

Durham E-Theses

*Simulations of high-field superconductors using
time-dependent Ginzburg–Landau theory with
accurate material parameters*

C.W.W. Haddon

How to cite:

Haddon, C.W.W. (2025) Simulations of high-field superconductors using time-dependent Ginzburg–Landau theory with accurate material parameters. Doctoral thesis, Durham University.

Use policy



This work is licensed under a [Creative Commons Attribution Share Alike 4.0](https://creativecommons.org/licenses/by-sa/4.0/)

Simulations of high-field superconductors
using time-dependent Ginzburg–Landau
theory with accurate material parameters

C.W.W. Haddon

Ph.D. Thesis

Durham University, Department of Physics

2025

Contents

Abstract	1
1 Introduction	3
1.1 Superconducting materials	3
1.2 Flux pinning and the critical current	5
1.3 Magnetic confinement fusion	7
1.4 Thesis overview	9
2 Superconductivity theory	11
2.1 Ginzburg–Landau theory	12
2.1.1 Free energy	13
2.1.2 Equilibrium equations	15
2.1.3 Homogeneous solution	18
2.1.4 London limit	18
2.1.5 Thermodynamic critical field	19
2.1.6 Flux quantisation	20
2.1.7 The coherence length and the proximity effect	21
2.1.8 Vortices	22
2.1.9 Abrikosov lattice	23
2.1.10 Surfaces	27
2.1.11 De-pairing current	27
2.1.12 Characteristic units	28
2.1.13 Time dependence	29
2.1.14 Pinning	30
2.2 BCS theory	31
2.2.1 The BCS system	31

2.2.2	Mean-field theory	32
2.2.3	Formal solution by Bogoliubov transformation	33
2.2.4	Expressions for Ginzburg–Landau parameters	34
3	Computational method	37
3.1	Discrete exterior calculus	38
3.1.1	Cellular complexes, chains, and discrete forms	38
3.1.2	Discrete exterior derivative	40
3.1.3	Dual complex and discrete Hodge star	40
3.1.4	Discrete tensors	42
3.1.5	Discrete bundle-valued forms	43
3.1.6	Discrete covariant derivative	44
3.1.7	Non-trivial bundles	45
3.1.8	Boundary conditions	47
3.2	Multigrid Ginzburg–Landau solver	48
3.2.1	Discrete Ginzburg–Landau equations	48
3.2.2	Gauss–Seidel relaxation	50
3.2.3	Non-linear multigrid	51
3.2.4	Rectilinear meshes	54
3.2.5	Transfer operators for real-valued forms	57
3.2.6	Transfer operators for the wavefunction	59
3.2.7	Fixing the average potential	60
3.2.8	Implementation summary	62
3.2.9	Critical current determination	63
3.2.10	Numerical experiments	65
3.2.11	Summary and future work	68
4	Preliminary simulations of REBCO with APCs	73
4.1	Computational method	74
4.2	Results	76
4.2.1	Effect of nanorod density	76
4.2.2	Effect of nanorod splay	78
4.2.3	Effect of effective mass	79
4.3	Conclusion	79

5 Simulations of niobium–titanium	81
5.1 Elementary pinning force	82
5.2 Metallurgy and microstructure	84
5.3 Ginzburg–Landau coefficients	85
5.4 Computational system	90
5.5 Anisotropy of the critical field	94
5.6 Critical current density	97
5.7 Conclusion	102
6 Conclusions and future work	111
Appendix	117
Supplementary material	117
Fusion CDT courses	118
Publications	118
Acknowledgements	119
Bibliography	121

Abstract

This thesis investigates the mechanisms which limit the critical current density in high-field superconductors using simulations based on time-dependent Ginzburg–Landau (TDGL) theory. The primary contributions are a fast and general method for solving the TDGL equations in large three-dimensional systems, a preliminary study of an anisotropic coated conductor with nanorod artificial pinning centres, and a study of niobium–titanium using a material model with Ginzburg–Landau (GL) coefficients which are derived from the measured properties of single-phase samples.

The solver uses a geometric multigrid method to solve the backward Euler discretised TDGL equations in the London gauge. It does not assume the frozen-field limit, all of the coefficients can vary in space, and the simulation time scales well with system size and does not depend strongly on the parameters. This allows a wider class of materials to be studied than previous solvers.

In the preliminary study of a system representing a REBCO coated conductor, it was assumed that the pinning landscape consists of resistive cylinders representing nanorod artificial pinning centres embedded in an anisotropic superconducting matrix. The effects of the density, resistivity, and splay of the nanorods were investigated. It was concluded that surface effects in the vicinity of matrix-pin interfaces can play an important role for dense, highly resistive nanorods, and that this might affect the observed critical fields in real materials.

For the model of niobium–titanium, the ratios of the GL coefficients α and β were calculated for the two phases using the measured properties of single-phase samples and GLAG theory. Microstructures were generated procedurally to match the precipitate thicknesses and titanium volume fractions seen in transmission electron microscope images of real wires. Critical field anisotropy simulations were run with a range of effective mass values and a comparison with experimental

data implied that the ratio of the effective mass in the titanium precipitates to that of the niobium–titanium matrix is about 2.

Critical current simulations were also run over a range of effective mass values. The magnitude of the high-field critical current density indicated that the effective mass ratio is closer to 4 than 2, but this comparison is sensitive to the precise values of the GL coefficients in a way that the critical field anisotropy is not. Given a 10% error in the calculated penetration depth for niobium–titanium, the critical current data are consistent with an effective mass of 2. It is concluded that the conventional understanding of pinning in niobium–titanium, that fluxons are pinned primarily by ribbon-shaped titanium precipitates, is consistent with a Ginzburg–Landau theoretic model where the pinning force primarily arises due to the higher carrier effective mass in titanium compared with that of niobium–titanium.

This work demonstrates that simulations based on TDGL theory can be used to study industrial superconducting materials, that using accurate coefficients is essential for properly capturing the mechanisms which limit the critical current, and that the carrier effective mass, which has been historically neglected in theoretical and computational studies of pinning, is a key parameter in niobium–titanium. This is a step in making TDGL simulations a more quantitative tool for critical current prediction.

Chapter 1

Introduction

1.1 Superconducting materials

Superconductivity is a phase of matter which occurs at low temperatures in materials with free charge carriers. It is characterised by the condensation of the charge carriers which leads to striking phenomena like the flow of electric current with no resistance, quantisation of magnetic flux, stable levitation due to flux pinning, and perfect diamagnetism of pure samples. Thirty-one of the sixty-two stable metallic elements are known to exhibit superconductivity at ambient pressure [1], and more do so at high pressures [2]. Even hydrogen is thought to be superconducting in its high-pressure metallic phase [3], [4]. Superconducting phases also occur in more extreme systems such as the interiors of neutron stars, where nuclear matter condenses and the supercurrent is carried by proton pairs [5]. Finally, the standard model itself contains a phase which exhibits characteristics of superconductivity. In the abelian Higgs model, the Higgs field is scalar field in a condensed state with a positive expected magnitude much like the order parameter of a superconducting phase transition [6]. In everyday metallic superconductors, the magnetic field acquires an effective mass and hence a finite range through its coupling to the condensate, which is exactly how the other elementary gauge fields acquire mass through their coupling to the Higgs field [7]. Topological defects like the vortices which are key to the performance of superconducting magnets may also be present in some of the elementary fields of nature [8].

We don't often encounter materials in the superconducting state because the

conditions on the surface of Earth happen to be unfavourable. Most metallic elements are only superconducting at temperatures lower than the boiling point of helium at atmospheric pressure. Of the pure elements, niobium has the highest transition temperature, 9.26 K [9], and certain alloys have slightly higher transition temperatures [1]. In these materials, the superconducting phase differs from the ordinary metallic phase in that the occupation of Landau quasiparticle orbitals with opposite spin and momentum are correlated. These correlations arise due to an effective attraction between quasiparticles of opposite spin and momentum which is mediated by phonons. Since this interaction is weak, only a small amount of thermal energy is required to disrupt the correlations, and therefore the transition temperatures are low. These materials are known as low-temperature superconductors (LTS).

Nuclear magnetic resonance imaging (MRI) is the main commercial application of superconductors [10] and almost all MRI devices use magnets wound from superconducting wire containing niobium–titanium, a ductile LTS alloy which can be used to build magnets reaching about 7 T at 4.2 K [11], or up to 11.7 T when cooled using superfluid helium [12]. Helium is the only suitable liquid cryogen for this material, since its transition temperature is less than the boiling point of hydrogen, neon, nitrogen, and every other substance. The second most important low-temperature superconductor, which is currently used in research magnets and will soon be used in particle accelerators [13] and fusion reactors [14], is Nb_3Sn , a brittle intermetallic compound. This material can reach higher current densities and magnetic fields but is more expensive and has worse mechanical properties than niobium–titanium. Because the superconductor is brittle, magnets are usually wound from a copper–niobium–tin composite and then heated so that the elements of the composite react to form Nb_3Sn .

Another, increasingly important, class of superconductors are the high-temperature superconductors (HTS), particularly the cuprates, materials with copper oxide planes in which the current preferentially flows, and more particularly the rare-earth barium copper oxides (REBCO), a family of anisotropic ceramics which share a crystal structure. In these materials the attractive interaction between electrons has a different mechanism which allows superconductivity to persist to temperatures of up to 93 K [15], higher than the boiling point of liquid nitrogen. Not only do HTS materials have higher transition temperatures, they also remain superconducting in higher magnetic fields and support higher current

densities than LTS materials. However, these materials are less widely used for a few practical reasons: they are expensive, they degrade easily, and they are hard to make into round wires. HTS conductors require specialist manufacturing procedures because the superconducting properties are sensitive to imperfections in the crystal structure and stoichiometry. For example, REBCO conductors are usually manufactured as a tape, where various buffer layers are deposited on a substrate to ensure that the top superconducting layer grows with a uniform crystallographic orientation. This is then coated in silver to reduce oxygen diffusion and preserve the optimal stoichiometry. Finally the whole assembly is coated in copper so that external electrical connections can be made. The resulting tape must not be heated for too long [16], put under strain [17], or exposed to water, lest the superconducting properties be permanently degraded.

In summary, although superconducting materials are many and varied, only a few are of present interest for applications in engineering. Pure elements like niobium, and certain alloys, are often used in electronic applications such as superconducting quantum interference devices¹ [20], which are used to make sensitive measurements of magnetic fields, and in superconducting quantum computers [21]. Niobium–titanium and Nb_3Sn are used in high-field magnets for nuclear magnetic resonance, particle accelerators, and magnetic confinement fusion. Other materials which have superior superconducting properties and are expected to see future application in high-field magnets include the LTS material MgB_2 [22], the HTS REBCO materials, and the cuprate bismuth–strontium–calcium–copper oxide [23].

1.2 Flux pinning and the critical current

When a superconductor is subject to a small magnetic field, lossless currents are induced to flow near the surface which prevent the field from penetrating into the interior. This is known as the Meißner–Ochsenfeld effect [24]. When the field is sufficiently large it becomes energetically favourable to form topological vortices in the superconducting condensate (as long as it is a so-called type-II

¹This technology was used in Gravity Probe B to measure the geodetic and Lense–Thirring effects in orbit of Earth [18], confirming the existence of a ‘gravitational effect which is analogous to electrodynamic induction’ [19].

superconductor, which all high-field superconductors are²). The condensate is depleted in the core of a vortex, which, in a three-dimensional superconductor, has the form of a long, thin tube whose axis is aligned with the magnetic field. A lossless electrical current flows around the core of each vortex so that each carries a quantum of magnetic flux. The formation of vortices is a mechanism by which a superconductor can accommodate some magnetic flux in its interior. The vortices are microscopic in radius and huge in number. When a flux density of 1 T penetrates a superconductor, about 500 million vortices pass through every square millimetre. The vortex radius depends on the material and temperature but it is typically on the order of nanometres.

If a superconductor in a uniform external magnetic field also carries an overall electrical transport current, this interacts with the circulating vortex currents resulting in a force on each vortex causing it to move in the direction which is perpendicular to the applied current and the magnetic field. Since this motion is dissipative, the superconductor develops an electrical resistance even though the superconducting condensate is still present throughout most of the material.

For a lossless current to flow, the vortices must be prevented from moving under the applied force. To this end, inhomogeneity is introduced into the material which results in variation of the vortex energy depending on its position relative to nearby defects. When a vortex sits in a position corresponding to a local minimum in energy, there is a pinning force which opposes the force exerted by the transport current. Defects which provide a pinning force include dislocations, grain boundaries, and second-phase inclusions. The point where the transport current becomes sufficiently large that the pinning force is exceeded and the vortices begin to move is called the critical current. A superconductor with too few defects has a low pinning force and therefore a low critical current. A superconductor with too many defects, or the wrong kinds of defects, has degraded superconducting properties. Optimised high-field superconducting materials have a distribution of defects which results in a high critical current under operating conditions.

²MgB₂ is not strictly 'type-II' [25], but it does support vortices.

1.3 Magnetic confinement fusion

One possible future application of high-field superconductors is to confine a fusion plasma for energy production. Whilst the economic viability of magnetic confinement fusion is in doubt in general [26], [27], superconducting magnets are certainly required.

The tokamak is the most developed fusion reactor concept for energy production. Most serious plans for demonstrator reactors are based on this concept [28]–[32]. The basis of the tokamak concept is to set up a magnetic field to confine a plasma in a toroidal vacuum vessel. The charged particles in a fusion plasma move helically in a magnetic field, where the axis of the helix is parallel to the field, so to a first approximation particles cannot move large distances transversely to the field direction. By taking a tube of magnetic flux and closing it in on itself to form a torus, the particles are constrained to orbit around the closed field lines, and are therefore confined. Strictly though, the orbits of charged particles are only helical in a uniform magnetic field. When the field is curved around the torus, particles drift in the transverse direction so they don't remain confined for very long. Confinement times are improved by introducing a twist to the axes of the helical orbits. This way, the particles move through the minor angles as they orbit the central axis, see Figure 1.1, and the drift effects cancel out, although confinement is still imperfect since the particles interact with each other. The primary component of the field, whose field lines encircle the major axis, is called the toroidal field. The twist is a result of a secondary component, called the poloidal field, whose field lines encircle the minor axis.

The toroidal field is generated by the toroidal field coils. These are the largest magnets in the system and they carry the most current because the toroidal field is usually the larger of the two field components. The poloidal field is generated mostly by an electrical current in the plasma itself, supplemented by poloidal field coils. Additional coils are also used to make fine adjustments to the shape of the magnetic field to improve control over the plasma. Usually, the plasma current is initiated inductively, by ramping up the current in a solenoid which lies along the major axis of the torus. Since the current can't be ramped up forever, other methods of driving a current through the plasma must then take over if the current is to be sustained for long periods of time. For example, directed particle beams can be used to simultaneously transfer heat and momentum to the plasma.

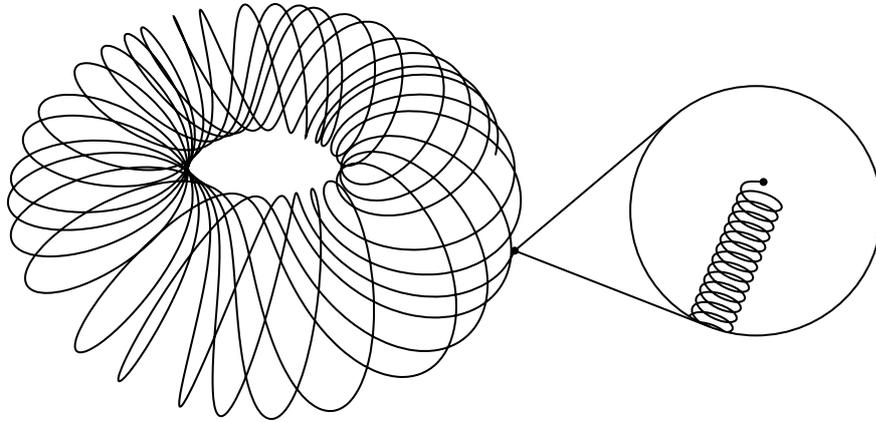


Figure 1.1: The path of a free charged particle in the magnetic field of a tokamak is a helix whose axis, the ‘guiding centre’, winds around the minor (poloidal) angles of the torus as it moves through the major (toroidal) angles. In a fusion plasma the path length is limited by collisions.

There are many difficulties in adapting superconducting materials so that they will be able to survive for any length of time in the environment of a fusion reactor. In particular, neutron irradiation from the plasma will disrupt the optimised microstructure so that the superconducting properties gradually degrade with increasing fluence [33]. Substantial radiation shielding is therefore required to maximise the lifetime of the superconducting magnets, since they represent a large fraction of the cost of a fusion reactor [27]. There are many other engineering issues like the requirement for low-resistance ‘re-mountable’ joints (i.e., joints where the superconductor can be separated and joined without damage) for ease of maintenance, the reduction of energy losses and heating under time-varying fields which is an intrinsic issue with high-field superconductors, and the design of cables and magnets which prevent excessive strain and transient load on the conductor itself. Some of these issues are common to all applications of superconducting magnets but are exacerbated by the scale of the magnet system of a tokamak. A detailed physical understanding of the mechanisms which limit the performance of high-field superconductors in this environment might allow us to design better superconducting materials for use in fusion reactors.

1.4 Thesis overview

The aim of this thesis is to study a practical high-field superconducting material, niobium–titanium, using simulations based on time-dependent Ginzburg–Landau theory with material parameters which are consistent with the measured properties of the particular phases which are seen in the actual microstructure. Some previous studies have included defects with realistic geometry in simulations, but the material properties of second-phase inclusions are usually set to arbitrary values representing a generic non-superconducting material, and consequently the magnitude of the critical current density in the simulation is not directly comparable with the measured value, although some qualitative features of the dependence on the magnetic field (magnitude and direction) can be reproduced.

Chapter 2 discusses relevant aspects of superconductivity theory. Most of the chapter is dedicated to the phenomenological Ginzburg–Landau theory, which forms the basis of the computational method, with a small section on the more fundamental BCS theory, from which the Ginzburg–Landau theory can be derived under certain conditions, thereby giving expressions for the phenomenological parameters in terms of normal state properties.

Chapter 3 describes the computational method used to solve the time-dependent Ginzburg–Landau equations. In order to model a wide class of materials accurately, a solver is required which solves the full system of equations and allows for spatial variation of all of the material parameters. Since existing three-dimensional solvers [34], [35] either don't allow for variation of some material parameters, or solve the Ginzburg–Landau equations only in a restrictive limit, a new solver is developed which uses a geometric multigrid method.

Chapter 4 describes the results of some preliminary simulations of REBCO with nanorod artificial pinning centres. It was found that for resistive nanorods, surface superconductivity in the vicinity of the matrix-pin interfaces can allow a lossless current to flow above the upper critical field of the matrix as long as the rods are sufficiently dense that the regions of surface superconductivity overlap allowing a supercurrent to flow across the whole system.

In Chapter 5 the conventional understanding of pinning in niobium–titanium, that ribbon-shaped titanium precipitates are primarily responsible for the high critical current density, is investigated. A computational model of niobium–titanium is developed using parameters which are constrained by experiments on

single-phase samples. Critical fields and critical currents are calculated for a range of microstructures and material parameters. It is concluded that the conventional explanation for the high critical current density in niobium–titanium can be justified, but that the pinning force must result primarily from a difference in the carrier effective mass between the precipitates and the matrix, rather than a difference in the condensation energy.

At least for the materials studied in this work, it appears that given sufficiently accurate material parameters for the various components of the system, simulations based on time-dependent Ginzburg–Landau theory can be used to obtain approximate critical current densities and critical fields. These methods might be used to guide the further optimisation of high-field superconducting materials.

Chapter 2

Superconductivity theory

Because macroscopic quantum coherence is unusual in terrestrial materials, some of the properties of superconductors can be surprising and unintuitive, but various theories exist which explain superconducting phenomena at different levels of detail.

The celebrated BCS theory describes the underlying mechanism using nothing more than many-body quantum mechanics, results from the theory of metals, and some well-chosen assumptions. Since it describes superconductivity at the finest level of detail it is often called a microscopic theory. Other microscopic theories include the generalisation of BCS theory to impure materials which is sometimes known as Gor'kov theory [36], and a greater generalisation which is important for first-principles modelling of low-temperature superconductors called Eliashberg theory [37], [38]. Whilst microscopic theory is not the main focus of this chapter, some discussion of BCS theory is given at the end.

The Ginzburg–Landau theory is a mesoscopic theory of superconductivity which can be obtained as a classical limit of BCS theory. It is the main subject of this chapter because it is the foundation of the computational method described in Chapter 3. Ginzburg–Landau theory predates BCS theory, being initially based on simple physical arguments rather than a detailed derivation from microscopic theory. For this reason it can be applied to superconductors for which there is no convincing microscopic description, such as the high-temperature cuprate superconductors.

There are macroscopic descriptions of superconductivity which are coarser still. One example is the critical state model [39], which can describe the irre-

versible magnetic response of high-field superconducting materials. These models are used in the design of superconducting magnets [40] to calculate the energy losses, heating, and forces when the superconductor is subjected to a changing external field or current. Whilst practically useful they have little to say about why some materials have better superconducting properties than others and they are not discussed in this chapter.

2.1 Ginzburg–Landau theory

Ginzburg–Landau theory [41] is a phenomenological theory of superconductivity based on Landau’s [42] theory of second-order phase transitions. In Landau theory, a second-order phase transition is characterised by a particular observable called the order parameter, which is zero in the less-ordered phase, above the transition temperature, and non-zero in the more ordered phase, below the transition temperature. It is assumed that the free energy can be written as a power series in the order parameter with coefficients which are simple functions of temperature, so that the equilibrium value of the order parameter, that is, the value which minimises the free energy, linearly approaches zero as the temperature approaches the transition temperature from below, leading to a discontinuity in the second derivative of the free energy as a function of temperature. An example of an order parameter is the magnetisation in a ferromagnet, which drops to zero at the Curie point. Ginzburg and Landau were motivated to make a particular choice for the order parameter based on some observed properties of superconductors.

The ability of a superconducting loop to carry a persistent current and therefore maintain angular momentum in spite of external perturbations is reminiscent of the ability of electrons in an atom to orbit the nucleus indefinitely without decay by electromagnetic radiation. Experiments show that the magnetic flux through such a loop is quantised just like the angular momentum of atomic orbitals [43]–[45]. These observations indicate the presence of macroscopic quantum coherence in the superconducting phase, suggesting that the charge carriers form some kind of condensate, an idea first proposed by London [46]. The key contribution of Ginzburg–Landau theory is that one can take the wavefunction of the state into which the charge carriers are condensed as an order parameter, expanding the free energy in terms of the squared-norm of the wavefunction, which in some way represents the fraction of the charge carriers in the condensed

state, and its spatial derivatives, which correspond to the kinetic energy of the condensate. This idea was proposed before the mechanism for low-temperature superconductivity was understood in detail, but the Ginzburg–Landau theory can also be derived as a limit of the more fundamental theories in the framework of many-body quantum mechanics under certain conditions [47].

2.1.1 Free energy

The state of the Ginzburg–Landau system on a base manifold M , representing the space occupied by the superconductor, consists of a section of a complex line bundle, E , and a connection on that bundle. In a particular local trivialisation (i.e., with a particular gauge choice) over a neighbourhood, U , the component in a unit basis of a section of E is expressed as a complex-valued zero-form, $\psi : U \rightarrow \mathbb{C}$, called the condensate wavefunction, and the local connection one-form, ω_A , is written in terms of a real one-form, called the magnetic vector potential, as

$$\omega_A = -i\frac{q}{\hbar}A \quad (2.1)$$

where q is the charge of the condensed particles and \hbar is the reduced Planck constant. The covariant derivative in the local trivialisation is then given by

$$D_{\omega_A}\psi = d\psi + \omega_A \wedge \psi \quad (2.2)$$

where d denotes the exterior derivative and \wedge denotes the exterior product (which will often be elided when one of the arguments is a zero-form).

There are several observables which can be calculated using ψ and A which are independent of the choice of local trivialisation. The curvature of the connection is proportional to the magnetic flux density, dA , the squared-norm of the condensate wavefunction, $\|\psi\|^2$, is interpreted as the superfluid density, and the supercurrent circulation¹ is given by

$$j_s = \hbar q m^{-1} \mathfrak{I}(\bar{\psi} D_{\omega_A} \psi) \quad (2.3)$$

where \mathfrak{I} denotes the imaginary part, $\bar{\cdot}$ denotes complex conjugation, and m^{-1} is the reciprocal effective mass for the charge carriers in the superconductor, which

¹I use the term circulation to refer to the one-form which is dual to the current density ($\dim M - 1$)-psuedoform.

is a symmetric² $(1, 1)$ -tensor field for an inhomogeneous and anisotropic material. Juxtaposition denotes tensor contraction, which is unambiguous in this case.

The Ginzburg–Landau free energy³, F , consists of three parts [41], [48],

$$F(\psi, A) = F_1(\psi) + F_2(\psi, A) + F_3(A). \quad (2.4)$$

The first part is the condensation energy. Since the order parameter (the condensate wavefunction) must be non-zero in the superconducting phase, it is necessary to include terms in the free energy which encourage the condensate to form when the temperature, T , is less than the transition temperature, T_c . Specifically, the condensation energy has the form

$$F_1(\psi) = \int_M \left(\frac{1}{2} \alpha \|\psi\|^2 + \frac{1}{4} \beta \|\psi\|^4 \right) \text{vol} \quad (2.5)$$

where vol is the volume pseudoform⁴ and α and β are phenomenological coefficients, which for generality we take to be real zero-forms rather than constants. Since ψ only exists in a local trivialisation, computing this integral involves choosing a partition of unity and a local trivialisation on each support. Even though ψ depends on the choice of local trivialisation, the value of the integral is well-defined because $\|\psi\|$ is independent of this choice, i.e., it is gauge invariant.

In the simplest case, E is trivial and the coefficients are constant, $\alpha(p) = \alpha_0$ and $\beta(p) = \beta_0$, where the condensation energy has a minimum corresponding to a constant non-zero wavefunction if $\alpha_0 < 0$ and $\beta_0 > 0$. These constants are parameters of the theory, and from a thermodynamic perspective they can be thought of as functions of temperature. If β_0 is negative then the free energy is unbounded from below, so β_0 must be positive for all temperatures. Since the minimum with a non-zero wavefunction exists only for $\alpha_0 < 0$, α_0 must be negative below T_c and positive above T_c .

²Symmetric as a linear map between one-forms with respect to the Hodge inner product.

³More properly, this is the free energy difference between the superconducting and normal states, but it is assumed that the normal state free energy is virtually constant in the vicinity of the transition.

⁴In a three-dimensional Euclidean space where the coordinates, x , y , and z have units of length and the orientation of the integration manifold agrees with the conventional right-handed orientation, the volume form (not pseudoform since we have picked out an orientation) can be written as $dx \wedge dy \wedge dz$.

The next part of the free energy is the kinetic part, given by

$$F_2(\psi, A) = \int_M \frac{\hbar^2}{2} \langle D_{\omega_A} \psi, m^{-1} D_{\omega_A} \psi \rangle \text{vol} \quad (2.6)$$

where $\langle \bullet, \bullet \rangle$ denotes the Hodge inner product, defined by $\langle \zeta, \eta \rangle \text{vol} = \bar{\zeta} \wedge \star \eta$ for possibly complex-valued forms ζ and η of the same degree where \star denotes the Hodge star operator⁵. Again, the integral is well-defined because the integrand is gauge invariant. This expression is essentially the usual kinetic energy from single-particle quantum mechanics, except it depends on the magnitude of the wavefunction because the total kinetic energy of the condensate depends on the number of condensed particles, and the reciprocal mass is a tensor rather than a scalar because we are dealing with quasiparticles in a material.

The final part of the free energy is (up to addition of a constant) the Gibbs free energy associated with the magnetic field,

$$F_3(A) = \int_M \frac{1}{2\mu_0} \|dA - \mu_0 \star^{-1} H\|^2 \text{vol}. \quad (2.7)$$

The magnetic induction is dA as always, so $dA - \mu_0 \star^{-1} H$ is the magnetisation, assuming the normal magnetic response of the material is much weaker than the magnetic induction due to the supercurrent, so that it can be neglected. Here, H is a pseudo-one-form which should be specified as a datum to impose an external magnetic field.

Since terms higher than second order in the superfluid density and the derivatives of the wavefunction are not included, the theory is only to be applied when the magnitude of the wavefunction is sufficiently small so that those terms can be reasonably neglected, which is the case when the temperature is close to the transition temperature because the phase transition is second order.

2.1.2 Equilibrium equations

The equilibria of the Ginzburg–Landau system are states for which the free energy is maximised or minimised. To obtain a set of equations whose solutions are equilibria, the free energy is varied with respect to the wavefunction and the

⁵I assume the version of the Hodge star which maps k -forms to $(\dim M - k)$ -pseudoforms, and vice-versa. Then no arbitrary handedness rule is required; the theory is parity-symmetric after all.

vector potential. The resulting equations are known as the Ginzburg–Landau equations, and they are given by

$$0 = (-\alpha - \beta\|\psi\|^2)\psi + \frac{\hbar^2}{2}\star^{-1}D_{\omega_A}\star m^{-1}D_{\omega_A}\psi, \quad (2.8)$$

$$0 = \frac{1}{\mu_0}\delta dA - \star^{-1}dH - \hbar q\mathfrak{I}(\bar{\psi}m^{-1}D_{\omega_A}\psi), \quad (2.9)$$

where δ is the codifferential, which is defined by $\delta\zeta = (-1)^k\star^{-1}d\star\zeta$ for a k -form ζ , subject to the boundary conditions

$$0 = i_{\partial}^*(\star m^{-1}D_{\omega_A}\psi), \quad (2.10)$$

$$0 = i_{\partial}^*(\star dA - \mu_0 H), \quad (2.11)$$

where \star denotes the pull-back and $i_{\partial} : \partial M \rightarrow M$ is the boundary inclusion map, so i_{∂}^* is the restriction to the boundary. The boundary conditions imply that the supercurrent through the boundary is zero and that the field on the boundary is equal to the applied field. These are the natural boundary conditions for the variation, which is to say that if they are not satisfied then the solutions the Ginzburg–Landau equations are not necessarily extrema of the free energy functional, but we will have cause to consider more general boundary conditions when modelling systems with an overall current flow, which correspond to minimising the free energy subject to a constraint on the total current.

The first Ginzburg–Landau equation can be interpreted as Schrödinger’s equation for a particle in a magnetic field, except there is an additional non-linear term which pushes the magnitude of the wavefunction towards its equilibrium value so the normalisation is fixed (since $\|\psi\|^2$ is an observable equal to the superfluid density) and the vector potential must be obtained self-consistently because there are a macroscopic number of condensed particles and the induction due to their motion can no longer be neglected. The second equation is essentially Ampère’s law where $\delta dA/\mu_0$ is the total current circulation, $\star^{-1}dH$ is the ‘free current’⁶, and $\hbar qm^{-1}\mathfrak{I}(\bar{\psi}D_{\omega_A}\psi)$ is the supercurrent which is induced to flow in the material. This equation tells us how the motion of the condensate affects the vector potential.

⁶The free current is an externally applied current which doesn’t interact with the condensate except via the magnetic field which it induces.

Internal boundary conditions can be obtained by varying the free energy for a superconductor which is divided into two regions, M_1 and M_2 , which share a common boundary, S , with inclusion maps $i_1 : S \rightarrow M_1$ and $i_2 : S \rightarrow M_2$. If $\psi = \psi_1$ in M_1 and $\psi = \psi_2$ in M_2 , and likewise for A , assuming that the dependent variables are continuous at the interface,

$$i_1^*(\psi_1) = i_2^*(\psi_2), \quad (2.12)$$

$$i_1^*(A_1) = i_2^*(A_2), \quad (2.13)$$

the natural boundary conditions are

$$i_1^*(\star m_1^{-1} D_{\omega_{A_1}} \psi_1) = i_2^*(\star m_2^{-1} D_{\omega_{A_2}} \psi_2), \quad (2.14)$$

$$i_1^*(\star dA_1) = i_2^*(\star dA_2), \quad (2.15)$$

where m_1^{-1} and m_2^{-1} are the reciprocal mass tensors on M_1 and M_2 respectively. These imply continuity of supercurrent and magnetic induction across the interface.

The Ginzburg–Landau equations are expressed in terms of the local condensate wavefunction and vector potential, and therefore depend on a local trivialisation. On a non-trivial bundle, where local trivialisations on multiple patches are required to cover the whole manifold, the Ginzburg–Landau equations must be satisfied by the local ψ and A on each patch. Where patches overlap, ψ and A are related by the transition map of the bundle. Given any local solution to the Ginzburg–Landau equations, (ψ, A) , by making a change of local trivialisation (i.e., a gauge transformation), one can obtain an infinite family of condensate wavefunctions and vector potentials which also satisfy the Ginzburg–Landau equations, corresponding to the same global section and therefore the same physical configuration. These are given by (ψ', A') where

$$\psi = e^{i\chi} \psi', \quad (2.16)$$

$$A = A' + \frac{\hbar}{q} d\chi \quad (2.17)$$

for an arbitrary zero-form χ , which is sometimes called a gauge function.

2.1.3 Homogeneous solution

The simplest solution to the Ginzburg–Landau equations is for a homogeneous superconductor, $\alpha = \alpha_0$ and $\beta = \beta_0$ ⁷, in zero external field, $H = 0$, on a trivial bundle. In this case a global wavefunction and vector potential exist. Any spatial variations only serve to increase the free energy so the stable equilibrium is the constant

$$(\psi, A) = (\sqrt{-\alpha_0/\beta_0}, 0) \quad (2.18)$$

when $\alpha_0 < 0$. When $\alpha_0 > 0$ the stable equilibrium is

$$(\psi, A) = (0, 0), \quad (2.19)$$

which remains a valid solution for all α_0 , although it is a maximum of the free energy and therefore unstable for $\alpha_0 < 0$. We call these the homogeneous superconducting solution and the normal state solution respectively. The former describes a stationary condensate with a homogeneous superfluid density of $n_0 = -\alpha_0/\beta_0$ and a free energy density of $f_0 = -\alpha_0^2/2\beta_0$. The latter describes a system with no condensate where the free energy is zero.

2.1.4 London limit

Assuming the domain is simply connected, the bundle, E , is trivial, and the wavefunction is non-zero everywhere, we can consider Ginzburg–Landau equations in the case where the wavefunction is real. This is because the phase can be eliminated from any solution by means of a gauge transformation. Letting $\psi = f$ where f is real-valued, the second Ginzburg–Landau equation becomes

$$\frac{1}{\mu_0} \delta dA = \star^{-1} dH - q^2 f^2 m^{-1} A. \quad (2.20)$$

Assuming a homogeneous and isotropic superconductor with scalar effective mass m_0 where the current density is small and neglecting second order terms, the first Ginzburg–Landau equation reduces to a pair of equations (from the real and imaginary parts) which are still satisfied by $f = \sqrt{n_0}$ as long as $\delta A = 0$, which

⁷The effective mass need not be homogeneous nor isotropic.

is a gauge condition for what is known as the London gauge⁸. Explicitly, we have

$$\delta A = 0, \quad (2.21)$$

$$\delta dA = \mu_0 \star^{-1} dH - \frac{1}{\lambda^2} A \quad (2.22)$$

subject to the boundary conditions

$$i_{\partial}^*(\star A) = 0, \quad (2.23)$$

$$i_{\partial}^*(\star dA - \mu_0 H) = 0 \quad (2.24)$$

where $\lambda = \sqrt{m_0/\mu_0 q^2 n_0} = \sqrt{-\beta_0 m_0/\mu_0 q^2 \alpha_0}$ is called the London penetration depth and the equations are known as the London equations.

Since $\delta dA/\mu_0$ is the current circulation, the London equations imply that A is proportional to the difference between the free current circulation, $\star^{-1} dH$, and the total current circulation, i.e., that $-A$ is proportional to the supercurrent circulation. Since $\delta \star^{-1} dH$ is automatically zero (all exact forms are closed), the condition $\delta A = 0$ expresses continuity of current, and the condition $i_{\partial}^*(\star dA - \mu_0 H) = 0$ expresses that the current flow into the boundary is just the free current, i.e., that the supercurrent into the boundary is zero.

It can be shown that for a superconductor obeying London's equations, an applied magnetic field is screened so that the magnetic induction decays exponentially with a length scale λ away from the free surface. This is due to a lossless supercurrent which is induced to flow near the surface. Likewise, if a superconducting wire carries a current, the current density flows only near the free surface of the wire, again decaying on a length scale λ in the interior. Since the penetration depth is usually on the scale of nanometres or micrometres, macroscopic superconductors in sufficiently weak external fields (so that the London approximation is valid) exhibit near-perfect diamagnetism. This is known as the Meißner–Ochsenfeld effect.

2.1.5 Thermodynamic critical field

According to London's equations, the homogeneous superconducting state persists in the interior of a superconductor when a magnetic field is applied. However,

⁸The boundary conditions differentiate the London gauge from the Coulomb gauge, which also has $\delta A = 0$, but where the boundary condition $i_{\partial}^*(\star A) = 0$ is replaced with the condition that the vector potential approaches zero at spatial infinity.

even if such a solution exists, it does not mean that it is thermodynamically preferable, and the Ginzburg–Landau free energy contains a penalty for the screening of the external field. If the superconductor in a constant external field is sufficiently large so that the contribution to the free energy from near the boundaries can be neglected, since in the interior we have $\psi = \sqrt{n_0}$ and $dA = 0$, the free energy is given by

$$F(\psi, A) = |M| \left(\frac{1}{2\mu_0} \|\mu_0 \star^{-1} H\|^2 - f_0 \right) \quad (2.25)$$

In the normal state, $\psi = 0$ and $dA = \mu_0 \star^{-1} H$, so the free energy is zero. Therefore the superconducting state is thermodynamically preferable if $\frac{\mu_0}{2} \|\star^{-1} H\|^2 < f_0$ and we call the critical field value for this scenario,

$$H_c = \sqrt{\frac{2f_0}{\mu_0}} = \sqrt{\frac{\alpha_0^2}{\mu_0 \beta_0}}, \quad (2.26)$$

the thermodynamic critical field. When the critical field is exceeded, the superconductor reverts to the normal state.

2.1.6 Flux quantisation

The magnetic flux through a superconducting loop is quantised. This can be understood by a simple argument which invokes the existence of the condensate wavefunction and the Meißner–Ochsenfeld effect. Consider a loop of superconducting wire carrying a persistent current. If the superfluid density is constant as in the London limit, the wavefunction is given by $\psi = \sqrt{n_0} e^{i\theta}$ for some phase function θ , and therefore the supercurrent circulation is given by

$$j_s = \frac{\hbar q}{m_0} \Im(\bar{\psi} D_{\omega_A} \psi) = \frac{\hbar q n_0}{m_0} \left(d\theta - \frac{q}{\hbar} A \right). \quad (2.27)$$

If γ is the closed curve along which the superconducting wire lies, integrating the circulation around this curve gives

$$\int_{\gamma} j_s = \frac{\hbar q n_0}{m_0} \int_{\gamma} \left(d\theta - \frac{q}{\hbar} A \right) = 2\pi n \frac{\hbar q n_0}{m_0} - \frac{q^2 n_0}{m_0} \Phi \quad (2.28)$$

where n is an integer, since the phase must wind by an integer multiple of 2π around the loop, and Φ is the flux through the loop. If the wire is sufficiently

thick and the curve passes through its axis, then, due to the Meißner–Ochsenfeld effect, the current circulation along the curve is zero. Hence

$$0 = 2\pi n \frac{\hbar q n_0}{m_0} - \frac{q^2 n_0}{m_0} \Phi \quad (2.29)$$

and therefore $\Phi = n\phi_0$ where ϕ_0 is called the magnetic flux quantum, and is given by

$$\phi_0 = \frac{2\pi\hbar}{q}. \quad (2.30)$$

Since the phase is a smooth function, a smooth deformation of the integration path cannot result in a discontinuous jump in the flux, and therefore sets of paths which are homotopic in the space where $\|j_s\| \approx 0$ must enclose the same flux.

2.1.7 The coherence length and the proximity effect

Whilst the first Ginzburg–Landau equation closely resembles the Schrödinger equation, it can also be viewed as a Helmholtz equation with an additional non-linear term. For a homogeneous and isotropic superconductor the natural length scale, called the coherence length, can be read off as

$$\xi = \sqrt{-\frac{\hbar^2}{2m_0\alpha_0}}. \quad (2.31)$$

In situations where the wavefunction is not constant, it usually changes on a length scale of ξ . For example, the Ginzburg–Landau equations for a superconductor occupying \mathbb{R}^1 with $\alpha = \alpha_0$ for $x > 0$ and $\alpha = -\alpha_0$ for $x < 0$ are solved by⁹

$$\psi = \begin{cases} \sqrt{n_0} \cdot \tanh \frac{x}{\sqrt{2}\xi} + \frac{\log 3}{2} & x > 0 \\ -\sqrt{2n_0} \cdot \operatorname{csch} \frac{x}{\xi} - \log 3 + 2\sqrt{2} & x < 0 \end{cases}. \quad (2.32)$$

This system models an interface between a superconductor and a normal metal. The condensate is depleted near the edge of the superconductor ($x > 0$) but it also survives some distance into the non-superconducting metal ($x < 0$). This is known as the proximity effect.

⁹A unit metric has been assumed so that the coordinate has units of length.

2.1.8 Vortices

There are spatially inhomogeneous solutions to the Ginzburg–Landau equations even for homogeneous materials. In an infinite homogeneous superconductor, the winding number of the phase of the wavefunction is topologically preserved, so there are equilibria for every integer winding number. For example, for a superconductor occupying the plane, \mathbb{R}^2 , suppose the wavefunction is given by $\psi = fe^{i\theta}$ where f is real-valued and only depends on the radial coordinate r , and θ is the polar coordinate. The phase winds by 2π around the origin so the winding number is 1. The winding of the phase is associated with a supercurrent which circulates around the origin, and since the winding becomes more and more rapid towards the phase singularity at $r = 0$, f must drop to zero so that the supercurrent remains bounded. We should also expect that f approaches $\sqrt{n_0}$ as r approaches infinity, since the derivative of the wavefunction is negligible far from the origin. So we have a circulating supercurrent around a core where the superfluid density drops to zero, and by the same argument as for the quantisation of flux in a superconducting loop, taking the contour of integration to infinity, the flux in this state is ϕ_0 . Such a solution is called a vortex or fluxon. More generally, a vortex with a winding number of n carries a flux of $n\phi_0$.

The thermodynamic stability of vortex solutions depends on the external magnetic field and the material in question. A superconductor containing vortices accommodates flux in the interior, allowing it to reclaim some of the free energy penalty for the expulsion of flux in the Meißner state, but they come at a cost in the condensation energy because of the normal core, as well in the kinetic energy due to the inhomogeneity of the wavefunction. It turns out that the normal core of a single-quantum vortex has a radius of roughly ξ , and the flux it carries is spread over a region with a radius of λ . Whether it is favourable to form vortices at any field depends on the ratio $\kappa = \lambda/\xi$ which is called the Ginzburg–Landau parameter. Formation of single-quantum vortices in high fields is disfavourable for $\kappa < 1/\sqrt{2}$ and favourable for $\kappa > 1/\sqrt{2}$, which separates superconductors into two classes called type-I and type-II respectively. Still, vortex solutions are metastable in type-I superconductors with no free surfaces because of the topological constraint. Furthermore, in type-I superconductors, amongst the solutions with a winding number of n , a single n -quantum vortex is preferable, whereas for type-II superconductors it is preferable to have n single-quantum vortices,

essentially because the free energy associated with a normal-superconductor interface, as between the normal core and the superconducting bulk, is positive for type-I superconductors and negative for type-II. Although there are no exact vortex solutions in terms of elementary functions, approximations of the superfluid density, current density, and magnetic flux density as a function of radius can be obtained [49].

2.1.9 Abrikosov lattice

For a homogeneous and isotropic type-II superconductor subject to a large external field, the global minimum in the free energy corresponds to a solution with closely-packed vortices with a density proportional to the amount of flux which penetrates the material [50]. There is an effective repulsion between vortices because there is a free energy cost associated with the superposition of the circulating currents around nearby vortex cores. Although the formation of vortices allows superconductivity to persist at fields higher than the thermodynamic critical field, eventually the density of vortices becomes so high that the vortex cores are forced to overlap and the system becomes fully normal. The magnetic field at which this occurs is called the upper critical field.

Near the upper critical field, the superfluid density is small and so is the magnitude of the supercurrent. The magnetic induction is therefore almost constant. The Ginzburg–Landau problem can be approximately solved in this regime by considering the wavefunction as a perturbation to the normal state solution at the upper critical field. In this case one considers the linearised form of the first Ginzburg–Landau equation (neglecting the $\|\psi\|^2\psi$ term), which is just the same as the problem of a charged particle in a magnetic field in quantum mechanics, and is solved by Landau orbitals in \mathbb{R}^2 . However, since the Ginzburg–Landau problem is not an eigenvalue problem, that is, the energy eigenvalue of the Schrödinger equation is replaced with the constant $-\alpha_0$, each Landau level is only a solution for a specific field value, and the lowest Landau level corresponds to the highest field, which is the upper critical field. For larger fields there is no solution.

The value of the upper critical field, H_{c2} , is given by

$$\mu_0 \|H\| = \mu_0 H_{c2} = \frac{\phi_0}{2\pi\xi^2} = -\frac{2m_0\alpha_0}{\hbar q} \quad (2.33)$$

and the associated family of solutions is

$$(\psi_k, A_k) = \left(e^{ikx} e^{-\frac{1}{2}(y/\xi - k\xi)^2}, -\mu_0 H_{c2} y \wedge dx \right) \quad (2.34)$$

where k is an arbitrary real parameter. However, arbitrary linear combinations of this solution with different values of k do not solve the full Ginzburg–Landau equations, except at precisely the upper critical field in which case the wavefunction is zero anyway. In order to find the precise solution of the linearised equations which is preferred close to the upper critical field, one must somehow include the non-linearity, as a perturbation or by minimising the free energy.

Abrikosov [50] chose to seek solutions which were periodic, writing the wavefunction as a linear combination of the form

$$\psi = \sum_{m \in \mathbb{Z}} c_m \cdot e^{imkx} \cdot e^{-\frac{1}{2}(y/\xi - mk\xi)^2} \quad (2.35)$$

where k is now taken to be constant and the coefficients c repeat with some periodicity N , so that $c_m = c_{m+N}$. He used $N = 1$ but it was later shown that the global minimum of the free energy is attained when $N = 2$, $c_0 = 1$, $c_1 = i$, and $k = 3^{1/4} \pi^{1/2} / \xi$, corresponding to a triangular lattice of vortices [51]. The normalised superfluid density for the resulting solution is shown in Figure 2.1. A similar triangular lattice solution is valid when the magnetic field is smaller, but it cannot be expressed as simply. Some numerical solutions are shown in Figure 2.2.

The Abrikosov solution remains valid when the domain is a flat torus, $M = T^2$, but the Stokes–Cartan theorem implies that the wavefunction bundle must be non-trivial in order to accommodate non-zero flux. The construction of the bundle is reminiscent of the Dirac monopole [53]. It implies a quantisation condition where the number of phase singularities in the wavefunction is a fixed integer, n , and the total magnetic flux is exactly $n\phi_0$.

For a flat torus of dimensions $w \cdot h$, given the obvious Cartesian coordinates, $x : M \rightarrow [0, w]$ and $y : M \rightarrow [0, h]$, one can take two patches, U and V , which are the pre-images of $[0, w] \times (0, h)$ and $[0, w] \times ([0, h/2) \cup (h/2, h])$ respectively. That is, each patch is a horizontal band, covering the whole torus except for a horizontal circle, and the missing circles from each patch are vertically offset by $h/2$ so that $U \cup V$ covers M .

The wavefunction is given by its components, ψ^U and ψ^V , in local trivialisations over U and V , and likewise the vector potential has a different expression

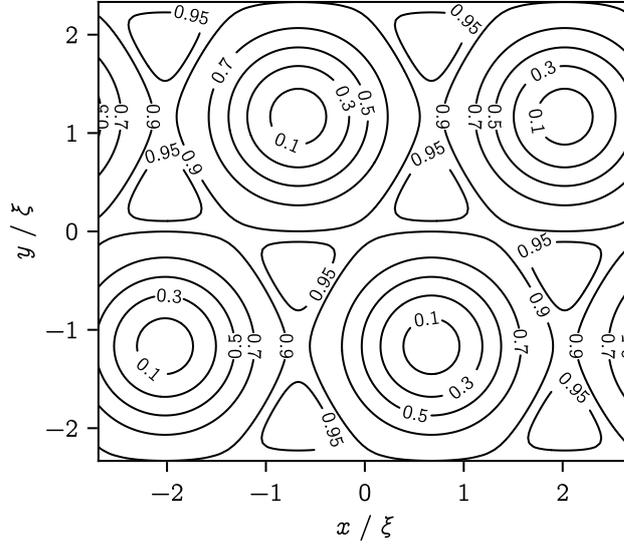


Figure 2.1: Superfluid density contours for the normalised triangular lattice solution to the linearised Ginzburg–Landau equations at the upper critical field.

in each patch. The patches U and V intersect in two disjoint regions, I_1 and I_2 , which are the pre-images under the coordinate map of $[0, w] \times (0, h/2)$ and $[0, w] \times (h/2, h)$ respectively. In each of these intersection regions, the values of ψ^U and ψ^V are components of a single section in two different unit bases. They are therefore related by a transition map which can be represented as multiplication by a complex-valued function with unit norm. For example, in I_1 , $\psi^U = \psi_1^U$ and $\psi^V = \psi_1^V$, so

$$\psi_1^V = e^{i\chi_1} \psi_1^U \quad (2.36)$$

for some smooth (circle-valued) zero-form χ_1 . The vector potentials which define the connection one-forms in the two local trivialisations are related by

$$A_1^V = A_1^U + \frac{\hbar}{q} d\chi_1. \quad (2.37)$$

The structure of the bundle is defined by the transition functions on the intersections of the patches which cover the manifold. For the Abrikosov lattice, the appropriate transition functions are

$$\chi_1 = +\pi n \frac{x}{w}, \quad (2.38)$$

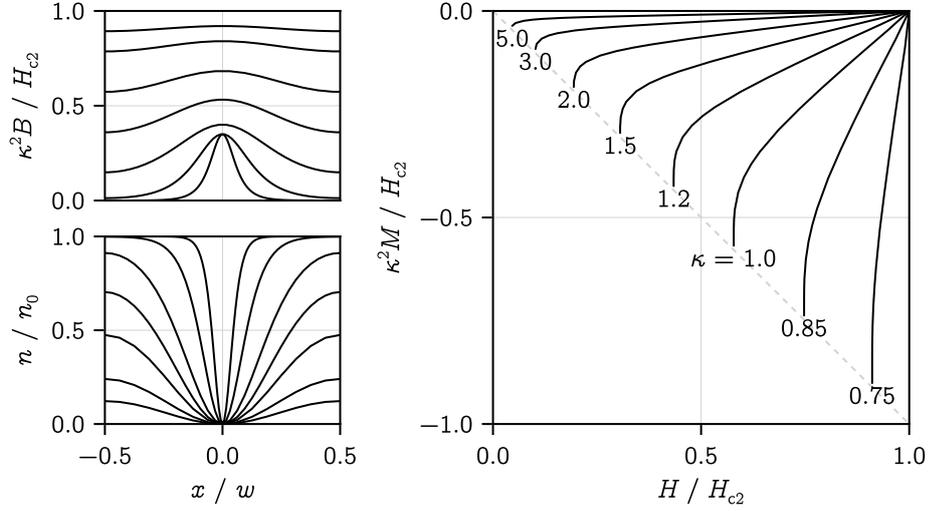


Figure 2.2: (Left) Profiles of the superfluid density and magnetic induction around a vortex core in the lattice with $\kappa = 3/2$ for magnetic flux densities 0.01, 0.05, 0.2, 0.4, 0.6, 0.8, and 0.9. (Right) Reversible magnetisation curves for lattice solutions. These plots are inspired by a figure of Brandt's [52], but were generated by numerically solving the Ginzburg–Landau equations using the method described in Chapter 3.

$$\chi_2 = -\pi n \frac{x}{w}. \quad (2.39)$$

When considering the wavefunction on one patch, say U , these essentially imply a jump¹⁰ in the phase between the bottom and the top edge of magnitude $2\pi n x / w$. There is a corresponding jump in the x -component of the vector potential. The magnetic flux is the integral of dA^U over U , since U almost covers M , missing only zero-measure subset. Its value is $n\phi_0$ as previously stated. The Abrikosov lattice solution on the patch U , which is valid for $n = 2$, is

$$\psi^U = \sum_{m \in \mathbb{Z}} c_m \cdot e^{imk(x-w/2)} \cdot \exp\left(-\frac{1}{2\xi^2} \left(y - \frac{h}{2} - mk\xi^2\right)^2\right), \quad (2.40)$$

$$A^U = -\frac{\phi_0}{2\pi} \cdot \frac{y - h/2}{\xi} \wedge \frac{dx}{\xi} \quad (2.41)$$

where $k = 3^{1/4} \pi^{1/2} / \xi$, $w = 2\pi/k$, $h = \sqrt{3}w$, and $c_m = 1$ if m is even and $c_m = i$ if m is odd.

¹⁰Technically, there is no discontinuity because the patch does not include the edge.

2.1.10 Surfaces

In the vicinity of the boundary of a superconductor, the translational symmetry of the Ginzburg–Landau system is broken, and solutions like the Abrikosov lattice are invalid. The boundary geometry can significantly affect the observed behaviour of superconducting samples.

A simple system with a boundary is the superconductor occupying a Euclidean half-space. In three-dimensions this system has a two-dimensional planar boundary. Such a system was considered by Saint-James and de Gennes [54] in the linearised high-field limit, where it was observed that although superconductivity is destroyed in the bulk if the constant field H parallel to the surface exceeds H_{c2} , superconductivity persists in a layer, approximately a coherence length wide, along the boundary surface up to a field of $H_{c3} \approx 1.695H_{c2}$. It is generally true that superconductivity persists above the upper critical field in the vicinity of free surfaces, or interfaces with resistive or insulating materials, which are parallel to the magnetic field. This has a negligible effect on the thermodynamic and transport properties of macroscopic bulk superconductors, but it is important in thin films, multilayers, and nanowires.

For superconductors with sharp corners, or parts which are so thin that pairs of boundaries are separated by less than a coherence length, superconductivity can survive beyond even H_{c3} . For example, the parallel critical field of a thin film diverges as the reciprocal of the film thickness [41], and a similar result holds for thin wires.

2.1.11 De-pairing current

The Ginzburg–Landau equations imply a limit to the amount of current which can be passed through a superconducting wire immersed in three-dimensional Euclidean space. For a very thin wire, with dimensions much smaller than the coherence length, the spatial variation of the current and superfluid density over the cross-section of the wire can be neglected. Although the critical field can be arbitrarily high for sufficiently thin wires, the current density cannot exceed [41]

$$J_D = \frac{2}{3\sqrt{3}} \cdot \frac{\phi_0}{2\pi\mu_0\lambda^2\xi}, \quad (2.42)$$

which we call the de-pairing current density. This illustrates that the Ginzburg–Landau equations do not respect Galilean relativity, the condensate motion is in

reference to the privileged rest frame of the crystal lattice¹¹.

2.1.12 Characteristic units

Since there are two natural length scales in the Ginzburg–Landau system, there are two approaches to writing the Ginzburg–Landau equations in dimensionless form. Some authors choose to express lengths in terms of the penetration depth [48], [50], [56], whereas in this work, they are expressed in terms of the coherence length since it is the smaller of the two length scales for the systems which we are interested in. The difference is where the factors of κ appear in the final equations. Firstly, the dimensionless metric, \hat{g} , is given in terms of the original metric, g , as

$$g = \xi^2 \hat{g}, \quad (2.43)$$

which induces dimensionless versions of the volume form and Hodge star. Likewise, writing the wavefunction in units of its equilibrium value, the vector potential in terms of the flux quantum, the magnetic field in terms of the natural current unit (assuming a three-dimensional domain¹²) and the material parameters in terms of some arbitrary reference values,

$$\psi = \sqrt{n_0} \hat{\psi}, \quad A = \frac{\phi_0}{2\pi} \hat{A}, \quad H = \frac{\xi \phi_0}{2\pi \mu_0 \lambda^2} \hat{H} = J_0 \xi^2 \hat{H}, \quad (2.44)$$

$$\alpha = \alpha_0 \hat{\alpha}, \quad \beta = \beta_0 \hat{\beta}, \quad m^{-1} = \hat{m}^{-1} / m_0, \quad (2.45)$$

we obtain the dimensionless Ginzburg–Landau equations,

$$0 = \left(\hat{\alpha} - \hat{\beta} \|\hat{\psi}\|^2 \right) \hat{\psi} + \star^{-1} D_{\hat{\omega}_{\hat{A}}} \star \hat{m}^{-1} D_{\hat{\omega}_{\hat{A}}} \hat{\psi}, \quad (2.46)$$

$$0 = \kappa^2 \delta d \hat{A} - \star^{-1} d \hat{H} - \mathfrak{F} \left(\hat{\psi} \hat{m}^{-1} D_{\hat{\omega}_{\hat{A}}} \hat{\psi} \right) \quad (2.47)$$

¹¹A rotating superconductor generates a magnetic field known as the London moment. Measurement of the London moment can be used to determine the electron mass (with small corrections), but naïve application of Ginzburg–Landau theory in the lab frame (in which the superconductor is rotating) implies that this measurement should give the effective mass of Landau quasiparticles. Tate et al. [55] measured the London moment in niobium and obtained an electron mass which is very close to the accepted value, $m \approx 1.000084 m_e$.

¹²The units of the norm of the current density and magnetic field depend on the dimension of the domain since the current density is a $(\dim M - 1)$ -pseudoform and the magnetic field is a $(\dim M - 2)$ -pseudoform. Once the appropriate units have been chosen, the dimensionless Ginzburg–Landau equations look the same regardless of the dimension.

subject to the boundary conditions

$$i_{\partial}^*(\star\hat{m}^{-1}D_{\hat{\omega}_{\hat{A}}}\hat{\psi}) = 0, \quad (2.48)$$

$$i_{\partial}^*(\kappa^2\star d\hat{A} - \hat{H}) = 0 \quad (2.49)$$

where the dimensionless connection form is given by $\hat{\omega}_{\hat{A}} = -i\hat{A}$ and the Hodge star and codifferential implicitly refer to the dimensionless metric.

2.1.13 Time dependence

The non-equilibrium behaviour of a superconductor can be described using time-dependent Ginzburg–Landau theory. The time-dependent Ginzburg–Landau equations are gradient-flow equations for the Ginzburg–Landau free energy, meaning that the wavefunction and vector potential gradually relax towards an equilibrium. They can also be derived from the microscopic theories under an even more restrictive set of assumptions than the ordinary Ginzburg–Landau equations.

The dimensionless time-dependent Ginzburg–Landau equations are given by [57]

$$\hat{\dot{\psi}} + i\hat{\varphi}\hat{\psi} = \left(\hat{\alpha} - \hat{\beta}\|\hat{\psi}\|^2\right)\hat{\psi} + \star^{-1}D_{\hat{\omega}_{\hat{A}}}\star\hat{m}^{-1}D_{\hat{\omega}_{\hat{A}}}\hat{\psi} \quad (2.50)$$

$$-\hat{\sigma}\left(\hat{\dot{A}} + d\hat{\varphi}\right) = \kappa^2\delta d\hat{A} - \star^{-1}d\hat{H} - \mathfrak{I}\left(\overline{\hat{\psi}}\hat{m}^{-1}D_{\hat{\omega}_{\hat{A}}}\hat{\psi}\right) \quad (2.51)$$

where $\hat{\varphi}$ is a real-valued zero-form representing the electrostatic scalar potential and $\hat{\sigma}$ is a new material parameter which is a symmetric $(1, 1)$ -tensor field. The solution now consists of the triple $(\hat{\psi}, \hat{A}, \hat{\varphi})$ where each entry is a time-parametrised family of differential forms of the appropriate degree and value where, for example, $\hat{\dot{\psi}}$ denotes differentiation of $\hat{\psi}$ with respect to the time parameter. The scalar potential is introduced to maintain the electromagnetic gauge symmetry and the dimensionless conductivity, $\hat{\sigma}$, represents the ratio of the time-scales for relaxation of the vector potential and the wavefunction. The second equation can now be read as Ohm's law, where $-\left(\hat{\dot{A}} + d\hat{\varphi}\right)$ is the dimensionless electric field, and the right-hand side represents the normal current, which is the difference between the bound current and the supercurrent. The magnetic field can vary with time but it must be slowly varying since we are working in the magnetic limit.

Time dependence introduces new characteristic units for time and for the electric field. The relaxation time for the wavefunction has to be introduced as a parameter, τ , whose value can be calculated in the microscopic theory, and the dimensionless conductivity represents the relative relaxation time for the current. The characteristic scale for the electric field is given by $E_0 = \phi_0/2\pi\xi\tau$.

2.1.14 Pinning

A perfect macroscopic type-II superconductor in an external magnetic field cannot carry any current without dissipation. This is because the transport current exerts a force on the vortices, causing them to move dissipatively through the material. The force can be viewed either as a Magnus force, where the transport current reinforces the circulating vortex current on one side of the core and counteracts it on the other, or as a Lorentz force due to the magnetic flux carried by the vortex interacting with the transport current. Any force causes the vortices to move immediately because a perfect superconductor is translationally symmetric, assuming it is sufficiently large that the effects at the boundary can be neglected¹³

In order to carry a lossless current density, defects must be introduced in the material. For example, consider a homogeneous type-II superconductor with an embedded cylinder of normal material, whose radius is roughly the coherence length, where the applied field is along the axis of the cylinder. Suppose there is a vortex near the cylinder. Since the condensation energy is reduced due to the decreased superfluid density in the vortex core and since the superfluid density is reduced in the normal cylinder anyway, it is preferable for the vortex core to lie inside the normal cylinder. We say that the vortex is pinned, since there is a force of attraction between the vortex and the cylinder which acts against the force due to the current, holding the vortex in place.

The situation becomes much more complicated when there are a variety of defect types, when they are randomly distributed, and when the density of vortices is high enough that their interaction can't be neglected. The vortices themselves act like elastic strings, but together they can form a rigid lattice with certain elastic properties, or a disordered glassy state, or a vortex liquid when thermal fluctuations are present. How the individual pinning forces of each defect lead

¹³This requires that the Ginzburg–Landau parameter is large, otherwise the current is confined to the boundary.

to a net pinning force on all the vortices taking into account the interactions is sometimes called the grand summation problem. In practice, strong assumptions are required to make analytic progress [58]–[61] but, in principle, quantitative predictions of the pinning force could be obtained by numerically solving the time-dependent Ginzburg–Landau equations, which is the subject of this thesis.

2.2 BCS theory

Whilst a detailed exposition of the microscopic theory of superconductivity is beyond the scope of this work, some discussion is required since some of the relations derived from microscopic theory will later be used to relate the coefficients of the Ginzburg–Landau free energy to more concrete measurable properties of the material.

The original mechanistic theory of superconductivity is due to Bardeen, Cooper, and Schrieffer [62] and it has been subsequently extended by many authors to deal with more complex and realistic materials [36], [37], [63]. A brief discussion of BCS theory will be given including the limits of its applicability and state some key results from the literature.

2.2.1 The BCS system

The starting point of the BCS theory is a model for metals based on Fermi liquid theory and the Born–Oppenheimer approximation. That is, the valence electrons interact with each other, as well as with the nuclei and core electrons, via the Coulomb interaction, but the nuclei and core electrons are treated only as a static external potential, whereas the valence electrons are treated as a strongly-interacting Fermi liquid. The fundamental excitations of the Fermi liquid are quasiparticles of the same number and charge as the electrons, but which interact with each other much more weakly. So weakly, in fact, that to a first approximation the interaction can be completely neglected. The model then consists of free particles subject to an external potential which, in the case of a non-magnetic pure metal, is periodic and independent of spin, so the orbitals are Bloch waves which are uniquely specified by a crystal momentum, $k \in K$, a band index, $i \in B$, and a spin projection, $\sigma \in S$, where $S = \{\uparrow, \downarrow\}$. The (grand canonical) Hamiltonian

is

$$H_0 = \sum_{i \in B} \sum_{k \in K} \sum_{\sigma \in S} \xi_i(k) c_{i,k,\sigma}^\dagger c_{i,k,\sigma} \quad (2.52)$$

where ξ is the energy of the Bloch waves measured from the chemical potential and c^\dagger and c are the creation and annihilation operators for the Bloch states of the quasiparticles, which from now on will just be called ‘electrons’.

The key idea in the BCS model is that the coupling between electrons and lattice vibrations in a metal can result in an effective interaction between electrons of opposite spin and momentum which is attractive. Fröhlich [64] was the first to suggest this, but BCS theory was the first fully developed theory of phonon-mediated superconductivity. In view of the success of the BCS theory, the correctness of this idea will not be justified to the same lengths as in the original work, and instead the effective interaction will simply be added into the Hamiltonian. The interaction is given by

$$H_{\text{int}} = \sum_{i \in B} \sum_{q \in K} \sum_{k \in K} V_i(q, k) b_{i,q}^\dagger b_{i,k} \quad (2.53)$$

where $b_{i,k} = c_{i,-k,\downarrow} c_{i,k,\uparrow}$ and V_i is the interaction matrix in momentum space. Clearly, this is just the general form of an interaction between pairs of electrons in the same band having opposite spin and momentum. The final BCS Hamiltonian is just the sum of the free particle Hamiltonian and the interaction,

$$H = H_0 + H_{\text{int}}. \quad (2.54)$$

Sometimes additional terms are added to represent, for example, randomly distributed impurities.

2.2.2 Mean-field theory

Although the original authors used a variational method, the modern approach to BCS theory begins with a mean-field ansatz. With the expectation that the charge carriers form a condensate, we assume that the b operators are well-approximated by their ensemble averages and rewrite the interaction term as

$$b_{i,q}^\dagger b_{i,k} \xrightarrow{\text{MF}} \langle b_{i,q}^\dagger \rangle b_{i,k} + \langle b_{i,k} \rangle b_{i,q}^\dagger - \langle b_{i,q}^\dagger \rangle \langle b_{i,k} \rangle. \quad (2.55)$$

Since it is assumed that the ensemble averages $\langle b_{i,k} \rangle$ are non-zero in the ground state, and since these operators do not conserve particle number unlike the full

Hamiltonian, we have assumed that the transition to the superconducting phase is a symmetry-breaking phase transition.

With the above assumption, the Hamiltonian can be re-written as

$$H^{\text{MF}} = H_0 + H_{\Delta}^{\text{MF}} + E^{\text{MF}} \quad (2.56)$$

where the interaction term has been replaced with the following non-interacting term,

$$H_{\Delta}^{\text{MF}} = - \sum_{i \in B} \sum_{k \in K} \left(\bar{\Delta}_i(k) b_{i,k} + \Delta_i(k) b_{i,k}^{\dagger} \right), \quad (2.57)$$

plus a constant term,

$$E^{\text{MF}} = \sum_{i \in B} \sum_{q \in K} \sum_{k \in K} V_i(q, k) \langle b_{i,q}^{\dagger} \rangle \langle b_{i,k} \rangle, \quad (2.58)$$

where Δ , known as the pair potential, is given by

$$\Delta_i(k) = - \sum_{q \in K} V_i(k, q) \langle b_{i,q} \rangle. \quad (2.59)$$

Δ can be non-zero only if some of the ensemble averages $\langle b_{i,k} \rangle$ are non-zero, so Δ itself is a suitable order parameter for the superconducting phase transition.

2.2.3 Formal solution by Bogoliubov transformation

Any quantum system where the Hamiltonian is at most quadratic in the ladder operators is solvable in principle, and a standard approach to solving the mean-field BCS system uses a transformation of the ladder operators known as a Bogoliubov [65] transformation. The new operators, γ^{\dagger} and γ , are defined by

$$c_{i,k,\uparrow} = u_i(k) \gamma_{i,k,\uparrow} - \bar{v}_i(k) \gamma_{i,-k,\downarrow}^{\dagger}, \quad (2.60)$$

$$c_{i,-k,\downarrow} = v_i(k) \gamma_{i,k,\uparrow} - \bar{u}_i(k) \gamma_{i,-k,\downarrow}^{\dagger}, \quad (2.61)$$

where u and v are given by

$$u_i(k) = \sqrt{\frac{1}{2} \left(1 + \frac{\xi_i(k)}{E_i(k)} \right)} \cdot \exp\left(+\frac{i}{2} \arg \Delta_i(k)\right), \quad (2.62)$$

$$v_i(k) = \sqrt{\frac{1}{2} \left(1 - \frac{\xi_i(k)}{E_i(k)} \right)} \cdot \exp\left(-\frac{i}{2} \arg \Delta_i(k)\right), \quad (2.63)$$

and where E is given by

$$E_i(k) = \sqrt{\xi_i(k)^2 + |\Delta_i(k)|^2}. \quad (2.64)$$

The new operators satisfy the fermionic anti-commutation relations, and they therefore describe fermionic quasiparticles which are sometimes known as bogoliubons. In the new basis the Hamiltonian has the form

$$H^{\text{MF}} = \sum_{i \in B} \sum_{k \in K} \sum_{\sigma \in S} E_i(k) \gamma_{i,k,\sigma}^\dagger \gamma_{i,k,\sigma} + \sum_{i \in B} \sum_{k \in K} (\xi_i(k) - E_i(k)) + E^{\text{MF}}, \quad (2.65)$$

which describes free bogoliubons with energy E , meaning that the ground state is the bogoliubon vacuum and the excited states are obtained by applying the bogoliubon creation operators. Since the minimum excitation energy is the minimal value of $|\Delta_i(k)|$, $|\Delta|$ is called the energy gap. The existence of an energy gap is partially responsible for the stability of the ground state in the presence of external perturbations.

2.2.4 Expressions for Ginzburg–Landau parameters

In order to calculate properties of the superconductor analytically, it is common to assume a free electron dispersion relation, that is, a single spherically symmetric band, and an approximate form of the phonon mediated interaction which gives a constant weak attraction between electrons with momenta near the Fermi surface. The assumption of a spherical Fermi surface is inaccurate for almost all superconductors [66]–[68], but scattering results in the averaging of most of the properties over the Fermi surface and the first order corrections from the true shape can be incorporated by correcting the Fermi surface area. In the case where the interaction is weak, it is reasonable to assume that it is constant, since the energy gap is in all cases much smaller than the Fermi energy, but there are certain materials where the interaction cannot be considered weak [25], in which case the phonons have to be considered as a separate species of particle rather than as an effective interaction between electrons.

Some of the quantities in the Ginzburg–Landau theory can be obtained without deriving the full free energy, for example, the thermodynamic critical field is essentially a measure of the free energy difference between the normal and superconducting states. With all the above assumptions it can be calculated in

the BCS theory as [69]

$$\mu_0 H_c^2 = N_F \frac{8\pi^2 (k_B T_c)^2}{7\zeta(3)} \left(1 - \frac{T}{T_c}\right)^2 \quad (2.66)$$

where N_F is the density of states *per spin direction* at the Fermi surface and ζ is the Riemann zeta function. This expression is the only one which we will need to constrain the Ginzburg–Landau coefficients for the materials studied in this work. Other measurements, such as those of the upper critical field as a function of temperature, can often be used to obtain the coefficients more directly. We will see that the most problematic parameter is the effective mass tensor, for which we will rely on fermiological methods like the de Haas–van Alphen effect.

There are many possible sources of inaccuracy of the determined Ginzburg–Landau coefficients. These include deviation from the Ginzburg–Landau theory due to application at temperatures far from the transition temperature, deviation from the BCS relations due to the failure of assumptions like weak coupling and a spherical Fermi surface, and inaccuracy in the measured properties themselves. These factors will lead to inconsistencies in the final observables. The choice of measured properties used to constrain the Ginzburg–Landau coefficients can affect the calculated values by up to 50%. Although this situation may seem dire we will see that there are conclusions which we can draw which are insensitive to errors in the coefficients of this magnitude.

Chapter 3

Computational method

In this chapter, the discretisation of the Ginzburg–Landau equations will be discussed, as well as how the discrete equations can be solved numerically. A major contribution of this work is the implementation of a solver for the time-dependent Ginzburg–Landau equations whose performance scales optimally with the size of the system, so that it can be used for three-dimensional critical current simulations. Unlike previous implementations, the frozen field limit is *not* assumed, all of the material parameters are allowed to vary in space, and both the effective mass and conductivity are allowed to be independently anisotropic.

The spatial discretisation, derived in the framework of discrete exterior calculus, is second-order accurate and valid for cylindrical, rectilinear, and well-centred simplicial complexes, although only rectilinear complexes are supported by the present implementation. This discretisation coincides exactly with a standard finite-differences approach for the Ginzburg–Landau equations, but the derivation provides a clearer geometric justification for the approach.

A fully implicit backward Euler temporal discretisation is used, leading to a system of nonlinear equations which are solved using a geometric multigrid method [70]. The multigrid method is an iterative technique which can be directly applied to the nonlinear equations and converges quickly regardless of the spatial resolution and value of the Ginzburg–Landau parameter unlike simple iterative methods (e.g., the Jacobi and Gauss–Seidel methods [71]). Multigrid methods also offer faster convergence than general-purpose linear system solvers applied to linearised equations or to the Jacobian system for Newton iteration.

In applying the geometric multigrid method to the fully implicit time-depen-

dent Ginzburg–Landau system for the first time, some special considerations were required. Firstly, in the transfer of the wavefunction between grids of different resolutions, care must be taken to parallel transport the components into the correct fibres so that the whole operation remains covariant. Secondly, it is essential for the convergence of the present implementation that the London gauge is used. This gauge is usually avoided in numerical Ginzburg–Landau theory [72] because it requires solving an additional equation, for the scalar potential, as opposed to the more common temporal gauge which simply has $\varphi = 0$.

A limited study of the accuracy and convergence of the solver is given towards the end of the chapter. The chapter concludes with a list of potential improvements in future work.

3.1 Discrete exterior calculus

Having expressed the Ginzburg–Landau equations in terms of the theory of exterior calculus, we can make use of the theory of *discrete* exterior calculus [73] to write down the corresponding discrete equations. Discrete exterior calculus includes the notion of a discrete differential form as well as discrete counterparts of the exterior derivative and Hodge star operators which preserve some important properties of the smooth versions exactly. After also defining the discrete wavefunction and covariant derivative we will obtain discrete equations which turn out to be identical to those of the standard ‘link variable’ technique [56].

3.1.1 Cellular complexes, chains, and discrete forms

The discrete counterpart to the base manifold is a cellular complex. This can be seen as a formalisation of a finite element mesh. An n -dimensional complex consists of a set of cells for each dimension 0 to n , each equipped with a specific reference orientation. A zero-dimensional cell is a vertex, a one-dimensional cell is an edge, a two-dimensional cell is a face, and a k -dimensional cell is called a k -cell. A k -cell can be any manifold which is homeomorphic to a k -ball, and popular choices for numerical applications include simplexes and hypercubes. The boundary of a k -cell consists of signed $(k - 1)$ -cells which are also part of the complex, the sign being positive where the induced orientation of a boundary cell agrees with its reference orientation, and negative otherwise. For example,

an edge has a boundary consisting of two vertices. An orientation on the edge distinguishes the start and end points. Since these are points they have the same (positive) reference orientation, so the start point is given a negative sign and the end point is given a positive sign.

A k -chain is a formal linear combination of k -cells. For example, if e_0 is an edge in a cellular complex, beginning at the vertex v_0 and ending at v_1 , then its boundary can be written as

$$\partial e_0 = v_1 - v_0, \quad (3.1)$$

which is a formal linear combination of zero-cells the v_0 and v_1 with coefficients -1 and $+1$ respectively, and is therefore a zero-chain. Furthermore, e_0 itself can be considered a one-chain, having a coefficient of 1 for e_0 and 0 for all other edges. The boundary operator can be extended from k -cells to all k -chains by linear continuation. That is, if χ is a chain defined by

$$\chi = \sum_i \chi_i \cdot c_i \quad (3.2)$$

where χ_i is a real coefficient and c_i is a k -cell for each i , then

$$\partial \chi = \sum_i \chi_i \cdot \partial c_i. \quad (3.3)$$

This definition has the nice property that applying the boundary operator to a sum of distinct cells with compatible orientations gives the correct boundary of their union. For instance, if e_1 is the edge from v_1 to v_2 , then

$$\partial(e_0 + e_1) = \partial e_0 + \partial e_1 = (v_1 - v_0) + (v_2 - v_1) = v_2 - v_0, \quad (3.4)$$

which is the correct boundary of $e_0 \cup e_1$ since the end point of e_0 and the starting point of e_1 are the same, they are both v_1 .

The discrete analogue of a differential k -form is a k -cochain, that is, a linear map from k -chains to the real numbers. The canonical map from smooth forms to discrete forms on a cellular complex is called the de Rham map, I [74]. It is defined by

$$I(\eta)(c) = \int_c \eta \quad (3.5)$$

for a smooth k -form, η , and a k -cell, c . For general k -chains the de Rham map is given by linear continuation. If a discrete k -form, η , is to be thought of as approximating a smooth form, then the de Rham map gives a means for interpreting the components, $\eta(c)$, on each cell. They are to be thought of as approximating the integral of the smooth form over the corresponding cells.

3.1.2 Discrete exterior derivative

The discrete exterior derivative, d , of a discrete $(k - 1)$ -form, ζ , is defined by

$$(d\zeta)(\chi) = \zeta(\partial\chi) \quad (3.6)$$

where χ is a k -chain. In other words, the discrete exterior derivative is dual to the discrete boundary operator. The discrete exterior derivative maps $(k - 1)$ -cochains to k -cochains like the smooth exterior derivative maps $(k - 1)$ -forms to k -forms. An important consequence of this definition is that the Stokes–Cartan theorem is preserved exactly under the de Rham map,

$$I(d\zeta)(c) = \int_c d\zeta = \int_{\partial c} \zeta = (dI(\zeta))(c) \quad (3.7)$$

for a smooth $(k - 1)$ -form, ζ .

3.1.3 Dual complex and discrete Hodge star

So far in the development of the discrete theory the metric has not been introduced. This reflects the fact that smooth forms, the exterior derivative, and integration do not depend on the metric. This is not the case for the Hodge star. Furthermore, the definition of the discrete Hodge star only approximates its smooth counterpart, unlike the previous definitions which are exact for injected smooth forms. To obtain an accurate discretisation we need to restrict the class of cellular complexes we are considering to those with a sufficiently nice dual complex.

The dual of a cellular complex is another cellular complex which has an $(n - k)$ -cell for each k -cell of the original, henceforth ‘primal’, complex. Given a cell, c , of the primal complex, let c^* be the corresponding cell of the dual complex, and let $c^{**} = c$. The dual complex is connected in a way which is compatible with the primal complex in that the boundary of a dual cell is the dual of the co-boundary of the corresponding primal cell (special consideration is required for cells on the boundary of the manifold, covered in Section 3.1.8). Since an orientation of the primal cells induces a transverse orientation of the dual cells, and since pseudoforms are integrated over submanifolds with transverse orientation, it is natural to represent the Hodge dual of a cochain on the primal complex as a cochain on the dual complex. Only oriented primal cells and transversely oriented dual cells, and therefore primal forms and dual pseudoforms, will be considered,

but in general one can also consider dual forms and the corresponding primal pseudoforms.

It is not possible for the discrete Hodge star to be exact for injected smooth forms like the exterior derivative is. This would require

$$I(\star\eta)(c^*) = (\star I(\eta))(c^*) \quad (3.8)$$

where \star is the discrete Hodge star, but this can't be satisfied for arbitrary η because the dual cells and primal cells are not identical, so the integrals of η over primal cells are not sufficient to determine the integral of $\star\eta$ over every dual cell. Instead, the discrete Hodge star is defined so that this requirement is satisfied for some specific set of smooth forms.

Often there is a natural choice of smooth forms which should be preserved under the discrete Hodge star. For example, consider a polar grid in two-dimensional Euclidean space. Suppose we require that for a radial edge, e , forms with a locally constant (that is, constant on a neighbourhood containing e and e^*) radial component are preserved exactly, which is guaranteed if dr itself is preserved,

$$I(\star dr)(e^*) = (\star I(dr))(e^*). \quad (3.9)$$

A simple calculation shows that if the radial and polar grid lines have equal spacing and the dual cells lie half-way between the primal ones (all according to the polar coordinates), then

$$(\star I(dr))(e^*) = \frac{|e^*|}{|e|} I(dr)(e) \quad (3.10)$$

where $|\cdot|$ denotes the measure of a primal or dual cell. This situation is illustrated in Figure 3.1. Extending this to arbitrary discrete forms, we write

$$(\star\eta)(e^*) = \frac{|e^*|}{|e|} \eta(e) \quad (3.11)$$

where e is a radial edge. It turns out that the same relationship holds for polar edges if $rd\theta$ is preserved (the intensity $\|rd\theta\|$ is constant under the Euclidean metric). In fact, the idea of preserving locally constant forms leads to the same formula for the discrete Hodge star for forms of any degree in any simplicial complex with the circumcentric dual (where the dual vertices are the circumcentres of n -cells)

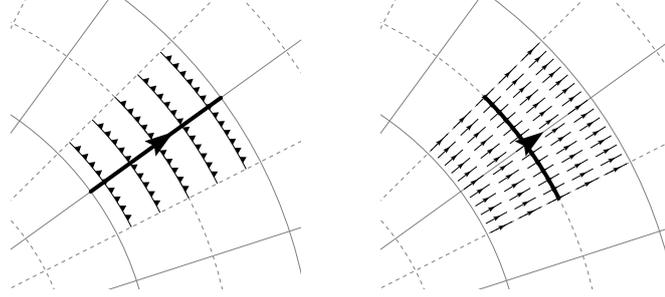


Figure 3.1: For a one-form with a constant radial component, the circulation along a primal radial edge is proportional to the flux of the dual form through the dual edge. The constant of proportionality is the ratio of the primal and dual edge lengths.

and in any rectilinear cubical complex with the barycentric dual. This discrete Hodge star is given by

$$(*\eta)(c^*) = \frac{|c^*|}{|c|} \eta(c). \quad (3.12)$$

This is the definition of the Hodge star which we will use. It is a good approximation of the smooth Hodge star for smooth forms which do not deviate much from constant in the neighbourhood of $c \cup c^*$.

3.1.4 Discrete tensors

The reciprocal mass and conductivity are symmetric (in the sense of linear operators) $(1, 1)$ -tensor fields. The reciprocal mass is essentially the diffusivity of the wavefunction, since it appears in the covariant Laplacian between the two derivatives. In the smooth case, the application of the reciprocal mass before the Hodge star can be combined into a new Hodge star with a new metric. This can be used to define the discrete covariant Laplacian with the reciprocal mass tensor, since we know how to construct a discrete Hodge star given a metric. However, the new metric doesn't necessarily preserve the structure of the complex. A rectilinear complex will not necessarily remain rectilinear under the new metric. A more approximate definition based on the view of a $(1, 1)$ -tensor field as a linear map between one-forms is therefore preferred.

Since tensor fields act linearly on forms multiplied by any smooth function, discrete tensors should be linear on each cell, in which case a discrete $(1, 1)$ -tensor is described by a single component on each edge [75, p. 84]. The discretisation of

a smooth tensor field is accomplished in much the same way as the discretisation of the Hodge star. Each edge has an associated form in its neighbourhood which should be preserved exactly, and the component is the integral along the edge of the result of applying the tensor to that unit form. This leads to a discretisation which has been used previously in discrete Ginzburg–Landau theory [76].

3.1.5 Discrete bundle-valued forms

Whilst the component of a bundle-valued form in some local trivialisation can be represented as a smooth form, it does not make sense to discretise this smooth form as a cochain because the integral of the component across a submanifold is not covariant. Instead, the integral of a bundle-valued form is defined by picking a particular reference point on the submanifold, transporting the components to the fibre over that point, and integrating the resulting form. In general, the result will depend on the paths chosen to transport over, but for a one-dimensional submanifold it does not.

The parallel transport from the end point of a curve, γ , to the start point is given in terms of the connection form, ω , as¹

$$\Gamma(\gamma) = e^{\int_{\gamma} \omega}. \quad (3.13)$$

The covariant integral of a one-form, ζ , with values in the bundle is then given by

$$\int_{\gamma} e^{\int_{\gamma_{[0,t]}} \omega} \cdot \zeta \quad (3.14)$$

where t is the parameter coordinate on γ and $\gamma_{[0,t]}$ is the part of γ with coordinate between 0 and t . When γ is an edge, e , of a cellular complex beginning at vertex v_0 and ending at vertex v_1 , this integral is the value of the covariant de Rham map, $J(\zeta)(e, v_0)$, on the edge e with the reference point v_0 . We represent the bundle-valued form as a map from a cell and a particular reference point on that cell to the complex numbers [77]. We can calculate the integral over multiple cells by adding the values of the discrete form as long as they have the same reference

¹This expression is only valid for bundles with an abelian structure group. When I mention bundle sections and bundle-valued forms I always mean the specific bundle E , whose structure group is $U(1)$.

point. The reference point can be changed from one end of an edge to the other by parallel transport, as

$$J(\zeta)(e, v_1) = e^{-\int_e \omega} \cdot J(\zeta)(e, v_0) = e^{-I(\omega)(e)} \cdot J(\zeta)(e, v_0). \quad (3.15)$$

3.1.6 Discrete covariant derivative

The discrete covariant derivative of a discrete form, ζ , representing the component of a bundle section, on an edge, e , is given by

$$(D_\omega \zeta)(e, v_0) = e^{+\omega} \cdot \zeta(v_1) - \zeta(v_0) \quad (3.16)$$

$$(D_\omega \zeta)(e, v_1) = \zeta(v_1) - e^{-\omega} \cdot \zeta(v_0) \quad (3.17)$$

where v_0 and v_1 are the start and end points of the edge and ω is the discrete connection form. This is just the same expression as for the ordinary exterior derivative except each term has been transported into the same fibre. It can be easily verified that this expression exactly matches the covariant integral of the smooth covariant derivative under the covariant de Rham map, just like the discrete exterior derivative matches its smooth counterpart.

The covariant Laplacian involves taking the covariant codifferential of a bundle-valued one-form. Since the resulting zero-form is evaluated at a point, by choosing that as the reference point for the evaluation of the bundle-valued one-form, all terms are already in the correct fibre so the application of the covariant codifferential becomes identical to the application of the ordinary codifferential. The discrete codifferential of a discrete bundle-valued one-form, η , is then given by

$$(-*^{-1}D_\omega*\eta)(v) = -\frac{1}{|v^*|}(*\eta)(\partial v^*, v). \quad (3.18)$$

The covariant laplacian of a bundle-valued zero-form, ζ , is then given by

$$(-*^{-1}D_\omega*D_\omega\zeta)(v) = -\frac{1}{|v^*|} \sum_{e^* \in \partial v^*} \frac{|e^*|}{|e|} (D_\omega \zeta)(e, v) \quad (3.19)$$

where it has been assumed that the bundle-valued form is linearly continued in the first argument and that the Hodge star has been extended to bundle-valued forms in a trivial way.

The expression for the supercurrent circulation contains a term of the form

$$\mathfrak{I}(\overline{\hat{\psi}} D_{\hat{\omega}_A} \hat{\psi}) \quad (3.20)$$

which cannot be discretised exactly. Knowing the integral of the covariant derivative along the edge and the values of the wavefunction at the end points is not sufficient to calculate the supercurrent circulation along the edge. We therefore resort to a central differences approximation,

$$\begin{aligned}\rho_s(\omega, \zeta)(e) &= \mathfrak{I}\left(\frac{1}{2}\left(\bar{\zeta}(v_0) + e^{-\omega(e)}\bar{\zeta}(v_1)\right)(D_\omega\zeta)(e, v_0)\right) \\ &= \mathfrak{I}\left(\bar{\zeta}(v_0)e^{\omega(e)}\zeta(v_1)\right).\end{aligned}\quad (3.21)$$

The central differences approximation has been evaluated in the fibre over v_0 but since the overall expression is gauge invariant it doesn't actually matter which fibre is chosen, they both result in the same expression.

3.1.7 Non-trivial bundles

The covariant derivative of the component of a bundle-valued form has been defined in a local trivialisation, which makes sense if all the cells involved lie in the patch over which the local trivialisation is defined. If the wavefunction bundle is non-trivial, more than one local trivialisation is required to cover the whole manifold, so there will be neighbouring vertices in the cellular complex where the values of the discrete wavefunction will be components in different local trivialisations.

In the computational method, each vertex and edge in the complex is assigned a local trivialisation so that the stored values of the wavefunction or the vector potential on those cells live in the corresponding local trivialisation, see Figure 3.2 for an illustration. Suppose an edge, e , in local trivialisation over a patch U originates at a vertex, v_0 , in U and ends at a vertex, v_1 , which is assigned a local trivialisation over another patch, V . The stored values of the wavefunction are $\psi_U(v_0)$ and $\psi_V(v_1)$ and the stored value of the vector potential is $A_U(e)$. Since e is entirely in U , the point v_1 lies in both U and V . Hence we can obtain $\psi_U(v_1)$ using the transition map as

$$\psi_U(v_1) = e^{+i\chi(v_1)}\psi_V(v_1) \quad (3.22)$$

where χ is the transition function on the intersection of U and V . The discrete covariant derivative in the fibre over v_0 in U is therefore given by

$$\left(D_{\omega_{A_U}}\psi_U\right)(e, v_0) = e^{-iA_U(e)}e^{+i\chi(v_1)}\psi_V(v_1) - \psi_U(v_0). \quad (3.23)$$

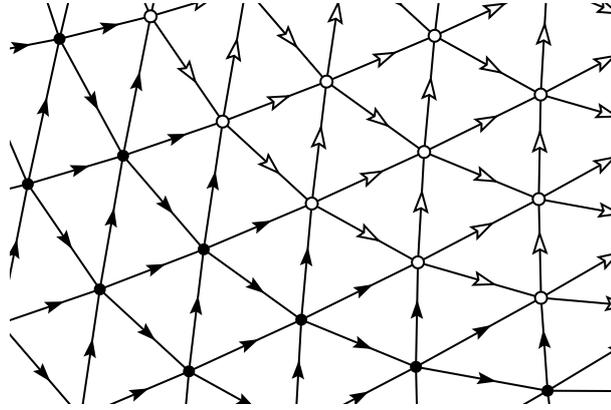


Figure 3.2: A triangular complex with two different local trivialisations, indicated by the vertex and edge colours. The edges which connect vertices of different colours carry a phase difference which is used to compute the covariant derivative along that edge and the magnetic induction on adjacent faces.

The covariant derivative at v_1 can be obtained by first parallel transporting to the fibre over v_1 in U and then using the transition map from U to V . It is given by

$$\left(D_{\omega_{A_V}} \psi_V\right)(e, v_1) = \psi_V(v_1) - e^{-i\chi(v_1)} e^{+iA_U(e)} \psi_U(v_0). \quad (3.24)$$

One obtains a similar expressions using $\chi(v_0)$ when an edge terminates at a vertex with which it shares a local trivialisation. The expression $D_{\omega_{(A+C)}} \psi$, where A and ψ are the components of the vector potential and wavefunction in the standard local trivialisation for each cell and $C(e)$ is the transition function for the start or end point of each edge, depending on which point shares the same local trivialisation, gives the correct component of the covariant derivative in the local trivialisation of each end point. Essentially, we have combined the parallel transport and change of local trivialisation operators in the definition of the covariant derivative.

For the derivative of the vector potential a similar argument holds. For a face, f , where all but one of its boundary edges lie in the patch U and one lies in the patch V , the magnetic induction is given by $dA_U(f) = (d(A + C))(f)$. Since dA_U is independent of the local trivialisation, it is also the case that $\delta dA_U = \delta d(A + C)$.

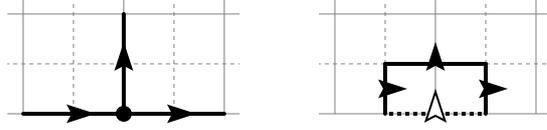


Figure 3.3: For vertices on the boundary of the complex, the dual of the coboundary is not equal to the boundary of the dual. In the case illustrated here, only three of the four dual edges which comprise the boundary of the dual cell can be obtained as the dual of a primal edge. Extra information is therefore required to compute the codifferential of a one-form at this vertex, namely the contribution from the fourth dual edge.

3.1.8 Boundary conditions

Earlier it was stated that the dual complex is connected in such a way that the dual of the coboundary of a primal cell is the boundary of the corresponding dual cell. This is true for internal cells but not for dual cells on the boundary of the complex. Dual cells on the boundary of the complex can have boundary cells which have no corresponding primal cell. An example of such a cell is shown in Figure 3.3. This means that the Hodge star and its inverse do not constitute a bijection between primal and the dual discrete forms and therefore to say that the discrete codifferential of a form, f , is $\delta f = *^{-1}d*f$ is to impose the artificial restriction that the value of $*f$ on the external boundaries of dual cells is zero. As a convention, the definition $\delta f = *^{-1}d*f$ is used but to discretise a smooth codifferential, one must add a correction to account for the flux through the dual boundary elements. This convention is used so that δ remains a linear, rather than an affine, operator.

In the vector potential equation, the expression $\delta d\hat{A}$ requires a boundary contribution. In the discrete theory, $d\hat{A}$ lives on the faces of the complex, so $*d\hat{A}$ lives on the dual vertices in two dimensions or the dual edges in three dimensions. In two dimensions, to compute $d*d\hat{A}$ on a dual edge which intersects the boundary of the complex we need to know the contribution from the dual vertex where it intersects the boundary, which can't be obtained as the dual of a primal face. According to the smooth boundary conditions, we have $\kappa^2*d\hat{A} = \hat{H}$ on the boundary, so there is a contribution from the dual boundary vertex which is equal to the value of the external field pseudoform on that vertex. Thus the discretisation of $\delta d\hat{A}$ is written as $\delta d\hat{A} + f_{\Lambda}$ where f_{Λ} is a primal one-cochain

given by

$$f_A(e) = \frac{1}{\kappa^2} \frac{|e|}{|e^*|} H(\partial e^* \cap \text{closure}(\partial M)) \quad (3.25)$$

where H is a dual zero cochain representing the external magnetic field, M is the full complex, and $\text{closure}(\partial M)$ is the boundary subcomplex of M .

3.2 Multigrid Ginzburg–Landau solver

3.2.1 Discrete Ginzburg–Landau equations

Using the definitions from the previous section, the dimensionless time-dependent Ginzburg–Landau equations can be written in discrete form using an backward Euler temporal discretisation [78] as

$$\frac{\psi_1 - e^{-ik\varphi_1}\psi_0}{k} = (\alpha - \beta\|\psi_1\|^2)\psi_1 + *^{-1}D_{\omega_{A_1+C}} *m^{-1}D_{\omega_{A_1+C}}\psi_1, \quad (3.26)$$

$$-\sigma\left(\frac{A_1 - A_0}{k} + d\varphi_1\right) = \kappa^2(\delta d(A_1 + C) + f_A) - *^{-1}dH_1 - m^{-1}\rho_s(\omega_{A_1+C}, \psi_1). \quad (3.27)$$

where k is the time step size and ψ_1 and ψ_0 denote the discrete wavefunction at the current and previous time steps respectively, and likewise for the vector potential, A , scalar potential, φ , and magnetic field, H . The phase factor multiplying ψ_0 in the left-hand side of the wavefunction equation results from writing the covariant time-derivative in the same way as the covariant spatial derivative, approximating the integral of φ along the temporal edge as $k\varphi_1$ [72]. We use a backward Euler discretisation because for large κ or high spatial resolution, the forward Euler method is only stable if the time step is extremely small [34]. The semi-implicit Crank–Nicolson is often used for diffusive equations like these [35], since it is unconditionally stable and second-order accurate in time, but the first-order backward Euler method is still preferred since taking k to infinity gives the discrete time-independent Ginzburg–Landau equations in the limit, and the object of critical current simulations is to determine the existence or non-existence of time-independent solutions.

Although the discrete Ginzburg–Landau equations are non-linear, it is convenient to express them in a common form which resembles an affine equation

where the linear operator is a function of the dependent variable. For example, the wavefunction equation can be expressed as

$$L^\psi(\psi_1, A_1, \varphi_1)(\psi_1) = F^\psi(\psi_0) \quad (3.28)$$

where the linear operator, L^ψ , is given by

$$L^\psi(\psi_1, A_1, \varphi_1) = e^{ik\varphi_1} \left(\frac{1}{k} - (\alpha - \beta \|\psi_1\|^2) - D_{\omega_{A_1+C}} * m^{-1} D_{\omega_{A_1+C}} \right) \quad (3.29)$$

and the right-hand side is given by $F^\psi(\psi_0) = \psi_0/k$. Similarly, the vector potential equation can be written as

$$L^A(A_1) + R^A(\psi_1, A_1, \varphi_1) = F^A(A_0) \quad (3.30)$$

where the linear operator is given by

$$L^A = \frac{\sigma}{k} + \kappa^2 \delta d, \quad (3.31)$$

the nonlinear part, R^A , is given by

$$R^A(\psi_1, A_1, \varphi_1) = \sigma d\varphi_1 - m^{-1} \rho_s(A_1 + C, \psi_1), \quad (3.32)$$

and the right-hand side is given by

$$F^A(A_0) = *^{-1} dH_1 + \frac{\sigma A_0}{k} - \kappa^2 (\delta dC + f_A). \quad (3.33)$$

Finally, an equation for the scalar potential can be obtained by taking the discrete codifferential of the vector potential equation, and it is given by

$$L^\varphi(\varphi_1) + R^\varphi(\psi_1, A_1) = F^\varphi(A_0) \quad (3.34)$$

where

$$L^\varphi = \delta \sigma d, \quad (3.35)$$

$$R^\varphi(\psi_1, A_1, \varphi_1) = \delta \left(\frac{\sigma A_1}{k} - m^{-1} \rho_s(A_1 + C, \psi_1) \right), \quad (3.36)$$

$$F^\varphi(A_0) = \frac{\delta \sigma A_0}{k}. \quad (3.37)$$

The natural boundary conditions of the Ginzburg–Landau equations specify $*d\hat{A}$ in terms of the external field \hat{H} . From Ampère’s law, this implies that the

total current in the system is the integral of the external current density $d\hat{H}$. This is because the free energy is always minimised when the net supercurrent is zero. However, we would like to constrain the current density in the superconductor to determine the critical current density. To do this, we modify the boundary condition on the vector potential, so

$$i_{\partial}^*(\kappa^2 * d\hat{A} - (\hat{H} + H')) = 0 \quad (3.38)$$

where dH' integrates to the desired current on the boundary of some cross-section of the system. In the discrete system this corresponds to a simple modification of the boundary term, f_A .

3.2.2 Gauss–Seidel relaxation

A general linear equation, $L(x) = f$, where L is a linear operator, x is an unknown vector, and f is a known vector, can be written in component-form in some basis as

$$\sum_j L_{i,j} x_j = f_i \quad (3.39)$$

where i and j are indices for the basis vectors and the sum ranges over the whole basis. Gauss–Seidel relaxation [71] is an iterative method for approximately solving such an equation. By splitting the linear operator into diagonal and off-diagonal parts, one can rearrange the equation to put the dependent variable on one side like so

$$L_{i,i} x_i + \sum_{j \neq i} L_{i,j} x_j = f_i, \quad (3.40)$$

$$x_i = \frac{f_i - \sum_{j \neq i} L_{i,j} x_j}{L_{i,i}}. \quad (3.41)$$

The exact solution must satisfy the above equation, and under certain conditions on the structure of L , one can iteratively approach a solution by repeatedly updating the components of x with the right-hand side,

$$x_i \leftarrow \text{iter}(L, f, x, i) = \frac{f_i - \sum_{j \neq i} L_{i,j} x_j}{L_{i,i}}. \quad (3.42)$$

Since the components are updated one-at-a-time, and the updated vector is immediately used to calculate the next updated component, the order in which the updates occur can affect the result.

The Gauss–Seidel method works equally well for non-linear systems of equations like the discrete Ginzburg–Landau equations. In this case the basis vectors are the unit cochains,

$$\delta_c(c') = \begin{cases} 1 & c = c' \\ 0 & c \neq c' \end{cases}, \quad (3.43)$$

and the dependent variables are updated according to

$$\psi_1(v) \leftarrow \text{iter}(L^\psi(\psi_1, A_1, \varphi_1), F^\psi(\psi_0), \psi_1, v), \quad (3.44)$$

$$A_1(e) \leftarrow \text{iter}(L^A, F^A(A_0) - R^A(\psi_1, A_1, \varphi_1), A_1, e), \quad (3.45)$$

$$\varphi_1(v) \leftarrow \text{iter}(L^\varphi, F^\varphi(A_0) - R^\varphi(\psi_1, A_1, \varphi_1), \varphi_1, v). \quad (3.46)$$

Unfortunately, for these equations, Gauss–Seidel relaxation converges extremely slowly. This is because the matrix row for each cell only has non-zero components for the cell’s direct neighbours, hence the updated value of each component only uses information from the nearest neighbours. When the diffusion coefficients are high compared with the mesh resolution, it takes many iterations for the necessary information to propagate across the mesh. However, fast convergence can be achieved by using a multigrid method combined with Gauss–Seidel relaxation.

3.2.3 Non-linear multigrid

Multigrid methods solve the issue of slow convergence of the Gauss–Seidel method on fine meshes by maintaining a hierarchy of approximations to the solution on several meshes of decreasing resolution [70]. It is quicker to solve the problem on a coarse mesh, interpolate the solution to a fine mesh as an initial estimate, then solve on the fine mesh, than it is to solve the problem on the fine mesh to begin with.

A discretised partial differential equation on a fine mesh can be represented as a system of non-linear equations of the form

$$F^0(u^0) = f^0 \quad (3.47)$$

where u^0 is a vector whose components are the the values of the solution on the vertices (or edges, faces, etc.) of the mesh, f^0 is the constant right-hand side vector, and F^0 is some non-linear function. Given an estimate, v^0 , of the solution, the residual, r^0 , is defined by

$$r^0 = f^0 - F^0(v^0) \quad (3.48)$$

so that if the estimate is perfect then the residual is zero. This leads to the so-called residual equation [79],

$$F^0(u^0) - F^0(v^0) = r^0. \quad (3.49)$$

If we have a true solution we can compute the error, that is, the difference between the solution and our estimate,

$$e^0 = u^0 - v^0. \quad (3.50)$$

The basic idea of non-linear multigrid is to solve a coarse version of the residual equation, compute the error, and interpolate the error to correct the fine grid estimate. In order to do this, we need to know how to convert the fine problem into a coarse problem and how to interpolate the coarse correction to the fine mesh.

If we have two meshes which approximate the same geometric domain at different levels of detail, the function which maps a form on the lower-detail coarse mesh to an ‘equivalent’ form on the higher-detail fine mesh is called interpolation, and the map in the other direction is called restriction. If the discrete form on the fine mesh approximates a smooth form, then the restricted form should just be a coarser approximation of that same smooth form, and vice versa for interpolation, although the approximation can’t be any better in this case because forms on the coarse mesh have fewer degrees of freedom than on the fine mesh.

Given a discrete form on a simplicial complex, the Whitney map is a general method for constructing a piecewise smooth form, that is, interpolating the values of the discrete form within each simplex. This might be used to define an interpolation operator from a coarse mesh to a fine mesh. Once we have a linear interpolation operator, a good choice for the restriction operator is simply the transpose of interpolation, scaled by a constant factor so that constant forms are restricted exactly [79] (at least in the case of zero forms). However, in the case of well-centred simplicial complexes, it is tricky to define a general method for constructing a coarser approximation to a given fine mesh. Instead, only regular rectilinear meshes will be considered, in which case constructing a coarser mesh is trivial: simply double the grid spacing along each axis.

Let ‘interpolate’ be the interpolation map from a coarse mesh to a fine mesh, and let ‘restrict’ be the restriction. The specific expressions for these operators for discrete forms on rectilinear meshes will be discussed shortly. Given an

approximation to the solution of our non-linear system of equations on the fine mesh, v^0 , we can produce an approximation on the coarse mesh given by $v^1 = \text{restrict}(v^0)$. If F^0 is expressed only in terms of operations from discrete exterior calculus then we can construct an analogous function on the coarse mesh, call it F^1 . We can therefore write the residual equation on the coarse grid as

$$F^1(u^1) = F^1(v^1) + \text{restrict}(r^0). \quad (3.51)$$

Once we solve this coarse equation for u^1 , which is easier than solving the fine grid equation, we can compute the correction, $u^1 - v^1$, and update our fine grid estimate as

$$v^0 \leftarrow v^0 + \text{interpolate}(u^1 - v^1) \quad (3.52)$$

which should be a better approximation depending on the properties of F^0 .

We can keep recursively applying this two-grid correction scheme on coarser and coarser meshes until we are left with a mesh consisting of only a few vertices, on which it's trivial to solve the non-linear equations. If we have an iterative method for improving the estimate on each level, call it 'relax' so that $\text{relax}(v^0)$ is a better estimate than v^0 , then a single multigrid cycle consists of the following procedure:

1. 'relax' a few times on the current level (initially level zero, the finest grid)
2. restrict the current estimate and residual to the next (coarser) level
3. compute the right-hand side of the coarse residual equation using the restricted estimate and residual
4. perform a multigrid cycle on the coarse level (recurse)
5. compute the correction on the coarser level
6. correct the estimate on the current level using the interpolated correction
7. 'relax' a few more times

On the finest level the right-hand side is a datum which is provided as part of the initial problem, and on the coarsest level we just 'relax' a few times since there's no coarser grid to transfer to. This procedure is known as a V-cycle because in each cycle the estimate is transferred to coarser and coarser, then finer and finer grids. An alternative is the W-cycle, which spends longer on the coarser grids by

replacing the recursive step with two multigrid cycles on the coarser grid instead of one.

There is much more to the theory of multigrid methods than is presented here [79], but the above forms the basis of the algorithm used in this work. The circumstances under which the multigrid method converges have not been explained, but it requires that the problem is well-posed, so that iterations at different levels converge to the same solution, and that the solution is smooth so that the interpolation and restriction operations retain sufficient accuracy that the correction is an improvement.

In the case of the discrete time-dependent Ginzburg–Landau equations, the solution, u , is the triple (ψ, A, φ) , the non-linear operator, F , is given by

$$F(\psi, A, \varphi) = \left(L^\psi(\psi, A, \varphi), L^A(A) + R^A(\psi, A, \varphi), L^\varphi(\varphi) + R^\varphi(\psi, A) \right) \quad (3.53)$$

and the right-hand side on the finest level is given by

$$f = \left(F^\psi(\psi_0), F^A(A_0), F^\varphi(\varphi_0) \right) \quad (3.54)$$

where (ψ_0, A_0, φ_0) is the state from the previous time step (although φ_0 is not used). The interpolation and restriction operators for the state will be defined shortly. In order to make the system well-posed, a particular gauge choice is required. A common choice for this system is the temporal gauge, $\hat{\varphi} = 0$, since this eliminates the scalar potential equation, which on its own is an elliptic boundary value problem and therefore particularly slow to solve with simple iterative methods. However, multigrid methods are good at solving Poisson's equation, and so we instead choose the London gauge, where $\delta \hat{A} = 0$ and $\star \hat{A} = 0$ on ∂M . This means that the expression $\delta d \hat{A}$ in the vector potential equation becomes $\Delta \hat{A}$, where Δ is the Laplace–de Rham operator. This makes the solutions smoother, which allows the multigrid method to converge. The equations are still under-determined in this form, and this will be discussed in Section 3.2.7.

3.2.4 Rectilinear meshes

Whilst many of the operations in exterior calculus can be defined for wider classes of complexes, the present implementation is specialised to rectilinear meshes. There are two reasons for this. Firstly, the interpolation and restriction operators described here can only be implemented for parallelepipeds and the definition of

the discrete Hodge star requires that those paralleletoes are in fact rectilinear so that the dual cells are orthogonal to their primal counterparts. Secondly, operators like the Laplacian, interpolation, and restriction can be computed on rectilinear meshes without explicitly storing their components as matrices. This drastically reduces the memory usage of the solver which means that larger systems can be studied.

The manifold we will discretise consists of an n -ary product of intervals and circles. For example, $M = S^1 \times [0, 1] \times [0, 1]$ is a three-dimensional manifold which is ‘periodic’ in the 1-direction and has boundaries in the 2- and 3-directions. There are natural dimensionless coordinates for S^1 and $[0, 1]$, and therefore for any product of these. One can define a constant diagonal metric (so the curvature is zero) in terms of these coordinates to give the manifold the desired dimensions, then the coordinates can then be rescaled in terms of the metric so that they have dimensions of length, and these are these rescaled coordinates, q^i , are the ones we will use.

A primal cell of a rectilinear mesh represents a k -dimensional submanifold of M . The points of this submanifold are

$$K(I, X) = \left\{ p \in M \mid \forall x \in 1 \dots n. \begin{cases} q^x(p) \in [I_x h_x, (I_x + 1)h_x] & x \in X \\ q^x(p) = I_x h_x & x \notin X \end{cases} \right\} \quad (3.55)$$

where $I = (I_1, \dots, I_n)$ is the spatial index of the cell, $X = (X_1, \dots, X_k)$ is a tuple of axis indices, $X_i \in 1 \dots n$, and h_x is the constant grid step size along axis x . The origin of the cell has coordinates $(I_1 h_1, \dots, I_n h_n)$ and it spans the axes X having a side length h_x along each $x \in X$. The orientation on K is the induced one from the orientation $\left(\frac{\partial}{\partial q^1}, \dots, \frac{\partial}{\partial q^n} \right)$ on M . See Figure 3.4 for an example.

One has to choose the dimensions of the manifold and the step size so that the manifold can be divided into $N_1 \cdot N_2 \cdot \dots \cdot N_n$ cells of degree n where $N_x \in \mathbb{N}$. If no directions are periodic then the number of vertices in the complex is then $(N_1 + 1) \cdot (N_2 + 1) \cdot \dots \cdot (N_n + 1)$. In general, the number of k -cells spanning the axes (X_1, \dots, X_k) along axis x is

$$\text{extent}(X)_x = N_x + (\text{if } x \notin X \wedge \neg\text{periodic}_x \text{ then } 1 \text{ else } 0). \quad (3.56)$$

For computational purposes, it’s convenient to have indices which wrap around on periodic axes. To this end, the index, I , of a cell spanning axes X

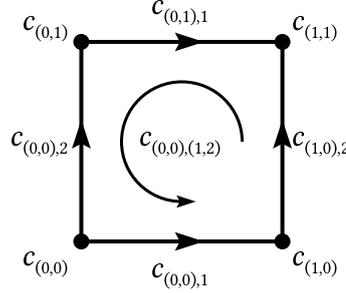


Figure 3.4: Indices and reference orientations of the cells of a small two-dimensional rectilinear mesh.

is defined to be in bounds if it is in the range $0 \dots \text{extent}(X)_x - 1$ along all non-periodic axes x .

$$\text{inbounds}(I, X) = \forall x \in 1 \dots n. \text{ periodic}_x \vee (I_x \geq 0 \wedge I_x < \text{extent}(X)_x) \quad (3.57)$$

If an index is in bounds, then it identifies a cell $K(\text{mod}(I, \text{extent}(X)), X)$ where $\text{mod}(I, \text{extent}(X))$ means taking the modulo of each index component with respect to the extent along each axis. Thus we define the primal cells as the chains

$$c_{I,X} = \text{if inbounds}(I, X) \text{ then } 1 \cdot K(\text{mod}(I, \text{extent}(X)), X) \text{ else } 0. \quad (3.58)$$

Returning the null chain for out-of-bounds indices will prove convenient.

The measure of a cell depends only on the axes which it spans. It is given by²

$$|c_{I,X}| = \prod_{x \in X} h_x. \quad (3.59)$$

The boundary operator is given by

$$\partial c_{I,X} = \sum_{i \in 1 \dots |X|} (-1)^{i-1} (c_{\text{inc}(X_i, I), \text{del}(i, X)} - c_{I, \text{del}(i, X)}) \quad (3.60)$$

where $|X|$ means the length of the tuple X , $\text{inc}(x, I)$ means the index I with the x -component incremented, and $\text{del}(i, X)$ means the tuple X with the i -th entry deleted. The signs reflect whether the reference orientation of the boundary cells agrees with the induced orientation.

²The definitions of the measure and boundary operators assume that the index of the cell which is being operated on is in bounds, otherwise they should return zero.

For each primal cell there is a corresponding dual cell. The boundary of the dual cell must be the dual of the coboundary of the primal cell. Assuming the axes X are unique and sorted, the boundary of a dual cell is given by

$$\partial c_{I,X}^* = \sum_{x \in \bar{X}} (c_{I,\text{ins}(X,x)}^* - c_{I,\text{ins}(\text{dec}(x,I),\text{ins}(X,x))}^*) \quad (3.61)$$

where \bar{X} is the complement of X , that is, the tuple of axes which are not in X , $\text{ins}(X, x)$ means the tuple X with the axis x inserted at the end, and $\text{dec}(x, I)$ means the index I with the x -component decremented. The geometry of the dual complex is defined by giving the locations of the dual vertices. In the rectilinear case the obvious choice is that the dual vertices should be the barycentres of the primal n -cells, in which case the measure of a dual cell is given by

$$|c_{I,X}^*| = \prod_{x \in \bar{X}} \left(\begin{array}{l} \text{if } \text{inbounds}(\text{inc}(x, I), X) \text{ then } \frac{h_x}{2} \text{ else } 0 \\ + \text{if } \text{inbounds}(\text{dec}(x, I), X) \text{ then } \frac{h_x}{2} \text{ else } 0 \end{array} \right). \quad (3.62)$$

From these definitions we can immediately write down the exterior derivative and the Hodge star, then we can write down the codifferential, then the Laplacian, as well as the covariant equivalents.

Evaluating the discrete Laplacian on a zero form at a particular vertex uses only the components of the zero form on that vertex and its nearest neighbours. In a rectilinear complex we can divide the vertices into two classes where the neighbours of each vertex are of the opposite class. This is called red-black colouring, and the vertices are coloured in an n -dimensional checkerboard pattern. When performing a Gauss–Seidel iteration of the discrete Laplacian on a rectilinear grid, if the updates are ordered so that all the vertices of one colour are updated before all the vertices of the other colour, then the update equation for a red vertex does not depend on other red vertices, and likewise for the black vertices. This means that all the red vertices can be updated in parallel, followed by all the black vertices. A similar argument holds for the edges. This is the ordering we will use for the Gauss–Seidel iterations, and this approach is called red-black Gauss–Seidel.

3.2.5 Transfer operators for real-valued forms

It is assumed that a coarse grid is refined by halving the grid step size along all axes. In this case alternating fine vertices, having index I , lie on top of a coarse

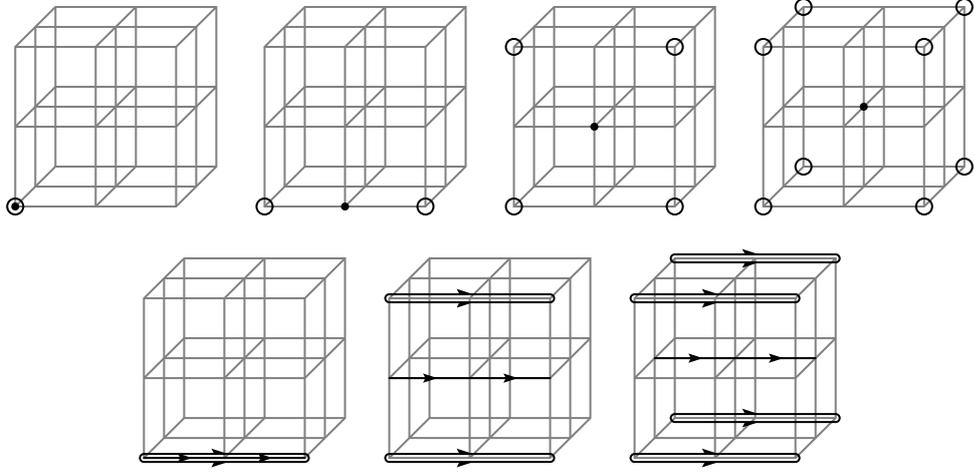


Figure 3.5: Coarse and fine vertices and horizontal edges for a small rectilinear mesh.

vertex with index $I/2$. A form defined on the coarse grid is transferred to the fine grid by standard linear interpolation. If c^0 are the cells of the fine grid and c^1 are the cells of the coarse grid, then a zero-form, f , defined on the coarse grid is interpolated as

$$l(f)(c_I^0) = \frac{1}{2^{\sum_{x \in 1 \dots n} \text{mod}(I_x, 2)}} \sum_{\delta \in \times_{x \in 1 \dots n} 0 \dots \text{mod}(I_x, 2)} f(c_{\lfloor I/2 \rfloor + \delta}^1). \quad (3.63)$$

The idea is that when a fine vertex lies between multiple coarse vertices, the values of f on those vertices are averaged, see Figure 3.5.

Since every coarse vertex has a corresponding fine vertex, transfer from the fine to the coarse grid could be defined by simply sampling the values from those fine vertices, a procedure known as injection. However, it turns out to be better to use the transpose of the interpolation operator, scaled by a constant factor so that constant forms are restricted exactly. Explicitly, the restriction of a zero-form, g , on the fine grid is

$$R(g)(c_I^1) = \frac{1}{2^n} \sum_{\delta \in \{-1, 0, +1\}^n} \frac{g(c_{\text{boundaries}(2I+\delta)}^0)}{2^{\sum_{x \in 1 \dots n} |\delta_x|}} \quad (3.64)$$

where special treatment of the boundary values is required, as given by

$$\text{boundaries}(I)_x = \begin{cases} I_x & \text{periodic}_x \vee (I_x \geq 0 \wedge I_x < \text{extent}_x) \\ I_x - 2 & I_x = \text{extent}_x \\ I_x + 2 & I_x = -1 \end{cases}. \quad (3.65)$$

Indices which are out of bounds are reflected into the interior, which is a standard central differences approximation of the boundary condition $\star dg = 0$.

For one-forms, interpolation is performed along the axes perpendicular to each edge, with injection being used in the parallel direction. For example, in a one-dimensional grid the coarse edge $c_{0,1}^1$ is equal to the sum of the fine edges $c_{0,1}^0 + c_{1,1}^0$. If f is a coarse one-form then the injected value is given by

$$I(f)(c_{0,1}^0) = I(f)(c_{1,1}^0) = \frac{1}{2}f(c_{0,1}^1), \quad (3.66)$$

and conversely, if g is a fine one-form, then

$$R(g)(c_{0,1}^1) = g(c_{0,1}^0) + g(c_{1,1}^0). \quad (3.67)$$

In the one-dimensional case there are no perpendicular directions, but the linear interpolation scheme is identical it's just performed with the injected edge values, see Figure 3.5.

3.2.6 Transfer operators for the wavefunction

The ordinary interpolation and restriction operators should not be used for the wavefunction because they are not covariant. The correct approach is to parallel transport the wavefunction values to the same point when interpolating. Linear interpolation on an n -dimensional grid can be viewed as nested one-dimensional interpolation and the fine-grid parallel transport operators are sufficient to define a covariant version of this operation.

In the two-dimensional case shown in Figure 3.6, there are three types of fine vertices: those that lie on top of a coarse vertex, those that lie between two coarse vertices, and those that lie between four coarse vertices. For fine vertices which lie on top a coarse vertex, the value of the wavefunction is injected:

$$I_{A^0}(\psi^1)(c_{2I}^0) = \psi^1(c_I^1) \quad (3.68)$$

where ψ^1 is the wavefunction on the coarse grid and A^0 is the vector potential on the fine grid. For those that lie between two coarse vertices, the values are averaged in a covariant way:

$$\begin{aligned} 2 \cdot I_{A^0}(\psi^1)(c_{2I+\delta_x}^0) &= e^{+iA^0(c_{2I,x}^0)} I_{A^0}(\psi^1)(c_{2I}^0) \\ &+ e^{-iA^0(c_{2I+\delta_x,x}^0)} I_{A^0}(\psi^1)(c_{2I+2\delta_x}^0) \end{aligned} \quad (3.69)$$

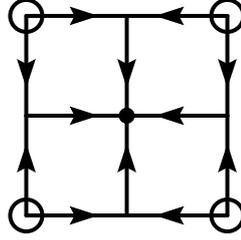


Figure 3.6: Covariant interpolation to the fine vertex (filled circle) at the centre of a coarse face. The values from the coarse vertices (open circles) are transported along the edges indicated by arrows. When several arrows meet, the values are averaged.

where $x \in 1..n$ and δ_x is an n -tuple which is zero for all components y not equal to x and one for $y = x$. For vertices which lie between four coarse vertices, the interpolated edge values are averaged:

$$\begin{aligned}
 4 \cdot l_{A^0}(\psi^1)(c_{2I+\delta_x+\delta_y}^0) &= e^{+iA^0(c_{2I+\delta_y,x}^0)} l_{A^0}(\psi^1)(c_{2I+\delta_y}^0) \\
 &+ e^{-iA^0(c_{2I+\delta_x+\delta_y,x}^0)} l_{A^0}(\psi^1)(c_{2I+2\delta_x+\delta_y}^0) \\
 &+ e^{+iA^0(c_{2I+\delta_x,y}^0)} l_{A^0}(\psi^1)(c_{2I+\delta_x}^0) \\
 &+ e^{-iA^0(c_{2I+\delta_x+\delta_y,y}^0)} l_{A^0}(\psi^1)(c_{2I+\delta_x+2\delta_y}^0)
 \end{aligned} \tag{3.70}$$

where $y \neq x$. Hopefully it's clear how this generalises to higher dimensions. Again, the restriction operator is given by the scaled transpose.

3.2.7 Fixing the average potential

The time-dependent Ginzburg–Landau equations in the London gauge subject to the variational boundary conditions are not well-posed because there is still some gauge freedom. A time-dependent homogeneous function can be added to the phase, and its time derivative to the scalar potential, to give another solution. To make the system well-posed one must fix the average scalar potential [79]. We therefore add the additional condition

$$\int_M \|\hat{\phi}\| = 0. \tag{3.71}$$

Computationally, this is implemented by subtracting the average value of the scalar potential after each multigrid cycle,

$$\varphi \leftarrow \varphi - \langle \varphi \rangle^0 \quad (3.72)$$

where $\langle \varphi \rangle^0 = \sum_{c_I} (*\varphi)(c_I^*) / \sum_{c_I} |c_I^*|$, the numerator being the integral of φ over the manifold and the denominator being the measure of the manifold.

In the smooth case the vector potential does not have the same issue. However, there are technical reasons why taking a similar approach to the scalar potential is beneficial. When taking the limit of large Ginzburg–Landau parameter, one can divide the vector potential equation by κ^2 and take $\kappa \rightarrow \infty$. This gives a vector Poisson’s equation where the spatial variation in the vector potential is determined only by the boundary conditions and not by the supercurrent. However, with the present statement of the vector potential equation, estimates with similar errors have residuals which scale like κ^2 . To get similar solution times for large κ , the residual for the vector potential is divided by κ^2 for the convergence check, which is equivalent to dividing the whole equation by κ^2 , but it is important to maintain accuracy in the time dependence of the average vector potential. This is not an issue for non-periodic directions, because those have Dirichlet boundary conditions which fix the average vector potential, but it is an issue for periodic directions where the time derivative is determined by the normal current densities.

The discrete vector potential equation is

$$-\sigma \left(\frac{A_1 - A_0}{k} + d\varphi_1 \right) = \kappa^2 (\delta d(A_1 + C) + f_A) - *^{-1} dH_1 - m^{-1} \rho_s(\omega_{A_1+C}, \psi_1). \quad (3.73)$$

Averaging the component along a periodic direction, x , gives

$$\frac{\langle \sigma A_0 \rangle_x^1 - \langle \sigma A_1 \rangle_x^1}{k} - \langle \sigma d\varphi_1 \rangle_x^1 = J_x - \langle m^{-1} \rho_s(\omega_{A_1+C}, \psi_1) \rangle_x^1 \quad (3.74)$$

where $\langle f \rangle_x^1 = \sum_{c_{I,x}} f(c_{I,x}) |c_{I,x}^*| / \sum_{c_I} |c_I^*|$ for a rectilinear one-form f and where J_x is the applied current density in the x -direction. Therefore the solution should have

$$\langle \sigma A_1 \rangle_x^1 - \langle \sigma A_0 \rangle_x^1 = k \left(\langle m^{-1} \rho_s(\omega_{A_1+C}, \psi_1) - \sigma d\varphi_1 \rangle_x^1 - J_x \right). \quad (3.75)$$

Given an estimate A_1 , we use this equation to iterate the estimate by first calculating $\langle \sigma A_1 \rangle_x^1$ and then updating the components of A_1 according to

$$A_1(c_{I,x}) \leftarrow A_1(c_{I,x}) + \frac{\left(k \left(\langle m^{-1} \rho_s(\omega_{A_1+C}, \psi_1) - \sigma d\varphi_1 \rangle_x^1 - J_x \right) - \left(\langle \sigma A_1 \rangle_x^1 - \langle \sigma A_0 \rangle_x^1 \right) \right)}{\sigma_{c_{I,x}}}. \quad (3.76)$$

The magnitude of this correction is also included in the convergence check. This approach is similar to that of Sadovskyy et al. [35], except the electric field is encoded in the time dependence of the vector potential instead of an additional gauge parameter.

3.2.8 Implementation summary

The solver described above is implemented as a Julia library which will be made available under a copyleft license. It can be run on a single processor core, suitable for two-dimensional simulations or short-running experiments, or a single graphics card, suitable for long-running three-dimensional critical current simulations. Most of the two-dimensional simulations in this work were performed on a single core of an AMD Ryzen 5 1600 or an Intel Core i9-10900K CPU. Three-dimensional simulations for purposes other than critical current determination were performed on an AMD RX 7600 and an Nvidia GeForce RTX 3080m GPU. Critical current simulations were performed on various Nvidia GPUs on Durham's NCC cluster.

The description of a system consists of the number of highest-degree cells along each axis, whether each axis is periodic, the grid step size along each axis, the time step size, the Ginzburg–Landau parameter, a condensation parameter and nonlinearity parameter value on each vertex, a reciprocal mass and conductivity value on each edge, a change-of-local-section phase difference on each edge, and the value of the exterior derivative of that phase difference on each face. The exterior derivative of the phase difference has to be stored because the phase is supposed to be circle-valued but it's stored as a real number, so the phase difference around a loop is ambiguous up to addition of $2\pi n$.

Discrete differential forms are represented as arrays of components, with some metadata for converting rectilinear cell indices into linear indices. The state of the system consists of the wavefunction, a complex-valued zero form, the vector potential, a one form, and the scalar potential, a zero form. In order

to perform Gauss–Seidel iteration and check for convergence, one must also store the right-hand side and the residual for each discrete equation, which both have the same shape as the state structure since there is one equation for each dependent variable. The parallel transport operators on each edge are also stored as a separate array and updated each time the vector potential estimate is updated.

In a multigrid system with several levels we store a version of all of the above data for each level, plus an additional copy of the state for each level which represents the restricted state from the finer level. This restricted state is used in constructing the right-hand side of the residual equation and computing the correction. The system is initialised by providing the mesh, parameters, material, and initial state for the finest level, along with the number of desired multigrid levels and the coarser levels are constructed by restricting these data.

The multigrid implementation performs fully coupled W-cycles of the backward Euler discrete equations in the London gauge. It performs one red-black Gauss–Seidel relaxation step before transferring to the coarser grid, and it performs two relaxation steps after correcting on the fine grid. The average vector potential is iterated before each multigrid cycle and the average scalar potential is fixed after each multigrid cycle. The infinity-norm of the residual on the finest grid level is used for the convergence criterion.

When stepping the system forward in time, additional data must be provided. These are a one form representing the boundary and field terms in the vector potential equation and the value of the applied current density in each direction, which is used in iterating the average vector potential. The field boundary term and the current density must agree in order for the iteration to converge.

A user of the library will typically construct the solver at the start of their program and then step the simulation forward in a loop. They might change the boundary conditions with each step to vary the current and field with time, and they may also compute and record some observables such as the average electric field, or the magnetic flux in some part of the system. There are some functions available for computing these observables from the state.

3.2.9 Critical current determination

The primary reason for the creation of the present solver was to determine the critical current density for large three-dimensional systems with full control of all

of the material parameters. This involves setting the current and field via specific expressions for the boundary term and phase jump and evaluating the electric field to decide if the system is in a dissipative state. By repeatedly changing the current density and checking whether the system reaches a stable configuration one can systematically search for the critical value.

Consider, for example, a three dimensional domain which is periodic in the 1- and 2-directions, but has insulating boundaries on the upper and lower surfaces in the 3-direction. If we want to apply a field in the 3-direction we have to introduce a non-trivial transition map because the system is doubly periodic in the (1, 2)-plane. If there are n fluxons in this plane then one valid choice of the transition maps has

$$C(c_{(i,\text{extent}(1)_{2-1,k},2)}) = \frac{n \cdot i}{2\pi \cdot \text{extent}(2)_1}, \quad (3.77)$$

$$dC(c_{(i,\text{extent}(1,2)_{2-1,k},(1,2))}) = \frac{n}{2\pi \cdot \text{extent}(2)_1}, \quad (3.78)$$

and all other components zero. If we then want to apply a current in the 1-direction, we can do so via the field boundary conditions on the insulating surfaces. Approximating the tangential field on the free surfaces as uniform, we use the boundary term

$$f_A(c_{(i,j,0),1}) = f_A(c_{(i,j,\text{extent}(1)_{3-1}),1}) = -\frac{J}{\kappa^2} \cdot \text{extent}_3 \cdot h_1 \quad (3.79)$$

which gives the appropriate tangential field difference to satisfy Ampère's law given a uniform current with density J .

An alternative to the above setup has periodic boundary conditions in the 3-direction as well. In this case there are no free surfaces on which to apply a field to satisfy Ampère's law. Instead, we superpose a uniform external current which cancels the first-order variation in the field which would be expected from our desired current. A current must flow in the superconductor to compensate for this externally applied current so that the field remains continuous. This achieved with the external field term

$$(*^{-1}dH_1)(c_{I,1}) = J \cdot h_1. \quad (3.80)$$

A common experimental method of determining the critical current of a real superconductor is to supply an increasing current to the sample and measure the voltage between two fixed points along its length, which is proportional to

the average electric field. When the electric field exceeds a given criterion, it is said that the critical current has been reached. A continuous current ramp is less practical in a simulation because there is a normal inductive contribution to the electric field due to the change in current, and the time scale of the simulation is tiny compared with that of a real experiment, as noted by Carty [80, p. 94]. Instead, we use the same stepped current ramp procedure as Blair [81] and Din [82].

The system is initialised with a small current, $10^{-8}J_0$, and allowed to relax for some time, 4000τ , then the following steps are taken repeatedly until the critical current is reached. The current is increased by some factor, 2%, causing a spike in the electric field, if this electric field drops below a criterion, $10^{-5}E_0$, then the current is increased again and the process repeats, but if it remains above the criterion for a long time, 8000τ , then the critical current has been reached. The criterion is in terms of the average electric field in the same direction as the current, the 1-direction, which is given by

$$E = -\left\langle \frac{A_1 - A_0}{k} + d\phi_1 \right\rangle_1^1 = \frac{\langle A_0 \rangle_1^1 - \langle A_1 \rangle_1^1}{k} \quad (3.81)$$

where, as before, the subscripts on the state variables denote the time step they are evaluated at, whereas the subscripts on the brackets indicate the component which is being averaged.

3.2.10 Numerical experiments

Having implemented the solver, it must be determined whether it accurately solves the discrete equations and whether the performance is as expected from the multigrid method. Unfortunately, exact non-trivial solutions to the Ginzburg–Landau equations in more than one dimension are not available, much less time-dependent solutions. Some simple checks are included as unit tests. These include checking that the flux in systems with non-trivial bundles matches the expected value, checking that the induced electric field depends approximately³ correctly on σ and the applied current in the normal state, and checking that the superconducting state is destroyed at the upper critical field or the de-pairing

³Due to the discretisation and the iterative nature of the solver, only approximate agreement is expected. In this test the grid step sizes along each axis were $(0.3\xi, 0.4\xi, 0.5\xi)$, the time step size was 0.1τ , the tolerance was 10^{-5} , and it was required that the electric field agreed with the expected value with an absolute error of less than $10^{-5}E_0$.

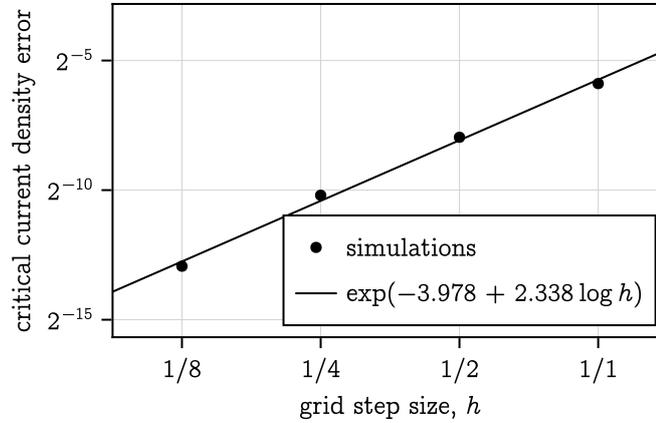


Figure 3.7: Error in the dimensionless critical current density versus grid step size for a Josephson junction system in the weak coupling limit. The system has dimensions $8\xi \cdot 2\xi$, a time step of $\frac{1}{2}\tau$, and a Ginzburg–Landau parameter of 20, and the junction has a thickness of 1ξ and a dimensionless condensation parameter of -20 .

current density. These tests can eliminate factor-of-two-type errors in the terms of the equations and in the boundary conditions.

Another important test for the multigrid London solver was to compare with simpler solvers using other gauges, namely an explicit solver in the temporal gauge and implicit solvers in the temporal and London gauges both using single-grid Gauss–Seidel instead of multigrid. These other solvers are included in the provided code. Since the explicit solver uses a different temporal discretisation, the level of agreement with the multigrid solver depends on the time step size, but the exact gauge invariance of the discrete equations implies that the implicit solvers should agree perfectly with the multigrid solver regardless of the gauge, apart from the small error in the solution due to the non-zero tolerance of the iterative method. This was confirmed by comparing gauge invariant quantities like the magnetic induction between time-dependent solutions like the one in Figure 3.8. Other tests of the accuracy of the solver include the comparison with Brandt’s [52] lattice solutions in Figure 2.2.

It can be seen that the solver accurately reproduces the hexagonal vortex lattice solutions, but the most relevant parameter for which we would like the solver to be accurate is the critical current density. Very accurate approximations

of the current density are available for weakly coupled, approximately one-dimensional Josephson junction systems. For a junction of thickness d in which the condensation parameter is decreased to $\alpha_n/\alpha_0 < 0$, the zero-field critical current density is approximately

$$J_c = J_0 \cdot \sqrt{\frac{\alpha_0}{\alpha_n}} \cdot e^{-\frac{d}{\xi}} / \sqrt{\frac{\alpha_0}{\alpha_n}}. \quad (3.82)$$

Since the critical current density is the boundary below which there exists stationary solutions, the accuracy of the numerical approximation should not depend on the time step. Since the spatial discretisation is effectively a second order central-differences approximation, the accuracy in the numerical critical current should depend on the grid step size as h^2 . Figure 3.7 shows the error in numerically determined values of J_c for a particular Josephson junction system versus the grid step size. The accuracy appears to scale as $h^{2.338}$, which is close to the expected order of accuracy. The difference may be due to the small number of samples.

The reason for implementing the multigrid approach was to improve performance for high κ and for small grid step sizes. To determine whether the implementation achieves this goal, several simulations were run to produce magnetisation loops for a square superconductor with various parameters and the average numbers of multigrid iterations per time step were recorded. One of the generated magnetisation loops is shown in Figure 3.8. Since there are no pinning sites in the system, it is a pure superconductor with insulating boundaries, the irreversibility of the magnetisation is due to the losses associated with the field ramp rate. The results of this experiment are illustrated in Figure 3.9. It can be seen that the number of cycles is virtually independent of κ over a wide range of values, and only weakly dependent on the grid step and time step sizes. For Jacobi and Gauss–Seidel methods the number of iterations drastically increases with decreasing h [35]. It may be concluded that the multigrid approach delivers on its promise of performance which is practically independent of these parameter choices. Figure 3.10 shows a comparison of the real run time of the simulation whose results are shown in Figure 3.8 for three different solvers, illustrating that the multigrid solver is faster in absolute terms as well. It should be noted however that the explicit method and the Gauss–Seidel method do not represent the state-of-the-art in Ginzburg–Landau solvers, but to my knowledge no existing solver exhibits the same scaling with grid size as the multigrid approach.

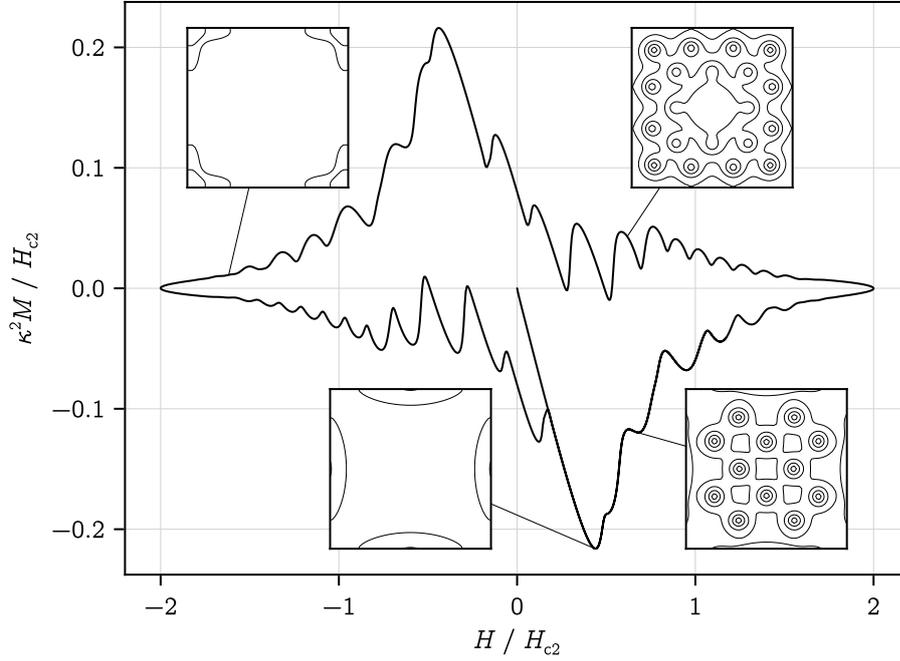


Figure 3.8: A magnetisation loop for a homogeneous square superconductor with dimensions $16\xi \cdot 16\xi$ and $\kappa = 2$. The applied field is given by $H = 2H_{c2} \sin 2\pi t / 2000\tau$. The insets show contours of the wavefunction magnitude (at intervals of $\frac{1}{5}\sqrt{n_0}$) at various points in the loop.

3.2.11 Summary and future work

A second-order finite-differences solver for the time dependent Ginzburg–Landau equations on rectilinear meshes has been implemented which allows for spatial variation of the condensation and nonlinearity parameters, and spatial variation and anisotropy of the the reciprocal mass and conductivity tensors. The solver operates in the London gauge using a backward Euler temporal discretisation. The multigrid approach offers favourable scaling with grid size which enables investigation of the dynamics of the condensate in large three-dimensional systems without assuming the high- κ limit.

To my knowledge, this is the first time the multigrid method has been applied to simultaneously solve both Ginzburg–Landau equations. Various authors have applied multi-level methods to solve various kinds of nonlinear Schrödinger equations in more restricted situations [83]–[86]. Only one of these prior works

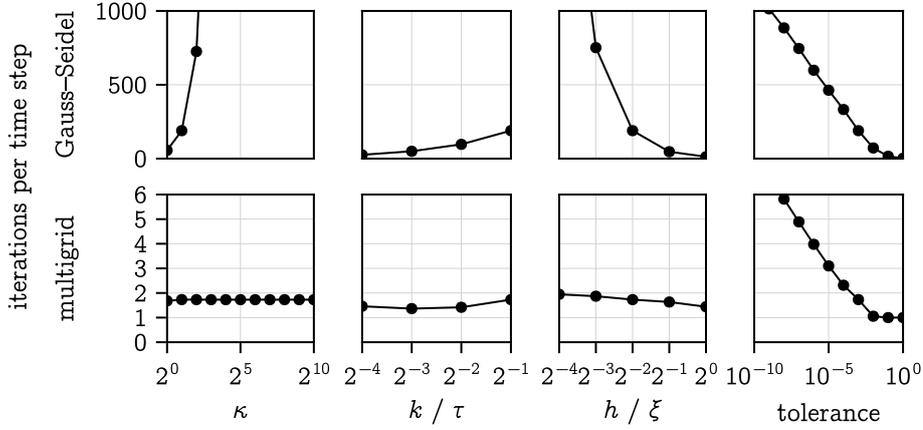


Figure 3.9: Dependence of the mean number of Gauss–Seidel iterations or multigrid cycles per time step on various parameters of the magnetisation loop simulation. Where unstated, $\kappa = 2$, $k = \frac{1}{2}\tau$, $h = \frac{1}{4}\xi$, and the tolerance is 10^{-3} . The Gauss–Seidel solver uses the temporal gauge whereas the multigrid solver uses the London gauge.

solves the wavefunction equation with an external field [85], and it is restricted to constant field (the high- κ limit) and homogeneous and isotropic materials. They use an algebraic multigrid method as a preconditioner for solving the Jacobian system for Newton iteration. None of these prior works simultaneously solves for the vector potential.

The present solver could be further generalised and improved in many ways, some of which are listed below in no particular order.

- Weight by reciprocal mass or conductivity in the interpolation and restriction operators. This could improve the accuracy of the interpolation and therefore further reduce the number of multigrid iterations required for anisotropic materials.
- Generalise to non-rectilinear meshes. For a well-centred simplicial complex, interpolation can be defined using the Whitney map and restriction is obtained from the transpose. However, mesh refinement is more complicated for a simplicial complex, red-black colouring is impossible in general, and storing the matrices of the linear operators explicitly uses a lot of memory.
- Treat the reciprocal mass and conductivity as different metrics which induce

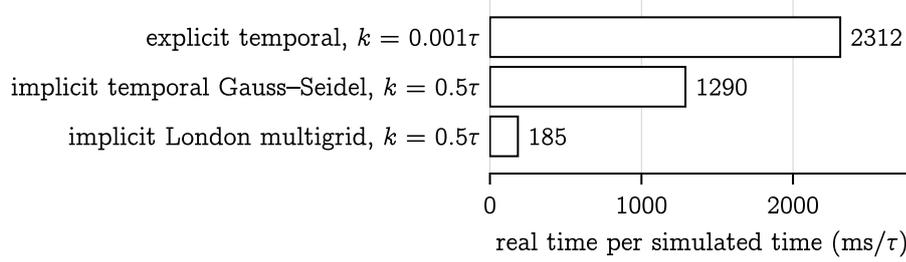


Figure 3.10: Simulation time comparison between three different solvers for the magnetisation loop simulation with $\kappa = 2$, $h = \frac{1}{4}\xi$, and a tolerance of 10^{-3} . Each simulation was run on a single thread on an AMD Ryzen 5 1600 CPU.

different Hodge star operators. For non-rectilinear meshes this would be both simpler and more accurate, as the off-diagonal components could be encoded accurately and the isotropic and anisotropic cases would be treated identically.

- Allow internal boundaries. Especially with non-rectilinear meshes, this could enable more accurate modelling of the geometry of defects. It could also let us couple the Ginzburg–Landau solver in the superconducting domain to a magnetostatic solver in the free space outside the superconductor, which would improve accuracy for certain materials like type-I superconductors and thin films.
- Support simulations in the limit of large Ginzburg–Landau parameter. The discrete equations can be written in such a way that literally setting $\kappa = \infty$, so that $1/\kappa = 0$, gives the correct limit, but the performance benefit is limited because expensive iterations of the vector potential are still performed whilst the other variables are converging. The condensate has no impact on the magnetic induction in this limit, so an optimised solver only needs to compute the induction once for each change in the external field.
- Solve the time-independent Ginzburg–Landau equations using Newton’s method. The time-dependent equations are effectively gradient descent of the free energy functional, so Newton’s method could be a faster way to reach a minimum, potentially speeding up critical current searches. Gradient descent can always be used as a fall-back if Newton’s method fails

to converge. Although multigrid methods are often used for boundary value problems, the method described in this work is not suitable for solving time-independent Ginzburg–Landau equations because they are not well posed due to the existence of multiple extrema of the free energy functional.

Chapter 4

Preliminary simulations of REBCO with APCs

In some work which was published as part of the 2022 Applied Superconductivity Conference [87], we studied pinning due to resistive nanorods in REBCO coated conductors using simulations based on time-dependent Ginzburg–Landau theory. That publication forms the basis of this chapter. The solver described in Chapter 3 did not exist when this work was completed, so a simpler method which is similar to that of Sadovskyy et al. [35] was used. We assumed the frozen field limit, but improved on prior work by considering the consequences of spatial variation in the effective mass.

The critical current densities of high-field superconductors are usually understood in terms of the pinning of fluxons by defects in the superconducting material. In REBCO coated conductors, these defects include natural pinning centers such as dislocations and twin boundaries, and artificial pinning centers such as oxide nanoparticles or nanorods. When, for example, yttrium barium copper oxide (YBCO) is made by pulsed laser deposition, nanorods are produced by doping the YBCO target with barium zirconate (BaZrO_3 , BZO) impurities [88]. In the deposited film, these impurities agglomerate into nanoparticles which self-assemble into columnar nanorods aligned with the crystal c -axis [89], [90]. The nanorods are thought to act as strong pinning centers because of the sharp boundary between the superconducting YBCO and the insulating BZO [91].

Aligned BZO nanorods result in c -axis correlated pinning, with a peak in the critical current density when the magnetic field is aligned with the c -axis. It

has been demonstrated that some degree of splay, so that the nanorods are not oriented perfectly along the c -axis, can broaden this peak and improve the critical current density by producing an entangled vortex matter ground state [92]–[94]. However, it has also been shown that films in which the nanorods are dense and highly aligned across the film thickness demonstrate exceptionally high critical current densities [95], [96].

Time-dependent Ginzburg–Landau theory can describe the interaction of fluxons with pinning centers, and simulations based on this theory have been used successfully to model pinning in coated conductors [97]. However, these simulations typically assume that pinning arises due to modulations in the critical temperature alone, whilst the other material parameters are constant throughout the matrix and the pins. In order to accurately model insulating oxide nanorods, modulation of the carrier effective mass, which is proportional to the normal state resistivity, must be considered.

For this work, a program was developed to solve the TDGL equations in the high- κ limit [98] whilst allowing all material parameters to vary spatially. Critical current data were produced for pinning landscapes with nanorods of varying density, splay, and normal-state conductivity. A simple geometric multigrid method [70] was adopted to solve the Poisson’s equation which constitutes part of the high- κ TDGL system. This significantly improved simulation times and was a stepping stone to the fully coupled nonlinear multigrid solver described in Chapter 3.

4.1 Computational method

The simulation setup used here is far more specialised than the one described in Chapter 3. It assumes the high- κ (frozen field) limit, and only supports three-dimensional fully periodic systems. The conductivity and reciprocal mass tensors are assumed to be equal, the field always points in the z -direction, and the current always flows in the x -direction. The equations are de-coupled by using a leapfrog method and linearised by replacing the $\|\hat{v}\|^2$ term with its value at the previous time step, and a linear multigrid method is applied to the scalar potential equation alone, with the Gauss–Seidel method being used to solve the wavefunction equation. See the original publication [87] for further details.

The distribution of material parameters in the computational domain is designed to model an anisotropic superconductor with embedded non-superconducting nanorods. Intrinsic defects are not included. The free parameters in Ginzburg–Landau theory are α , β , and m , called the condensation parameter, the nonlinearity parameter, and the effective mass. For the superconducting matrix, the z -axis of the system coincides with the crystalline c -axis and is taken to have an anisotropy factor of $\gamma = \xi_{ab}/\xi_c = 3$, which is representative of commercial REBCO tapes [99]. The nanorods are considered circular and isotropic, with a relative condensation parameter of $\hat{\alpha} = -1$ and an effective mass 10 times that in the a - b plane unless otherwise stated. When the maximum splay angle is $\theta \neq 0$, each nanorod is sheared in a random direction in the x - y plane, by a random amount uniformly distributed in $[0, \arctan \theta]$, about its base at $z = 0$. When the maximum splay angle is zero, each nanorod is a cylinder of radius $r = \xi_{ab}$ aligned with the c -axis. The nanorods are distributed by dividing the x - y plane into rectangles and placing a nanorod at a randomly chosen point in each rectangle. A non-superconducting layer is included at the bottom and top of the system to mimic insulating boundary conditions in the z -direction so that the nanorods can be splayed by arbitrary angles without introducing artificial breaks at the system boundary. The dimensionless nonlinearity parameter is set to 1 throughout the entire domain. An example of a computational domain with splayed nanorods is shown in Figure 4.1. The critical current is determined by a well-established upward current ramp procedure [100] using an electric field criterion of $10^{-5}E_0$.

The grid resolution used in the simulation is $h = \frac{1}{2}\xi_{ab}$ and the system dimensions are $64\xi_{ab} \cdot 64\xi_{ab} \cdot 32\xi_{ab}$. Since the grid resolution is fairly coarse, the material parameters are first sampled on a high-resolution grid, four-times the linear resolution of the final simulation, which is then down-sampled to the final grid resolution. In order to implement the boundary condition correctly, it is important for each neighbouring pair of vertices with different condensation parameters ($\hat{\alpha} = \pm 1$) to be connected by an edge with a large effective mass ($\hat{m} = 10$), otherwise the order parameter is artificially suppressed at the edge of the superconducting region. To ensure this, when testing whether each edge is inside a nanorod, a larger radius of $r + \frac{1}{2}h = \frac{5}{4}\xi_{ab}$ is used, since the test is performed at the midpoint of the edge¹.

¹This is approximately correct, but a better approach is to interpolate the effective masses using the fraction of the edge which is inside the cylinder.

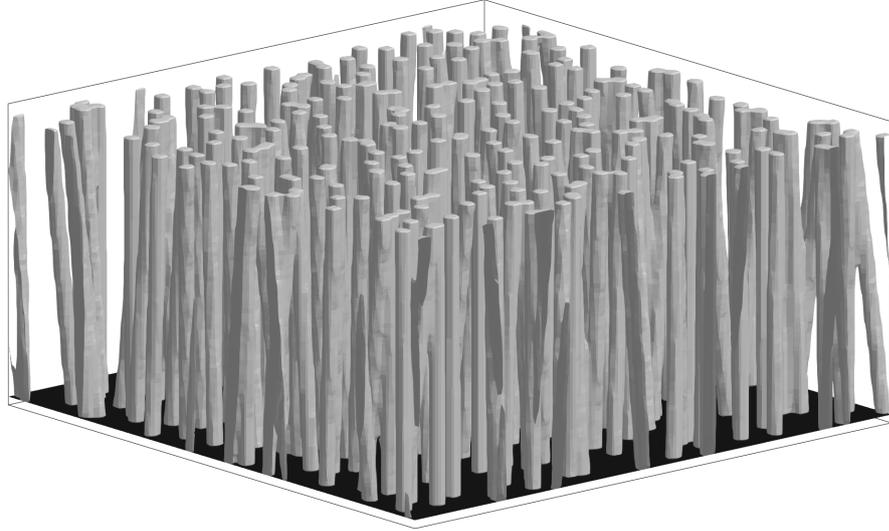


Figure 4.1: An isosurface of the condensation parameter in a computational domain representing a REBCO system with splayed nanorods.

4.2 Results

4.2.1 Effect of nanorod density

Figure 4.2 shows the critical current density versus applied field for highly resistive ($\hat{\alpha} = -1$, $\hat{m} = 10$) nanorods aligned with the c -axis with varying density (for YBCO, $J_0 = J_{0,ab} \approx 1.1 \text{ GA/cm}^2$ [101]).

In low fields and for higher pin densities, each fluxon occupies its own pin, and each pin contains at most one fluxon, resulting in a field-independent critical current density. When the number of fluxons approaches the number of pins (i.e., the matching field), fluxons begin to enter the superconducting matrix. Due to the irregular distribution of pins, where some are very close or overlap, this happens just below the matching field. At fields greater than the matching field, the mechanism which limits J_c switches from the de-pinning of single fluxons to the coherent motion of the vortex matter past the pins. For low pin densities this results in a rapid drop in J_c because the vortex matter must shear only slightly to flow past the sparse matrix of pinned fluxons. At high pin densities the fluxons in the matrix are effectively trapped by the pinned fluxons. In this case, the decrease in critical current density is primarily driven by the decrease

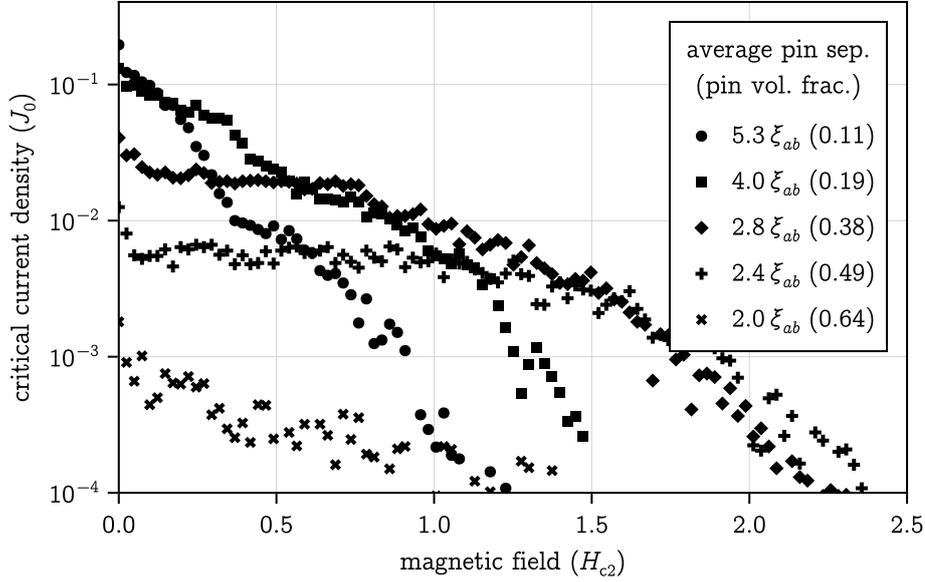


Figure 4.2: Critical current density versus field for aligned highly resistive nanorods with varying density. The separation is the distance between the rod centres.

in the condensate fraction as more fluxons enter the superconducting matrix. Furthermore, there is no significant change in the critical current at H_{c2} , the upper critical field of the matrix, since surface superconductivity persists in the vicinity of the matrix-pin interfaces. In fact, superconductivity can persist to very high fields as the thickness of the superconducting layer between neighboring pins decreases [41], [102], [103]. For the system with very high pin density, an average pin separation of $2\xi_{ab}$, the pins, whose radii are ξ_{ab} , are sufficiently dense to block the supercurrent flow entirely, resulting in a negligible critical current density for all fields.

Figure 4.3 shows the fraction of an a - b cross section which has a superfluid density greater than 1% of the equilibrium value versus field for different pin densities. The dashed vertical lines indicate ‘effective upper critical field’, extracted by linearly extrapolating the high field pinning force versus field data derived from the data in Figure 4.2. Although the effective upper critical field is greater than the upper critical field of the matrix, it does not represent the point at which the superconducting state is destroyed, or even that the critical current density is zero. When the field is such that the suppression of the superconducting state is

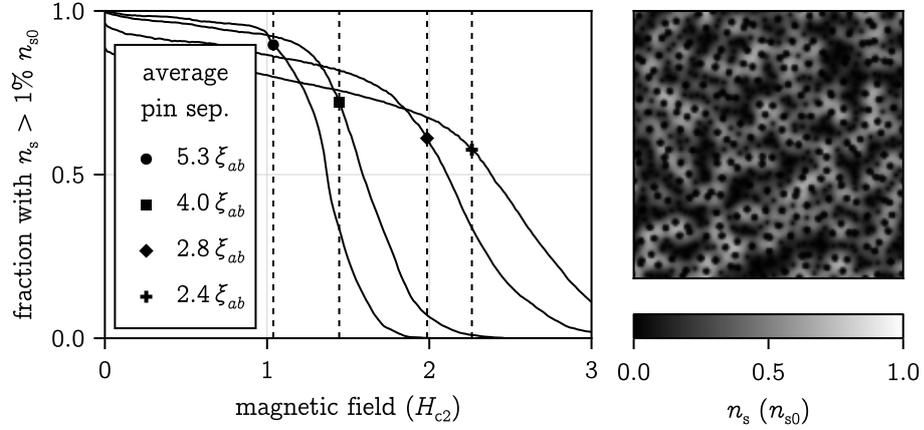


Figure 4.3: (Left) Superconducting fraction versus field for aligned highly resistive nanorods with varying density. Dashed lines indicate the ‘effective upper critical fields’ extracted from a linear fit of the high-field tail of the pinning force data derived from Figure 4.2. (Right) Cross-section in the a - b plane of the superfluid density with an average pin separation of $4\xi_{ab}$ at the effective upper critical field of $1.44H_{c2}$.

sufficient to block the last percolative current path there is a precipitous drop in J_c , which is not visible in Figure 4.2.

4.2.2 Effect of nanorod splay

Figure 4.4 shows the critical current density versus applied field for highly resistive nanorods with average spacing $4\xi_{ab}$ and varying maximum splay angle. Introducing splay reduces the critical current across the whole field range and smooths out the transition from the low-field to the high-field behavior. We attribute this to the reduction in the pinned fluxon length fraction. In high-current HTS wire with BZO APCs, continuous c -axis-aligned nanorods are associated with increased critical current densities [95], [96], but splayed nanorods have also been associated with increased J_c due to the entanglement of fluxons in the vortex matter ground state [92]–[94]. The fact that a small amount of splay reduces the critical current density, and that J_c is substantially independent of splay for larger angles, indicates that some aspects of the real pinning landscape are not yet captured.

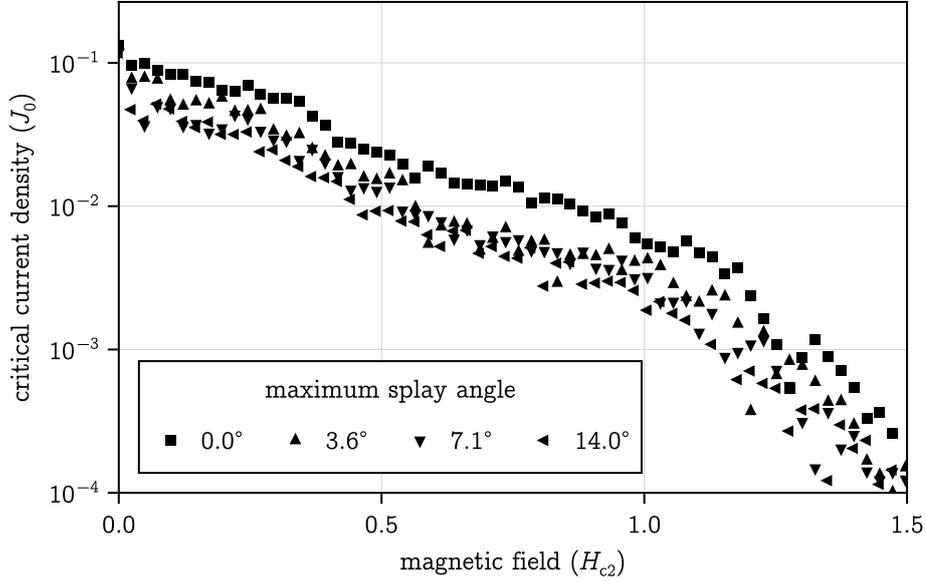


Figure 4.4: Critical current density versus field for highly resistive nanorods with average spacing $4\xi_{ab}$ and varying maximum splay angle.

4.2.3 Effect of effective mass

Figure 4.5 shows the critical current density versus applied field for nanorods with average spacing $4\xi_{ab}$, maximum splay angle 7.1° , and varying effective mass, which is proportional to the normal state resistivity. The data for the material with conducting rods ($\hat{m} = 1$) show the expected precipitous drop in J_c at H_{c2} , whereas the other two materials have non-zero critical currents beyond the upper critical field of the matrix. It is clear that the mechanism which limits J_c in high fields for systems with highly resistive pins is fundamentally different from that of conducting pins. TDGL models which incorporate only variation of the condensation parameter cannot describe this mechanism.

4.3 Conclusion

A TDGL model of REBCO with highly resistive nanorod APCs has been implemented showing high-field critical current density limited by percolative current flow between regions of internal surface superconductivity in the vicinity of

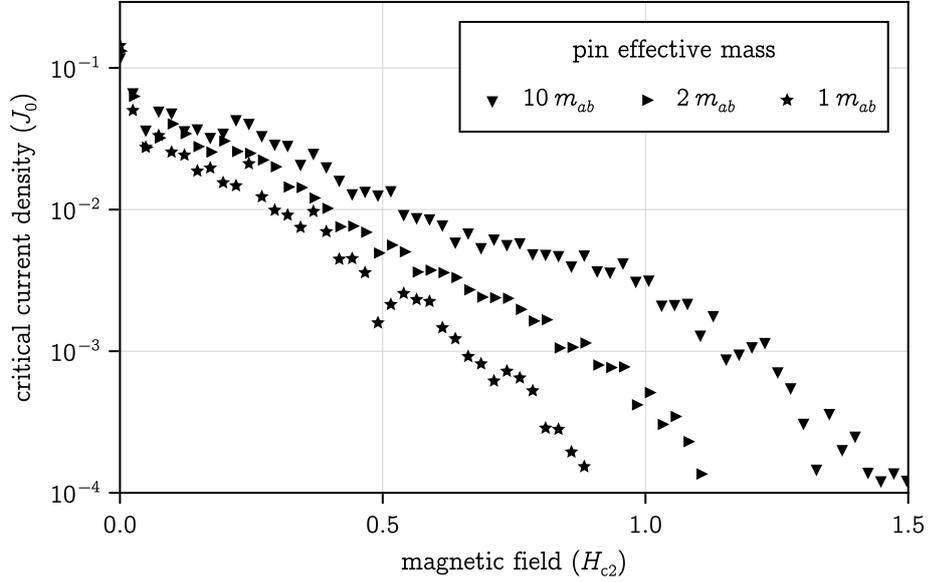


Figure 4.5: Critical current density versus field for nanorods with average spacing $4\xi_{ab}$, maximum splay angle 7.1° , and varying effective mass.

matrix–pin interfaces. Future work might include the strain field which arises due to the lattice parameter mismatch between the matrix and pin material, which is not included in the present model. TDGL simulations were accelerated by use of a multigrid method to solve the Poisson’s equation for the scalar potential. Since the publication of this work, we have developed a new solver which applies the multigrid method to the full TDGL system.

Chapter 5

Simulations of niobium–titanium

Niobium–titanium is the most widely used industrial superconductor. A typical magnet coil is wound from a multifilamentary wire consisting of strands which are made from a billet of niobium–titanium alloy by a procedure of drawing and heat treatments. The initial composition, the number of drawing stages, the final strain, and the heat treatment temperatures and times have been optimised to produce material with a high critical current density at the typical operating field of 5 T at 4.2 K [104]. It is thought that the optimised procedure leads to high critical current densities because it produces an inhomogeneous microstructure which effectively pins vortices, consisting of laminar α -phase (hexagonal close-packed) titanium precipitates embedded in the β -phase (body-centred cubic) niobium–titanium matrix.

Amongst high-field superconducting materials, this system is uniquely suited to modelling with Ginzburg–Landau theory since both phases are superconductors in their own right, so their superconducting properties can be measured and their Ginzburg–Landau coefficients can be calculated. This is in contrast to, for example, nanorod artificial pinning centres in REBCO, which are made of insulating materials which can't be effectively described within Ginzburg–Landau theory, requiring ad hoc choices to be made for their Ginzburg–Landau coefficients. This chapter describes an attempt to simulate pinning in niobium–titanium using time-dependent Ginzburg–Landau theory with parameters which are determined from experimental measurements.

5.1 Elementary pinning force

Pinning in realistic materials is a complex problem because of the many competing forces acting on the vortices. An isolated defect exerts an elementary pinning force on an isolated vortex, but when there are many vortices and many pinning sites the equilibrium configuration also depends on the line tension of each vortex and the interactions between vortices. The simplest part of this problem is determining the elementary pinning force but even this can't be calculated exactly in Ginzburg–Landau theory. In fact, there is no exact closed-form solution even for an isolated vortex in a homogeneous superconductor. Before proceeding directly to the model of niobium–titanium, it will be useful to understand how each of the Ginzburg–Landau coefficients affects the force exerted by a thin planar pin on nearby vortex.

For a homogeneous pin with a fixed geometry embedded in an otherwise homogeneous superconductor, there are three parameters which can affect the dimensionless pinning force. These are the ratios of the condensation parameter, nonlinearity parameter, and effective mass in the pin and the superconductor, $\hat{\alpha}$, $\hat{\beta}$, and \hat{m} . Depending on the values of these parameters, the pinning force can be attractive or repulsive (or repulsive at short range and attractive at long range, for example). In the literature, some authors refer to pinning due to changes in the thermodynamic critical field, as a proxy for the condensation energy, or changes in the Ginzburg–Landau parameter, κ , rather than changes in the Ginzburg–Landau coefficients themselves [60]. A change in the condensation energy is interpreted as core pinning whereas a change in κ is interpreted as magnetic pinning. This categorisation neglects a third type of pinning, due to changes in the effective mass, which results from the redistribution of the vortex currents, which prefer to flow in parts of the system with low effective mass.

Since the precipitates in niobium–titanium are laminar, we will consider the elementary pinning force due to a thin normal layer, with perfectly parallel planar boundaries, embedded in a superconducting matrix. The maximum pinning force occurs when the transport current direction and the magnetic field are perpendicular to each other and both are parallel to the planar surfaces. In this case the problem can be treated in two dimensions. When the properties of the normal layer are such that the pinning force is attractive, there is a minimum in the free energy when the vortex core lies inside the pin, which leads to a positive

critical current density, since a positive Lorentz force is required to push the vortex out of this minimum. When the properties are such that the pinning force is repulsive, there is a maximum in the free energy when the vortex lies inside the pin, and so, in this geometry, the critical current density is again positive, since a positive Lorentz force is required to push the vortex through the pin. This is not the case for all possible pin geometries but it is generically true; especially symmetric arrangements are required to avoid any local minima in the pinning landscape.

If $\hat{\alpha} = \hat{\beta} = \hat{m} = 1$ in the pin then the pinning force is zero, since the pin has exactly the same properties as the superconductor. We can therefore view small changes in the properties as perturbations, in which case the pinning force should depend linearly on those parameters, and we should be able to add the contributions from the changes in each parameter. However, the case where $\hat{\alpha} = \hat{\beta} = \hat{m} = 1$ is singular since the local minimum in the free energy as a function of the vortex position disappears. We can, a priori, have different minima for $\hat{\alpha} < 1$ than for $\hat{\alpha} > 1$ for example, so we should avoid linear extrapolation through this point.

Figure 5.1 shows the dependence of the pinning force for a planar pin of unit thickness on the various material parameters in isolation, as obtained by critical current simulations. Despite the caution against extrapolating through the zero-point, the derivative of the pinning force does appear to be continuous through $\hat{\alpha} = \hat{\beta} = \hat{m} = 1$ in every case. The pinning force data show that, generally speaking, worse superconductors (lower $\hat{\alpha}$, higher $\hat{\beta}$, higher \hat{m}) attract vortices, and better superconductors repel vortices. The fact that reducing $\hat{\alpha}$ has an opposite effect to reducing $\hat{\beta}$ is significant for pinning in niobium–titanium. Changes in κ do affect the pinning force but this effect is diminishing for large κ , and κ is large in most practical high-field superconductors. Once we calculate the Ginzburg–Landau coefficients for niobium–titanium and pure titanium, these data will allow us to estimate the elementary pinning force. However, since this calculation is for an isolated vortex, the overall pinning force in high field will be substantially different.

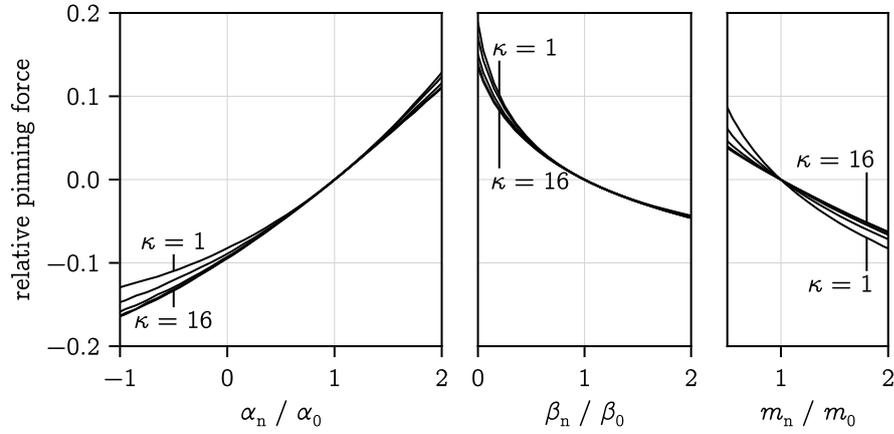


Figure 5.1: Pinning force for a single fluxon on a single planar pin of thickness 1ξ with current parallel to the plane as a function of the relative condensation parameter, nonlinearity parameter, and mass inside the planar pin. The system is $16\xi \cdot 16\xi$ and the different lines are for Ginzburg–Landau parameters $\kappa \in \{1, 2, 4, 8, 16\}$. Unspecified parameters have a value of 1. Positive force represents repulsion and negative force represents attraction.

5.2 Metallurgy and microstructure

Optimised niobium–titanium alloy has a fraction of 0.47 titanium by weight [104], but heat treatments cause titanium to precipitate so that the volume fraction of pure titanium in optimised material is about 0.17 [105]. The conventional lattice parameter for β -phase niobium is 0.33 nm, depending slightly on the composition [106]. For pure α -phase titanium, the lattice parameters are 0.295 nm and 0.468 nm [107]. Calculating the molar density, using the relative atomic masses of niobium and titanium, and the volume fraction of titanium, the atomic fraction of titanium in the β -phase matrix is then 0.56.

Titanium initially precipitates at grain boundary triple points. The precipitates are initially equiaxed and homogeneously and isotropically distributed but when the wire is drawn down, the precipitates become elongated in the drawing direction and distorted into ‘ribbon-like’ structures in the perpendicular plane [108]. Transmission electron microscopy shows that the titanium ribbons become thinner and their separation becomes smaller as the wire is drawn further. In optimised material the ribbon thickness is approximately 1.5 nm [105].

5.3 Ginzburg–Landau coefficients

Both niobium–titanium and titanium exhibit phonon-mediated superconductivity. This means that we should be able to use theoretical formulae to calculate the coefficients in their Ginzburg–Landau free energy functionals but we must decide which parameters have been most reliably determined by experiment, and which theoretical approximations can be justified for these materials.

Assuming the niobium–titanium matrix is purely β -phase with no variation in composition, use measurements of single-phase samples will be used. Muller [109] made resistive measurements of the upper critical field and resistive, inductive, and calorimetric measurements of the transition temperature for single-phase samples with various compositions. Debye frequencies and electronic specific heat coefficients were also extracted from heat capacity measurements. The relevant data are reproduced in Table 5.1. It was found that the values of $H'_{c2}(T_c)$ predicted by the Ginzburg–Landau–Abrikosov–Gor'kov theory¹ using the measured residual resistivity and electronic specific heat coefficient did not agree very well with the measurements of the upper critical field, and that the more general theory (with additional parameters to fit) of Werthamer, Helfand, and Hohenberg [110] was required. However, since our simulations are based on Ginzburg–Landau theory we must somehow extract Ginzburg–Landau coefficients from the measured data, so we assume that the differences from the simple microscopic theory arise from corrections to the coherence length and penetration depth rather than a complete breakdown of the theory. Such corrections can arise from the breaking of certain assumptions, like the assumptions of a spherical Fermi surface and of weak electron-phonon coupling. We will therefore rely mostly on the relations between quantities within Ginzburg–Landau theory, in preference to those derived from microscopic theory.

Three parameters are required to characterise a homogeneous and isotropic phase in a multi-phase system within Ginzburg–Landau theory. For example, the condensation parameter, α , nonlinearity parameter, β , and carrier effective mass, m , are sufficient, but an alternative set are the coherence length, ξ , penetration depth, λ , and effective mass, or the thermodynamic and upper critical fields and the effective mass. These parameters have different temperature dependences so

¹This is the version of Ginzburg–Landau theory with parameters derived from BCS theory [36], [63].

at.% Nb	T_c K	$-H'_c(T_c)$ T/ μ_0 K	$H_{c2}(4.2\text{ K})$ T/ μ_0	γ_{mol} J/mol/K ²
1				
20	7.19	2.59	7.0	0.0051
25	8.27	3.22	10.6	0.0055
30	8.71	3.25	11.4	0.0086
35	9.07	2.98	11.6	0.0063
40	9.23	2.87	11.5	0.0082
50	9.47	2.45	10.7	0.0056
60	9.88	2.09	9.3	0.0042
70	9.86	1.50	7.6	0.0032

Table 5.1: Measured properties of β -phase niobium–titanium from Muller [109].

they are usually quoted at some reference temperature, say $T = 0$, and some conventional temperature dependence is assumed, or they may be directly measured at the particular temperature of interest, which in our case is 4.2 K, the typical operating temperature of a niobium–titanium superconducting magnet.

In Ginzburg–Landau theory, the coherence length can be obtained directly from the upper critical field as

$$\xi = \sqrt{\frac{\phi_0}{2\pi\mu_0 H_{c2}}}. \quad (5.1)$$

The penetration depth is difficult to directly measure, and the expression in Ginzburg–Landau theory depends on the thermodynamic critical field,

$$\lambda = \frac{\phi_0}{2\pi\sqrt{2}\mu_0 H_c \xi}. \quad (5.2)$$

The thermodynamic critical field is related to the free energy difference between the normal and superconducting states, and in BCS theory it can be approximated near the transition as

$$\mu_0 H_c = \sqrt{\mu_0 N_F \frac{8\pi^2 (k_B T_c)^2}{7\zeta(3)} \left(1 - \frac{T}{T_c}\right)^2}. \quad (5.3)$$

where ζ is the Riemann zeta function and N_F is the density of states per spin at

the Fermi level. The Sommerfeld expression for the density of states is

$$N_F = \frac{\frac{3}{2}\gamma}{(\pi k_B)^2} \quad (5.4)$$

where γ is the electronic specific heat coefficient. These formulae can be used together to obtain the characteristic length scales from the electronic specific heat coefficient and the upper critical field.

Notice that the thermodynamic critical field depends linearly on the temperature, and this is true of the upper critical field as well within Ginzburg–Landau theory, due to the assumption that the condensation parameter depends linearly on temperature and the nonlinearity parameter and effective mass are constant in the vicinity of the transition. This is clearly not true of the measured upper critical field, so we must choose whether to use standard Ginzburg–Landau theory, extrapolating linearly from the transition temperature, or a more ad hoc approach of directly applying the Ginzburg–Landau formulae using the measured upper critical field at 4.2 K. Either way we are assuming a linear temperature dependence of the thermodynamic critical field. The calculated length scales for both of these options are shown in Table 5.2. The Ginzburg–Landau coefficients for the linear extrapolation are higher because the measured upper critical field is strictly lower than the linear extrapolation.

For titanium, the superconducting properties are known with less precision than those of niobium–titanium. This is because the transition temperature is very low (less than 1 K) and superconducting titanium has no applications. Still, some measurements have been performed mainly for the purpose of testing BCS theory. For example, Thiemann, Dressel, and Scheffler [111] studied the microwave response of a sample of pure titanium over a frequency range which included the energy gap and were able to obtain the size of the gap, the transition temperature, the penetration depth, and the critical field². Titanium is referred to in their work as a type-I superconductor, but the stated values of the penetration depth and critical field do not support this interpretation according to Ginzburg–Landau theory. Taking their zero-temperature penetration depth and the lower limit on the critical field gives a Ginzburg–Landau parameter of $\kappa = 1.64$, which is greater than the critical value of $1/\sqrt{2}$.

²Although, they say that their measurement of the critical field is not representative of ‘the conventional definition of B_c ’.

at.% Nb	ξ	λ	κ	ξ	λ	κ
1	nm	nm	1	nm	nm	1
20	6.86	391	57.0	6.52	411	63.0
25	5.57	340	61.0	5.01	378	75.4
30	5.37	254	47.4	4.74	289	60.9
35	5.33	278	52.1	4.76	311	65.2
40	5.35	235	43.9	4.77	263	55.1
50	5.55	261	47.1	5.05	287	56.9
60	5.95	261	43.9	5.27	295	56.0
70	6.58	271	41.2	6.23	287	46.1

Table 5.2: Calculated Ginzburg–Landau length scales for β -phase niobium–titanium at 4.2 K. The first set directly use the measured upper critical field at 4.2 K whereas the second set use a linear extrapolation from the data near the transition temperature.

The transition temperature in titanium is sensitive to impurity concentration, differing by a factor of two between a sample with an impurity concentration (by weight) of 10 ppm and one with an impurity concentration of 2000 ppm [112]. This introduces a large uncertainty in the Ginzburg–Landau length scales. The purity of the titanium precipitates in niobium–titanium is not likely to be high, so the coefficients are likely to differ substantially from those of high-purity titanium.

Since niobium–titanium magnets are usually operated at 4.2 K, we need the Ginzburg–Landau coefficients of titanium at this temperature, and since this is above the transition temperature of titanium we have no choice but to extrapolate the coefficients from near the transition temperature. According to the theory of thermal fluctuations, the Ginzburg–Landau coefficients follow approximately the same temperature dependence above the transition temperature as they do below [113], but thermal fluctuations are so weak in titanium due to the low transition temperature that this prediction cannot be verified. Even if the fluctuations were detectable, we have to extrapolate the coefficients so far beyond the transition that they would not provide any assurance of the validity of the extrapolation. This is another source of possible error in the determined coefficients.

Collings and Ho [114] give the electronic specific heat coefficient of high-

purity titanium as $\gamma_{\text{mol}} = 3.36 \text{ mJ/mol/K}^2$. Peruzzi et al. [112] give the transition temperature of a high purity sample as $T_c = 511.3 \text{ mK}$. Using the molar density calculated in Section 5.2 to obtain the volumetric electronic specific heat coefficient we can again calculate the *zero-temperature* thermodynamic critical field as $\mu_0 H_c = 12.2 \text{ mT}$. The critical field temperature derivative is $\mu_0 H'_c = -0.091 \text{ T/K}$ which gives a zero-temperature coherence length of $\xi = 84.1 \text{ nm}$. The penetration depth can then be calculated as $\lambda = 227.2 \text{ nm}$, which agrees reasonably well with the value obtained by Thiemann et al. [111].

In the dimensionless Ginzburg–Landau equations, the parameters of the niobium–titanium matrix are used as the reference values for the condensation parameter, nonlinearity parameter, and effective mass. The quantities we require are then the ratios of these parameters in the normal precipitates to those in the superconducting matrix, α_n/α_s , β_n/β_s , and m_n/m_s . Leaving the effective mass for now, the condensation parameter and nonlinearity parameter ratios can be expressed in Ginzburg–Landau theory with the usual simple temperature dependence (α linear in temperature, β and m constant) as

$$\frac{\alpha_n}{\alpha_s} = \left(\frac{\xi_{s0}}{\xi_{n0}} \right)^2 \cdot \frac{1 - \frac{T}{T_{cn}}}{1 - \frac{T}{T_{cs}}} \cdot \frac{m_s}{m_n} = -0.0232 \frac{m_s}{m_n} \quad (5.5)$$

$$\frac{\beta_n}{\beta_s} = \left(\frac{\xi_{s0}}{\xi_{n0}} \right)^2 \cdot \left(\frac{\lambda_{n0}}{\lambda_{s0}} \right)^2 \cdot \left(\frac{m_s}{m_n} \right)^2 = 0.00240 \left(\frac{m_s}{m_n} \right)^2 \quad (5.6)$$

where the values of the Ginzburg–Landau length scales for pure titanium and β -niobium–titanium with 40 at.% Nb have been substituted. Over a wide mass range the condensation parameter and nonlinearity parameter are small in the pins compared with those in the matrix.

Since the ratio of the coefficients is so small, the large uncertainty in the values for titanium does not lead to a large uncertainty in the pinning force, since according to Figure 5.1, the pinning force depends smoothly on these ratios, so even if the coefficients in titanium were wrong by a factor of two, the ratios would still be close to zero and the pinning force would be essentially the same. In case where $\hat{\alpha} = \hat{\beta} = 0$, the condensation energy in the pins is zero, and the distribution of the condensate is determined only by the values on the boundary, that is, by the proximity coupling. Since the pinning forces resulting from the differences in $\hat{\alpha}$ and $\hat{\beta}$ in this case are opposite in sign and similar in magnitude,

their effects partially cancel. That niobium–titanium has a high critical current density despite this cancellation suggests that most of the pinning force arises from a difference in effective mass, although, as previously noted, this may not hold in high fields.

There is no simple measurement of a macroscopic Ginzburg–Landau observable in a single-phase sample which can be used to determine the effective mass, but the effective mass in Ginzburg–Landau theory is just twice the Landau quasiparticle effective mass averaged (due to collisions) over the region of the Fermi surface where the superconducting gap is present [69], and it can therefore be obtained, at least in principle, using fermiological methods like the de Haas–van Alphen effect. However, it is not possible to obtain a clear de Haas–van Alphen signal for a random alloy like niobium–titanium and, even if it was, it is not simple to take the effective masses for specific orbits on the Fermi surface obtained by de Haas–van Alphen measurements and use them to calculate an appropriate average for use in the Ginzburg–Landau model. For this, a detailed model of the Fermi surface is required, whose parameters can be constrained by the de Haas–van Alphen data.

The most relevant de Haas–van Alphen data in the literature are for pure niobium [67] and pure titanium [115]. The measured effective masses for electron cyclotron orbits in titanium range from $1m_e$ to $3.4m_e$. For niobium the effective masses range from $1.4m_e$ to $5.56m_e$, ignoring the mass stated for ‘ β oscillations’ which the authors say are ‘not from the dHvA effect, but rather a result of oscillatory magnetoresistance’. Based on these data, it will be assumed that the ratio of the averaged effective mass in titanium to that of niobium is greater than (roughly) $\frac{1}{5}$ and less than 3. Assuming the values for niobium–titanium are similar to those of niobium, simulations were run with effective mass values which cover this range to determine which effective mass values are consistent with the observed critical current density in niobium–titanium wires.

5.4 Computational system

The diameter of a real niobium–titanium filament is on the order of $10\ \mu\text{m}$ [104], or 1000ξ , and the length is on the order of kilometres. Since these scales are impractical to capture in a simulation where the grid step must be a fraction of a coherence length, the simulation domain is set up to model a small volume

inside the filament to obtain the local critical current density subject the local magnetic field. When a magnetic field is applied perpendicular to the filament axis and the filament is fully saturated with flux then, neglecting the pinning due to the surface of the filament, the critical current density for the whole filament is the same as the local critical current density over some part of the cross section whose scale is larger than the scale of the inhomogeneity of the microstructure.

Since the microstructure is elongated during the drawing process, the length scale for variation in the drawing direction is thousands of times longer than the length scales in the perpendicular plane [116]. It is therefore also impractical to capture this scale in the simulation and we consider the microstructure to be homogeneous in the drawing direction. This means that the system is translationally invariant in the drawing direction so the component of the pinning force in this direction is always zero in the computational system, that is, the flux line lattice is free to slide along the precipitates. Since the transport current flows in the drawing direction, the average Lorentz force in this direction is also zero. The drawing direction will be called the x -direction. In all critical current simulations the applied magnetic field will lie in the plane perpendicular to the drawing direction and, since the microstructure is isotropic in this plane, the direction is arbitrary. The direction of the applied field in critical current simulations will be called the z -direction. The y -direction is the direction perpendicular to x and z so (x, y, z) are the three Cartesian coordinate functions. We align the grid to these axes and the same symbols are used as labels to index into triples like the grid step, $h = (h_x, h_y, h_z)$.

The grid step must be made small enough to resolve the smallest length scales in the system but not so small that the number of mesh elements is too large and the simulation times become too long. The smallest intrinsic length scale is the coherence length of the niobium–titanium matrix, being much smaller than the penetration depth of the matrix material, and the coherence length and penetration depth of the titanium precipitates. However, the titanium ribbons have a thickness which is a fraction of a coherence length, the smallest thickness used here being 0.3ξ , so this length scale must also be resolved. It was decided that a grid step of $h_x = h_y = h_z = \frac{1}{3}\xi$ was a good compromise between accuracy and simulation time for most of the microstructures studied, and this is the value used for all of the simulations of the niobium–titanium analogue system. The microstructure is rasterised to the simulation grid, weighting by the fraction of

each dual cell which is occupied by the ribbons, so although the grid step is larger than the smallest ribbon thickness, it is found that the results follow the general trends from the thicker ribbons.

We would like to simulation domain as large as possible to model the macroscopic behaviour but we are again limited by simulation time. The cross-section in the y – z -plane should be sufficiently large as to represent a suitable average over the possible distributions of titanium ribbons, and the size in the x -direction should be large enough to include at least a few vortex-lattice-spacings for the value of the critical current density at a particular field to be representative of a macroscopic sample. It is convenient for the dimensions to be powers-of-two, so we use a mesh with dimensions $64 \cdot 128 \cdot 128$ (cells), which corresponds roughly to a volume of $100 \text{ nm} \cdot 200 \text{ nm} \cdot 200 \text{ nm}$.

To allow the current to flow in the x -direction, this direction must be periodic. The system must also be periodic in the y -direction so that there are no free surfaces in the direction of flux flow, whose pinning force would otherwise dominate over the bulk. The z -direction may or may not be periodic. The method for applying a transport current depends on whether it is so, as described in Section 3.2.9. The critical current simulations in this work used insulating boundary conditions in the z -direction with a thin normal coating on each surface to reduce the effects of the artificial boundaries.

All simulations use a time step of $\frac{1}{2}\tau$. The size of the time step should not affect the accuracy of the critical current and critical field values since they are not dynamical properties. Similarly, a relatively loose tolerance of 10^{-3} is used and the conductivity is set to the same value as the reciprocal mass. The parameters of the critical current search procedure are the exemplary values given in Section 3.2.9.

Procedural generation is used to produce microstructures which share similar features to the observed microstructure. Using procedural generation rather than tracing the microstructure from images³ will allow us to systematically vary the parameters of the microstructure, like the ribbon thickness and the overall volume fraction, to see how they affect the critical current, and to guarantee that the microstructure is continuous across the periodic ‘boundaries’ of the

³A previous study of pinning in a high-temperature superconductor, YBCO, used positions and radii of nanoparticles derived from a three-dimensional scanning transmission electron tomogram [97]. The authors were able to reproduce some of the qualities of the experimental field and angular dependence of the critical current in their simulations.

system. Other authors have used Voronoi tessellation for procedural generation of microstructures which resemble grain boundary networks [100]. A drawback of procedural generation is that there may be features which are important for pinning in the real microstructure which we are not aware of, and so are not reproduced in the computational microstructure, or vice versa, features in the computational microstructure which are not present in the real microstructure but which affect the critical current.

Since the system is translationally symmetric in the x -direction, we need only generate the distribution of ribbons in the perpendicular plane. In order to produce a random microstructure which is periodic in both directions, we use a smooth periodic noise function and place the ribbons along the contours of this function. The thickness of the ribbons is controlled directly and their separation is controlled by changing the interval between the function values at which contours are drawn.

To generate a smooth periodic noise function, we begin with a function in frequency space, F , and take the discrete inverse Fourier transform to obtain a function in real space, $\mathfrak{R} \circ \mathcal{F}^{-1}(F)$. The frequency representation is given by

$$F(k) = \sum_i w_i \cdot e^{i\varphi(k)} \cdot e^{-|k|^2 / f_i^2} \quad (5.7)$$

where w_i and f_i are positive real numbers for each i in some finite set, and $\varphi(k)$ is a random phase for each k . Figure 5.2 shows the Fourier transform of two terms in the sum whose f values differ by a factor of 4. The f values chosen are fractions of the ‘base’ value, $f_i = f_0/2^i$, up to some maximum f_n , and the corresponding weights are $w_i = 1/2^{2n}$. With $n = 2$ this was found to produce microstructures which resembled the microscope images to a reasonable degree, although there are of course many other choices which produce similar results.

The contours of the noise function are closed, whereas the precipitates in niobium–titanium have ends. To produce precipitates with a certain length we draw the contours with a ‘dashed’ pattern, lowering the pen and drawing until the ribbon reaches a certain length, then raising the pen and continuing along the contour for some gap size, then lowering the pen again.

An example microstructure is shown in Figure 5.3. The ribbon thickness is controlled directly and the separation between the ribbons is controlled by varying the interval between the isovalues at which the contours are drawn. One could in principle vary the length and smoothness of the ribbons, or draw some

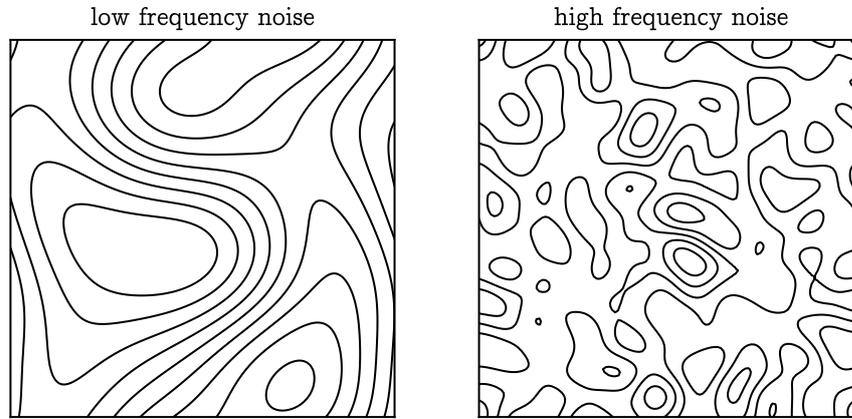


Figure 5.2: Smooth periodic noise functions with different frequency content.

of these variables from probability distributions, but for the purposes of this work only the thickness and separation are varied. The code which implements the microstructure generation is included in the supplementary material.

5.5 Anisotropy of the critical field

Despite niobium–titanium being an isotropic superconductor, it has been observed that the properties of optimised niobium–titanium wire are anisotropic [117]. This could arise from some intrinsic anisotropy which is introduced in the matrix by the drawing process, but it could also arise due to the geometric anisotropy in the microstructure, that is, the fact that the normal precipitates are very long and straight in the drawing direction but shorter and distorted in the perpendicular plane. Since we are assuming that the properties of the niobium–titanium matrix are the same as those of single-phase samples, we should investigate the latter possibility, and how the anisotropy depends on the effective mass ratio between the matrix and the pins.

Chislett–McDonald et al. [117] used resistive transport and calorimetric measurements to determine the critical field for niobium–titanium filaments with the magnetic field oriented in either the longitudinal or the transverse direction at a few temperatures both for filaments as supplied by the manufacturer, and for filaments with the outer copper etched away. They found that in all cases the

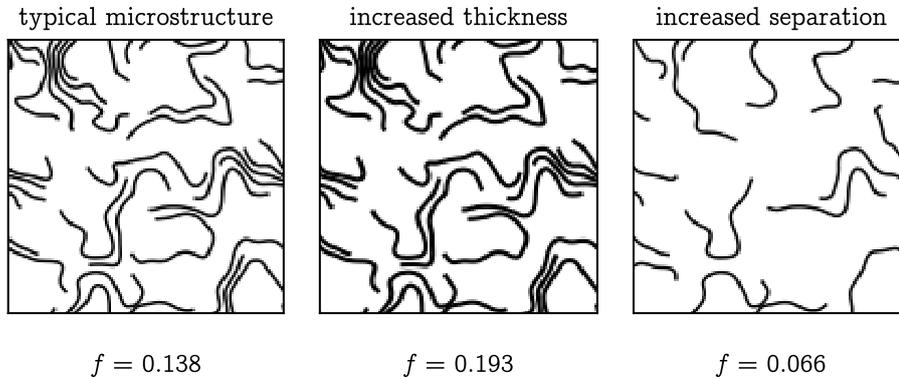


Figure 5.3: Varying two parameters of the microstructure generation procedure and their effect on the precipitate volume fraction. Each image represents part of the cross-section of a strand, having dimensions of roughly $200 \text{ nm} \cdot 200 \text{ nm}$. Black represents the titanium precipitates and white represents the niobium–titanium matrix.

critical field was higher for the longitudinal orientation. This was also the case for critical current measurements at 4.2 K where the critical field in the longitudinal direction was 1.10 times higher than the critical field in the transverse direction. Assuming the difference in the critical field is due to surface superconductivity (see Section 2.1.10) at the matrix-pin interface allowing a lossless current to flow above the critical field of the matrix, it should be possible to reproduce the anisotropy with Ginzburg–Landau simulations if the material properties of the precipitates are correct and the geometry is sufficiently accurate.

To determine the critical field, we set up a computational system as described in Section 5.4 with all three axes being periodic. To vary the field in the longitudinal direction, we vary the number of fluxons in the y – z -plane, and to vary the field in the transverse direction, we vary the number of fluxons (by changing the phase differences between patches in our non-trivial bundle) in the x – y -plane. In either case, we begin at a flux density where the system is completely normal and remove fluxons one-at-a-time, perturbing the wavefunction slightly to encourage the condensate to nucleate, and then allowing the system to equilibrate for some time. After each equilibration period the average superfluid density is recorded. The effective upper critical field is then obtained as the highest field where the

superfluid density exceeds some small criterion. This approach is similar to the experimental calorimetric technique as it measures the bulk superfluid fraction, which is related to the condensation energy.

The effective upper critical field should depend on the ribbon thickness and average separation. For a thin superconducting slab in parallel field the effective upper critical field is inversely proportional to the thickness [41], so it is expected that the critical field for the ribbon microstructure will increase as the separation between the ribbons is reduced. However, when the superconducting layers are bordered not by a perfect insulator but by a normal metal, we expect the critical field to go to zero as the fraction of superconducting material approaches zero [82, p. 66].

Figure 5.4 shows the critical field ratios for a set of microstructures with ribbon thicknesses of 0.3ξ , 0.5ξ , and 0.7ξ , and various separations using a superfluid density criterion of $1\% \cdot n_0$. The dependence of the ratio on the effective mass is roughly independent of the criterion in the range from $10^{-2}n_0$ to $10^{-9}n_0$, although the absolute values of the critical fields do depend on the criterion. As in the real wire, the critical field in the direction parallel to the wire axis is always greater than in the perpendicular direction. This is because the precipitates are elongated along the drawing axis, so in the parallel orientation their interfaces are parallel with the field over much greater lengths than in the perpendicular orientation, and surface superconductivity arises when the field is parallel to an interface.

As the effective mass in the precipitates approaches from above that of the matrix, the critical field in the parallel direction becomes nearly constant with a value slightly less than that of the matrix, and the critical field in the perpendicular direction continues to decrease. As the effective mass in the precipitates increases, there is a regime where the critical fields in both directions depend almost linearly on the effective mass ratio. The critical field ratio is greater than one regardless of whether the effective mass of the precipitates is larger or smaller than the effective mass of the matrix. When the effective mass is larger, the parallel critical field is enhanced more than the perpendicular field, whereas when it is smaller, the perpendicular critical field is suppressed more than the parallel critical field. For all reasonable values of the microstructural parameters, the effective upper critical field ratio of 1.1 is consistent with an effective mass ratio between the pins and the matrix of about 2.

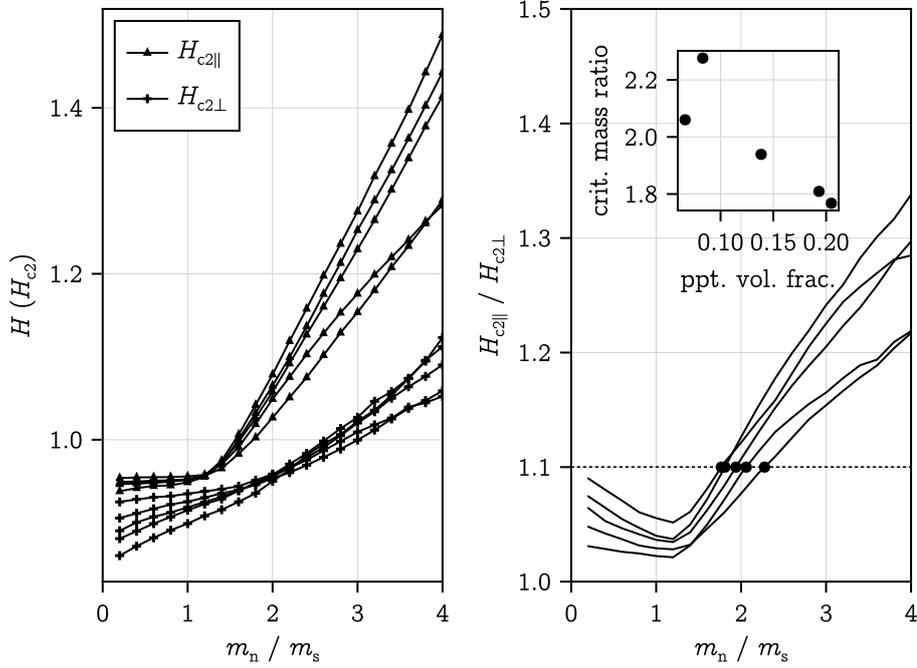


Figure 5.4: The ratio of the effective upper critical fields parallel to, and perpendicular to, the wire drawing direction versus the effective mass in the normal precipitates for a set of ribbon microstructures with various thicknesses and separations. The upper critical field is defined as the field magnitude where the superfluid density drops below 1% of the homogeneous zero-field equilibrium value. Inset: the value of the effective mass which gives a critical field ratio of 1.1 versus the precipitate volume fraction of each microstructure.

5.6 Critical current density

Critical current density simulations have been run in a parameter range which encompasses the range of reasonable values for the precipitate thickness, separation, and effective mass ratio. Since the microstructure is random to some extent, two different samples were also used to check that the qualitative behaviour is similar when the parameters of the microstructure are the same. The collected data are shown in Figure 5.5.

Several observations can be made from the raw data. Firstly, the behaviour of all systems is erratic in low fields, with the critical current density showing

rapid variation with field which is not seen in real materials. Secondly, all of the systems with higher effective mass in the precipitates, and some of those with lower effective mass, show an approximately linear decrease in the critical current density in high fields, which is seen in real materials. Finally, in all cases where the effective mass is the same in precipitates and the matrix, the critical current density is comparatively tiny over most of the field range. Each of these observations will be discussed in more detail.

Figure 5.6 shows an example of some critical current data for a particular system with ribbon microstructure. In this case the effective mass in the precipitates is 4 times that of the matrix, the thickness of the ribbons is half a coherence length, and the volume fraction of second phase material is 0.14. These data exhibit a linear dependence of the critical current density on the applied field close to the upper critical field and erratic variation in low field. This linear dependence is not seen in other pinning landscapes such as those with nanorods, nanoparticles, or grain boundaries (see Chapter 4, Sadovskyy et al. [97], and Blair [81]), so its presence is an indication that the generated microstructure is capturing some important features of the real pinning landscape. However, the erratic behaviour in low-field is not seen in real materials. It is believed that this is an artefact due to the scale of the simulation domain, a possibility which is now explored via a simpler analogue system.

A cross-section in the x - y -plane of a system with ribbon microstructure consists of a series of bands of normal material embedded in a superconducting matrix, in other words, an irregular multilayer. The full system can therefore be considered as a stack of coupled multilayer systems where the thickness of each normal and superconducting layer depends on the z -coordinate of the slice. To simplify the system, we can consider a regular two-dimensional multilayer, which would correspond to a series of ribbons with equal spacing which are aligned with the z -axis. This two-dimensional system is much simpler to simulate and to understand, but still retains some of the unusual (compared to real materials) behaviour of the three-dimensional simulations.

Figure 5.7 shows some critical current data from simulations of a multilayer system where the material parameters of the normal layers are similar to those of the titanium precipitates in niobium–titanium. The system is a doubly-periodic square superconductor with evenly spaced normal planes. The plot shows similar erratic variation of the critical current density in low fields as is seen in the

three-dimensional simulations, despite the vastly simpler microstructure, and the variation remains when the system size and the grid resolution are increased.

After watching the flow of vortices in the multilayer system, it's clear the erratic variation in the critical current density is due to a matching effect. Certain numbers of fluxons produce highly stable and symmetric lattices which are commensurate with the underlying microstructure. These stable configurations result in high critical current densities. Adding or removing a single fluxon from such a configuration results in a defect in the lattice, a vacancy or an interstitial, which can move through the system much more easily, resulting in a lower critical current density, see Figure 5.8. It is not as simple as the pins being saturated with fluxons, as with an array of cylindrical pins, since the fluxons can slide along the pins. In the state of flux motion, the whole lattice does not move as a solid, but the rows of the lattice de-pin one after another, so the lattice simultaneously slides in the vertical direction. Similar behaviour is seen in the niobium–titanium analogue system. It is expected that these matching effects disappear in high fields since the local hexagonal order dominates when the density of vortices is high, and it is not necessarily possible to establish a globally ordered state. Defects in the initial state like vortex lattice grain boundaries are frozen in for the duration of the simulation.

In a regular multilayer system, matching effects persist independent of the system size, but for a macroscopic system like niobium–titanium where the spacing and orientation of the titanium ribbons varies over the cross-section, and where the number of fluxons is much greater than in our simulations, it is expected that the vortex matter will be able to equilibrate such that is highly stable in each local region. In this case the variation disappears and the critical current density is high. From this point onwards, the low-field critical current density will be ignored and we focus only on the high-field behaviour, which appears to vary systematically with the microstructural parameters, and is reliable across different microstructures with the same parameters.

Figure 5.9 shows the critical current density as a function of the applied magnetic field for a particular system with ribbon microstructure where the effective mass in the precipitates is equal to that of the matrix. The values of the magnetic field and critical current density have been converted to SI units using the *unextrapolated* length scales for niobium–titanium with 40 at.% Nb from Table 5.2. Although the decision was made to use the extrapolated length scales in

calculating the relative Ginzburg–Landau coefficients, since otherwise the difference parameters and two phases are not treated consistently, the unextrapolated lengths (that is, with the coherence length obtained directly from the measured upper critical field at 4.2 K) are used to compare with experimental data since otherwise there is a substantial discrepancy in the critical field. It is clear from this plot, and from all of the data with equal effective masses in the precipitates and the matrix, that, as expected, a difference in condensation energy alone cannot explain pinning in niobium–titanium. Therefore, if the important details of the microstructure are captured, the strong pinning in niobium–titanium must arise from a difference in effective mass.

Since the steady linear decrease of the critical current density in high-field is reliable despite the small system size, we can compare different microstructures by their gradients in that region. To do so, the critical current data for fields greater than half the upper critical field for each microstructure are fit with a straight line of the form

$$J_c = a \left(1 - \frac{\|H\|}{b} \right) \quad (5.8)$$

where a and b are free parameters, as in Figure 5.6. The parameter b is an effective upper critical field, which is usually close to the critical field of the matrix, slightly larger for high-mass precipitates and slightly lower for low-mass precipitates. The magnitude of the critical current density is proportional to the parameter a .

Figure 5.10 shows the values of these fit parameters for the three different mass ratios and the full range of microstructures. The horizontal lines show the values from an identical fit to the experimental data shown in Figure 5.9, converted to dimensionless form using the same, unextrapolated, Ginzburg–Landau coefficients. It can be seen that precipitates having a dimensionless effective mass of $\hat{m} = 0.2$ or of $\hat{m} = 1$ never result in critical current densities which are as high as those seen in experiments. It can also be seen that the critical current for $\hat{m} = 0.2$ does not depend strongly on the precipitate volume fraction, whereas for $\hat{m} = 1$, $\hat{m} = 2$, and $\hat{m} = 4$, the critical current density scales with the precipitate volume fraction. This latter behaviour is consistent with experiments, although the experimental volume fractions for the final wires all lie in a small range [108]. An effective mass of $\hat{m} = 4$ can account for the experimentally observed critical current densities with reasonable precipitate volume fractions. However, we previously found that the anisotropy of the critical field is consistent

with an effective mass of $\hat{m} = 2$, and precipitates with this lower effective mass require unrealistically high precipitate volume fractions to reach the observed critical current density.

The effective upper critical fields shown in Figure 5.10, which were extracted from critical current, data are qualitatively consistent with the transverse critical fields shown in Figure 5.4, which were obtained via the superfluid density. That is, high-mass precipitates result in a small enhancement in the transverse critical field, whereas low-mass precipitates result in a small suppression. The experimental data have an effective upper critical field which is less than the matrix value, which seems to suggest the presence of low-mass precipitates.

Whilst the magnitude of the critical field and critical current density in experimental measurements do not seem to be consistent with an effective mass ratio of 2, which was the value obtained from the critical field anisotropy, it should be noted that these quantities are sensitive to the precise values of the calculated Ginzburg–Landau coefficients for the matrix, whereas the anisotropy factor of the critical field is not. The coherence length can be determined relatively reliably from the measured upper critical field, but the penetration depth is calculated using the BCS theory, which makes many assumptions which are unjustified in the niobium–titanium system, and it depends on the electronic specific heat coefficient, which is difficult to measure accurately. If the penetration depth is substantially inaccurate, the resulting error in κ would lead to an error in J_0 and therefore an error in the dimensionless critical current density calculated from the experimental data. In calculating the dimensionless critical current density, a value for the Ginzburg–Landau parameter of $\kappa = 43.9$ was used, which was based on the directly measured upper critical field at 4.2 K and the extrapolated BCS thermodynamic critical field obtained from the electronic specific heat. If an error in the penetration depth meant that the real Ginzburg–Landau parameter was 10% lower, $\kappa = 39.5$, then the natural current density unit would be 23% larger, which would mean that a mass of $m_n = 2m_s$ would be consistent with the experimental data, as shown by the dotted horizontal lines in Figure 5.10.

Only the critical current density for high-mass precipitates shows a systematic variation with precipitates thickness. The other parameters only seem to depend on the overall precipitate volume fraction. It can be clearly seen in Figure 5.10 that for equal volume fractions, thicker precipitates result in higher critical current densities in the $\hat{m} = 4$ case. Experiments have revealed an opposite dependence

in the real material, where the pinning force per volume fraction decreases with ribbon thickness [108]. However, one cannot directly control the thickness in the drawing process, so it's possible that the improvement in the critical current density is due to some other change in the microstructure.

Is it reasonable to conclude that pinning in niobium–titanium is *primarily* due to a difference in effective mass? It's clear that the pinning force is much smaller if there is no difference in effective mass between the precipitates and the matrix, but on the other hand, Figure 5.1 shows that the elementary pinning forces due to an increased mass or a decreased condensation parameter are attractive whereas the pinning force due to a decreased nonlinearity parameter is repulsive. Maybe the pinning force is small because of a cancellation between the two terms in the condensation energy, in which case the magnitudes of the pinning force due to the effective mass and condensation parameter may be similar, and we could not conclude that the effective mass term is the most important factor.

Figure 5.1 gives the *elementary* pinning force, so it is only expected to reflect the critical current density in very low field, where the interactions between vortices can be neglected. The pinning forces due to the different terms in the condensation energy may be different in high field. Figure 5.11 shows the high-field critical current density for a particular niobium–titanium analogue system, and the effects of ‘turning off’ the condensation parameter or nonlinearity parameter. The high-field critical current density is almost entirely accounted for by the difference in the condensation parameter alone. The reason for this is that in high field, the magnitude of the order parameter is small, and hence the nonlinear term in the free energy can be neglected. We therefore conclude that the high-field pinning force resulting from differences in the condensation energy is in fact only due to the difference in $\hat{\alpha}$, and therefore that there is no significant cancellation due to the competing terms. It is then reasonable to say that a difference in effective mass is the primary reason for the high critical current density in niobium–titanium.

5.7 Conclusion

According to the present model, strong pinning in niobium–titanium arises primarily due to the increased carrier effective mass in the titanium precipitates. There are three pieces of evidence which support this conclusion. Firstly, high-

effective-mass precipitates can explain the enhancement of the effective upper critical field in the longitudinal direction versus the transverse direction. Secondly, the high-field critical current density scales as expected with the precipitate volume fraction with high-mass precipitates. Thirdly, only high-mass precipitates lead to high-field critical current densities which are similar in magnitude to those measured in experiments on niobium–titanium filaments. Pinning due to a reduced effective mass or condensation energy fail to explain these phenomena.

That high-effective-mass precipitates can account for the critical field anisotropy, the magnitude of the critical current density, and its dependence on the titanium volume fraction indicates that other microstructural features like crystal defects, concentration gradients, and disordered interfaces between phases do not significantly contribute to pinning in high-field niobium–titanium conductor. However, there is an inconsistency with the effective masses required to reproduce the critical field anisotropy and the magnitude of the critical current density, the former indicating an effective mass ratio of $\hat{m} = 2$, and the latter being consistent with an effective mass ratio of $\hat{m} = 4$, although, the magnitude of the critical current density is sensitive to errors in the calculated Ginzburg–Landau coefficients, whereas the critical field anisotropy is not.

There are other ways in which the model is inconsistent with experiments on real wires. For example, images of the microstructure and critical current data from various stages in the drawing process indicate that thinner and more closely spaced precipitates result in higher critical current densities even when the precipitate volume fraction is roughly the same, whereas the simulations show the opposite behaviour, that for the same volume fraction, thicker precipitates lead to higher critical current densities. Finally, the perpendicular critical field for real wires, as obtained by critical current measurements, is about 10% less than the measured critical field for single-phase niobium–titanium with the appropriate composition whereas, in the simulations, the upper critical field for high-mass precipitates is always higher than that of the matrix. Depending on the criterion used for the critical field, low-mass precipitates can result in a 10% reduction. However, the critical field is again sensitive to the values of the Ginzburg–Landau coefficients, so this difference could also arise from errors in the material properties.

In summary, the computational model suggests that the high critical current density in niobium–titanium wires cannot be explained by pinning due to a differ-

ence in condensation energy between the matrix and titanium precipitates. It can be explained by a difference in the carrier effective mass, but some inconsistencies between the model and the observed properties of real wires remain. These may arise from assumptions like the perfect separation of the two phases or errors in the calculated Ginzburg–Landau parameters due to extrapolation far from the transition temperatures.

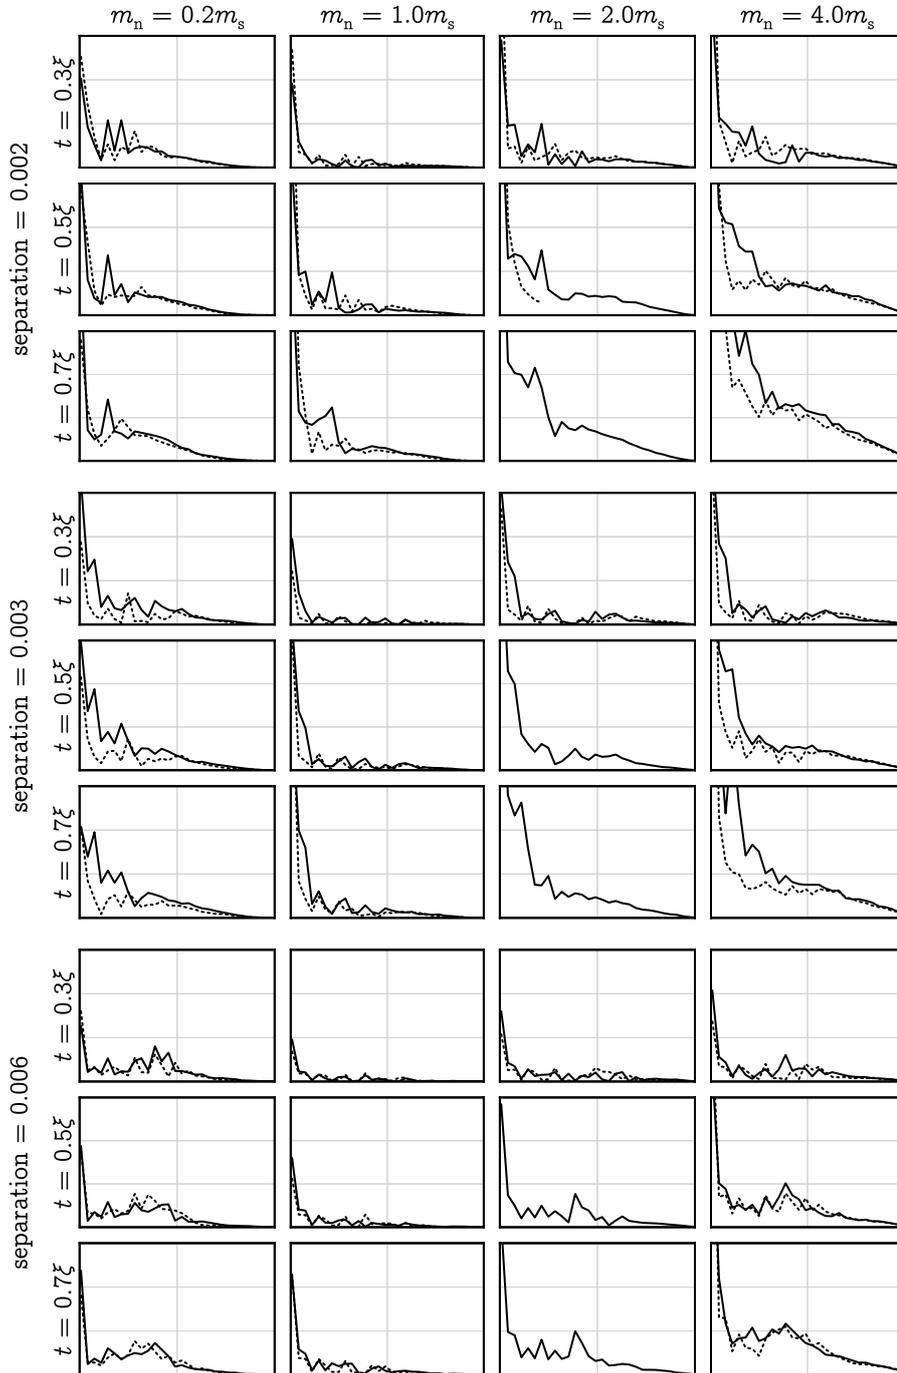


Figure 5.5: Compilation of critical current data. In each plot the x-axis represents the field, ranging from 0 to H_{c2} , and the y-axis represents the critical current density, ranging from 0 to $0.015J_0$. For most of the configurations, two different randomly generated microstructures were used, and the dotted line is for the second of these samples.

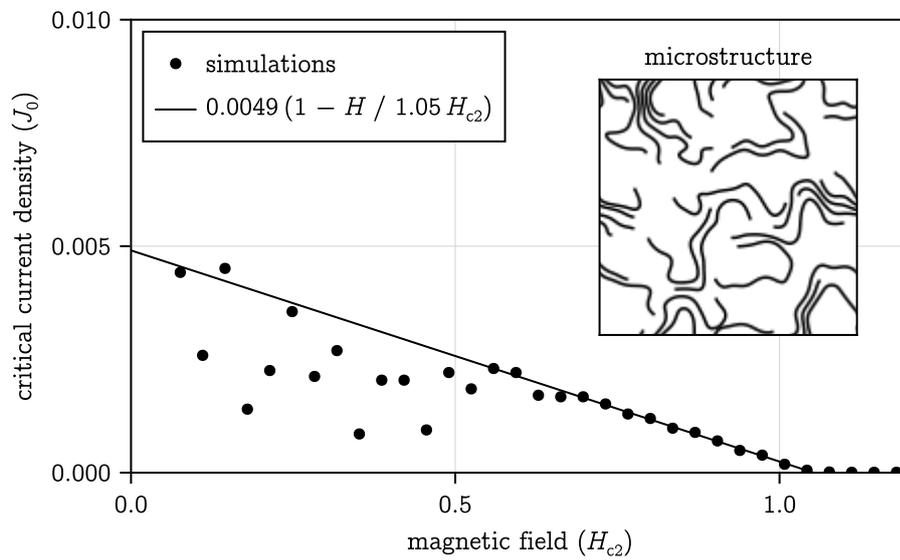


Figure 5.6: Simulated critical current density as a function of applied field for a typical system with ribbon microstructure where $m_n = 4m_s$, the ribbon thickness is $\xi/2$, and the precipitate volume fraction is 0.14.

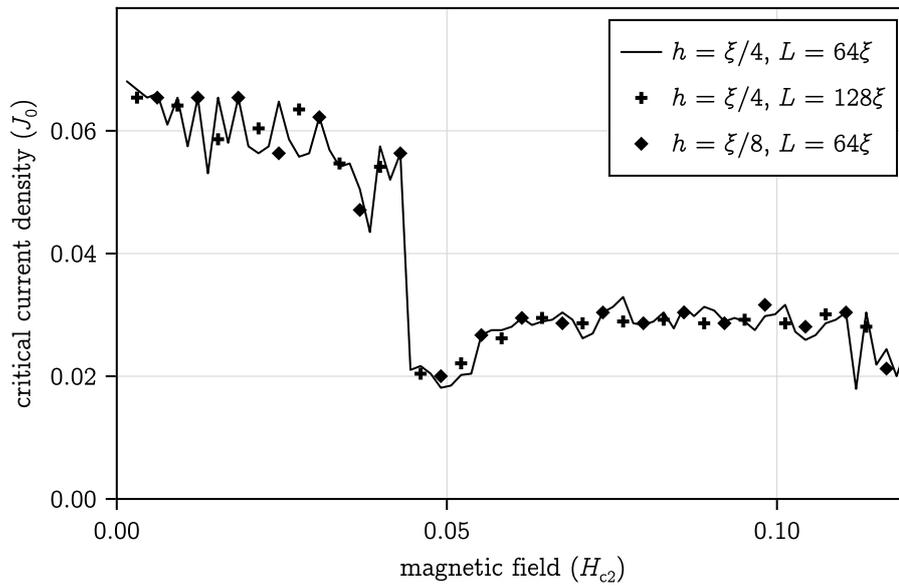


Figure 5.7: Critical current density for two-dimensional multilayer systems consisting of a doubly-periodic homogeneous superconducting matrix evenly spaced normal planes. The normal planes have $\alpha = \beta = 0$, $m_n = 2m_s$, and a thickness of $\xi/2$, the separation between the centres of the planes is 16ξ , and $\kappa = 50$. The current flows along the interfaces. The solid line is for a system of dimensions $64\xi \cdot 64\xi$ and the other data are for similar systems, one with twice the linear dimensions and one with half the grid spacing.

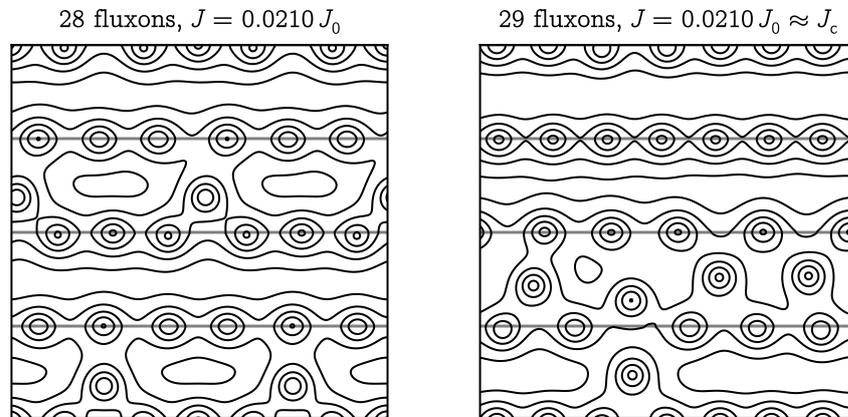


Figure 5.8: Current streamlines for two snapshots of the multilayer system with different numbers of fluxons. Although the first system has fluxons which sit between the pins, it has a much higher critical current density because there are simply too many fluxons in the second system to comfortably sit either on the pins or directly between them.

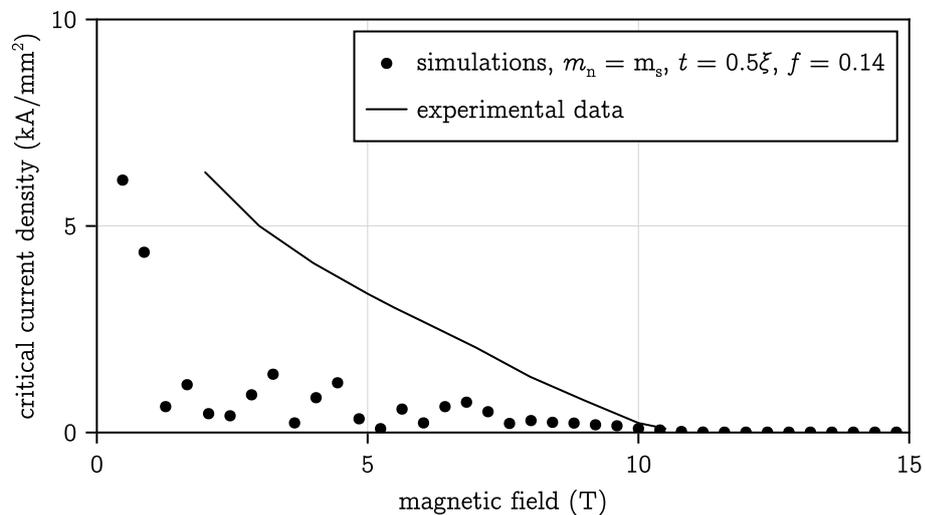


Figure 5.9: Simulated critical current data for a particular system with ribbon microstructure having the same effective mass in the precipitates as in the matrix, converted to SI units and compared with experimental data from Meingast and Larbalestier [118].

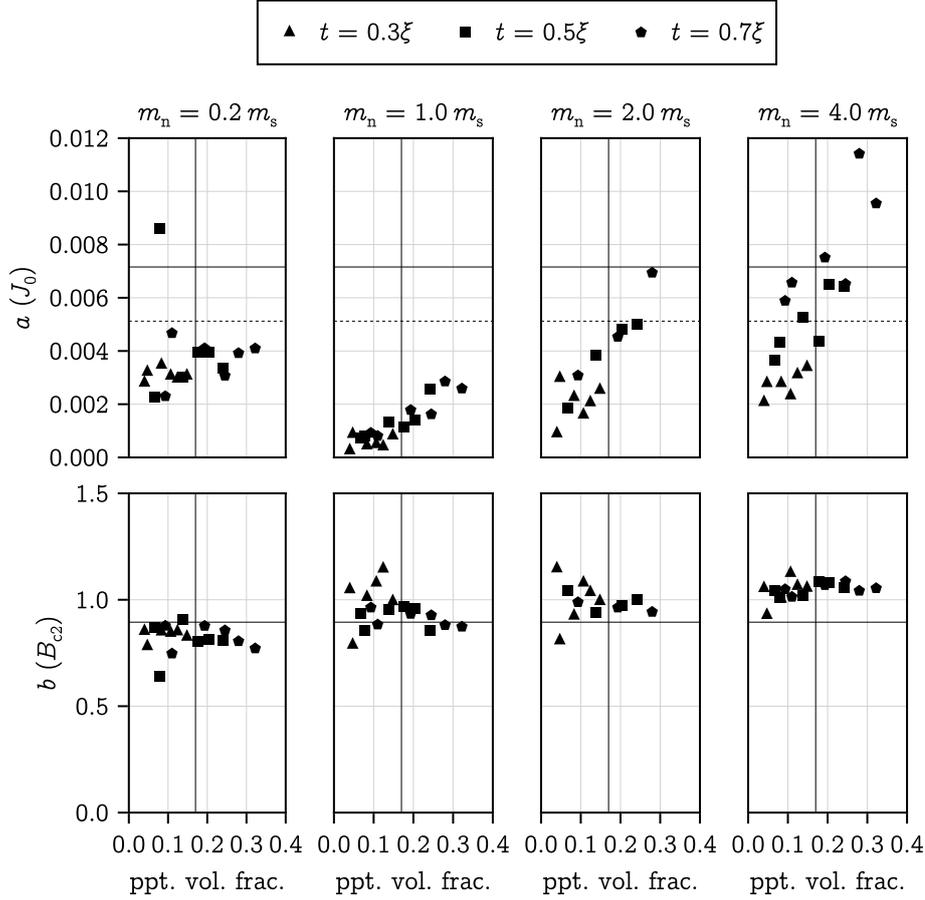


Figure 5.10: Plots of the fit parameters characterising the high-field critical current density as a function of the precipitate volume fraction for various microstructures and effective mass ratios. The horizontal lines across all panels shows the values of the same parameter for an identical fit to the experimental data shown in Figure 5.9, converted to dimensionless form using the same reference values. The dotted lines use a κ value which is 10% less. The vertical lines indicate the precipitate volume fraction in a real wire.

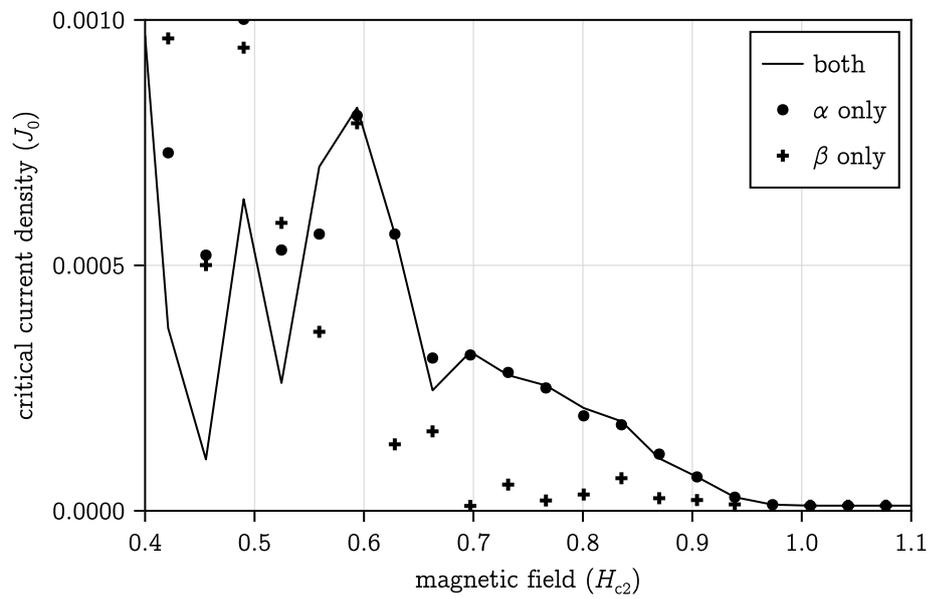


Figure 5.11: High-field critical current density for a niobium–titanium analogue system with no mass difference between the precipitates and the matrix. The ribbon thickness is $\xi/2$ and the precipitate volume fraction is 0.14. The points show the effects of ‘turning off’ the condensation parameter or nonlinearity parameter (artificially setting the value in the precipitates to be the same as in the matrix).

Chapter 6

Conclusions and future work

The aim of this work was to demonstrate that it is practical to simulate the dynamics of the condensate in high-field superconductors like niobium–titanium using Ginzburg–Landau theory with microstructures and material parameters which are consistent with real microstructures and measured properties of single phase samples. Most prior work makes use of overly idealised microstructures and ad hoc choices of material parameters. This is partly because of the nature of pinning sites in, for example, REBCO tapes. It is not obvious how to model oxide nanoparticles in Ginzburg–Landau theory since the oxide is not superconducting in the first place, so the transition temperature, critical field, and penetration depth do not exist to be measured. It is also because of the limitations of existing solvers for the Ginzburg–Landau equations. Solvers for large-scale three-dimensional systems assume the frozen field limit and most assume the homogeneity of all material parameters except for the condensation parameter, and the more general solvers do not scale well with system size, making them prohibitively slow for large systems.

In some preliminary simulations, we investigated pinning by resistive nanorods in REBCO tape, still assuming the frozen field limit and using ad hoc material parameters but lifting the assumption of the homogeneity of the effective mass, which is proportional to the resistivity. Not only do resistive rods pin more strongly than conductive ones, but the interpretation of the mechanism limiting the critical current density is completely different in high fields. An interface with a resistive material supports surface superconductivity above the upper critical field of the matrix, and when the pins are sufficiently close together, these regions

of surface superconductivity overlap so that a supercurrent can flow across the whole system.

Since a generic inclusion will have an effective mass which is different from the matrix, this shows that the field at which the critical current density drops to zero is not necessarily the same as the upper critical field of the matrix, even neglecting flux line lattice melting. More impure superconductors tend to have higher critical fields, since more scattering leads to a shorter coherence length, but there may be an additional increase in the effective upper critical field due to surface superconductivity in the vicinity of interfaces in inhomogeneous samples which is usually not considered. It also shows that there are effects which are of fundamental importance in determining the critical current density which cannot be captured in simulations which only allow spatial variation in the condensation parameter.

A major contribution of this work is a computational method for solving the fully coupled time-dependent Ginzburg–Landau equations which allows all of the material parameters to vary in space and which scales well with system size. This is required for accurate treatment of systems consisting of multiple superconducting and normal materials, some of which might not have very large Ginzburg–Landau parameters. A geometric multigrid method was used to solve the nonlinear system of equations which arises from a backwards Euler discretisation of the time-dependent Ginzburg–Landau equations in the London gauge. To my knowledge, this is the first time that such a method has been applied to simultaneously solve the full system of equations.

We applied the new solver to investigate the most widely used high-field superconductor, niobium–titanium. In this material, it is thought that the primary pinning centres are titanium precipitates. Since titanium is itself a superconductor at very low temperatures, we were able to use the measured properties of both phases to obtain estimates of the ratios of the Ginzburg–Landau coefficients in the titanium precipitates versus the niobium–titanium matrix. We generated microstructures with random distributions of ribbon-like precipitates with varying thicknesses and separations which are similar to the microstructures which are observed by transmission electron microscopy of niobium–titanium.

Measurements of the macroscopic superconducting properties of single phase samples cannot constrain the carrier effective mass but, as demonstrated in the REBCO simulations, the ratio of the effective masses can have a drastic effect on

the pinning. However, because of the anisotropy in the microstructure due to the wire drawing, there is an anisotropy in the effective upper critical field due to a differences in the surface superconducting state at the interfaces depending on the field direction. This anisotropy depends on the effective mass ratio, so we were able to calculate the upper critical field using simulations and compare the results with experimental data to constrain the effective mass ratio.

From the simulations, we found that the effective upper critical field is always higher when the field is parallel to the drawing direction, compared to when it is perpendicular, regardless of the effective mass ratio and microstructural parameters. When the effective mass is larger in the precipitates, this is attributed to enhanced surface superconductivity when the field is parallel to the drawing direction because the precipitates are highly aligned in that direction, so there is a large interfacial area parallel to the field. We find that the experimentally observed critical field anisotropy factor is consistent with an effective mass ratio of 2.

Critical current simulations were also run with a range of effective mass values and microstructural parameters. Since the small size of the system compared with a real wire led to unpredictable variation in the critical current density in low fields, we compared the effective upper critical field and high-field critical current density for each system with experimental data. We found that the effective mass in the titanium precipitates must be higher than that of the matrix in order to explain the magnitude of the critical current density. Using the calculated Ginzburg–Landau coefficients, an effective mass ratio of 4 is required, but given some uncertainty in the measured properties, an effective mass of 2 could be consistent.

Whilst high-mass precipitates can consistently explain the critical field anisotropy and high-field critical current density in niobium–titanium, the dependence of the critical current density on the microstructural parameters was not as expected. The conventional understanding of why the optimised heat treatment and drawing process results in high critical current densities is that it produces titanium precipitates which are thin and closely spaced. Thinner and more closely spaced precipitates have been shown to result in higher pinning forces. However, in the case of high-mass precipitates, which were the only systems which showed systematic variation with ribbon thickness in our simulations, we saw the opposite behaviour. That is, for the same titanium volume fraction, thicker

precipitates resulted in higher pinning forces.

Although high mass precipitates are required to explain the pinning force, since the differences in the condensation parameter and nonlinearity parameter are of a similar magnitude to the difference in mass, it does not necessarily imply that the effective mass is the dominant factor. One of the advantages of simulation is that we can test scenarios which would be impossible to test in an experiment, and in this case we can simply ‘turn off’ the pinning due to the condensation and nonlinearity parameters. By doing so we observe that the difference in nonlinearity parameter barely contributes to the high-field critical current density because in high-field the nonlinear term in the free energy is small anyway. Therefore in the case where the mass is the same in the matrix and the precipitates, the pinning force is only due to the difference in condensation parameter, and the critical current density is small. Therefore, we conclude that the pinning in the high-mass case is predominantly due to the difference in mass. If high mass precipitates are responsible for pinning the real material, the condensation parameter must provide about 15% of the high-field pinning force, the effective mass about 85%, and the nonlinearity parameter does not contribute at all.

In conclusion, pinning in niobium–titanium was modeled satisfactorily, approximately reproducing the observed critical field anisotropy and high-field critical current density. It is likely that titanium precipitates are indeed the dominant pinning centres in niobium–titanium, and that the pinning force arises because the carrier effective mass in the precipitates is about twice that of the matrix. This demonstrates that time-dependent Ginzburg–Landau simulations with accurate material parameters are practical and can provide insights which would be difficult to discover by other means.

In future work, these methods will hopefully be applied to other materials. For pinning sites which are not superconducting, accurate simulations may allow effective Ginzburg–Landau coefficients to be obtained in the same way that we were able to constrain the effective mass using the critical field anisotropy. One particularly relevant example are the defects which are created by neutron irradiation in fusion magnets. The solver described here should be able to simulate superconductivity in most low-temperature and high-temperature superconductors. Those with multiple condensates, like MgB_2 , can’t be simulated at present but the inclusion of a second wavefunction would be a simple extension of the

existing code.

There are many other potential improvements which could be made to the solver including support for non-rectilinear meshes, internal boundary conditions, and special treatment for the frozen field and steady state limits. These are discussed more in Section 3.2.11.

In the long term, Ginzburg–Landau simulations might be used to inform the optimisation of superconducting materials [119]. Although these simulations take a significant amount of computational time, they are much quicker and cheaper than a full production cycle of a real conductor. If the errors in the Ginzburg–Landau coefficients and simulated critical current can be understood for many different superconductors, one might be able to check whether some modification of the material would lead to higher critical currents, and thereby improve upon locally optimal designs.

Appendix

Supplementary material

The following materials are provided to supplement the main text. The updated code will be made available under a copyleft license.

- `MultDGL_nbti.jl`, the Julia library which implements the multigrid method for the TDGL system as used in the simulations of niobium–titanium. There are some minor differences from the method described in Chapter 3 because the implementation has been improved since the niobium–titanium simulations were performed.
- `nbti_material.jl`, a script which generates niobium–titanium analogue microstructures.
- `nbti_run.jl`, a script which drives the critical current simulations for niobium–titanium.
- `nbti_data.tar.xz`, an archive containing input parameters and critical current data for niobium–titanium analogue systems.
- `nbti_vid_side.avi` and `nbti_vid_top.avi`, videos of the flux motion in a niobium–titanium analogue system with low-mass precipitates showing different cross-sections of the wavefunction magnitude.
- `figures_code_data.tar.xz`, an archive of all the code and data required to generate the figures in this work. This is a large file and some of the code is not very pretty!

Fusion CDT courses

As a student of the Fusion CDT, I completed the following courses.

- Introduction to fusion plasmas (*dist.*)
- Introduction to materials (*dist.*)
- Computational techniques (*dist.*)
- Materials applications (*dist.*)
- Fusion technology (*dist.*)
- Integrated systems and project management (*dist.*)
- Radiation damage (*n.a.*)
- Nuclear power (*n.a.*)
- Finite element modelling (*pass*)
- Characterisation and analytical tools (*dist.*)
- Frontiers of fusion (*n.a.*)
- Collaboratory (*n.a.*)

The Collaboratory is a collaborative project with another institution. I spent eight weeks at Oxford Instruments developing a method to estimate AC losses in superconducting solenoids during a quench.

Publications

The following works were published as part of this research.

- C.W.W. Haddon, A.I. Blair, F. Schoofs, and D.P. Hampshire, ‘Computational Simulations Using Time-Dependent Ginzburg–Landau Theory for Nb–Ti-Like Microstructures’, *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 4, pp. 1–5, Jun. 2022, doi: 10.1109/TASC.2022.3156916.
- C.W.W. Haddon and D.P. Hampshire, ‘Fast Multigrid Simulations of Pinning in REBCO With Highly Resistive Nanorods’, *IEEE Transactions on Applied Superconductivity*, vol. 33, no. 5, pp. 1–5, Aug. 2023, doi: 10.1109/TASC.2023.3253065.

Acknowledgements

Thank you to E.C., without whose love and support I wouldn't be where I am today, to my colleagues in the superconductivity group at Durham for the many enjoyable and interesting discussions, and to my supervisor for his patience, constructive feedback, and passion for physics and superconductivity.

This work made use of the excellent computing facilities of the NCC at Durham University, and of many Julia libraries including `CUDA.jl` [120] and `Makie.jl` [121]. I am grateful to the developers and maintainers of free software, whose work is indispensable and often undertaken without any expectation of financial reward.

Finally, I would like to thank the people who generously offer their physics-related writings, talks, and videos to the public for free. I have been particularly inspired by the writings of John Baez and the podcasts of Jim Rantschler and Randy Morrison.

This work was funded by EPSRC grant EP/L01663X/1 and the UK Government Department for Business, Energy, and Industrial Strategy.

Bibliography

- [1] G.W. Webb, F. Marsiglio, and J.E. Hirsch, “Superconductivity in the elements, alloys and simple compounds,” *Physica C: Superconductivity and its Applications*, vol. 514, pp. 17–27, Jul. 2015, doi: 10.1016/j.physc.2015.02.037.
- [2] J.J. Hamlin, “Superconductivity in the metallic elements at high pressures,” *Physica C: Superconductivity and its Applications*, vol. 514, pp. 59–76, Jul. 2015, doi: 10.1016/j.physc.2015.02.032.
- [3] J.M. McMahon and D.M. Ceperley, “High-temperature superconductivity in atomic metallic hydrogen,” *Phys. Rev. B*, vol. 84, no. 14, p. 144515, Oct. 2011, doi: 10.1103/PhysRevB.84.144515.
- [4] D. van der Marel and C. Berthod, “Superconductivity in metallic hydrogen,” *Newton*, vol. 1, no. 1, p. 100002, Mar. 2025, doi: 10.1016/j.newton.2024.100002.
- [5] N. Chamel, “Superfluidity and Superconductivity in Neutron Stars,” *J Astrophys Astron*, vol. 38, no. 3, p. 43, Sep. 2017, doi: 10.1007/s12036-017-9470-9.
- [6] F. Englert, “Nobel Lecture: The BEH mechanism and its scalar boson,” *Rev. Mod. Phys.*, vol. 86, no. 3, pp. 843–850, Jul. 2014, doi: 10.1103/RevModPhys.86.843.
- [7] T. Xiang and C. Wu, *D-wave superconductivity*. Cambridge, United Kingdom; New York, NY, USA: Cambridge University Press, 2022. doi: 10.1017/9781009218566.
- [8] T.W.B. Kibble, “Topology of cosmic domains and strings,” *J. Phys. A: Math. Gen.*, vol. 9, no. 8, p. 1387, Aug. 1976, doi: 10.1088/0305-4470/9/8/029.
- [9] V. Novotny and P.P.M. Meincke, “Single superconducting energy gap in pure niobium,” *J Low Temp Phys*, vol. 18, no. 1, pp. 147–157, Jan. 1975, doi: 10.1007/BF00116976.
- [10] T.C. Cosmos and M. Parizh, “Advances in Whole-Body MRI Magnets,” *IEEE Transactions on Applied Superconductivity*, vol. 21, no. 3, pp. 2104–2109, Jun.

- 2011, doi: 10.1109/TASC.2010.2084981.
- [11] R.M. Scanlan, A.P. Malozemoff, and D.C. Larbalestier, “Superconducting materials for large scale applications,” *Proceedings of the IEEE*, vol. 92, no. 10, pp. 1639–1654, Oct. 2004, doi: 10.1109/JPROC.2004.833673.
- [12] N. Boulant *et al.*, “Commissioning of the Iseult CEA 11.7 T whole-body MRI: current status, gradient–magnet interaction tests and first imaging experience,” *Magn Reson Mater Phy*, vol. 36, no. 2, pp. 175–189, Apr. 2023, doi: 10.1007/s10334-023-01063-5.
- [13] E. Todesco *et al.*, “The High Luminosity LHC interaction region magnets towards series production,” *Supercond. Sci. Technol.*, vol. 34, no. 5, p. 053001, Mar. 2021, doi: 10.1088/1361-6668/abdba4.
- [14] N. Mitchell *et al.*, “The ITER Magnet System,” *IEEE Transactions on Applied Superconductivity*, vol. 18, no. 2, pp. 435–440, Jun. 2008, doi: 10.1109/TASC.2008.921232.
- [15] M.K. Wu *et al.*, “Superconductivity at 93 K in a new mixed-phase Y-Ba-Cu-O compound system at ambient pressure,” *Phys. Rev. Lett.*, vol. 58, no. 9, pp. 908–910, Mar. 1987, doi: 10.1103/PhysRevLett.58.908.
- [16] J. Lu, Y. Xin, B. Jarvis, and H. Bai, “Oxygen out-diffusion in REBCO coated conductor due to heating,” *Supercond. Sci. Technol.*, vol. 34, no. 7, p. 075004, May 2021, doi: 10.1088/1361-6668/abfd0c.
- [17] P. Gao, J. Mao, J. Chen, X. Wang, and Y. Zhou, “Electromechanical degradation of REBCO coated conductor tapes under combined tension and torsion loading,” *International Journal of Mechanical Sciences*, vol. 223, p. 107314, Jun. 2022, doi: 10.1016/j.ijmecsci.2022.107314.
- [18] C.W.F. Everitt *et al.*, “Gravity Probe B: Final Results of a Space Experiment to Test General Relativity,” *Phys. Rev. Lett.*, vol. 106, no. 22, p. 221101, May 2011, doi: 10.1103/PhysRevLett.106.221101.
- [19] A. Einstein, “Gibt es eine Gravitationswirkung, die der elektrodynamischen Induktionswirkung analog ist?,” *Vierteljahrschrift für Gerichtliche Medizin und Öffentliches Sanitätswesen*, vol. 44, pp. 37–40, Jan. 1912.
- [20] J. Clarke, “SQUIDS,” *Scientific American*, vol. 271, no. 2, pp. 46–53, 1994, Accessed: Aug. 28, 2025. [Online]. Available: <https://www.jstor.org/stable/24942801>
- [21] M. Kjaergaard *et al.*, “Superconducting Qubits: Current State of Play,” *Annual Review of Condensed Matter Physics*, vol. 11, no. Volume 11, 2020, pp.

- 369–395, Mar. 2020, doi: 10.1146/annurev-conmatphys-031119-050605.
- [22] J. Nagamatsu, N. Nakagawa, T. Muranaka, Y. Zenitani, and J. Akimitsu, “Superconductivity at 39 K in magnesium diboride,” *Nature*, vol. 410, no. 6824, pp. 63–64, Mar. 2001, doi: 10.1038/35065039.
- [23] H. Maeda, Y. Tanaka, M. Fukutomi, and T. Asano, “A New High- T_c Oxide Superconductor without a Rare Earth Element,” *Jpn. J. Appl. Phys.*, vol. 27, no. 2A, p. L209, Feb. 1988, doi: 10.1143/JJAP.27.L209.
- [24] W. Meissner and R. Ochsenfeld, “Ein neuer Effekt bei Eintritt der Supraleitfähigkeit,” *Naturwissenschaften*, vol. 21, no. 44, pp. 787–788, Nov. 1933, doi: 10.1007/BF01504252.
- [25] H.J. Choi, D. Roundy, H. Sun, M.L. Cohen, and S.G. Louie, “The origin of the anomalous superconducting properties of MgB_2 ,” *Nature*, vol. 418, no. 6899, pp. 758–760, Aug. 2002, doi: 10.1038/nature00898.
- [26] T.E.G. Nicholas *et al.*, “Re-examining the role of nuclear fusion in a renewables-based energy mix,” *Energy Policy*, vol. 149, p. 112043, Feb. 2021, doi: 10.1016/j.enpol.2020.112043.
- [27] B. Lindley, T. Roulstone, G. Locatelli, and M. Rooney, “Can fusion energy be cost-competitive and commercially viable? An analysis of magnetically confined reactors,” *Energy Policy*, vol. 177, p. 113511, Jun. 2023, doi: 10.1016/j.enpol.2023.113511.
- [28] A.J. Creely *et al.*, “Overview of the SPARC tokamak,” *Journal of Plasma Physics*, vol. 86, no. 5, p. 865860502, Oct. 2020, doi: 10.1017/S0022377820001257.
- [29] M. Siccino *et al.*, “Development of the plasma scenario for EU-DEMO: Status and plans,” *Fusion Engineering and Design*, vol. 176, p. 113047, Mar. 2022, doi: 10.1016/j.fusengdes.2022.113047.
- [30] Y. Song *et al.*, “Engineering design of the CFETR machine,” *Fusion Engineering and Design*, vol. 183, p. 113247, Oct. 2022, doi: 10.1016/j.fusengdes.2022.113247.
- [31] I.T. Chapman, S.C. Cowley, and H.R. Wilson, “The Spherical Tokamak for Energy Production: theme issue introduction,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 382, no. 2280, p. 20230416, Aug. 2024, doi: 10.1098/rsta.2023.0416.
- [32] J.S. Kang, G. Jo, J.-M. Kwon, and B.G. Hong, “Physics and technology maturity level required for the K-DEMO design points,” *Fusion Engineering and*

- Design*, vol. 211, p. 114741, Feb. 2025, doi: 10.1016/j.fusengdes.2024.114741.
- [33] D.X. Fischer, R. Prokopec, J. Emhofer, and M. Eisterer, “The effect of fast neutron irradiation on the superconducting properties of REBCO coated conductors with and without artificial pinning centers,” *Supercond. Sci. Technol.*, vol. 31, no. 4, p. 044006, Mar. 2018, doi: 10.1088/1361-6668/aaadf2.
- [34] T. Winiecki and C.S. Adams, “A Fast Semi-Implicit Finite-Difference Method for the TDGL Equations,” *Journal of Computational Physics*, vol. 179, no. 1, pp. 127–139, Jun. 2002, doi: 10.1006/jcph.2002.7047.
- [35] I.A. Sadovskyy, A.E. Koshelev, C.L. Phillips, D.A. Karpeyev, and A. Glatz, “Stable large-scale solver for Ginzburg–Landau equations for superconductors,” *Journal of Computational Physics*, vol. 294, pp. 639–654, Aug. 2015, doi: 10.1016/j.jcp.2015.04.002.
- [36] A.A. Abrikosov and L.P. Gor’kov, “On the Theory of Superconducting Alloys; I. The Electrodynamics of Alloys at Absolute Zero,” *Journal of Experimental and Theoretical Physics*, vol. 35, no. 8, pp. 1090–1098, 1959, [Online]. Available: <http://jetp.ras.ru/cgi-bin/e/index/e/8/6/p1090?a=list>
- [37] G.M. Eliashberg, “Interactions between Electrons and Lattice Vibrations in a Superconductor,” *Soviet Physics JETP*, vol. 11, no. 3, pp. 696–702, 1960, [Online]. Available: <http://www.jetp.ras.ru/cgi-bin/e/index/e/11/3/p696?a=list>
- [38] G.M. Eliashberg, “Temperature Green’s Function for Electrons in a Superconductor,” *Soviet Physics JETP*, vol. 12, no. 5, pp. 1000–1002, 1961, [Online]. Available: <http://www.jetp.ras.ru/cgi-bin/e/index/e/12/5/p1000?a=list>
- [39] C.P. Bean, “Magnetization of Hard Superconductors,” *Phys. Rev. Lett.*, vol. 8, no. 6, pp. 250–253, Mar. 1962, doi: 10.1103/PhysRevLett.8.250.
- [40] Y. Iwasa, *Case studies in superconducting magnets: design and operational issues*. in Selected topics in superconductivity. New York: Plenum Press, 1994.
- [41] V.L. Ginzburg and L.D. Landau, “On the Theory of Superconductivity,” *Journal of Experimental and Theoretical Physics*, vol. 20, no. 12, pp. 1064–1082, 1950.
- [42] L.D. Landau, “On the Theory of Phase Transitions,” *Ukrainian Journal of Physics*, vol. 53, pp. 26–35, 2008.
- [43] B.S. Deaver and W.M. Fairbank, “Experimental Evidence for Quantized Flux in Superconducting Cylinders,” *Phys. Rev. Lett.*, vol. 7, no. 2, pp. 43–46, Jul.

- 1961, doi: 10.1103/PhysRevLett.7.43.
- [44] R. Doll and M. Näbauer, “Experimental Proof of Magnetic Flux Quantization in a Superconducting Ring,” *Phys. Rev. Lett.*, vol. 7, no. 2, pp. 51–52, Jul. 1961, doi: 10.1103/PhysRevLett.7.51.
- [45] R.D. Parks, “Quantized Magnetic Flux in Superconductors,” *Science*, vol. 146, no. 3650, pp. 1429–1435, Dec. 1964, doi: 10.1126/science.146.3650.1429.
- [46] F. London, “A New Conception of Supraconductivity,” *Nature*, vol. 140, no. 3549, pp. 793–796, Nov. 1937, doi: 10.1038/140793a0.
- [47] L.P. Gor’kov, “Microscopic Derivation of the Ginzburg–Landau Equations in the Theory of Superconductivity,” *Journal of Experimental and Theoretical Physics*, vol. 36, no. 9, pp. 1364–1367, 1959, [Online]. Available: <http://www.jetp.ras.ru/cgi-bin/e/index/e/9/6/p1364?a=list>
- [48] Q. Du, M.D. Gunzburger, and J.S. Peterson, “Analysis and Approximation of the Ginzburg–Landau Model of Superconductivity,” *SIAM Review*, vol. 34, no. 1, pp. 54–81, 1992, Accessed: Dec. 12, 2024. [Online]. Available: <https://www.jstor.org/stable/2132785>
- [49] J.R. Clem, “Simple model for the vortex core in a type II superconductor,” *J Low Temp Phys*, vol. 18, no. 5, pp. 427–434, Mar. 1975, doi: 10.1007/BF00116134.
- [50] A.A. Abrikosov, “On the Magnetic Properties of Superconductors of the Second Group,” *JETP*, vol. 32, no. 6, p. 1442, 1957, Accessed: Sep. 11, 2023. [Online]. Available: <http://jetp.ras.ru/cgi-bin/e/index/e/5/6/p1174?a=list>
- [51] W.H. Kleiner, L.M. Roth, and S.H. Autler, “Bulk Solution of Ginzburg–Landau Equations for Type II Superconductors: Upper Critical Field Region,” *Phys. Rev.*, vol. 133, no. 5A, pp. A1226–A1227, Mar. 1964, doi: 10.1103/PhysRev.133.A1226.
- [52] E.H. Brandt, “Precision Ginzburg–Landau Solution of Ideal Vortex Lattices for Any Induction and Symmetry,” *Phys. Rev. Lett.*, vol. 78, no. 11, pp. 2208–2211, Mar. 1997, doi: 10.1103/PhysRevLett.78.2208.
- [53] P.A.M. Dirac, “Quantised singularities in the electromagnetic field,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 133, no. 821, pp. 60–72, Sep. 1931, doi: 10.1098/rspa.1931.0130.
- [54] D. Saint-James and P.G. de Gennes, “Onset of superconductivity in decreasing fields,” *Physics Letters*, vol. 7, no. 5, pp. 306–308, Dec. 1963, doi:

- 10.1016/0031-9163(63)90047-7.
- [55] J. Tate, B. Cabrera, S.B. Felch, and J.T. Anderson, “Precise Determination of the Cooper-Pair Mass,” *Phys. Rev. Lett.*, vol. 62, no. 8, pp. 845–848, Feb. 1989, doi: 10.1103/PhysRevLett.62.845.
- [56] W.D. Gropp, H.G. Kaper, G.K. Leaf, D.M. Levine, M. Palumbo, and V.M. Vinokur, “Numerical Simulation of Vortex Dynamics in Type-II Superconductors,” *Journal of Computational Physics*, vol. 123, no. 2, pp. 254–266, Feb. 1996, doi: 10.1006/jcph.1996.0022.
- [57] Q. Du and P. Gray, “High-Kappa Limits of the Time-Dependent Ginzburg–Landau Model,” *SIAM Journal on Applied Mathematics*, vol. 56, no. 4, pp. 1060–1093, 1996, Accessed: Dec. 12, 2024. [Online]. Available: <https://www.jstor.org/stable/2102603>
- [58] A.M. Campbell and J.E. Evetts, “Flux vortices and transport currents in type II superconductors,” *Advances in Physics*, vol. 21, no. 90, pp. 199–428, Mar. 1972, doi: 10.1080/00018737200101288.
- [59] E.J. Kramer, “Scaling laws for flux pinning in hard superconductors,” *Journal of Applied Physics*, vol. 44, no. 3, pp. 1360–1370, Mar. 1973, doi: 10.1063/1.1662353.
- [60] D. Dew-Hughes, “Flux pinning mechanisms in type II superconductors,” *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, vol. 30, no. 2, pp. 293–305, Aug. 1974, doi: 10.1080/14786439808206556.
- [61] T. Matsushita, *Flux pinning in superconductors*, Second edition. in Springer series in solid-state sciences, no. 178. Heidelberg: Springer, 2014.
- [62] J. Bardeen, L.N. Cooper, and J.R. Schrieffer, “Theory of Superconductivity,” *Phys. Rev.*, vol. 108, no. 5, pp. 1175–1204, Dec. 1957, doi: 10.1103/PhysRev.108.1175.
- [63] A.A. Abrikosov and L.P. Gor’kov, “Superconducting Alloys at Finite Temperatures,” *Soviet Physics JETP*, vol. 36, no. 1, pp. 220–221, 1959, [Online]. Available: <http://www.jetp.ras.ru/cgi-bin/e/index/e/9/1/p220?a=list>
- [64] H. Fröhlich, “Theory of the Superconducting State. I. The Ground State at the Absolute Zero of Temperature,” *Phys. Rev.*, vol. 79, no. 5, pp. 845–856, Sep. 1950, doi: 10.1103/PhysRev.79.845.
- [65] N.N. Bogoliubov, “A New Method in the Theory of Superconductivity. I,” *Soviet Physics JETP*, vol. 7, no. 1, pp. 41–46, 1958, [Online]. Available:

- <http://www.jetp.ras.ru/cgi-bin/e/index/e/7/1/p41?a=list>
- [66] A.V. Gold, “An experimental determination of the fermi surface in lead,” *Philos Trans A Math Phys Eng Sci*, vol. 251, no. 989, pp. 85–112, Nov. 1958, doi: 10.1098/rsta.1958.0010.
 - [67] D.P. Karim, J.B. Ketterson, and G.W. Crabtree, “A de Haas–van Alphen study of niobium: Fermi surface, cyclotron effective masses, and magnetic breakdown effects,” *J Low Temp Phys*, vol. 30, no. 3, pp. 389–423, Feb. 1978, doi: 10.1007/BF00114959.
 - [68] N. Doiron-Leyraud *et al.*, “Quantum oscillations and the Fermi surface in an underdoped high- T_c superconductor,” *Nature*, vol. 447, no. 7144, pp. 565–568, May 2007, doi: 10.1038/nature05872.
 - [69] N.R. Werthamer, “The Ginzburg–Landau equations and their extensions,” in *Superconductivity*, R.D. Parks, Ed., New York: Marcel Dekker, Inc., 1969.
 - [70] A. Brandt, “Multi-Level Adaptive Solutions to Boundary-Value Problems,” *Mathematics of Computation*, vol. 31, no. 138, pp. 333–390, 1977, doi: 10.2307/2006422.
 - [71] R. Barrett, Ed., *Templates for the solution of linear systems: building blocks for iterative methods*. in Software, environments, tools. Philadelphia: SIAM, Society for Industrial and Applied Mathematics, 1994.
 - [72] Q. Du, “Discrete Gauge Invariant Approximations of a Time Dependent Ginzburg–Landau Model of Superconductivity,” *Mathematics of Computation*, vol. 67, no. 223, pp. 965–986, 1998, Accessed: Mar. 12, 2024. [Online]. Available: <https://www.jstor.org/stable/2585166>
 - [73] A.N. Hirani, “Discrete Exterior Calculus,” Ph.D. Thesis, California Institute of Technology, 2003. doi: 10.7907/ZHY8-V329.
 - [74] T. Tarhasaari, L. Kettunen, and A. Bossavit, “Some realizations of a discrete Hodge operator: a reinterpretation of finite element techniques,” *IEEE Transactions on Magnetism*, vol. 35, no. 3, pp. 1494–1497, May 1999, doi: 10.1109/20.767250.
 - [75] L.J. Grady and J.R. Polimeni, *Discrete calculus: applied analysis on graphs for computational science*. London: Springer, 2010.
 - [76] G.J. Carty, M. Machida, and D.P. Hampshire, “Numerical studies on the effect of normal-metal coatings on the magnetization characteristics of type-II superconductors,” *Phys. Rev. B*, vol. 71, no. 14, p. 144507, Apr. 2005, doi: 10.1103/PhysRevB.71.144507.

- [77] T. Braune, Y. Tong, F. Gay-Balmaz, and M. Desbrun, “A Discrete Exterior Calculus of Bundle-valued Forms.” arXiv, Jun. 08, 2024. doi: 10.48550/arXiv.2406.05383.
- [78] M.T. Heath, *Scientific computing: an introductory survey*, Second edition. Boston, Mass.: McGraw-Hill, 2002.
- [79] W.L. Briggs, V.E. Henson, and S.F. McCormick, *A Multigrid Tutorial*, 2nd ed. in Other Titles in Applied Mathematics, no. 72. Philadelphia, Pa: Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 2000. doi: 10.1137/1.9780898719505.
- [80] G.J. Carty, “Studies of coated and polycrystalline superconductors using the time dependant Ginzburg-Landau equations,” Doctoral, Durham University, 2006. Accessed: Jun. 24, 2025. [Online]. Available: <https://etheses.dur.ac.uk/2656/>
- [81] A. Blair, “Simulations of Critical Currents in Polycrystalline Superconductors Using Time-Dependent Ginzburg–Landau Theory,” Doctoral, Durham University, 2021. Accessed: Mar. 26, 2025. [Online]. Available: <https://etheses.dur.ac.uk/13993/>
- [82] B.P. Din, “Computational and Analytic Time-Dependent Ginzburg-Landau Theory for High-Resistivity High-Field Superconducting Josephson Junctions,” Doctoral, Durham University, 2024. Accessed: Jun. 24, 2025. [Online]. Available: <https://etheses.dur.ac.uk/15489/>
- [83] Q. Chang and G. Wang, “Multigrid and adaptive algorithm for solving the nonlinear Schrödinger equation,” *Journal of Computational Physics*, vol. 88, no. 2, pp. 362–380, Jun. 1990, doi: 10.1016/0021-9991(90)90184-3.
- [84] M.-C. Lai and Y.-H. Tseng, “A fast iterative solver for the variable coefficient diffusion equation on a disk,” *Journal of Computational Physics*, vol. 208, no. 1, pp. 196–205, Sep. 2005, doi: 10.1016/j.jcp.2005.02.005.
- [85] N. Schlömer and W. Vanroose, “An optimal linear solver for the Jacobian system of the extreme type-II Ginzburg–Landau problem,” *Journal of Computational Physics*, vol. 234, pp. 560–572, Feb. 2013, doi: 10.1016/j.jcp.2012.10.013.
- [86] F. Xu, Y. Guo, and M. Xie, “A novel multilevel finite element method for a generalized nonlinear Schrödinger equation,” *Journal of Computational and Applied Mathematics*, vol. 457, p. 116280, Mar. 2025, doi: 10.1016/j.cam.2024.116280.
- [87] C.W.W. Haddon and D.P. Hampshire, “Fast Multigrid Simulations of Pin-

- ning in REBCO With Highly Resistive Nanorods,” *IEEE Transactions on Applied Superconductivity*, vol. 33, no. 5, pp. 1–5, Aug. 2023, doi: 10.1109/TASC.2023.3253065.
- [88] J.L. MacManus-Driscoll *et al.*, “Strongly enhanced current densities in superconducting coated conductors of $\text{YBa}_2\text{Cu}_3\text{O}_{7-x} + \text{BaZrO}_3$,” *Nature Mater*, vol. 3, no. 7, pp. 439–443, Jul. 2004, doi: 10.1038/nmat1156.
- [89] Y. Yamada *et al.*, “Epitaxial nanostructure and defects effective for pinning in $\text{Y}(\text{RE})\text{Ba}_2\text{Cu}_3\text{O}_{7-x}$ coated conductors,” *Appl. Phys. Lett.*, vol. 87, no. 13, p. 132502, Sep. 2005, doi: 10.1063/1.2061874.
- [90] A. Goyal *et al.*, “Irradiation-free, columnar defects comprised of self-assembled nanodots and nanorods resulting in strongly enhanced flux-pinning in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ films,” *Supercond. Sci. Technol.*, vol. 18, no. 11, p. 1533, Oct. 2005, doi: 10.1088/0953-2048/18/11/021.
- [91] K. Matsumoto and P. Mele, “Artificial pinning center technology to enhance vortex pinning in YBCO coated conductors,” *Supercond. Sci. Technol.*, vol. 23, no. 1, p. 014001, Dec. 2009, doi: 10.1088/0953-2048/23/1/014001.
- [92] T. Hwa, P. Le Doussal, D.R. Nelson, and V.M. Vinokur, “Flux pinning and forced vortex entanglement by splayed columnar defects,” *Phys. Rev. Lett.*, vol. 71, no. 21, pp. 3545–3548, Nov. 1993, doi: 10.1103/PhysRevLett.71.3545.
- [93] L. Civale *et al.*, “Reducing vortex motion in $\text{YBa}_2\text{Cu}_3\text{O}_7$ crystals with splay in columnar defects,” *Phys. Rev. B*, vol. 50, no. 6, pp. 4102–4105, Aug. 1994, doi: 10.1103/PhysRevB.50.4102.
- [94] F.J. Baca, P.N. Barnes, R.L.S. Emergo, T.J. Haugan, J.N. Reichart, and J.Z. Wu, “Control of BaZrO_3 nanorod alignment in $\text{YBa}_2\text{Cu}_3\text{O}_{7-x}$ thin films by microstructural modulation,” *Appl. Phys. Lett.*, vol. 94, no. 10, p. 102512, Mar. 2009, doi: 10.1063/1.3097234.
- [95] V. Selvamanickam, M.H. Gharahcheshmeh, A. Xu, E. Galstyan, L. Delgado, and C. Cantoni, “High critical currents in heavily doped $(\text{Gd},\text{Y})\text{Ba}_2\text{Cu}_3\text{O}_x$ superconductor tapes,” *Appl. Phys. Lett.*, vol. 106, no. 3, p. 032601, Jan. 2015, doi: 10.1063/1.4906205.
- [96] V. Selvamanickam, M.H. Gharahcheshmeh, A. Xu, Y. Zhang, and E. Galstyan, “Critical current density above 15 MA cm^{-2} at 30 K, 3 T in 2.2 μm thick heavily-doped $(\text{Gd},\text{Y})\text{Ba}_2\text{Cu}_3\text{O}_x$ superconductor tapes,” *Supercond. Sci. Technol.*, vol. 28, no. 7, p. 072002, May 2015, doi: 10.1088/0953-2048/28/7/072002.
- [97] I.A. Sadovskyy, A.E. Koshelev, A. Glatz, V. Ortalan, M.W. Rupich, and M.

- Leroux, "Simulation of the Vortex Dynamics in a Real Pinning Landscape of $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ Coated Conductors," *Phys. Rev. Appl.*, vol. 5, no. 1, p. 014011, Jan. 2016, doi: 10.1103/PhysRevApplied.5.014011.
- [98] H.G. Kaper and H. Nordborg, "The frozen-field approximation and the Ginzburg-Landau equations of superconductivity." arXiv, Jul. 14, 2000. doi: 10.48550/arXiv.math/0007092.
- [99] A.P. Smith, M.J. Raine, E. Surrey, S. Awaji, T. Okada, and D.P. Hampshire, "3-D Properties in (RE)BCO Tapes Measured in Fields up to 35T," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, pp. 1–5, Aug. 2019, doi: 10.1109/TASC.2019.2895218.
- [100] A.I. Blair and D.P. Hampshire, "Modeling the Critical Current of Polycrystalline Superconducting Films in High Magnetic Fields," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, pp. 1–5, Aug. 2019, doi: 10.1109/TASC.2019.2895213.
- [101] G. Wang, M.J. Raine, and D.P. Hampshire, "How resistive must grain boundaries in polycrystalline superconductors be, to limit J_c ?" *Supercond. Sci. Technol.*, vol. 30, no. 10, p. 104001, Aug. 2017, doi: 10.1088/1361-6668/aa7f24.
- [102] A.I. Buzdin, "Multiple-quanta vortices at columnar defects," *Phys. Rev. B*, vol. 47, no. 17, pp. 11416–11419, May 1993, doi: 10.1103/PhysRevB.47.11416.
- [103] A. Bezryadin and B. Pannetier, "Edge superconducting states," *Phys. Scr.*, vol. 1996, no. T66, p. 225, Jan. 1996, doi: 10.1088/0031-8949/1996/T66/041.
- [104] D.C. Larbalestier and P.J. Lee, "New developments in niobium titanium superconductors," in *Proceedings Particle Accelerator Conference*, May 1995, p. 1276-1281 vol.2. doi: 10.1109/PAC.1995.505199.
- [105] C. Meingast, P.J. Lee, and D.C. Larbalestier, "Quantitative description of a high J_c Nb-Ti superconductor during its final optimization strain. I. Microstructure, T_c , H_{c2} , and resistivity," *Journal of Applied Physics*, vol. 66, no. 12, pp. 5962–5970, Dec. 1989, doi: 10.1063/1.343624.
- [106] J.L. Murray, "The Nb–Ti (Niobium-Titanium) system," *Bulletin of Alloy Phase Diagrams*, vol. 2, no. 1, pp. 55–61, Jun. 1981, doi: 10.1007/BF02873704.
- [107] R.M. Wood, "The Lattice Constants of High Purity Alpha Titanium," *Proc. Phys. Soc.*, vol. 80, no. 3, p. 783, Sep. 1962, doi: 10.1088/0370-1328/80/3/323.
- [108] P.J. Lee and D.C. Larbalestier, "Determination of the flux pinning force of α -Ti ribbons in Nb46.5wt% Ti produced by heat treatments of varying temperature, duration and frequency," *J Mater Sci*, vol. 23, no. 11, pp. 3951–3957,

- Nov. 1988, doi: 10.1007/BF01106819.
- [109] H.J. Muller, “The Upper Critical Field of Niobium–Titanium,” University of Wisconsin–Madison, 1989. [Online]. Available: https://fs.magnet.fsu.edu/~lee/asc/pdf_papers/theses/hjm89phd.pdf
- [110] N.R. Werthamer, E. Helfand, and P.C. Hohenberg, “Temperature and Purity Dependence of the Superconducting Critical Field, H_{c2} . III. Electron Spin and Spin-Orbit Effects,” *Phys. Rev.*, vol. 147, no. 1, pp. 295–302, Jul. 1966, doi: 10.1103/PhysRev.147.295.
- [111] M. Thiemann, M. Dressel, and M. Scheffler, “Complete electrodynamics of a BCS superconductor with μeV energy scales: Microwave spectroscopy on titanium at mK temperatures,” *Phys. Rev. B*, vol. 97, no. 21, p. 214516, Jun. 2018, doi: 10.1103/PhysRevB.97.214516.
- [112] A. Peruzzi, E. Gottardi, I. Peroni, G. Ponti, and G. Ventura, “The influence of impurity concentration and magnetic fields on the superconducting transition of high-purity titanium,” *Nuclear Physics B - Proceedings Supplements*, vol. 78, no. 1, pp. 576–580, Aug. 1999, doi: 10.1016/S0920-5632(99)00607-6.
- [113] A.I. Larkin and A.A. Varlamov, “Fluctuation Phenomena in Superconductors,” in *Superconductivity: Conventional and Unconventional Superconductors*, K.H. Bennemann and J.B. Ketterson, Eds., Berlin, Heidelberg: Springer, 2008, pp. 369–458. doi: 10.1007/978-3-540-73253-2_10.
- [114] E.W. Collings and J.C. Ho, “Magnetic Susceptibility and Low-Temperature Specific Heat of High-Purity Titanium,” *Phys. Rev. B*, vol. 2, no. 2, pp. 235–244, Jul. 1970, doi: 10.1103/PhysRevB.2.235.
- [115] G.N. Kamm and J.R. Anderson, “De Haas–van Alphen Oscillations in Titanium,” in *Low temperature physics - LT 13*, 1974.
- [116] P.J. Lee and D.C. Larbalestier, “Development of nanometer scale structures in composites of Nb–Ti and their effect on the superconducting critical current density,” *Acta Metallurgica*, vol. 35, no. 10, pp. 2523–2536, Oct. 1987, doi: 10.1016/0001-6160(87)90149-0.
- [117] S.B.L. Chislett-McDonald, Y. Tsui, E. Surrey, M. Kovari, and D.P. Hampshire, “The magnetic field, temperature, strain and angular dependence of the critical current density for Nb–Ti,” *J. Phys.: Conf. Ser.*, vol. 1559, no. 1, p. 012063, Jun. 2020, doi: 10.1088/1742-6596/1559/1/012063.
- [118] C. Meingast and D.C. Larbalestier, “Quantitative description of a very high critical current density Nb–Ti superconductor during its final optimization

- strain. II. Flux pinning mechanisms,” *Journal of Applied Physics*, vol. 66, no. 12, pp. 5971–5983, Dec. 1989, doi: 10.1063/1.343625.
- [119] I.A. Sadovskyy *et al.*, “Toward Superconducting Critical Current by Design,” *Advanced Materials*, vol. 28, no. 23, pp. 4593–4600, 2016, doi: 10.1002/adma.201600602.
- [120] T. Besard, C. Foket, and B. De Sutter, “Effective Extensible Programming: Unleashing Julia on GPUs,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 4, pp. 827–841, Apr. 2019, doi: 10.1109/TPDS.2018.2872064.
- [121] S. Danisch and J. Krumbiegel, “Makie.jl: Flexible high-performance data visualization for Julia,” *Journal of Open Source Software*, vol. 6, no. 65, p. 3349, Sep. 2021, doi: 10.21105/joss.03349.