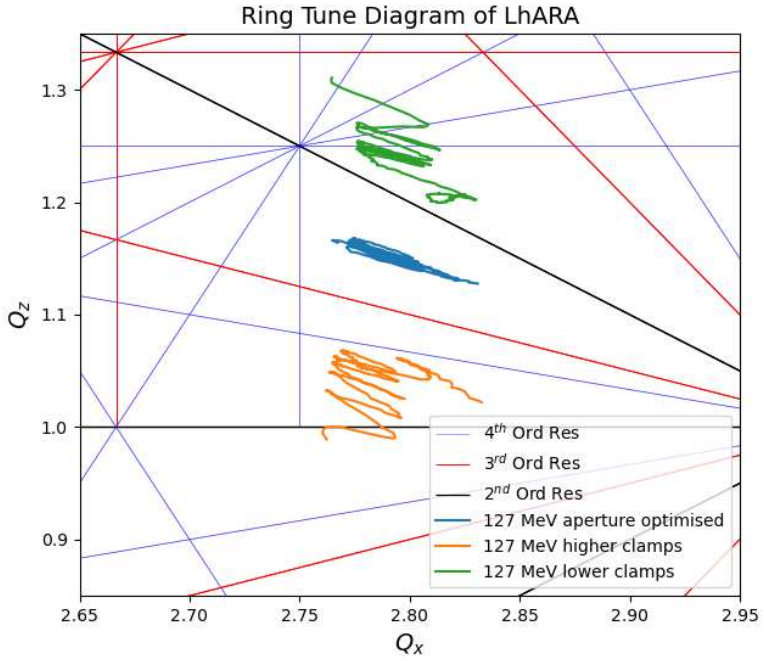


Figure 4.27: Variation in ring tune when the optimised clamp aperture is displaced by a fixed amount in the upward (orange) and downward (green) directions.

Chapter 5

Prototype magnet

A prototype magnet was planned for the FETS-FFA project; however, due to project constraints, producing a full-scale prototype within the available timeframe was not feasible. Instead, a scaled-down version was designed for manufacture, a single radial sector magnet with parameters reduced by a factor of one third. The objectives of this prototype are threefold: to investigate the manufacturing process of magnets with radially distributed coils, to benchmark measurements against the OPERA model in order to establish a control methodology, and to develop and verify correction schemes for locally correcting closed orbit distortion and higher order field errors.

It should be noted that this prototype is not suited to studying the spiral characteristics of the magnet, including the manufacturing process of spiral nested coils and spiral focusing effects. Similarly, saturation effects will be difficult to examine under normal operating currents, and cross talk between neighbouring magnets will not be observable. Nevertheless, the benchmarked OPERA simulations could still provide valuable insight into the behaviour of the full scale doublet.

5.1 Lattice and magnet parameters

To form a lattice with sufficient focusing strength, focusing and defocusing settings were developed for the prototype magnet. By operating the magnet at these two settings, tracking studies can be carried out by overlapping the resulting field maps. The lattice parameters are shown in Table 5.1, the magnet parameters are summarised in Table 5.2.

Table 5.1: Parameters of the prototype FFA lattice

Parameter	Value
Cell opening angle (deg)	30
Number of cells in the ring	12
F/D ratio	1.4339
Field index k	7.66
Target horizontal ring tune	4.66
Target vertical ring tune	2.30
Momentum range (keV)	160 – 650

Table 5.2: Parameters of the prototype FFA magnet

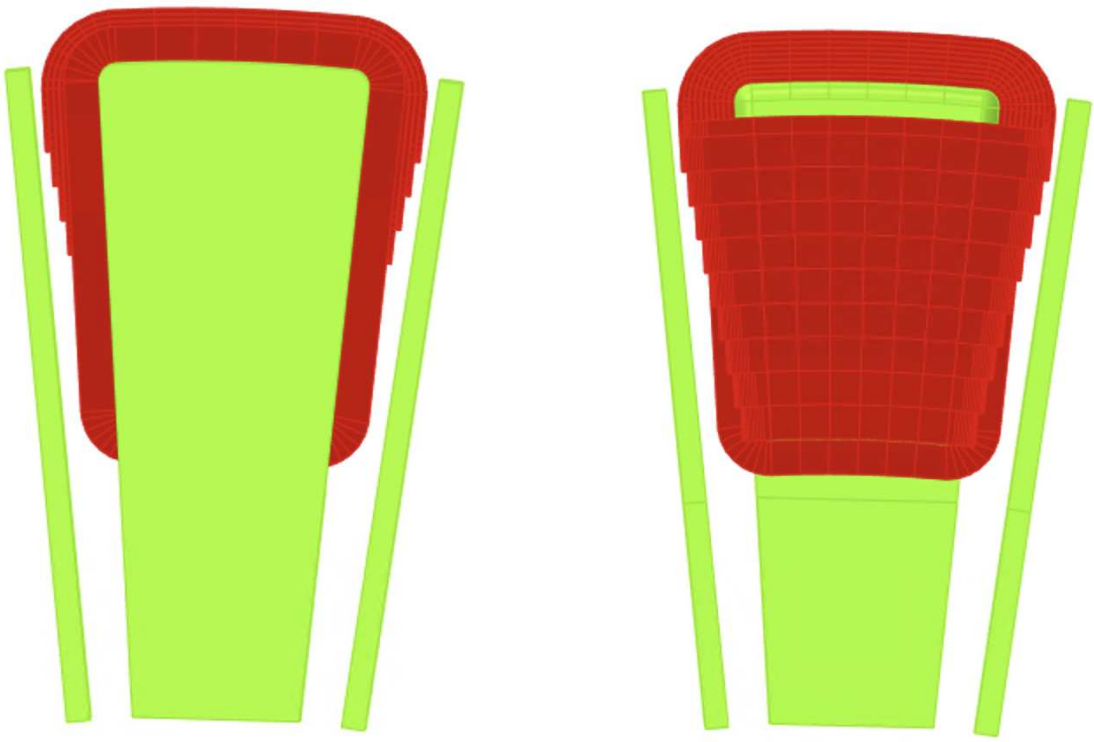
Parameter	Value
Pole opening angle (deg)	9.25
Opening angle between pole and clamp (deg)	2.36
Gap height of iron pole (mm)	$\pm$21
Min/Max radius of the trim coils (m)	1.143 / 1.428
Min/Max radius of the GFR (m)	1.189 / 1.411
Number of trim coils per magnet	18
Number of power supplies per magnet	9

The coil configuration follows a similar arrangement to that employed in the FETS-FFA and LhARA projects, but without a spiral angle. The trim coils are of saddle type and are nested, wound across two overlapping layers on the pole face and returning through two vertical layers. To maximise power efficiency, all trim coil returns are directed towards the higher radius region. Deliberate spacing has been introduced between the return layers to accommodate the coil return assembly, as shown in Fig. 5.1. The main coil is represented as a solid conductor with cross sectional dimensions of 30 mm $\times$ 30 mm, while the trim coils are modelled with dimensions of 4 mm $\times$ 30 mm.

5.2 Optimisation

The magnet model was optimised under two different settings to represent the focusing and defocusing configurations. The corresponding parameters are: $r0 = 1.3$ m, BL_{0_F} = -0.067916 Tm, BL_{0_D} =0.047364 Tm. Following the procedures introduced in Chap. 2, the currents were optimised for the integrated field, with the final results shown in Fig. 5.2.

The field quality in terms of the field index k is very similar for both configurations, indicating that the magnet is operating in the unsaturated regime where the magnetic response is linear.



(a) Top view

(b) Pole face view

(c) Side view without clamps

Figure 5.1: OPERA 3D model of the prototype magnet, with the coils (red) and iron (green).

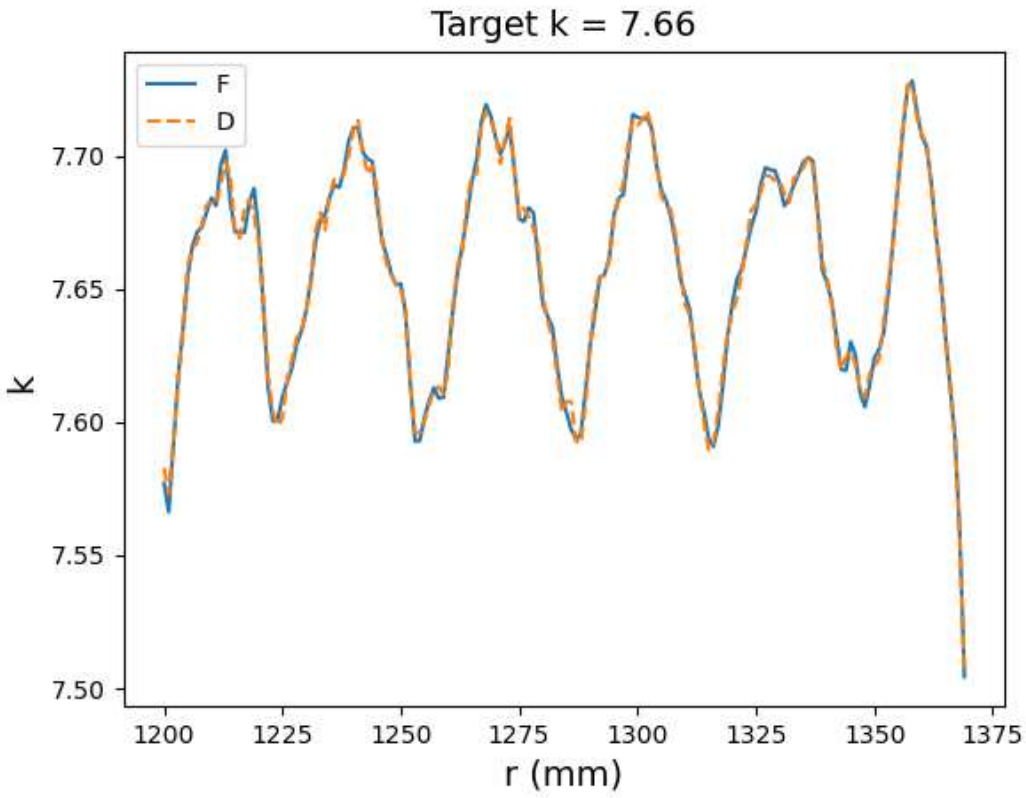


Figure 5.2: Field index k evaluated from the OPERA-3D model of the prototype magnet.

5.3 Future work

Tesla Engineering Ltd [74] was selected as the manufacturer of the prototype magnet. They were responsible for the mechanical design and fabrication of the magnet, shown in Fig. 5.3, which has been delivered to RAL.

Simulations indicate that the field clamps are positioned too close to the magnet, producing an unphysical fringe field decay characterised by a slight bump. This effect can be avoided by magnetically connecting the clamps to the magnet return yoke, which changes the direction of the magnetic flux, as illustrated in Fig. 5.4, yielding a considerably smoother fringe field fall off, as shown in Fig. 5.5.

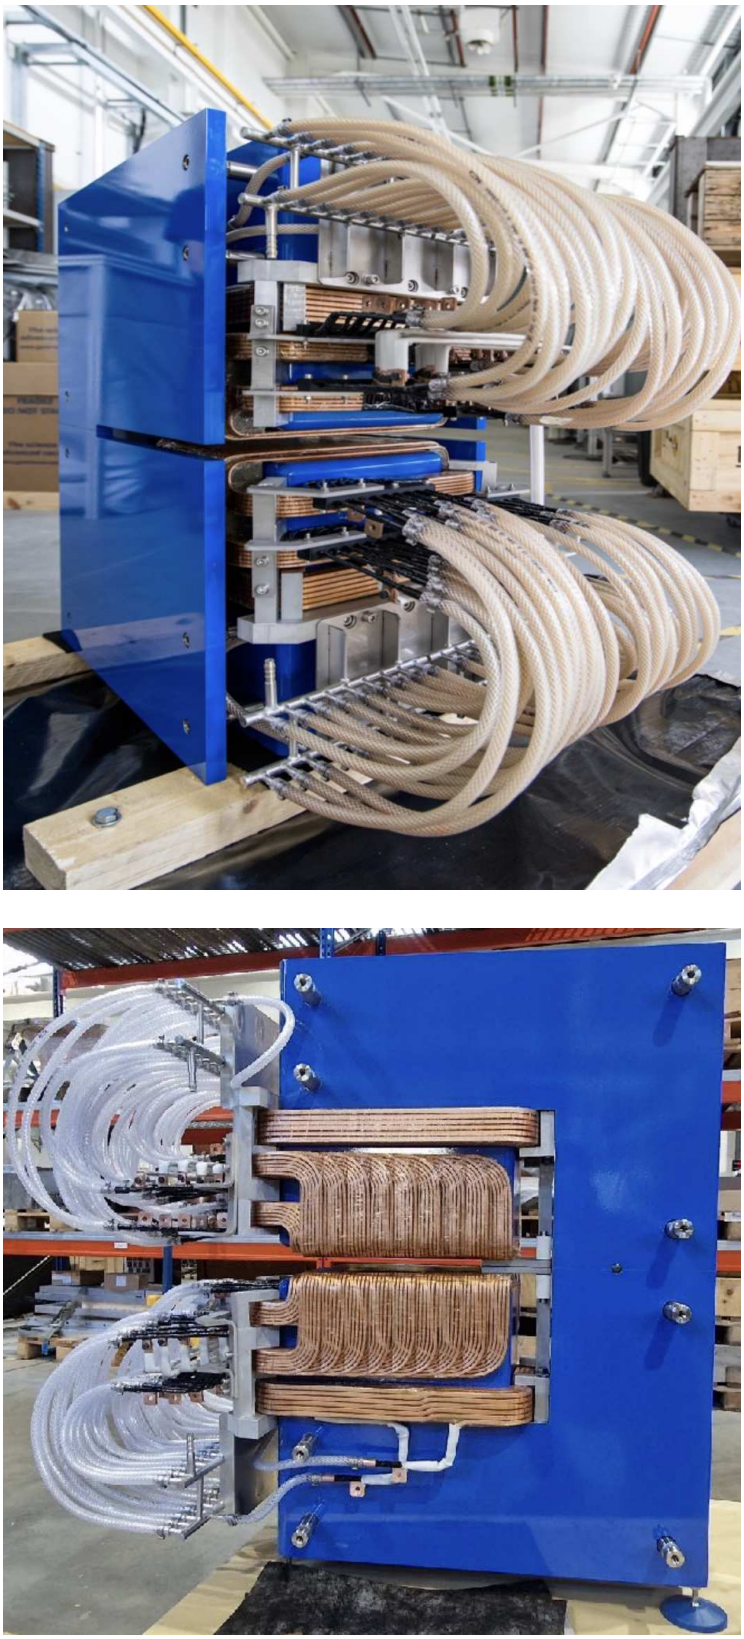


Figure 5.3: The coils (bronze) and iron yoke (blue) of the prototype magnet, with the electrical and water connection assembly for the coils located towards the outer radius of the magnet. The lower figure shows the magnet without the field clamps to highlight the coil structure.

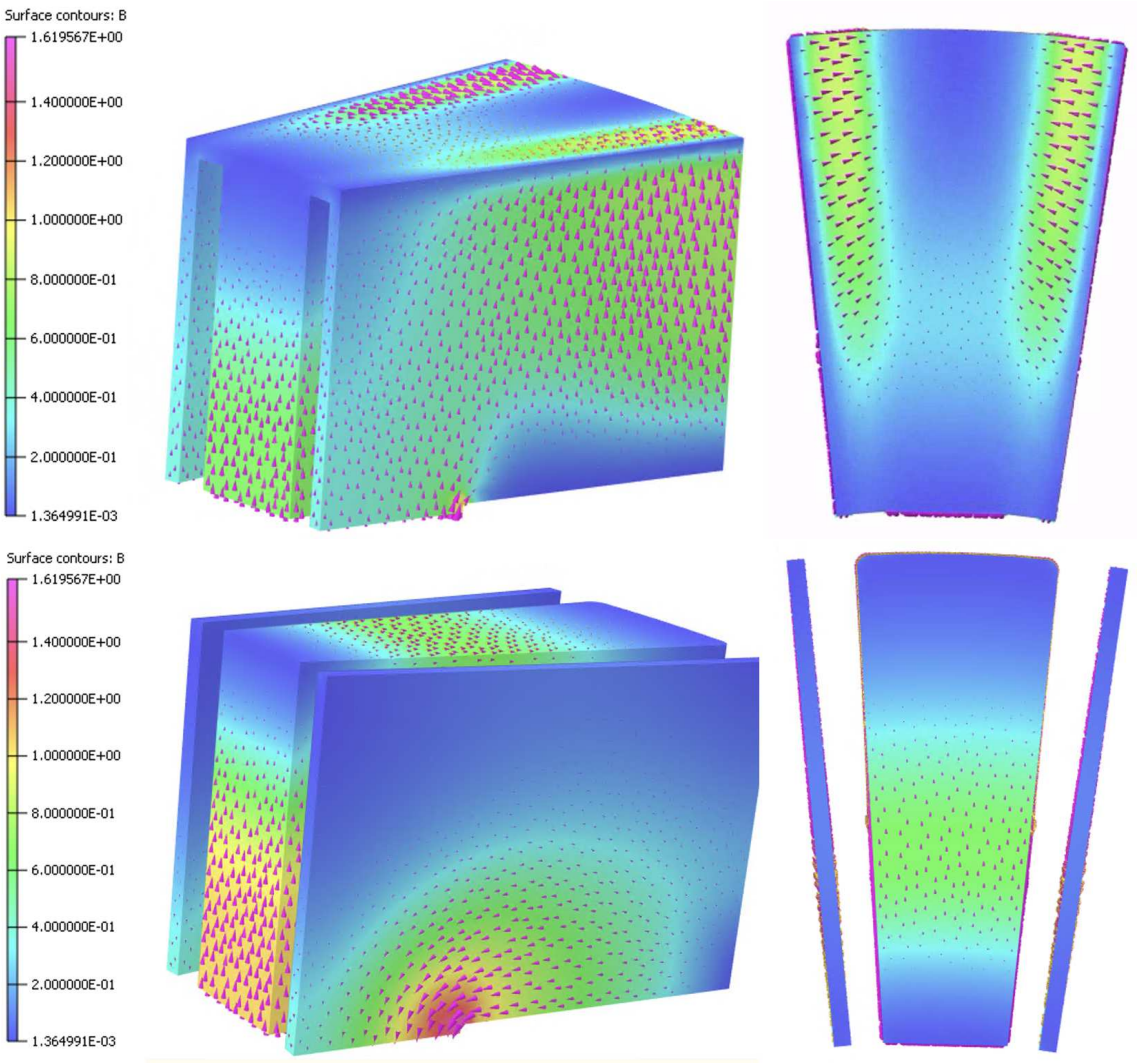


Figure 5.4: Flux direction in the magnet model shown with the saturation level, shown without the coils for clarity. (Top) Clamps connected to the return yoke. (Bottom) Separated clamps. The left panel shows a side view and the right panel shows a top view of the magnet model.

A planned upgrade to the magnet will incorporate an iron plate bridging the clamps and the return yoke, replicating the effect of a direct clamp-to-yoke connection. In the interim, field measurements will proceed using the existing separated configuration.

The first stage involves working with the ISIS Controls group to establish a control system for powering the magnet, after which field maps will be measured using the nominal current settings obtained from OPERA simulations. These measurements will benchmark the real world magnet behaviour against the OPERA model predictions, providing experimental validation of the simulations. Concurrently, a Jacobian matrix will be derived from experimental data to characterise the sensitivity of the magnetic field to variations in each power supply current. This empirical Jacobian will subsequently be employed to iteratively refine the current settings until the desired field configuration is achieved. With the currents optimised, separate three dimensional field maps will be measured for the F and D magnet configurations and superimposed to construct the complete field map of a single cell, which will serve as the basis for particle tracking studies.

The second stage is dedicated to verifying the operating principles of a variable k magnet through parameter optimisation across a range of k values. Since the good field region shifts with varying k, the number of active trim coils changes correspondingly, necessitating a systematic approach to jointly optimise the trim coil configuration alongside the required power supply currents. The three dimensional field maps generated for each configuration will then be used in tracking studies to confirm that the magnet delivers the expected performance across the full range of k values.

Finally, the effectiveness of local correction schemes will be investigated, including closed orbit distortion correction and local gradient correction, to evaluate the degree to which field errors can be mitigated in practice. This will demonstrate the versatility of radially distributed trim coils as a correction tool, highlighting their potential to address both orbit and gradient errors within a single, flexible coil arrangement.

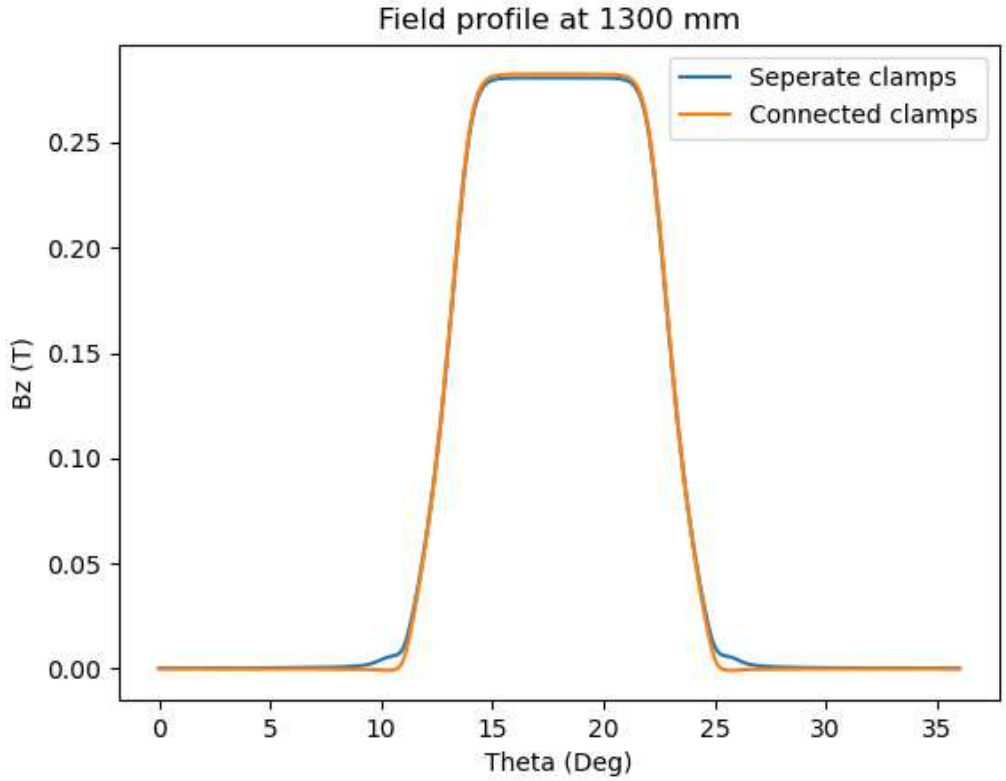


Figure 5.5: Magnetic field distribution in the median plane for the two clamp connection configurations.

Chapter 6

Conclusion

A comprehensive design framework has been developed for scaling Fixed Field Alternating Gradient (FFA) accelerator magnets incorporating radially distributed trim coils, whose magnetic field gradient can be varied radially through adjustment of the individual coil current settings.

The magnetic field distribution for a scaling FFA was derived, and the focusing properties decomposed into longitudinal and radial components for systematic analysis. This analysis establishes the geometric requirements of the magnet: the longitudinal boundaries must conform to a spiral profile, with the radial sector representing the special case of zero spiral angle. Trim coils crossing the pole face should approximate arcs centred on the machine axis as closely as possible. Within this geometry, the trim coil currents are optimised such that the integrated field BL follows a radial distribution of the form:

$$BL(r)_i = BL_{0i}\left(\frac{r}{r_0}\right)^{k+1} \tag{6.1}$$

where the subscript i denotes each individual magnet within a cell, each independently satisfying this distribution, and BL_0 is the integrated magnetic field at reference radius r_0.

Field clamps were incorporated on both longitudinal sides of the magnet, positioned at a fixed opening angle relative to the magnet and conforming to the same spiral angle, to confine and control the fringe field. Connecting the clamps to the magnet return yoke produced two beneficial effects: an increase in the peak field at the magnet centre for equivalent current settings, and improved fringe field containment characterised by a sharper fall off. For a flat pole magnet, field clamps of constant aperture were found to be insufficient for the fringe field to satisfy the scaling condition; their apertures must therefore be optimised simultaneously with the current settings. Since the entrance clamp exerts a greater influence on the tune, and to limit the complexity of the optimisation problem, only the entrance clamp aperture was optimised in the present work.

FFA accelerators are capable of reaching relativistic energies, and their fixed field nature frees them from the magnet ramping constraints. These characteristics make FFAs particularly attractive for high intensity applications. Operational experience from high intensity machines indicates

that lattice optics should offer flexibility in tune space to identify optimal working points [61, 62]. Motivated by this, the FETS-FFA prototype lattice was designed to demonstrate the viability of FFAs for high intensity operation as a candidate for an ISIS-II upgrade [54]. A novel doublet spiral lattice was introduced [63], which allows the magnetic field ratio to be adjusted to modify the tune whilst keeping the machine circumference smaller than that of conventional doublet lattices with no spiral angle. When combined with radially distributed trim coils [64], variation of the field index k provides an additional degree of freedom in the accessible tune space. Magnetically connecting the clamps to the doublet return yoke improved overall magnet efficiency, reducing the yoke height from 700 mm to 300 mm, with a consequent weight reduction of 20% and a 13% decrease in power consumption.

As demonstrated in Chapter 3, field index variation between $k = 6.20$ and $k = 8.75$ was achieved, enabling exploration of both horizontal and vertical ring tunes in the range 3–4. For a fixed momentum range, the radial beam excursion varies with k, and different numbers of trim coils are energised and optimised for each configuration. The entrance clamp aperture was optimised to achieve the best overall tune across all working points, yielding a maximum ring tune variation of ± 0.04. The decision to optimise only the entrance clamp aperture was motivated by the reduction in the number of free parameters during optimisation; furthermore, since no field clamp is present between the two magnets of the doublet, the fringe field in that region does not fully conform to the scaling law regardless, and optimisation of the exit clamp aperture was therefore not pursued. The achieved tune variation confirms both the flexibility of the FFA lattice in tune space and its ability to maintain small tune excursions throughout acceleration, supporting its potential for high intensity operation.

The discovery of FLASH radiotherapy [11] has further broadened the appeal of FFAs for medical applications. Their high repetition rates and ability to vary output energy without degraders make them strong candidates for next generation radiotherapy facilities. The proposed LhARA facility [70, 71] aims to couple a laser-driven particle source with a singlet spiral FFA ring to deliver variable energy beams in support of radiobiological research. Currently constrained to 10 Hz by the laser system, the FFA ring is designed to be compatible with future repetition rates of hundreds of Hz.

As detailed in Chapter 4, coil current settings were optimised for two energy ranges: a maximum energy case spanning 15 MeV injection to 127 MeV extraction, and a lower energy case spanning 7 MeV injection to 57 MeV extraction. Differences in magnetisation behaviour due to magnetic saturation were observed between the two regimes. The entrance clamp aperture was optimised to maximise the separation between the working tune at maximum extraction energy and nearby resonances. Following optimisation, the working tune points for both energy ranges were found to overlap and to avoid resonances up to fourth order. Since field clamps are present on both sides of the singlet magnet, further improvements may be achieved by optimising the apertures of both the entrance and exit clamps such that the resulting fringe fields conform to the scaling law.

To validate the magnet behaviour experimentally, a scaled down prototype without spiral angle has been fabricated by Tesla Engineering Ltd [74]. This prototype is intended to demonstrate

the operating principles of the trim coils in establishing the radial field gradient. Measurement results will be used to benchmark the magnet against OPERA computational models [13], and a methodology for experimentally evaluating the Jacobian will be established to enable current setting optimisation from measured data. Three dimensional field maps will be acquired in both focusing and defocusing configurations, to be superimposed and used in particle tracking studies. Future work will address closed orbit distortion and local radial gradient correction, with the broader aim of fully demonstrating the versatility of radially distributed trim coils.

Bibliography

[1] S. Brooks. "Non-scaling fixed-field proton accelerator with constant tunes". In: *JACoW* IPAC2023 (2023), TUPM031. DOI: 10.18429/JACoW-IPAC2023-TUPM031.

[2] K. R. Symon et al. "Fixed-Field Alternating-Gradient Particle Accelerators". In: *Physical Review* 103.6 (1956), pp. 1837–1859.

[3] T. Ohkawa. "Two-beam fixed field alternating gradient accelerator". In: *Review of Scientific Instruments* 29.2 (1958), pp. 108–110.

[4] A. A. Kolomensky and A. N. Lebedev. *Theory of Cyclic Accelerators*. Amsterdam: North-Holland, 1966, pp. 77–81.

[5] Alfred Gray. *Modern Differential Geometry of Curves and Surfaces with Mathematica*. 2nd. CRC Press, 1997.

[6] Ernest D. Courant, M. Stanley Livingston, and Hartland S. Snyder. "The Strong-Focusing Synchroton—A New High Energy Accelerator". In: *Phys. Rev.* 88 (5 Dec. 1952), pp. 1190–1196. DOI: 10.1103/PhysRev.88.1190. URL: https://link.aps.org/doi/10.1103/PhysRev.88.1190.

[7] G. H. Rees. "NON-ISOCHRONOUS & ISOCHRONOUS, NON-SCALING FFAG DESIGNS". In: 2007. URL: https://api.semanticscholar.org/CorpusID:55632668.

[8] F. T. Cole et al. "Electron model fixed field alternating gradient accelerator". In: *Rev. Sci. Instrum.* 28 (1957), pp. 403–420. DOI: 10.1063/1.1715895.

[9] D. W. Kerst et al. "Electron model of a spiral sector accelerator". In: *Rev. Sci. Instrum.* 31 (1960), pp. 1076–1106. DOI: 10.1063/1.1716818.

[10] Francis T. Cole et al. "MURA 50-MeV Electron AcceleratorGeneral Description and Review. I". In: 1964. URL: https://api.semanticscholar.org/CorpusID:101173164.

[11] V. Favaudon et al. "Ultrahigh Dose-Rate FLASH Irradiation Increases the Differential Response Between Normal and Tumor Tissue in Mice". In: *Science Translational Medicine* 6.245 (July 2014). Erratum published in Sci. Transl. Med. 11(523): eaba4525 (2019), 245ra93. DOI: 10.1126/scitranslmed.3008973.

[12] COMSOL AB. *COMSOL Multiphysics® Software*. https://www.comsol.com. Accessed: 2026-04-13.

[13] Dassault Systemes. *OPERA FEA Software*. URL: http://operafea.com.

[14] J.-B. Lagrange. "The Particle Tracking Code Fixfield". In: *Proc. IPAC'21* (Campinas, SP, Brazil). International Particle Accelerator Conference 12. https://doi.org/10.18429/JACoW-IPAC2021-TUPAB209. JACoW Publishing, Geneva, Switzerland, Aug. 2021, TUPAB209, pp. 1905–1906. ISBN: 978-3-95450-214-1. DOI: 10.18429/JACoW-IPAC2021-TUPAB209. URL: https://jacow.org/ipac2021/papers/tupab209.pdf.

[15] Shinji Machida. "S-code for fixed field alternating gradient accelerator design and particle tracking". In: *Int. J. Mod. Phys. A* 26 (2011). Ed. by Carol Johnstone, Martin Berz, and Pavel Snopok, pp. 1794–1806. DOI: 10.1142/S0217751X11053195.

[16] Francois Meot. “The Ray tracing code Zgoubi”. In: *Nucl. Instrum. Meth. A* 427 (1999), pp. 353–356. DOI: 10.1016/S0168-9002(98)01508-3.

[17] Andreas Adelmann et al. *OPAL a Versatile Tool for Charged Particle Accelerator Simulations*. 2019. arXiv: 1905.06654 [physics.acc-ph]. URL: https://arxiv.org/abs/1905.06654.

[18] M. Aiba et al. “Development of a FFAG Proton Synchrotron”. In: *Proc. EPAC’00* (Vienna, Austria). European Particle Accelerator Conference 7. JACoW Publishing, Geneva, Switzerland, Aug. 2000, pp. 581–583. URL: https://jacow.org/e00/papers/MOP1B21.pdf.

[19] Shinji Machida et al. “COMMISSIONING OF 150MeV FFAG SYNCHROTRON”. In: 2004. URL: https://api.semanticscholar.org/CorpusID:56077664.

[20] T. Adachi et al. “A 150-MeV FFAG Synchrotron with ’Return-Yoke Free’ Magnet”. In: *Conf. Proc. C* 0106181 (2001). Ed. by Peter W. Lucas and Sara Webber, pp. 3254–3256.

[21] M. Tanigaki et al. “Construction of FFAG Accelerators in Kurri for ADS Study”. In: June 2005, pp. 350–352. ISBN: 0-7803-8859-3. DOI: 10.1109/PAC.2005.1590431.

[22] C. H. Pyeon et al. “First Injection of Spallation Neutrons Generated by High-Energy Protons into the Kyoto University Critical Assembly”. In: *Journal of Nuclear Science and Technology* 46.12 (2009), pp. 1091–1093. DOI: 10.1080/18811248.2009.9711620.

[23] Yoshiharu Mori. “Development of FFAG accelerators and their applications for intense secondary particle production”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 562 (June 2006), pp. 591–595. DOI: 10.1016/j.nima.2006.02.044.

[24] Kota Okabe et al. “Design and Construction of FFAG Magnets for the ERIT System at KURRI”. In: *IEEE Trans. Appl. Supercond.* 20.3 (2010), pp. 740–743. DOI: 10.1109/TASC.2010.2042696.

[25] Jaroslaw Pasternak et al. “Accelerator and Particle Physics Research for the Next Generation Muon to Electron Conversion Experiment - the PRISM Task Force”. In: *Conf. Proc. C* 100523 (2010). Ed. by Akira Noda et al., WEPE056.

[26] A. Sato et al. “Six-sector FFAG Ring to Demonstrate Bunch Rotation for PRISM”. In: *Conf. Proc. C* 0806233 (2008). Ed. by Ivan Andrian and Christine Petit-Jean-Genaz, THPP007.

[27] Thomas Planche et al. “Design of a prototype gap shaping spiral dipole for a variable energy prototherapy FFAG”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 604 (June 2009), pp. 435–442. DOI: 10.1016/j.nima.2009.02.026.

[28] Marie-Julie Leray et al. “Magnetic Measurements of the RACCAM Prototype FFAG Dipole”. In: *Particle Accelerator Conference (PAC 09)*. 2010, FR5REP110.

[29] R. Barlow et al. “EMMA: The world’s first non-scaling FFAG”. In: *Nucl. Instrum. Meth. A* 624 (2010), pp. 1–19. DOI: 10.1016/j.nima.2010.08.109.

[30] S. Machida et al. “Acceleration in the linear non-scaling fixed-field alternating-gradient accelerator EMMA”. In: *Nature Physics* 8.3 (2012), pp. 243–247. DOI: 10.1038/nphys2179.

[31] A. Bartnik et al. “CBETA: First Multipass Superconducting Linear Accelerator with Energy Recovery”. In: *Phys. Rev. Lett.* 125 (4 July 2020), p. 044803. DOI: 10.1103/PhysRevLett.125.044803. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.044803.

[32] S Bogacz et al. “20-24 GeV FFA CEBAF ENERGY UPGRADE”. In: Aug. 2021. DOI: 10.18429/JACoW-IPAC2021-MOPAB216.

[33] K. Peach et al. “PAMELA - A Model for an FFAG based Hadron Therapy Machine”. In: *Conf. Proc. C* 070625 (2007). Ed. by C. Petit-Jean-Genaz, p. 2880. DOI: 10.1109/PAC.2007.4440607.

[34] {J. M.} Garland et al. “NORMA - A normal-conducting, scaling racetrack FFAG”. English. In: *IPAC 2014*. Ed. by Christine Petit-Jean-Genaz et al. 5th International Particle Accelerator Conference, IPAC 2014 ; Conference date: 15-06-2014 Through 20-06-2014. Switzerland: JACoW, Jan. 2014, pp. 2198–2201. DOI: 10.18429/JACoW-IPAC2014-WEPRO100.

[35] Brookhaven National Laboratory et al. *Electron Acceleration for eRHIC with the Non-Scaling FFAG*. Technical Report. Brookhaven National Laboratory, 2004. URL: http://www.osti.gov/servlets/purl/15007992-eL7LVo/native/.

[36] T. Ohkawa. "FFAG Electron Cyclotron". In: *Bulletin of the American Physical Society* 30 (1985), p. 20.

[37] G. Leleux, J. Proy, and M. Salvat. *F.F.A.G. Hélicoïdal: Étude de la stabilité bétatron*. Tech. rep. Rapport O.C. 70. Service de Physique Appliquée, Section d'Optique Corpusculaire, 1959.

[38] J. Teichmann. "Accelerators with Vertically Increasing Field". In: *Soviet Journal of Atomic Energy* 12 (1963), pp. 507–512.

[39] Stephen Brooks. "Vertical orbit excursion fixed field alternating gradient accelerators". In: *Phys. Rev. ST Accel. Beams* 16 (8 Aug. 2013), p. 084001. DOI: 10.1103/PhysRevSTAB.16.084001. URL: https://link.aps.org/doi/10.1103/PhysRevSTAB.16.084001.

[40] Y. Mori, Y. Yonemura, and H. Arima. "A Proposal of Harmonictron". In: *Memoirs of the Faculty of Engineering, Kyushu University* 77.2 (Dec. 2017).

[41] Kyosuke Adachi et al. "An electron model of vertical FFA accelerator for Harmonytron". In: *JACoW* IPAC2023 (2023), TUPA043. DOI: 10.18429/JACoW-IPAC2023-TUPA043.

[42] S. Machida et al. "Optics Design of Vertical Excursion Fixed-Field Alternating Gradient Accelerators". In: *Phys. Rev. Accel. Beams* 24.2 (2021), p. 021601. DOI: 10.1103/PhysRevAccelBeams.24.021601. arXiv: 2011.10783 [physics.acc-ph].

[43] J.-B. Lagrange et al. "FETS-FFA Ring Study". In: *Proc. IPAC'21* (Campinas, SP, Brazil). International Particle Accelerator Conference 12. https://doi.org/10.18429/JACoW-IPAC2021-TUPAB208. JACoW Publishing, Geneva, Switzerland, Aug. 2021, TUPAB208, pp. 1901–1904. ISBN: 978-3-95450-214-1. DOI: 10.18429/JACoW-IPAC2021-TUPAB208. URL: https://jacow.org/ipac2021/papers/tupab208.pdf.

[44] M. Tanigaki et al. "Present status of FFAG accelerators in KURRI for ADS study". In: *Conf. Proc. C* 060626 (2006). Ed. by Christopher Prior, pp. 2367–2369.

[45] M. Tanigaki et al. "Construction of FFAG Accelerators in Kurri for ADS Study". In: *Proceedings of the 2005 Particle Accelerator Conference*. 2005, pp. 350–352. DOI: 10.1109/PAC.2005.1590431.

[46] Masahiro Yoshimoto et al. "The magnet design study for the FFAG accelerator". In: *IEEE Transactions on Applied Superconductivity* 14 (2004), pp. 397–401. URL: https://api.semanticscholar.org/CorpusID:27670819.

[47] Damien Neuvéglise and Francois Méot. "An Alternative Design for the RACCAM Magnet with Distributed Conductors". In: *Particle Accelerator Conference (PAC 09)*. 2010, FR5REP095.

[48] D. W. Kerst et al. *A FFAG accelerator with spirally ridged poles*. Tech. rep. MURA-DWK/KMT/LWJ/KRS-3. Midwestern Universities Research Association (MURA), 1954.

[49] H. A. Enge. "Deflecting Magnets". In: *Focusing of Charged Particles, Vol. II*. Ed. by A. Septier. Academic Press Inc., 1967, pp. 203–264.

[50] J.W.G. Thomason. "The ISIS Spallation Neutron and Muon Source—The first thirty-three years". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 917 (Nov. 2018). DOI: 10.1016/j.nima.2018.11.129.

[51] S. Machida. "FFA Options for ISIS Upgrade and the Feasibility Study". In: *JPS Conf. Proc.* 33 (2021). Ed. by Naohito Saito, p. 011005. DOI: 10.7566/JPSCP.33.011005.

[52] J.-B. Lagrange et al. "Progress on Design Studies for the ISIS II Upgrade". In: *Proc. 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019* (Melbourne, Australia). International Particle Accelerator Conference 10. https://doi.org/10.18429/JACoW-IPAC2019-TUPTS068. Geneva, Switzerland: JACoW Publishing, June 2019, pp. 2075–2078. ISBN: 978-3-95450-208-0. DOI: doi:10.18429/JACoW-IPAC2019-TUPTS068. URL: http://jacow.org/ipac2019/papers/tupts068.pdf.

[53] A. Letchford et al. “Status of the RAL Front End Test Stand”. In: *6th International Particle Accelerator Conference*. 2015, THPF105. DOI: 10.18429/JACoW-IPAC2015-THPF105.

[54] S. Machida. “FFA design study for a high intensity proton driver”. In: *JACoW* IPAC2023 (2023), TUPA044. DOI: 10.18429/JACoW-IPAC2023-TUPA044.

[55] D. W. Posthuma de Boer et al. “ISIS-II Update”. In: *Proceedings of the 20th Annual Meeting of the Particle Accelerator Society of Japan*. Japan, 2023. URL: https://www.pasj.jp/web_publish/pasj2023/proceedings/PDF/WEOB/WEOB3.pdf.

[56] J. Wei et al. “Low-loss design for the high-intensity accumulator ring of the Spallation Neutron Source”. In: *Phys. Rev. ST Accel. Beams* 3 (2000), p. 080101. DOI: 10.1103/PhysRevSTAB.3.080101.

[57] H. Hotchi et al. “Achievement of a low-loss 1-MW beam operation in the 3-GeV rapid cycling synchrotron of the Japan Proton Accelerator Research Complex”. In: *Phys. Rev. Accel. Beams* 20 (2017), p. 060402. DOI: 10.1103/PhysRevAccelBeams.20.060402.

[58] H. Hotchi et al. “Beam commissioning of the 3-GeV rapid cycling synchrotron of the Japan Proton Accelerator Research Complex”. In: *Phys. Rev. ST Accel. Beams* 12 (2009), p. 040402. DOI: 10.1103/PhysRevSTAB.12.040402.

[59] H. Hotchi. “Effects of the Montague resonance on the formation of the beam distribution during multiturn injection painting in a high-intensity proton ring”. In: *Phys. Rev. Accel. Beams* 23 (2020), p. 050401. DOI: 10.1103/PhysRevAccelBeams.23.050401.

[60] S. Henderson et al. “The Spallation Neutron Source accelerator system design”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 763 (2014), pp. 610–673. DOI: 10.1016/j.nima.2014.03.067.

[61] V. Danilov et al. “SNS Ring Optics Tuning”. In: *Proceedings of PAC 2001*. 2001. URL: https://proceedings.jacow.org/p01/.

[62] H. Hotchi et al. “J-PARC 3-GeV RCS: 1-MW Beam Operation and Beyond”. In: *Journal of the Physical Society of Japan (Conference Proceedings) / HB2018* (2020). URL: https://www.researchgate.net/publication/343143300_J-PARC_3-GeV_RCS_1-MW_beam_operation_and_beyond.

[63] S. Machida. “Scaling Fixed-Field Alternating-Gradient Accelerators with Reverse Bend and Spiral Edge Angle”. In: *Phys. Rev. Lett.* 119 (6 Aug. 2017), p. 064802. DOI: 10.1103/PhysRevLett.119.064802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.119.064802.

[64] J.-B. Lagrange et al. “FFA Magnet for Pulsed High Power Proton Driver”. In: *JACoW* HB2023 (2024), THBP02. DOI: 10.18429/JACoW-HB2023-THBP02.

[65] S. Machida et al. *Conceptual Design Report of FETS-FFA (Fixed Field Alternating Gradient Accelerator): A Demonstrator of a High Intensity FFA*. Tech. rep. ISIS/TECH. STFC Rutherford Appleton Laboratory, 2025. URL: http://purl.org/net/epubs/work/62105376.

[66] F. Méot and F. Lemuet. “Developments in the ray-tracing code Zgoubi for 6-D multiturn tracking in FFAG rings”. In: *Nuclear Instruments and Methods in Physics Research Section A* 547.3 (2005), pp. 638–651.

[67] N. R. Datta, S. Rogers, and S. Bodis. “Challenges and Opportunities to Realize “The 2030 Agenda for Sustainable Development” by the United Nations: Implications for Radiation Therapy Infrastructure in Low- and Middle-Income Countries”. In: *International Journal of Radiation Oncology, Biology, Physics* 105.5 (Dec. 2019), pp. 918–933. DOI: 10.1016/j.ijrobp.2019.04.033.

[68] Berry R J. “EFFECTS OF RADIATION DOSE-RATE: From Protracted, Continuous Irradiation to Ultra-High Dose-Rates from Pulsed Accelerators”. In: *British Medical Bulletin* 29 (1973), pp. 44–47. ISSN: 0007-1420. DOI: 10.1093/oxfordjournals.bmb.a070955.

[69] Y. Prezado and G. R. Fois. “Proton Minibeam Radiation Therapy: A Proof of Concept”. In: *Medical Physics* 40.3 (Mar. 2013), p. 031712. DOI: 10.1118/1.4791648.

[70] Galen Aymar et al. “LhARA: The Laser-hybrid Accelerator for Radiobiological Applications”. In: *Frontiers in Physics* Volume 8 - 2020 (2020). ISSN: 2296-424X. DOI: 10.3389/fphy.2020.567738. URL: https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2020.567738.

[71] William Shields et al. “Status update of the laser-hybrid accelerator for radiobiological applications”. In: *JACoW* IPAC2025 (2025), TUPB103. DOI: 10.18429/JACoW-IPAC2025-TUPB103.

[72] D. Gabor. “A Space-Charge Lens for the Focusing of Ion Beams”. In: *Nature* 160 (1947), pp. 89–90. DOI: 10.1038/160089b0.

[73] LhARA Collaboration. *ITRF/LhARA Conceptual Design Report*. STFC Technical Report STFC-TR-2026-007. 275 p. UKRI Science and Technology Facilities Council, 2026. DOI: 10.5286/stfctr.2026007. URL: https://doi.org/10.5286/stfctr.2026007.

[74] Tesla Engineering Limited. *Tesla Engineering Ltd Group Website*. Accessed 29 April 2026. 2026. URL: https://www.tesla.co.uk/.

Appendix A

Numerical procedure to produce trim coil geometry

To determine the precise geometry of the trim coils on the pole face, a numerical procedure was developed that accounts for the full three dimensional spiral geometry of the coil, including its thickness and minimum bending radius. Since the cross section of the trim coil must remain constant throughout, once the shape of a particular edge has been defined, all remaining edges can be computed accordingly. Each trim coil is parametrised by the radius at which it crosses the pole face.

A key requirement of the trim coil design is to maximise the length of the constant radius section on the pole face. To ensure that the coils can be routed around the pole as closely as possible in a smooth and physically realisable manner, the design procedure begins at the spiral edges of the pole. By accounting for the geometry of the bends required to route the coils around the pole, it is possible to determine the starting radius at the pole edge such that the desired radius is achieved at the centre of the pole face.

By convention, θ is measured as positive in the clockwise direction, therefore:

$$\begin{aligned} x &= r\sin\theta \\ y &= r\cos\theta \end{aligned} \tag{A.1}$$

The procedure begins from point b in Fig. A.1, which lies on the spiral edge of the pole. The values of r and θ at this point must be determined subject to the constraint imposed by Eq. 1.70. The local gradient of the spiral is obtained from:

$$\frac{dr}{d\theta} = \left(\frac{d\theta}{dr}\right)^{-1}, \frac{dy}{dx} = \frac{-r\sin\theta + (dr/d\theta)\cos\theta}{r\cos\theta + (dr/d\theta)\sin\theta} \tag{A.2}$$

The tangent angle ϕ is then given by:

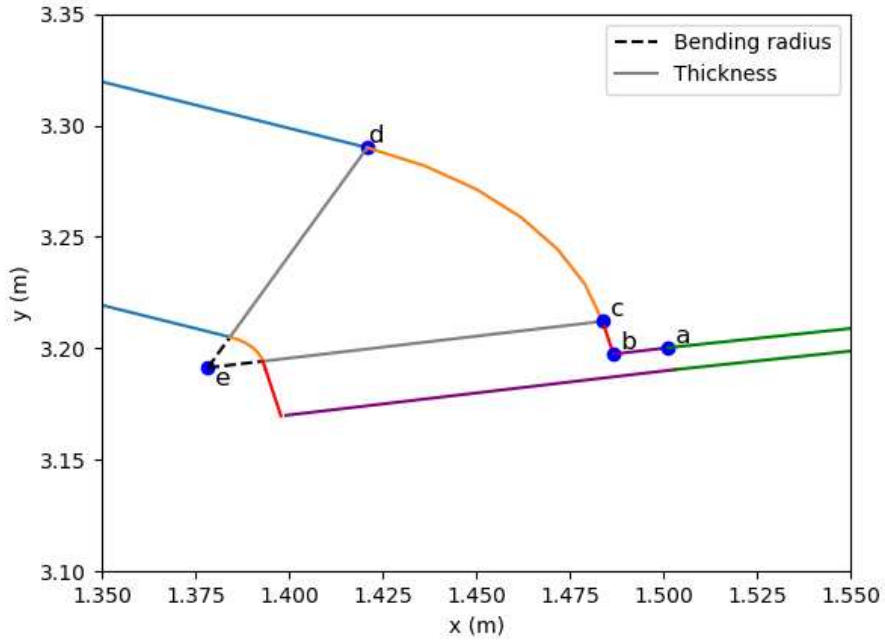


Figure A.1: Detailed schematic indicating the labelled points introduced in the trim coil coordinate calculation procedure.

$$\phi = \arctan\left(\frac{dy}{dx}\right). \tag{A.3}$$

The coil must be bent perpendicularly to the spiral edge by an angle of $\phi + \frac{\pi}{2}$ as illustrated in Fig 2.11. The coordinates of point c, denoted (c_x, c_y), are given by $(r\sin(\theta) + r_b\cos(\phi + \frac{\pi}{2}), r\cos(\theta) + r_b\sin(\phi + \frac{\pi}{2}))$, where r_b is the bending radius of the coil. Point d is the location at which the trim coils begins to follow a constant radius from the origin; it is connected to the point c via a circular arc whose centre of curvature is the point e. The coordinates of the point e, denoted (e_x, e_y), are given by $(c_x + (r_t + r_b)\cos(\phi + \pi), c_y + (r_t + r_b)\sin(\phi + \pi))$, where r_t is the thickness of the trim coil. Point e lies on the line connecting the origin to point d, thereby imposing a constraint on the radius of point d, which defines the parametrising radius of the trim coil:

$$e_x^2 + e_y^2 = (r_d - r_t - r_b)^2 \tag{A.4}$$

This equation is solved numerically to determine r, from which θ is obtained via the spiral angle equation Eq. 1.70. This uniquely determines the coordinates of point b such that the trim coil reaches the desired radius at the centre of the pole face. The edge $\overline{ba}$ corresponds to the bend introduced in Fig. 2.12 and follows the spiral edge of the pole. The coordinates of point a are given by $(b_x + r_b cos(\phi), b_y + r_b sin(\phi))$. An analogous calculation is performed on the entrance side of the magnet, and with the remainder of the trim coil conforming to the pole geometry, the complete coil geometry is fully determined.

Appendix B

Metaopera

Listing B.1: Structure of MetaOpera, the Python code used to generate command files for constructing OPERA models.

```
# ============================================================
# Imports
# ------------------------------------------------------------
# Import operating system tools
# Import file handling and subprocess tools
# Import threading and timing tools
# Import configuration parser
# Import numerical and plotting libraries
# Import metaopera interface
# ============================================================

import os
import shutil
import subprocess
import sys
import threading
import time
from configparser import ConfigParser

import matplotlib.pyplot as plt
import metaopera as mo
import numpy as np


# ============================================================
# Utility functions for OPERA COMI file generation
# ============================================================

def stop(f):
    # Force OPERA execution to halt for debugging
    print("SELECT AUTOUPDATE=YES",file=f)
```

```
    print("STOP",file=f)

def header(f):
    # Write standard COMI header
    print("SELECT AUTOUPDATE=YES",file=f)
    print("CLEAR LOGFILES=yes",file=f)
    print("GUIOPTIONS SET CONSOLEVIEW=YES CONSOLEDOCK=YES   ",file=
       f) #show the OPERA terminal consol
    print("WCS OPTION=UNSET                                 ",file=
       f) #reset coordinate system
    print("SELECT AUTOUPDATE=NO                             ",file=
       f) #turn off display update: makes the construction of the
       model much faster
    print("$CONST #MAXEDGEHDLPTS 512",file=f)

    print("",file=f)

def footer(f):
    # Restore standard COMI closing commands
    print("SELECT AUTOUPDATE=YES                            ",file=
       f)




# ============================================================
# Coil construction routines
# ============================================================

# Generate coil from four bounding trajectories using the 8 node
   brick option in OPERA
# coilshape is a file containing the coordinates of the bounding
   trajectories
# coilshape[:,0],coilshape[:,1],coilshape[:,2] are the x,y,z
   coordinates of the first trajectory etc




def coil_four_trajectory(f,coilshape,drive_label='Default',N=1,cur
   =0.,closed_loop=True):

    pass
    l = len(x)
    if len(y) != l:
        raise RuntimeError("the length of the x and y array don't
           match: these two arrays that define the shape of your
           coil must be of equal length")
```

```
node = []
for i in range(4):
    node.append({'x':np.zeros(l), 'y':np.zeros(l), 'z':np.zeros
        (l)})
node[0]['x'] = coilshape[:,0]
node[0]['y'] = coilshape[:,1]
node[0]['z'] = coilshape[:,2]

node[1]['x'] = coilshape[:,3]
node[1]['y'] = coilshape[:,4]
node[1]['z'] = coilshape[:,5]

node[2]['x'] = coilshape[:,6]
node[2]['y'] = coilshape[:,7]
node[2]['z'] = coilshape[:,8]

node[3]['x'] = coilshape[:,9]
node[3]['y'] = coilshape[:,10]
node[3]['z'] = coilshape[:,11]

if closed_loop:
    lend = l
else:
    lend = l-1

for i in range(lend):
    print("BRICK8 OPTION=NEW -KEEP XP1={:g} YP1={:g} ZP1={:g}
        XP2={:g} YP2={:g} ZP2={:g} XP3={:g} YP3={:g} ZP3={:g}
        XP4={:g} YP4={:g} ZP4={:g} XP5={:g} YP5={:g} ZP5={:g}
        XP6={:g} YP6={:g} ZP6={:g} XP7={:g} YP7={:g} ZP7={:g}
        XP8={:g} YP8={:g} ZP8={:g} INCIRCUIT=NO CIRCUITELEMENT=
        CURD={}/AREA TOLERANCE=TOLERANCE DRIVELABEL={} THETA2=0
        PHI2=0 PSI2=0 XCEN2=0 YCEN2=0 ZCEN2=0 LCNAME='Global
        coordinate system' RXY=1 RYZ=0 RZX=0 SYMMETRY={}
        MODELCOMPONENT=NO".format(
    node[0]['x'][i],node[0]['y'][i],node[0]['z'][i],
    node[1]['x'][i],node[1]['y'][i],node[1]['z'][i],
    node[2]['x'][i],node[2]['y'][i],node[2]['z'][i],
    node[3]['x'][i],node[3]['y'][i],node[3]['z'][i],
    node[0]['x'][np.mod(i+1,l)],node[0]['y'][np.mod(i+1,l)],
        node[0]['z'][np.mod(i+1,l)],
    node[1]['x'][np.mod(i+1,l)],node[1]['y'][np.mod(i+1,l)],
        node[1]['z'][np.mod(i+1,l)],
    node[2]['x'][np.mod(i+1,l)],node[2]['y'][np.mod(i+1,l)],
        node[2]['z'][np.mod(i+1,l)],
```

```
        node[3]['x'][np.mod(i+1,l)],node[3]['y'][np.mod(i+1,l)],
            node[3]['z'][np.mod(i+1,l)],
        cur,drive_label,
        N),file=f)
        print('',file=f)


def coil_box(f,xy,z0,Dz,w,name='coilbox'):
    # Construct a reduced potentail air region, all the conductors
       need to be enclosed in this region in which OPERA solves the
        magnetic field via the Biot-Sacart law

    # xy is the x and y coordinates of the region, z0 the minimum z
        coordinate and Dz the height

    box_xy = np.array(xy)
    l = len(box_xy)

    #Now the following variables can either be scalar or array:
    dum=0.
    list = [w,dum]
    for i in range(len(list)):
        if not (hasattr(list[i], "__len__")): #if z is a scalar,
           turn it into an array of length = len(x)
            list[i] = np.ones(l)*list[i]
        elif len(list[i]) != l:
            raise RuntimeError("the length of your z array doest
               not make sense: z should either be a float (planar
               coil) or have the same length than the x and y
               arrays")
    w,dum = list #cast the list back into the original variables

    for i in range(l):
        v = [xy[np.mod(i+1,l)][0]-xy[i-1][0], xy[np.mod(i+1,l)][1]-
           xy[i-1][1]]
        v = v / (np.linalg.norm(v) + 1e-16)
        vperp = (v[1],-v[0])

        box_xy[i][0] = xy[i][0]-vperp[0]*w[i]
        box_xy[i][1] = xy[i][1]-vperp[1]*w[i]

    xy_surface_z_sweep(f,box_xy,z0,Dz,name=name)


# ============================================================
# Background and wire geometry
# ============================================================
```

```
def set_background(f):
    # Set normal symmetry for xy plane, Translation symmetry for xz
        and yz planes
    print("BACKGROUND OPTION=SET SHAPE=TRIM XYSYMMETRYPLANE=YES "
      "YZSYMMETRYPLANE=YES ROTZNUM=1 ZXSYMMETRYPLANE=YES "
      "EMRXY=NORMMAGN EMRYZ=TANGMAGN EMRZX=TANGMAGN", file=f)

def wireedge(f,x0,y0,z0,x1,y1,z1):
    #Produce wireedge object in OPERA
    print("WIREEDGE Name='WireEdge' X0={:.10f} Y0={:.10f} Z0={:.10f
       } X1={:.10f} Y1={:.10f} Z1={:.10f} EDGETYPE=STRAIGHT".format
       (x0,y0,z0,x1,y1,z1), file=f)



def cover(f,name='surf'):
    # Convert closed wire loop into surface
    print("PICK OPTION=TOGGLE, | PICK PROPERTY=Name LABEL=WireEdge"
       ,file=f)
    print("COMBINE OPERATION=UNION +REGULAR",file=f)
    print("PICK PROPERTY=Name LABEL=WireEdge",file=f)
    print("COVER +REGULAR",file=f)
    print("PICK PROPERTY=Name LABEL=WireEdge",file=f)
    print("RENAME NAME='{}' UNIQUENAME='{}'".format(name,name),file
       =f)

def xy_surface_z_sweep(f,xy,z0,Dz,name='surf'):
    # Create 3D solid by sweeping 2D xy surface in z
    l = len(xy)
    for i in range(0,l-1):
        wireedge(f,xy[i][0],xy[i][1],z0,xy[i+1][0],xy[i+1][1],z0)
    #close the contour and cover
    wireedge(f,xy[-1][0],xy[-1][1],z0,xy[0][0],xy[0][1],z0)
    cover(f,name)
    print("PICK PROPERTY=System LABEL='Sheet face' ",file=f)
    print("SWEEP KEEP=NO  TYPE=DISTANCE REGULAR=NO RIGID=NO
       DISTANCE={} DRAFTTYPE=NONE".format(-Dz),file=f)
    print("",file=f)

def xz_surface_r_sweep(f,xz,angle_deg,name='surf'):
    # Create 3D solid by rotational sweep of 2D surface
    print("WCS OPTION=UNSET",file=f) #reset coordinate system
    print("TRANSFORM OPTION=WCS KEEP=YES TYPE=ROTATE ROTU=1 ROTV=0
       ROTW=0 ANGLE=90",file=f) #reset coordinate system

    l = len(xz)
    for i in range(0,l-1):
        wireedge(f,xz[i][0],xz[i][1],0,xz[i+1][0],xz[i+1][1],0)
```

```
    #close the contour and cover
    wireedge(f,xz[-1][0],xz[-1][1],0,xz[0][0],xz[0][1],0)
    cover(f,name)
    print("WCS OPTION=UNSET",file=f) #reset coordinate system

    print("PICK PROPERTY=System LABEL='Sheet face' ",file=f)
    print("SWEEP KEEP=NO TYPE=ROTATE REGULAR=NO RIGID=NO ROTU=0
        ROTV=0 ROTW=1 POSU=0 POSV=0 POSW=0 ANGLE={} DRAFTTYPE=NONE".
        format(angle_deg),file=f)
    print("",file=f)


# ============================================================
# Boolean operations and material assignments
# ============================================================

def combine(f,list_names,name=None,operation="UNION",option="+
    REGULAR"):
    # Boolean union/intersection/subtraction of objects in OPERA
    if(name is None): #if no body name is given, use the name of
        the first body in the list
        name = list_names[0]
    print("FILTER TYPE=BODY",file=f)
    for body in list_names:
        print("PICK OPTION=ADD PROPERTY=UniqueName LABEL='{}'".
            format(body),file=f)
    print("COMBINE OPERATION={} {}".format(operation,option),file=f
        )
    print("PICK PROPERTY=UNIQUENAME LABEL='{}'".format(list_names
        [0]),file=f)
    print("RENAME NAME='{}' UNIQUENAME='{}'".format(name,name),file
        =f)

def celldata(f,name,material='Air',potential='Default',elementtype=
    'quadratic',meshtype='NONE',DSL=0,MS=1000.0):
    # Assign material, mesh and solver properties
    print("PICK OPTION=ADD PROPERTY=UniqueName LABEL='{}'".format(
        name),file=f)
    print("CELLDATA OPTION=MODIFY MATERIALLABEL='{}' POTENTIAL={}
        ELEMENTTYPE={} LEVEL={} SIZE={:g} ELEMSHAPEPREF={}".format(
        material,potential,elementtype,DSL,MS,meshtype),file=f)


def cylinder(f,height,radius,thickness=0.,name="Cylinder"):
    # Generate cylindrical solid, useful for setting for air
        regions around the model
    # If positive thickness given, thickness goes inwards. If
        negative, thickness goes outwards
```

```
    print("CYLINDER Name='{}' X0=0 Y0=0 Z0=0 X1=0 Y1=0 Z1={} +TUBE
        THICKNESS={} SHAPECONTROL=SIMPLE MAJORRADIUS={} MINORRADIUS
        ={} TOPRADIUS={} SIDES=2".format(name,height,thickness,
        radius,radius,radius),file=f)
    print("",file=f)

def rotate_all(f,angle_deg,option='APPLY'):
    # Rotate full model
    print("SELECT AUTOUPDATE=YES",file=f)
    print("FILTER TYPE=BODY",file=f)
    print("PICK ALL",file=f)
    print("TRANSFORM OPTION={} KEEP=NO TYPE=ROTATE ROTU=0 ROTV=0
        ROTW=1 ANGLE={} COUNT=1".format(option,angle_deg),file=f)
    print("SELECT AUTOUPDATE=NO",file=f)


# ============================================================
# OPERA model setup and solver execution
# Important the save the model with mesh details
# ============================================================

def setup_model(f, filename='temp', bhfile=None, UNITS = "MM",
    TOL_ANALYSIS = 1.0E-4, TOL_GEOMETRY = 1.0E-4, TOL_FILL = 1.0E-4,
    MS=1000.):
    # Define BH data
    # Set nonlinear solver options
    # Create mesh
    # Save model with mesh
    if(os.path.exists(f"{filename}.opc")):
        os.remove(f"{filename}.opc")
    if(os.path.exists(f"{filename}.op3")):
        os.remove(f"{filename}.op3")

    config_parser = ConfigParser()
    config_parser.read(os.path.join(os.path.dirname(mo.__file__),"
        config.ini"))
    opera_config = config_parser['OPERA']

    print("MATERIALS PICK 'steel'",file=f)

    if(bhfile is None): #use OPERA default BH data
        print("MATERIALS OPTION=MODIFY MULINEARITY=NONLINEAR
            MUANISOTROPY=ISOTROPIC BH='Default'",file=f)
    else:
        print("BHDATA OPTION=NEW  LABEL='BHDATA'",file=f )
        print(f"BHDATA OPTION=LOAD LABEL=BHDATA FILE={bhfile}",
            file=f)
```

```
        print("MATERIALS OPTION=MODIFY MULINEARITY=NONLINEAR
            MUANISOTROPY=ISOTROPIC BH='BHDATA'",file=f)

    print("ANALYSISDATA OPTION=LOAD PROGRAM=TOSCAMAGN",file=f)

    print(f"ANALYSISDATA OPTION=SET PROGRAM=TOSCAMAGN LINEAR=NO
        NLITERTYPE=NEWTON NITERATIONS=51 TOLERANCE={TOL_ANALYSIS} HX
        =0 HY=0 HZ=0 RHS=ADAPTIVE POTENTIALCUT=YES",file=f)



    print(f"SOLVERS SOLVENOW=NO SAVEMODEL=YES, | SOLVERS OPTION=
        TEST FILE='{filename}' UNITS={UNITS} ELEMENT=MIXED SURFACE=
        CURVED | COMMENT CLEAR=YES TYPE=DBTITLE  | MODEL CREATE |
        MESH SIZE={MS} NORMALTOL=30.0 SURFACETOL=0.0 TOLERANCE={
        TOL_GEOMETRY} TYPE=PREFERTETRA | FILL TOL={TOL_FILL} |
        SOLVERS OPTION=OVERWRITE",file=f)

    print(f"SAVE OPTION=MESH FILE='{filename}'",file=f)
    # Saves the meshed model into opcb files

    if(os.environ['OPDEBUG'].upper() != 'DEBUG'):

        print("//close the modeller",file=f)
        print("END",file=f)
        print("yes",file=f)



class Spinner:
    # Terminal progress indicator, spins and indicates the progress
        is happening
    busy = False
    delay = 0.1

    @staticmethod
    def spinning_cursor():
        while 1:
            for cursor in '|/-\\': yield cursor

    def __init__(self, delay=None):
        self.spinner_generator = self.spinning_cursor()
        if delay and float(delay): self.delay = delay

    def spinner_task(self):
        while self.busy:
            sys.stdout.write(next(self.spinner_generator))
            sys.stdout.flush()
```

```
            time.sleep(self.delay)
            sys.stdout.write('\b')
            sys.stdout.flush()

    def __enter__(self):
        self.busy = True
        threading.Thread(target=self.spinner_task).start()

    def __exit__(self, exception, value, tb):
        self.busy = False
        time.sleep(self.delay)
        if exception is not None:
            return False

def run(exec,file,flag="",live_term_output=False):
    # Generic subprocess launcher
    ### exec can be the path to the 3D modeller, 3D postprocessor,
       the 2D modeller, etc
    DEVNULL = open(os.devnull, 'wb')
    print(f'Now runing {exec} {flag} {file}', end=' ')
    with Spinner():
        if(live_term_output):
            with subprocess.Popen("{} {} {}".format(exec,flag,file)
                .split(), stdout=subprocess.PIPE, stderr=subprocess.
                PIPE, bufsize=1, universal_newlines=True) as p:
                 for line in p.stdout:
                     print(line, end='')
        else:
            #I chose to set stdout=DEVNULL since the modeller
                outputs nothing, and to silence the annoying
                warnings I keep getting from Qt
            with subprocess.Popen("{} {} {}".format(exec,flag,file)
                .split(), stdout=DEVNULL, stderr=subprocess.PIPE,
                bufsize=1, universal_newlines=True) as p:
                 for line in p.stderr:
                     print(line, end='')

    print("") #skip a line in the terminal




# ============================================================
# End of code
# ============================================================
```

Listing B.2: A seperate configuration file is required to define executable paths (where OPERA is installed in the machine)

```
[OPERA]
# command to execute the different modules of OPERA
MODELLER = path\operafea-modeller
SOLVER   = path\operafea-op3solve
POST     = path\operafea-post
OPDEBUG  = no
#OPDEBUG options are 'no' to run without visual UI or 'debug' to
   run the visual UI for debugging purposes
```

Appendix C

Perpedicular distance between spirals

Given two spirals defined by distinct values of θ_0,

$$\begin{aligned} \text{Spiral 1} : \theta^1 &= \theta_0^1 + \tan\xi \ln\frac{r}{r_0} \\ \text{Spiral 2} : \theta^2 &= \theta_0^2 + \tan\xi \ln\frac{r}{r_0} \end{aligned} \tag{C.1}$$

the objective is to determine the opening angle ϕ between the two spirals, given a perpendicular separation a at an arbitrary point along one of them. It is convenient to decompose the opening angle ϕ into two separate contributions $\varepsilon + \alpha$; the full geometry of this problem is illustrated in Fig. C.1. Without loss of generality, the reference point may be chosen at r_0 on Spiral 2. By evaluating the distance $\bar{12}$ between the two spirals at this point, the following relation can be obtained:

$$\begin{aligned} a\cos\xi &= r_0 \sin\alpha \\ \rightarrow \alpha &= \arcsin\frac{a}{r_0}\cos\xi \end{aligned} \tag{C.2}$$

The angle ε is simply $\theta - \theta_0^1$,and may therefore be written as:

$$\varepsilon = \tan\xi \ln\frac{r}{r_0} \tag{C.3}$$

The total opening angle between the two spirals, $\varepsilon + \alpha$ is then given by:

$$\varepsilon + \alpha = \tan\xi \ln\frac{r}{r_0} + \arcsin\frac{a}{r_0}\cos\xi \tag{C.4}$$

Making use of Eq. C.2, the following identity is obtained:

$$\begin{aligned} a\cos\xi &= r_0\sin\alpha \\ \rightarrow \frac{a}{r_0}\cos\xi &= \sin\alpha \\ \rightarrow \frac{a^2}{r_0^2}\cos^2\xi &= \sin^2\alpha \\ \rightarrow 1-\frac{a^2}{r_0^2}\cos^2\xi &= 1-\sin^2\alpha \\ \rightarrow \sqrt{1-\frac{a^2}{r_0^2}\cos^2\xi} &= \cos\alpha \end{aligned} \tag{C.5}$$

The radius r may be expressed in terms of a, r_0 and ξ as:

$$\begin{aligned} r &= a\sin\xi + r_0\cos\alpha \\ \rightarrow r &= a\sin\xi + r_0\sqrt{1-\frac{a^2}{r_0^2}\cos^2\xi} \end{aligned} \tag{C.6}$$

Substituting Eq. C.6 into Eq. C.7, the final expression for the opening angle is obtained:

$$\phi = \varepsilon + \alpha = \tan\xi \ln\left(\frac{a}{r_0}\sin\xi + \sqrt{1-\frac{a^2}{r_0^2}\cos^2\xi}\right) + \arcsin\left(\frac{a\cos\xi}{r_0}\right) \tag{C.7}$$

Figure C.1: Labeled diagram of the geometry between two spirals.

Appendix D

Set up model with Metaopera

Listing D.1: Model generation script utilising functions defined in MetaOpera to perform geometry construction, coil definition, and solver execution, producing the corresponding OPERA database file.

```
import os
import sys
import numpy as np
from configparser import ConfigParser
import metaopera as mo
import glob
import re

# ============================================================
# INITIALISATION
# ============================================================

sys.path.insert(0, os.path.abspath(os.path.join(os.getcwd(), os.
    pardir)))
sys.path.insert(0, os.path.abspath(os.path.join(os.path.join(os.
    getcwd(), os.pardir), os.pardir)))

config_parser = ConfigParser()
config_parser.read(os.path.join(os.path.dirname(mo.__file__), "
    config.ini"))
opera_config = config_parser["OPERA"]

os.environ["OPDEBUG"] = opera_config["OPDEBUG"]
modeller = opera_config["MODELLER"]
solver = opera_config["SOLVER"]
post = opera_config["POST"]

if os.environ["OPDEBUG"] != "debug":
    os.environ["VFGRAPHICS"] = "none"
```

```
def getKey(filename):
    file_text_name = os.path.splitext(os.path.basename(filename))
    file_last_num = os.path.basename(file_text_name[0]).split('_')
    return int(file_last_num[-1])


# ============================================================
# PARAMETERS
# ============================================================


#Different mesh sizes for different regions of the model

ms_yoke = 50
ms_gfr = 15
ms_gfr_gfr = 3
ms_coilbox = 50
ms_pole = 15
ms_clamp = 25
ms_magbox = 200
ms_background = 500


bkgd_h = 2.0 * 1000
gap_height = 47.5 + 0.5
pole_height = 405
yoke_height = pole_height + 600
GFR_height = gap_height / 2


#Rotates the magnet model into the middle of the defined OPERA
    region

rotate_by = -31.5


# ============================================================
# GEOMETRY INPUTS (EXTERNALLY DEFINED)
# ============================================================
# All arrays below must be defined externally:
# pole_top, F_pole, midyoke, outer_midyoke, inyoke,
# outer_inyoke, coil_box, GFR_of_GFR,
# clamp_en, clamp_ex, clamp extensions, clamp cuts,
# r_pole
#Current_in_coil must also be defined externally
```

```
# The entire return yoke is separated into different sections made
    out of pole_top, midyoke, outer_midyoke, inyoke,
# outer_inyoke etc...

#clamp extensions connects the clamps magnetically with the rest of
     the magnet

#r_pole defined the radius of the GFR

#Closer to the magnet, a region with an even smaller mesh size is
    defined: GFR_of_GFR



#Example on loading in the pole geometry

pole_top=np.loadtxt('pole_top_geometry.txt')









# ============================================================
# FILE HANDLING
# ============================================================

filename = "LhARA_example"

#Where the coordinate files of the coils are stored
path = "LhARA/coils"

#Names of the coordinate files of the coils end with their
    corresponding number. For example, LhARA_coil_0, LhARA_coil_1
    etc...

#Sorts the coordinate file of the coils in ascending order so the
    correct label and current can be specified later



all_coordinates_file = glob.glob(path + "/*.txt")
all_coordinates_file.sort(key=getKey)
```

```
# ============================================================
# BUILD COMI FILE
# ============================================================

with open(f"{filename}.comi", "w") as f:

    mo.header(f)

    # ------------------------------------------------------------
    # MAGNET STRUCTURE
    # ------------------------------------------------------------

    mo.xy_surface_z_sweep(f, midyoke, pole_height, yoke_height -
        pole_height, name="midyoke")
    mo.xy_surface_z_sweep(f, outer_midyoke, pole_height,
        yoke_height - pole_height, name="outer_midyoke")
    mo.xy_surface_z_sweep(f, inyoke, 0, yoke_height, name="inyoke")
    mo.xy_surface_z_sweep(f, outer_inyoke, 0, yoke_height, name="
        outer_inyoke")
    mo.xy_surface_z_sweep(f, pole_top, pole_height, yoke_height -
        pole_height, name="pole_top")

    mo.xy_surface_z_sweep(f, pole_top, gap_height, pole_height -
        gap_height, name="F_pole")

    mo.coil_box(f, coil_box, 0, pole_height, 0, name="coilbox")


    # ------------------------------------------------------------
    # CLAMPS
    # ------------------------------------------------------------

    mo.xy_surface_z_sweep(f, clamp_en, 70, -(yoke_height - 70),
        name="clamp_en")
    mo.xy_surface_z_sweep(f, clamp_en_connection, pole_height,
        yoke_height - pole_height, name="clamp_en_connection")

    mo.xy_surface_z_sweep(f, clamp_ex, 52, yoke_height - 52, name="
        clamp_ex")
    mo.xy_surface_z_sweep(f, clamp_ex_extension, 52, -(pole_height
        - 52), name="clamp_ex_extension")
    mo.xy_surface_z_sweep(f, clamp_ex_connection, pole_height, (
        yoke_height - pole_height), name="clamp_ex_connection")


    # ------------------------------------------------------------
    # CUTS to modify the clamp aperture of entrance and exit clamp
        separately
```

```
# ------------------------------------------------------

mo.xz_surface_r_sweep(f, en_clamp_gap_shape, 720 / N, name="
    en_cut_clamp")
mo.xz_surface_r_sweep(f, ex_clamp_gap_shape, 720 / N, name="
    ex_cut_clamp")

mo.combine(f, ["clamp_en", "en_cut_clamp"], operation="SUBTRACT
    ")
mo.combine(f, ["clamp_ex", "ex_cut_clamp"], operation="SUBTRACT
    ")


# ------------------------------------------------------
# MATERIAL AND MESH SIZE ASSIGNMENT
# ------------------------------------------------------

mo.celldata(f, "inyoke", material="steel", potential="Total",
    elementtype="quadratic", DSL=30, MS=ms_pole)
mo.celldata(f, "outer_inyoke", material="steel", potential="
    Total", elementtype="quadratic", DSL=30, MS=ms_pole)
mo.celldata(f, "outer_midyoke", material="steel", potential="
    Total", elementtype="quadratic", DSL=30, MS=ms_pole)

mo.celldata(f, "midyoke", material="steel", potential="Total",
    elementtype="quadratic", DSL=30, MS=ms_yoke)
mo.celldata(f, "pole_top", material="steel", potential="Total",
     elementtype="quadratic", DSL=30, MS=ms_yoke)

mo.celldata(f, "coilbox", material="Air", potential="Reduced",
    elementtype="quadratic", DSL=40, MS=ms_coilbox)

mo.celldata(f, "clamp_en", material="steel", potential="Total",
     elementtype="quadratic", DSL=100, MS=ms_clamp)
mo.celldata(f, "clamp_ex", material="steel", potential="Total",
     elementtype="quadratic", DSL=100, MS=ms_clamp)


# ------------------------------------------------------
# COILS

# ------------------------------------------------------
number_of_conductors = len(all_coordinates_file)

drive_label=[0,1,2,3,4,5,6,7,8,9,1,2,3,4,5,6,7,8,9]


for i in range(number_of_conductors):
```

```
    LhARA_drive_label = np.loadtxt(all_coordinates_file[i])
    #Labels the coils based on the power supply connections.
    #Here the two pole face layers share the same current
       setting


    coilshape = np.loadtxt(all_coordinates_file[i])

    mo.coil_four_trajectory(
        f,
        coilshape[:,0], coilshape[:,1], coilshape[:,2],
        coilshape[:,3], coilshape[:,4], coilshape[:,5],
        coilshape[:,6], coilshape[:,7], coilshape[:,8],
        coilshape[:,9], coilshape[:,10], coilshape[:,11],
        LhARA_drive_label,
        N=1,
        cur=current_in_coil[i],
        closed_loop=True
    )


print(f"TRANSFORM OPTION=APPLY TYPE=ROTATE ROTW=1 ANGLE={
   rotate_by}", file=f)


# ------------------------------------------------------
# AIR REGIONS
# Several air regions needs to be defined, smaller meshes where
    the field values matters
# Differences between the mesh sizes of surrounding regions
   cannot be too large, so several air boxes are needed
# ------------------------------------------------------

mo.cylinder(f, GFR_height, "r_pole", name="GFR")
mo.xy_surface_z_sweep(f, GFR_of_GFR, 0, 10, name="GFR_of_GFR")

mo.combine(f, ["GFR_of_GFR", "GFR"], operation="INTERSECTION")

mo.cylinder(f, yoke_height + 160, 5000, thickness=3000, name="
   magbox")
mo.cylinder(f, bkgd_h, 7000, thickness=5500, name="background")


print("BACKGROUND OPTION=SET SHAPE=TRIM XYSYMMETRYPLANE=YES
   YZSYMMETRYPLANE=YES ZXSYMMETRYPLANE=YES", file=f)
```

```
    mo.footer(f)

    mo.setup_model(f, filename=filename, bhfile="../BH_data/
        COMSOL_BH.bh", MS=ms_background)


# ============================================================
# RUN OPERA modeller to set up the model
# RUN OPERA solver to solve for the B field
# ============================================================

mo.run(modeller, os.path.join(os.getcwd(), f"{filename}.comi"),
    flag="-local")
mo.run(solver, os.path.join(os.getcwd(),f'{filename}.op3'), flag ="
    -nthreads=16",live_term_output=True)
```

Appendix E

Current optimisation procedure

Listing E.1: Coil current optimisation script based on the SciPy nonlinear least-squares algorithm. Additional simulations are appended to the existing OPERA database rather than creating a new one, significantly reducing computation time by reusing previously computed information.

```
import os
import sys
import numpy as np
import pandas as pd
import glob
import re
from configparser import ConfigParser
from scipy.optimize import least_squares
import metaopera as mo


# ============================================================
# INITIALISATION
# ============================================================

sys.path.insert(0, os.path.abspath(os.path.join(os.getcwd(), os.
    pardir)))
sys.path.insert(0, os.path.abspath(os.path.join(os.path.join(os.
    getcwd(), os.pardir), os.pardir)))

config_parser = ConfigParser()
config_parser.read(str(os.path.dirname(mo.__file__) + "/config.ini"
    ))
opera_config = config_parser["OPERA"]

os.environ["OPDEBUG"] = opera_config["OPDEBUG"]

modeller = opera_config["MODELLER"]
solver = opera_config["SOLVER"]
```

```
post = opera_config["POST"]

coil_current = np.loadtxt("coil_current.txt")

if os.environ["OPDEBUG"] != "debug":
    os.environ["VFGRAPHICS"] = "none"


# ============================================================
# UTILITY FUNCTIONS
# ============================================================

def getKey(filename):
    file_text_name = os.path.splitext(os.path.basename(filename))
    file_last_num = os.path.basename(file_text_name[0]).split('_')
    return int(file_last_num[-1])

### Jacobian needs to be externally computed
### Can be done using the same code as below but manually varing
   the current by a desired amount
def jacobian_callable(params):
    return np.loadtxt('LhARA_jacobian.txt')


def scaling_B(BL0, r0, r, k):
    return BL0 * (r / r0) ** k+1


# ============================================================
# FIELD ANALYSIS FUNCTION
# ============================================================

def Integrated_field(df_constant_R):

    column_names = ['R','integrated_B_F']
    empty_df = pd.DataFrame(columns=column_names)

    number_of_radii = df_constant_R['R'].nunique()
    number_of_theta = df_constant_R['Theta'].nunique()

    for i in range(number_of_radii):

        r_constant = df_constant_R.iloc[number_of_theta*i:
           number_of_theta*i+number_of_theta][::-1].reset_index(
           drop=True)
        radius = r_constant['R'][0]
```

```
        total_integral = np.trapz(r_constant['Bz'], x=np.radians(
            r_constant['Theta']))

        data_to_append = {
            'R': [radius],
            'integrated_B_F': [total_integral]
        }

        empty_df = pd.concat([empty_df, pd.DataFrame(data_to_append
            )], ignore_index=True)

    return empty_df


# ============================================================
# OPTIMISATION
# ============================================================

itnumber = 2


def cost_function(params):

    global itnumber

    filename = f"LhARA_example"

    # ------------------------------------------------------
    # Open the previously created OPERA model, vary the currents
       accordingly
    # ------------------------------------------------------
    with open(f"{filename}.comi", "w") as f:

        print(f"LOAD OPTION=NEW FILE='path/{filename}.opcb'", file=
            f)

        drive_label = [0,1,2,3,4,5,6,7,8,9]

        #Scale the currents in the coils accordingly to the
            residual using the Jacobian externally computed

        for i in range(len(drive_label)):

            Drive_label = f"Coil_{drive_label[i]}"
            scale = params[i] / coil_current[i]

            print(f"DRIVE PICK '{Drive_label}'", file=f)
            print(f"DRIVE OPTION=MODIFY SCALE={scale}", file=f)
```

```
        print("DRIVE UNPICK", file=f)

    print(f"SOLVERS SOLVENOW=NO SAVEMODEL=YES OPTION=ADD |
       COMMENT TEXT='Current_optimised_it_{itnumber}' | SOLVERS
        FILE='path/{filename}.op3'", file=f)


# ------------------------------------------------------------
# RUN OPERA
# ------------------------------------------------------------

mo.run(modeller, os.path.join(os.getcwd(), f"{filename}.comi"),
    flag="-local")
mo.run(solver, os.path.join(os.getcwd(), f"{filename}.op3"),
   flag="-nthreads=16", live_term_output=True)


# ------------------------------------------------------------
# POST-PROCESSING (B FIELD MAP on the midplane)
# Define externally Fieldmap_coordinates.txt which contains the
    coordinates where the field should be evaluated #r #theta #
   x #y.

# ------------------------------------------------------------

with open('btheta.comi', 'w') as f:
    print(f"ACTIVATE CASE={itnumber} MODELSYMMETRY=DATABASE
       FILE='{filename}.op3'| LOAD", file=f)
    print(f"$open 1 path/Field_map_{filename}_{itnumber}.txt
       overwrite", file=f)
    print("$open 2 Fieldmap_coordinates.txt read", file=f)
    print("$format 1 fixed 15 5", file=f)
    print("$assign 1 1 1 1 1", file=f)
    print("$do #i 1 row_of_Fieldmap_coordinates", file=f)
    print("$read 2 #r #theta #x #y", file=f)
    print("POINT XP=#x  YP=#y  ZP=0  COMP=Bz", file=f)
    print("$write 1 #r      #theta  #x #y Bz", file=f)
    print("$end do", file=f)
    print("$close 1", file=f)
    print("$close 2", file=f)
    print("END", file=f)
    print("YES", file=f)


# ------------------------------------------------------------
# ANGULAR PROCESSING
# ------------------------------------------------------------
```

```
    mo.run(post, os.path.join(os.getcwd(), 'btheta.comi'), flag="-
        local")

    df_optimise = pd.read_csv(f"Field_map_{filename}_{itnumber}.txt
        ", header=None, delim_whitespace=True)
    df_optimise.columns = ['R', 'Theta', 'x', 'y', 'Bz']

    df_effective = Integrated_field(df_optimise)


    # ------------------------------------------------------------
    # RESIDUALS
    # ------------------------------------------------------------

    residuals = (df_effective['integrated_B_F'] * df_effective['R']
        / 1000)

    print(params, np.sum(residuals**2))

    itnumber += 1

    return residuals


# ================================================================
# OPTIMISATION CALL
# ================================================================

params = np.loadtxt('coil_current.txt')



#Using least_squares functions from scipy to do non-linear least
    squares optimisation

result = least_squares(cost_function, params, jac=jacobian_callable
    , method='lm', ftol=1e-9)



optimized_params = result.x

print(optimized_params)
```