Symmetries in the Kinematic Dynamos and Hydrodynamic Instabilities of the ABC Flows

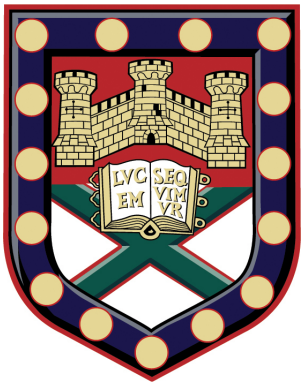

Submitted by

Samuel Edward Jones

to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Mathematics, November 2013.

I certify that all material in this thesis which is not my own work has been identified and that no material is included for which a degree has previously been conferred upon me.

. . . . . .

Samuel Edward Jones

Abstract

This thesis primarily concerns kinematic dynamo action by the 1:1:1 ABC flow, in the highly conducting limit of large magnetic Reynolds number R_m. The flow possesses 24 symmetries, with a symmetry group isomorphic to the group O_{24} of orientation-preserving transformations of a cube. These symmetries are exploited to break up the linear eigenvalue problem into five distinct symmetry classes, which we label I–V. The thesis discusses how to reduce the scale of the numerical problem to a subset of Fourier modes for a magnetic field in each class, which then may be solved independently to obtain distinct branches of eigenvalues and magnetic field eigenfunctions.

Two numerical methods are employed: the first is to time step a magnetic field in a given symmetry class and obtain the growth rate and frequency by measuring the magnetic energy as a function of time. The second method involves a more direct determination of the eigenvalue using the eigenvalue solver ARPACK for sparse matrix systems, which employs an implicitly restarted Arnoldi method. The two methods are checked against each other, and compared for efficiency and reliability. Eigenvalue branches for each symmetry class are obtained for magnetic Reynolds numbers R_m up to 10^4 together with spectra and magnetic field visualisations. A sequence of branches emerges as R_m increases and the magnetic field structures in the different branches are discussed and compared. All symmetry classes are found to contain a dynamo, though dynamo effectiveness varies greatly between classes, suggesting that the symmetries play an important role in the field amplification mechanisms.

A closely related problem, that of linear hydrodynamic stability, is also explored in the limit of large Reynolds number R_e. As the same symmetry considerations apply, the five symmetry classes of the linear instability can be resolved independently, reducing the size of the problem and allowing exploration of the effects of the symmetries on instability growth rate. Results and visualisations are obtained for all five classes for R_e up to 10^3, with comparisons drawn between the structures seen in each class and with those found in the analogous magnetic problem. For increasing R_e, multiple mode crossings are observed within each class, with remarkably similar growth rates seen in all classes at $R_e = 10^3$, highlighting a lack of dependence on the symmetries of the instability, in contrast with the magnetic problem.

This thesis also investigates the problem of large-scale magnetic fields in the 1:1:1 ABC flow through the introduction of Bloch waves that modify the periodicity of the magnetic field relative to the flow. Results are found for a field with increased periodicity in a single direction for R_m up to 10^3; it is established that the optimal scale for dynamo action varies as R_m increases, settling on a consistent scale for large R_m. The emerging field structures are studied and linked with those of the original dynamo problem. On contrasting this method with a previous study in which the flow is instead rescaled, it is shown that the use of Bloch waves drastically increases the range of possible scales, whilst cutting required computing time. Through a multiple-scale analysis, the contribution from the α-effect is calculated for the 1:1:1 ABC flow and is seen in growth rates for $R_m \ll 1$.

Acknowledgements

I am pleased to acknowledge everyone that supported me through the research process and in the writing of this thesis. I would like to thank Prof. Andrew Soward for useful advice and for helping me to access relevant literature for my research. Additionally, I am grateful to Dr. Ismaël Bouya and Dr. Emmanuel Dormy for helpful comments and interesting discussions, and for sending me their paper in advance of its publication. I thank the referees of my co-authored paper for their constructive criticism, which has proven useful in writing this thesis.

As the investigations in this thesis are heavily computational, I have relied on various computing facilities at the University of Exeter. The larger calculations reported in this paper were performed using the University of Exeter Supercomputer and I wish to thank Dr. David Acreman for providing me access to this parallel processing environment. I would like to express my gratitude to Dr. Peter Leggett and Robert O'Neale for extensive technical support and numerous time-saving coding tips, in addition to repeatedly fixing problems caused by my resource-intensive simulations on the Linux cluster. I am grateful for the support of NAG in furthering my expertise in high-performance computing and to the authors of the package VAPOR, used in producing the three-dimensional visualisations.

I would like to offer genuine thanks to my supervisors Prof. Andrew Gilbert and Prof. Mitchell Berger, who jointly envisaged this project and guided me through my research. I wish to express my enormous gratitude to Prof. Gilbert for sharing his knowledge, expertise and enthusiasm for all aspects of my research and in the wider subject area. I also owe Prof. Gilbert a special thanks for his never-ending

support, boundless patience and constant encouragement, without which I would never have come this far.

Finally, I want to acknowledge the financial support of the Leverhulme Trust for this research project, and to thank them for making this work possible.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Everything starts somewhere, although many physicists disagree.

Terry Pratchett

We begin this thesis with a comprehensive introduction to the general subject area within which this research lies. Starting with the key concepts and equations of fluid dynamics, we move on to the subfield of magnetohydrodynamics (MHD), which concerns the interaction of electrically conducting fluids and magnetic fields, and lay out the broad context in which these investigations take place. Within MHD, we focus on dynamo theory, which seeks to explain how astrophysical bodies generate and maintain magnetic fields against ohmic decay. We outline several fundamental concepts that are relevant to the investigations in this thesis before moving on to develop the immediate context of our study through discussion of the background literature. The specific objectives of this thesis are then set out and finally a breakdown of the layout of the thesis is given.

1.1 Governing equations of Fluid Dynamics

In 1687, Isaac Newton published his *Philosophiæ Naturalis Principia Mathematica* and its contents have since been critical in the development of the physical sciences. Of great significance are the three fundamental laws of motion (taken from Newton *et al.*, 1999):

1. Every body perseveres in its state of being at rest or of moving uniformly straight forward, except in so far as it is compelled to change its state by forces impressed.

2. A change in motion is proportional to the motive force impressed and takes place along the straight line in which that force is impressed.

3. To any action, there is always an opposite and equal reaction; in other words, the action of two bodies upon each other are always equal and always opposite in direction.

In using these laws, it has been possible to explain many physical phenomena, from the macroscopic orbits of the planets (when combined with Newton's law of universal gravitation) to the relatively microscopic friction forces caused by an object being dragged along a surface. Of course, the situations described were highly idealised, for example the objects were considered point masses, but the ideas of inertia, conservation of momentum and reaction forces had been established. These ideas, when applied (with some modifications) to fluids, were to become the basis of Fluid Dynamics as it is known today.

The fundamental properties of fluids are radically different to those of solids. For instance fluids are deformable, can be compressed and have the ability to flow and so, describing the motion and behaviour of fluids therefore requires an approach that accounts for these properties.

Core to the modelling of fluids are two assumptions: continuity of the fluid and the conservation of fluid properties. The first of these, continuity, is the idea that

a fluid is a continuous entity and, as such, properties such as velocity, pressure and density (amongst others) are well-defined at all infinitesimal points in the fluid. Of course, in reality this is not the case, as all fluids consist of numerous discrete molecules and so on small enough scales, this assumption will not hold. In the typical macroscopic scales dealt with when modelling fluids, the assumption of continuity holds well. The second assumption, conservation, is actually a set of assumptions, dealing with individual properties of the fluids. With regards to fluid dynamics, the properties considered important are mass, energy and momentum and it is sensible to assume that, within a closed system, these properties are conserved.

The momentum equation (1.1) and continuity equation (1.2) were originally introduced by Leonhard Euler in 1757 and described the evolution of a non-viscous (inviscid) flow through its velocity field. One form of the momentum equation is given as

$$\rho \left(\partial_t + \boldsymbol{u} \cdot \nabla\right) \boldsymbol{u} + \nabla p = \mathbf{0}, \tag{1.1}$$

with the continuity equation as

$$\partial_t \rho + \nabla \cdot (\rho \boldsymbol{u}) = 0, \tag{1.2}$$

and with $\boldsymbol{u}$ being the velocity, p as the pressure term and ρ as the fluid density. Further approximations can be made, a typical one being that the fluid is incompressible, and (1.2) can be reduced to

$$\nabla \cdot \boldsymbol{u} = 0. \tag{1.3}$$

This is referred to as the divergence-free condition and is interpreted as the fluid having no sources or sinks. The Euler equations are not capable of describing the behaviour of viscous fluids, however, and alternative systems of equations can be used, such as the Navier–Stokes equations.

The Navier–Stokes equations (named after Claude–Louis Navier and George Stokes) describe the motion of fluids which are viscous, differing from the Euler equations, and as such include a term for the viscosity; a measure of a fluid's resistance to shearing flows. The equations were derived to describe fluid behaviour whilst allowing for conservation of mass, energy and momentum and can, like the Euler equation, describe both compressible and incompressible fluids. The incompressible form of the Navier–Stokes equation can be written as

$$\rho\left(\partial_t + \boldsymbol{u}\cdot\nabla\right)\boldsymbol{u} = -\nabla\mathcal{P} + \mu\nabla^2\boldsymbol{u} + \mathcal{F}, \tag{1.4}$$

with $\mathcal{P}$ being the pressure term and $\mathcal{F}$ representing body forces applied to the fluid with velocity profile $\boldsymbol{u}$. The viscosity component is represented by the diffusion of momentum, shown by $\nabla^2\boldsymbol{u}$, with a corresponding parameter μ; the dynamic viscosity. The non-dimensionalised form of (1.4) is given by multiplying both sides through by $L/\rho U^2$, giving

$$\partial_t\boldsymbol{u} + \boldsymbol{u}\cdot\nabla\boldsymbol{u} = -\nabla p + \varepsilon\nabla^2\boldsymbol{u} + \boldsymbol{f}, \tag{1.5}$$

where L is the characteristic length-scale of the flow and with $U = \bar{\boldsymbol{u}}$ being the mean velocity. Now p and $\boldsymbol{f}$ represent the scaled pressure and forcing terms, with $\varepsilon = \mu/\rho UL = R_e^{-1}$ being the inverse Reynolds number. As a further consequence of incompressibility, (1.3) also holds.

With regards to the physical interpretation of terms in (1.5), the first term $\partial_t\boldsymbol{u}$ is the local acceleration, that is change of velocity at a given point. This is very different from the second term, $\boldsymbol{u}\cdot\nabla\boldsymbol{u}$ which is termed the advective acceleration, which describes the acceleration of a particle due to having moved with the flow. These two terms constitute the material acceleration (material derivative, with respect to time, of $\boldsymbol{u}$). The fourth term $\varepsilon\nabla^2\boldsymbol{u}$ is related to the viscosity and behaves as a dissipation term, in this case acting to dissipate kinetic energy. The final term $\boldsymbol{f}$ represents any external forcing, such as gravity. Lastly, the inverse Reynolds number ($\varepsilon = R_e^{-1}$) can be interpreted as the scale of the diffusive (viscous) forces

to the advective forces. It can then be determined, for a flow $\boldsymbol{u}$, how important these respective forces are in the evolution of the flow.

The Navier–Stokes equations are used to solve a huge range of problems in Applied Fluid Dynamics, from the classic model of flow through a pipe to modelling atmospheric processes for weather and climate forecasts. In Theoretical Fluid Dynamics, the evolution of less physical flows may be studied and understanding how these flows behave for varying R_e is often of interest.

The Navier–Stokes equations, when combined with Maxwell's equations, can be used in the study of Magnetohydrodynamics.

1.2 Magnetohydrodynamics

As a discipline of Fluid Dynamics, Magnetohydrodynamics (abbreviated as MHD) concerns the interaction of magnetic fields and electrically-conducting fluids. The term MHD was initially coined by Hannes Alfvén in 1942, though experiments that are generally considered to be of this field pre-date this term's usage by more than a century. For instance, in 1832, Michael Faraday made an attempt to measure the voltage in the River Thames induced by the motion of its water through the Earth's magnetic field, however his attempt was considered a failure.

Despite the isolated studies of the 19th Century, it wasn't until the middle of the 20th Century that MHD emerged as a field of its own. Often cited is the hypothesis put forward by Joseph Larmor in 1919 (see Larmor, 1919), where the process by which the Sun generates and maintains its magnetic field is through the flow of an electrically conducting fluid, termed Dynamo Theory. This explanation was not well received at the time and took more accurate equipment, detailed observations and significant theoretical developments for scientists to realise its actual value as a theory. An experiment detailed in Hartmann (1937), but carried out much earlier, involved using magnetic fields to drive the flow of liquid mercury; a simple electromagnetic pump, though by his own admission, the pump was not at all

efficient. These two initial examples, however, mark the first steps in what is now a very active and important field of study.

Current MHD research can be divided, as most areas of mathematical physics can, into two categories: applied and theoretical. Theoretical MHD, with which Larmor's theory resides, is concerned with describing complex phenomena through 'chunked', or idealised, models using the governing equations of MHD. Much research in theoretical MHD concerns attempts to describe naturally occurring self-sustaining magnetic fields, or dynamos. Initially, the dynamos of the Earth and Sun were considered, but as technology and the quality of observations improved, it became known that other astrophysical phenomena, such as accretion discs and galaxies, contain dynamos. It is well-accepted that the flow of the respective fluids – for instance, liquid metal in the outer core of the Earth and plasma in stars – generates and maintains a magnetic field through rotation and convection, but accurately and fully modelling these processes has still not been achieved. The difficulty being faced is that current technology does not allow us to observe and measure the flows that are driving these dynamos in addition to the fact that computing power limits the number of processes included in any given model. The current challenges in this area lie in identifying the most important processes and flow features for dynamo action to occur. Some of these processes will be discussed in section 1.3.

In addition to solving the "dynamo problem", descriptions of other observable phenomena are also being sought. For example, a highly active area is in modelling the processes behind the creation and demise of sunspots, which are seen as a consequence of the complex plasma flows in the Sun's convective zone. Although it is now known that sunspots are formed in pairs by emerging loops of magnetic field, with a sunspot surrounding each "end" of the loop, the dynamics of how these coronal loops originally form and are forced out into the photosphere is not fully understood. In recent years, there has been a spate of interest in the causes of solar flares and how magnetic reconnection in the coronal loops can produce such events, in part due to the damage that these high-energy events can cause to

man-made satellites.

In parallel to the theoretical understanding of magnetohydrodynamical problems such as natural dynamos, real physical applications have also been realised. One major area of research is in power generation, more specifically nuclear fission and nuclear fusion. Fast (breeder) reactors are one such example in which the heat is transferred away from the fissile material through the electromagnetic pumping of a liquid metal, often sodium. In the development of nuclear fusion as a viable method of power generation, magnetic confinement, which restricts the motion of high-energy plasma through strong magnetic fields, is seen as a promising field of study. There are also many industrial applications, particularly in the handling of molten metals. Magnetic fields are routinely used to stir liquid metals, drive flows or even inhibit flows in them. For instance in the casting of metals, homogeneity of the resulting object is often desired and so stirring the liquid metal without the use of a physical stirrer is required. The use of magnetofluids in propulsion has also been considered with magnetohydrodynamic principles applied to water-based and space-based propulsion systems. Such examples include the well-known "Yamato 1" boat (see Takezawa *et al.*, 1995) and the Magnetoplasmadynamic (MPD), space-based, thrusters in development by several independent institutes. Many practical problems, such as the magnitude of current required to power such devices, still need to be overcome.

The theoretical and applied areas of MHD are not completely separate and are often complementary fields of study. In more recent years, due to improvements in technology, laboratory-scale dynamos have become mechanically and economically possible. Through ventures such as the Riga Dynamo experiment (see Gailitis *et al.*, 2003) and the Von Kármán Sodium experiment (Monchaux *et al.*, 2009), dynamo action has been shown to be possible on small scales and in different geometries and the important processes, at least on the laboratory scale, are beginning to be understood. This allows the theories of the geodynamo, as well as astrophysical dynamos, to be backed up by actual empirical data.

The key equations of MHD can be derived from Maxwell's equations, originally

published between 1861 and 1862, and which expanded on the work of Faraday, Ampère and Gauss. Maxwell unified prior theories on electricity and magnetism into a set of four fundamental laws which, when used in combination, can describe nearly all electromagnetic phenomena. The equations, listed in differential form are given by

$$\nabla \cdot \boldsymbol{E} = \rho/\varepsilon, \tag{1.6}$$

$$\nabla \cdot \boldsymbol{B} = 0, \tag{1.7}$$

$$\nabla \times \boldsymbol{E} = -\partial_t \boldsymbol{B}, \tag{1.8}$$

$$\nabla \times \boldsymbol{B} = \mu \left(\boldsymbol{J} + \varepsilon \partial_t \boldsymbol{E} \right), \tag{1.9}$$

with $\boldsymbol{E}$ being the electric field, $\boldsymbol{B}$ the magnetic field and $\boldsymbol{J}$ being the (electric) current density. Differing from their earlier usage with respect to the Navier–Stokes equations, ρ, ε and μ now represent the quantities of charge density, permittivity (also referred to as the dielectric constant) and magnetic permeability, respectively. As a simplification, or rather as part of a series of approximations, the latter two parameters often take their free-space values ε_0 and μ_0.

Gauss' law of magnetism (1.7), may be interpreted as non-existence of magnetic monopoles. It can be seen as the magnetic equivalent of the continuity equation (1.3), as it requires the magnetic field to contain no sources and sinks. In (1.8), Faraday's law of induction is given, describing the interaction between a changing magnetic field $\boldsymbol{B}$ and the corresponding electric field $\boldsymbol{E}$. Ampère's law, suitably modified by Maxwell and given in (1.9), describes how magnetic fields can be produced by either (or both) the flow of electrical current or a change in the electric field, and completes the link between the magnetic fields and electric fields. As with the Navier–Stokes equations, these are classical relations and break down in specific circumstances, such as when dealing with intensely strong fields or exceptionally short distances. Maxwell's equations, however, are seen as near-perfect approximations in describing macroscopic-scale phenomena, such as those of MHD.

In the MHD equations, Ampère's original law, stating $\nabla \times \boldsymbol{B} = \mu \boldsymbol{J}$, is used and

results in the requirement that $\nabla \cdot \boldsymbol{J} = 0$. This is not unreasonable, as the term $\partial_t \boldsymbol{E}$ in (1.9), often referred to as the displacement current, has a corresponding coefficient $\mu\varepsilon$ with a free-space value, $\mu_0\varepsilon_0 = 1/c^2$ by definition, where c is the speed of light. The typical fluid velocities (of MHD problems) are of the order of V, with $V = L/T$ a ratio of the typical spatial and temporal scales L and T (respectively), and are far smaller than the speed of light c. Therefore the displacement current can be neglected: it is a factor of V^2/c^2 smaller than $\nabla \times \boldsymbol{B}$ and thus negligible for $V^2 \ll c^2$. Equations (1.8) and (1.9), without the displacement current, are combined with Ohm's Law, given by

$$\boldsymbol{J} = \sigma \left(\boldsymbol{E} + \boldsymbol{u} \times \boldsymbol{B} \right), \tag{1.10}$$

to produce the induction equation

$$\partial_t \boldsymbol{B} = \nabla \times (\boldsymbol{u} \times \boldsymbol{B}) - \nabla \times (\eta \nabla \times \boldsymbol{B}) . \tag{1.11}$$

The parameter σ is the conductivity (of the fluid) and $\eta = (\mu\sigma)^{-1}$ is the magnetic diffusivity. Although η is dependent on the temperature and density, it is common to assume that it is constant (that is, incompressible and constant temperature fluid) and so (1.11) when combined with (1.7) gives

$$\partial_t \boldsymbol{B} = \nabla \times (\boldsymbol{u} \times \boldsymbol{B}) + \eta \nabla^2 \boldsymbol{B}. \tag{1.12}$$

In parallel with the Navier–Stokes equation, to which the induction equation is similar in form, the Magnetic Reynolds number $R_m = UL/\eta$ is the ratio of the advection, $\nabla \times (\boldsymbol{u} \times \boldsymbol{B})$, to the diffusion, $\nabla^2 \boldsymbol{B}$. Of course, if the Lorentz force is significant enough to affect the fluid flow, the Navier–Stokes equation (1.5) also requires solving, with $\boldsymbol{f} = \boldsymbol{J} \times \boldsymbol{B}$ (in a closed system) as the Lorentz force acting back on the flow. In this case, (1.12) is non-linear in $\boldsymbol{B}$, as the resulting $\boldsymbol{u}$ depends on $\boldsymbol{B}$. Due to flow and magnetic field being divergence free, it is often worth

reducing the advective term to give a simplified form of the induction equation:

$$\partial_t \boldsymbol{B} + \boldsymbol{u} \cdot \nabla \boldsymbol{B} = \boldsymbol{B} \cdot \nabla \boldsymbol{u} + \eta \nabla^2 \boldsymbol{B}. \tag{1.13}$$

Dependent on the phenomenon studied, various approximations can be made. One of these is that the Lorentz force is negligible relative to the strength of the flow. In this case, the Navier–Stokes equation need no longer be solved, as the flow $\boldsymbol{u}$ is no longer dependent on $\boldsymbol{B}$. This is the core assumption in the study of kinematic dynamos, in which the effects of particular flows on the evolution of certain seed fields are investigated.

It is clear that the dimensionless parameter $R_m = UL/\eta$ is very important in determining the evolution of the flow. If $R_m \ll 1$, it is seen that diffusion is the major contributing term in the evolution of the field, whereas if $R_m \gg 1$, advection dominates. In laboratory dynamos, R_m is low, owing to the low velocities and short length scales; for example, in the Riga dynamo experiment the magnetic Reynolds number was calculated to be $R_m \approx 23$ at the maximum power level. For astrophysical objects, however, R_m is far larger; $O\left(10^3\right)$ in the Earth's core, $O\left(10^8\right)$ for the solar convection zone and $O\left(10^{12}\right)$ in the solar corona. With such extreme values of R_m, it is often of interest to study the limiting case of $R_m \to \infty \, (\eta \to 0)$ which implies zero diffusivity (i.e. that the fluid is perfectly conducting) and the magnetic field is essentially 'frozen-in' or Lie dragged by the flow. Ideal MHD uses this assumption, which is accurate on short timescales or over long length scales.

1.3 Dynamo Theory

Dynamo Theory is one of the most important theories to emerge from the developing field of MHD. It describes how a magnetic field is generated and perpetuated through the rotation and convection of electrically conducting fluids. Dynamo theory is applicable to any situation in which dynamo action is possible, but is most often applied to astrophysical objects such as the Earth and Sun. In the Earth, it is

believed that convection is driven by both temperature gradients in the fluid outer core (thermal convection) and compositionally buoyant fluid formed through the solidification of heavier elements, such as iron, at the inner core boundary (compositional convection). In the Sun, it is the motion of plasma in the convective zone that is important. The earliest such ideas can be dated back to 1919, when Joseph Larmor proposed that the Sun's magnetic field may be generated and maintained against ohmic decay through the internal motions of its fluids, this marking the beginning of an important paradigm shift in the scientific community of the time.

The leap between the proposal of such an idea and formulating a formal mathematical description of the processes is long and convoluted. Only in the past twenty years or so have numerical models successfully demonstrated some of the more prominent characteristics of the geodynamo (see Glatzmaier and Roberts, 1995), thanks to extensive research on the fundamental processes at work and massive improvements in computing power. The difficulty with such models is that the physical systems are extremely complex with a multitude of processes occurring on a huge range of scales.

This section provides an outline of some of the basic ideas and concepts within dynamo theory, but focuses more on the considerations of kinematic dynamo theory that are most relevant to this thesis.

1.3.1 Antidynamo theorems

Within MHD, one of the best ways to help identify potential dynamo mechanisms is to establish situations where dynamo action cannot occur. One of the best examples of an *antidynamo* theorem was demonstrated by Cowling (1933), in which he proves analytically that an axisymmetric magnetic field cannot be self-sustaining, and so it must be impossible for dynamo action to produce a purely axisymmetric field, as originally assumed by Larmor. Cowling's antidynamo theorem was critical, as it identified that other mechanisms must be at work to maintain the seemingly axisymmetric magnetic fields of the Earth and Sun.

Another form of Cowling's theorem, but stated in Cartesian coordinates, shows that a magnetic field which is a function of two spatial variables in three-dimensional space cannot be maintained by dynamo action. In such a case, the flow must also be a function of the same two spatial variables, meaning that there is no source for two of the three magnetic field components. Moffatt (1978) (see also Gilbert, 2003) proves that this is the case by showing that the growth rate γ of such dynamos cannot be greater than that of the slowest decaying mode, which is negative by definition.

Furthermore, it is possible to show that planar flows, those that are functions of all three spatial variables but only have two non-zero components, are unable to amplify magnetic fields and thus are not dynamos. This is demonstrated by Zeldovich (1957) and later generalised to spherical geometries (i.e. flows restricted to spherical shells) by other authors.

One result of the development of antidynamo theorems is that effective dynamos tend not to be highly symmetric, lacking, in particular, reflectional symmetries.

1.3.2 Mean-field theory and the α-effect

In many physical systems, the need to reduce the complexity of these systems in order to model them is essential. Mean field theory, which is not just confined to MHD, requires replacing small-scale or short-term processes with their averages or combining many complex interacting processes by a single process describing their average effects. The α-effect is the result of one such application to modelling of dynamos in flows with stochastic elements. The effect was initially realised by Parker (1955), though was given its current name when described by Steenbeck *et al.* (1966).

When modelling dynamos in astrophysical objects, mathematical descriptions of the flows are often unattainable because of the multitude of different scales present, making them too complex to describe analytically, and because observation of the motions themselves is not possible with current observational equipment. In models

of the solar dynamo, for example, the turbulent flows may be resolvable (at the surface of the photosphere) to a reasonably fine scale but subsurface flow profiles may be impossible to discern.

Moffatt (1978) details the method, and we give a brief summary here. A flow $\boldsymbol{U}$ can be described as a sum of a mean flow $\boldsymbol{U}_0$, a known function of space and time, and a fluctuating (i.e. random or turbulent) flow $\boldsymbol{u}$, which averages out, over some scale, to zero ($\langle \boldsymbol{u} \rangle = 0$). Similarly a magnetic field $\boldsymbol{B}$ can be written as a sum of a mean field $\boldsymbol{B}_0$ and a fluctuating field $\boldsymbol{b}$, which averages to zero ($\langle \boldsymbol{b} \rangle = 0$). The induction equation (1.12) can then be written as its mean and fluctuating parts by

$$\partial_t \boldsymbol{B}_0 = \nabla \times (\boldsymbol{U}_0 \times \boldsymbol{B}_0) + \nabla \times \mathcal{E} + \eta \nabla^2 \boldsymbol{B}_0, \tag{1.14}$$

$$\partial_t \boldsymbol{b} = \nabla \times (\boldsymbol{U}_0 \times \boldsymbol{b}) + \nabla \times (\boldsymbol{u} \times \boldsymbol{B}_0) + \nabla \times (\boldsymbol{u} \times \boldsymbol{b} - \mathcal{E}) + \eta \nabla^2 \boldsymbol{b}, \tag{1.15}$$

with $\mathcal{E} = \langle \boldsymbol{u} \times \boldsymbol{b} \rangle$ being the mean *electromotive force* (e.m.f.). If the e.m.f. can be described as a function of $\boldsymbol{U}_0$ and $\boldsymbol{B}_0$, then (1.14) can be solved without ever explicitly needing to know $\boldsymbol{b}$. By taking the leading term in the expansion of $\mathcal{E}$, it is seen that the contribution of the e.m.f. to the mean field $\boldsymbol{B}_0$ is

$$\mathcal{E}_i^{(0)} = \alpha_{ij} B_{0j}, \tag{1.16}$$

that is, a contribution proportional to the strength of the mean field itself. If the fluctuating field $\boldsymbol{u}$ is statistically isotropic, i.e. rotationally invariant, then α can also be approximated as symmetric with the contribution as

$$\mathcal{E}^{(0)} = \alpha \boldsymbol{B}_0, \tag{1.17}$$

with α a scalar. This formulation is central to much of dynamo theory, as it offers a way to handle small-scale turbulence, by using its average properties. Additionally, the poloidal $\rightleftharpoons$ toroidal field cycle is completed as the α-effect describes a mechanism in which poloidal field can be generated from toroidal field. (On a sphere, poloidal field would be parallel to lines of longitude, whereas toroidal field would

be parallel to lines of latitude.) Critically, the α-effect shows that it is possible for small-scale motions (such as convective turbulence) to drive large-scale magnetic fields.

Importantly, the α-effect cannot arise in small-scale flows $\boldsymbol{u}$ that contain reflective symmetry, as it is a pseudo-scalar and so changes sign under a coordinate inversion (reflection). One way in which these flows can avoid containing reflective symmetry is if they contain helicity $H = \boldsymbol{u} \cdot \nabla \times \boldsymbol{u}$, which does not necessarily destroy the required homogeneity of $\boldsymbol{u}$. This requirement is not unphysical, as for flows with large-scale differential rotation, such as those of the Sun, the Coriolis force leads to the generation of helicity.

Following from the concept of the α-effect, α^2- and $\alpha\omega$-dynamos are those for which the α-effect is important in different ways. If the differential rotation is not significant (and so is negligible as a source of toroidal field), the α-effect serves as a source for both poloidal and toroidal field, with such dynamos termed as 'α^2-dynamos'. Alternatively, if the dominant source of toroidal field is through differential rotation, referred to as the ω-effect, and the α-effect is a source of poloidal field, the dynamos are described as '$\alpha\omega$-dynamos'. The α-effect will be revisited in chapter 7, as part of the investigation into the effect of large-scale magnetic fields on dynamo growth.

1.3.3 Kinematic dynamo theory

In the previous section, the main assumption for kinematic dynamo theory was given by explaining that the Lorentz force remains negligible as the magnetic field evolves. This identifies that the magnetic field is not strong enough to change the flow, meaning that the flow is prescribed and only the kinematic induction equation (1.13) requires solving. Realistically this arrangement is unphysical, though it does help to explain why astrophysical bodies possess magnetic fields and is useful in describing the initial stages of non-linear dynamo evolution. On a theoretical basis this allows for interesting analysis of the given flow $\boldsymbol{u}$ such as identifying the

particular mechanisms responsible for effective field amplification.

Kinematic (linear) dynamo theory is useful for ascertaining whether specific flows are capable of dynamo action, that is for a range of R_m, the growth rate γ of the magnetic field is positive. Unlike in non-linear dynamos, where the magnetic field reaches saturation after sufficiently changing the flow, the magnetic energy can grow exponentially, with γ as the rate of exponential growth. Due to the linearity of (1.13) in $\boldsymbol{B}$, this is equivalent to solving an eigenvalue problem with many eigenvalues λ, and for which $\gamma_0 = \mathrm{Re}(\lambda_0)$ is the growth rate of the fastest growing eigenmode. Additionally, kinematic dynamo theory is useful for stability analysis of equilibria states of fully non-linear dynamo problems, as it is able to identify whether equilibria (in which there is no magnetic field) are stable states ($\gamma \leq 0$) or unstable states ($\gamma > 0$).

1.3.4 Fast and slow dynamos

In the kinematic regime, one important distinction (originally made by Vainshtein and Zeldovich, 1972) between particular (laminar) flows is that of their respective dynamos being 'fast' or 'slow'. A fast dynamo is one which for which $\gamma \to \gamma_0 > 0$ for $R_m \to \infty$, that is in the perfectly conducting limit, the growth rate has reached an asymptotic limit which is greater than zero. Alternatively, a slow dynamo is one for which $\gamma \to \gamma_0 \leq 0$ as $R_m \to \infty$, that is in the perfectly conducting limit, the flow is unable to support dynamo action. It is impossible to resolve growth rates in the limit of $R_m \to \infty$, owing to the ever-smaller diffusive scales needing to be resolved as R_m increases. Even for large R_m, particularly for three-dimensional flows, the growth rates can be difficult to obtain. However, since fast dynamos reach an asymptotic limit, the application of asymptotic theory to such dynamos is often possible.

The need to find fast dynamos is driven by our understanding of astrophysical dynamos. The solar dynamo is considered to have features of a fast dynamo, as its magnetic field evolves on the time scale of months or years. Though the evolution of

the Sun's magnetic field is not kinematic, the mechanisms for field amplification are considered to be similar (see Childress and Gilbert, 1995). The Stretch-Twist-Fold mechanism is discussed in section 1.3.5.

Conclusive proof of fast dynamo action in physically realistic flows is not possible through current mathematical theories and so relies on numerical evidence and asymptotic analysis for classification. The main difficulty in classifying dynamos is that the smallest scales (diffusive) to be resolved are $O\left(R_m^{-1/2}\right)$ and so in the high-R_m regimes, numerical analysis is limited by computational power, owing to the high resolutions needed, particularly in fully three-dimensional dynamos.

The Ponomarenko dynamo (Ponomarenko, 1973) is perhaps the most basic example of a dynamo, and is described as solid body rotation of a cylinder of fluid, with a vertical velocity component, and with zero fluid velocity outside this cylinder: the discontinuity of the fluid flow at the boundary is critical in the generation and amplification of magnetic fields. Through high-R_m numerical analysis (and asymptotic approximations), it is seen that the fastest growing modes appear to be independent of R_m, implying that this may be a fast dynamo (see Gilbert, 1988). By contrast, the smooth Ponomarenko dynamo, which is a modification of the Ponomarenko dynamo that avoids the discontinuity, has been identified as being slow.

The existence of slow dynamos has been analytically and numerically proven, with one example being the spatially-periodic Roberts dynamo. As the flow is integrable, it is more amenable for asymptotic analysis, which identifies a decreasing growth rate with increasing R_m (see Soward, 1987), confirming the findings of the original numerical analysis (see Roberts, 1970). Soward (1987) argues that, although it is a slow dynamo, it is actually close to being fast, with the growth rate $\gamma = O\left(\log\log R_m / \log R_m\right)$. The Roberts flow is a special case of the ABC flows, with $A = B, C = 0$. These flows are discussed in section 1.4.

1.3.5 Stretch, twist and fold

The stretch-twist-fold (STF) mechanism builds on the original models put forward by Alfvén (1950) for dynamo action in the presence of strong advection and weak diffusion (i.e. high-R_m regimes). The mechanism describes how field lines are advected with the flow which, if the flow is sufficiently complicated, can stretch and twist these field lines in a constructive sense, resulting in tubes of magnetic flux of the same polarity that are aligned. STF is perhaps the best known of the proposed fast dynamo mechanisms, and was described by Vainshtein and Zeldovich (1972), with further development by Childress and Gilbert (1995). Figure 1.1 illustrates

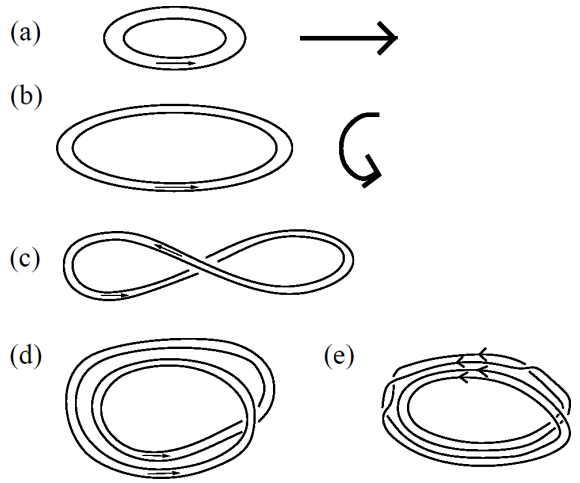

Figure 1.1: The stretch-twist-fold process of magnetic field amplification, from Gilbert (2003).

the idealised STF mechanism, which relies on the conservation of magnetic flux within the flux tubes. The process involves the stretching of a flux tube depicted in (a) to (b). The flow then twists the field, shown in (b) to (c), and then folds it back on itself in such a way that the flux tubes are aligned, polarity-wise, seen in (d). On repeating the process, the two flux tubes are doubled to four and in continued repetition of the process, the effective number of flux tubes increases exponentially. The process depicted does not necessarily consider diffusion, but in such an example, diffusion would serve to smooth field without cancellation of opposing fields, as folding is constructive. This, of course, relies on the very weak

diffusion seen in the high-R_m regimes of fast dynamo mechanisms.

The magnetic flux in a single flux tube is defined as the flux density B that passes through the cross-sectional area A of the tube and is calculated as BA, this quantity being conserved for a given flux tube. In being stretched by the flow, the flux tube's cross-sectional area decreases and so to conserve flux, the flux density increases accordingly. For instance, a reduction of the area to $A/2$ will result in a flux density of $2B$ passing through this reduced area. Through constructive folding of the tubes, assuming that the process results in two (connected) flux tubes of the same length as the original, each flux tube has half the cross-sectional area and twice the flux density, resulting in twice the total flux of the original tube. With repeated iterations, the number of flux tubes increases exponentially and the magnetic flux density (and thus magnetic energy) also experiences exponential growth. In non-perfectly conducting regimes, flux is lost through diffusion and so the growth rate of an STF dynamo is dependent on the efficiency of these amplifying mechanisms relative to the diffusive losses.

The whole STF process relies on the ability of the flow to exponentially stretch magnetic field lines, thus a chaotic flow is required; these are discussed in the next subsection.

1.3.6 Chaotic flows

To emulate the stochastic nature of small-scale flows in the solar dynamo, kinematic models can be applied to chaotic flows. In chaotic flows, two fluid particles placed infinitesimally close may experience exponential divergence on being advected through the flow. For such flows, the streamlines will be space-filling in the regions that are chaotic, though it may only be small areas of the flow that actually contain 'chaos', with the rest of the flow being integrable. One possible way to achieve this is through turbulence, in which the flows either contain an element of randomness or are time-dependent (i.e. unsteady) flows. As detailed above, the ability to stretch magnetic field lines admits the possibility of mechanisms (includ-

ing STF) that are considered essential for fast dynamos, thus chaos is considered an important ingredient of a fast dynamo.

It is possible to create chaotic trajectories in flows that only depend on two spatial variables by introducing time-dependence. The CP and LP flows of Galloway and Proctor (1992) are such an example. A Poincaré section identifies that the flows indeed have chaotic trajectories and by numerical simulations it is seen that the CP flow may be fast, while the LP flow shows indications of being slow.

It is also possible for steady three-dimensional flows to be chaotic and thus amenable for fast dynamo action. In an artificial 3-D flow defined on a Riemannian manifold, Arnold *et al.* (1981) demonstrate that exponential growth of magnetic energy is possible, concluding that streamline chaos, in such flows, favours magnetic field amplification, with the growth rate controlled by the largest Lyapunov exponent (maximum rate of line stretching possible in the flow). The flows which are the subject of this thesis are perhaps the best-known example of steady three-dimensional chaotic flows; the ABC flows. These flows are interesting as the (Lagrangian) chaos is restricted to thin bands, surrounded by integrable flow, and may be candidates for fast dynamos. These flows are the subject of the next section.

1.4 ABC flows and ABC dynamos

In this section, we review the literature surrounding the ABC flows, which are the subject of the investigations in this thesis. We refer to studies both in hydrodynamic and dynamo contexts and aim to provide a synopsis of prior studies, to bring the reader up to speed with the current level of understanding of these flows. Although the literature on ABC flows is not particularly extensive, we have endeavoured to include those studies that have direct relevance to this thesis.

The ABC flows are a family of flows, defined on $\mathbb{T}^3$ (the 3-torus), that are given as

$$\boldsymbol{u} = (C\sin z + B\cos y, A\sin x + C\cos z, B\sin y + A\cos x), \tag{1.18}$$

with $A, B, C \in \mathbb{R}$ and x, y and z periodic, with period 2π. Although the flows were originally introduced in Beltrami (1889) (see figure 1.2) as an example of a class of flows with the Beltrami property, given by

$$\nabla \times \boldsymbol{u} = k\boldsymbol{u}, \tag{1.19}$$

the flows have, much more recently, become the subject of many studies, owing to their interesting properties and features.

$$u = T_2 \cos 2Ty + T_3 \operatorname{sen} 2Tz,$$
$$v = T_3 \cos 2Tz + T_1 \operatorname{sen} 2Tx,$$
$$w = T_1 \cos 2Tx + T_2 \operatorname{sen} 2Ty,$$

Figure 1.2: The original flow profile in Beltrami (1889), satisfying the $\nabla \times \boldsymbol{u} = k\boldsymbol{u}$ condition. Note that T_1, T_2 and T_3 are equivalent to A, B and C.

It is noted that in (1.18), the parameters A, B and C have been assigned, respectively, to the functions of x, y and z. This differs to the original usage seen in Arnold (1965), Hénon (1966) up to Bouya and Dormy (2013) but coincides with the form given in Gilbert (1992) and Childress and Gilbert (1995) and seems a more logical definition.

This section aims to explore the relevant literature for the ABC flows, with the objective of providing the reader with enough information to understand the context of the work carried out for this thesis.

1.4.1 Early investigations of the ABC flows

The flows were first introduced in their current form by Arnold (1965), due to them being exact three-dimensional steady solutions to the Euler equations (1.1,1.2) and Navier–Stokes equations (1.5), provided that, for the latter, the correct forcing term (to counter viscous dissipation) is included. One flow, $A = \sqrt{3}, B = \sqrt{2}, C = 1$, is

identified by Hénon (1966) as chaotic through numerical calculations.

With respect to dynamo action, ABC flows are first considered by Childress (1970), in which the $A = B = C = 1$ (1:1:1) flow

$$\boldsymbol{u} = (\sin z + \cos y, \sin x + \cos z, \sin y + \cos x), \tag{1.20}$$

is introduced as a potential model for a space-periodic three-dimensional kinematic dynamo, however all calculations are purely analytical. It is not until Arnold and Korkina (1983) that the first numerical calculations are undertaken using the 1:1:1 flow. In their investigation, they identify that kinematic dynamo action is indeed possible for (1.20) and find that the magnetic field only grows for a limited range of $\eta = R_m^{-1}$. The values are given as $R_m \approx 8.93$ and $R_m \approx 17.5$ as the respective lower and upper limits of this 'window' (with a maximum at $R_m \approx 12.5$), and with the field decaying for values outside this range. They also note the form of the growing field in this window as

$$\boldsymbol{b} = (\cos y - \sin z, \cos z - \sin x, \cos x - \sin y), \tag{1.21}$$

similar to the 1:1:1 flow, except with the sines negated. As an extension to this investigation, Arnold (1984) discusses the numerical work of the previous paper in more detail but admits that the inherent symmetry of the flow (1.20) is clearly important, as in his Galerkin approximation, he finds certain spectral modes to be zero-valued. Arnold presents the generating symmetries as

$$\mathsf{g}_4(x, y, z) \mapsto (x + \pi/2, z - \pi/2, \pi/2 - y), \quad \mathsf{g}_3(x, y, z) \mapsto (y, z, x), \tag{1.22}$$

with $\mathsf{g}_4^4 = \mathsf{g}_3^3 = \mathsf{i}$. The full group, G, of 24 orientation-preserving symmetries is then generated by these symmetries. A description of how this can be achieved for the different ABC flows will be given in chapter 2. Also noted by Arnold is that the group G is isomorphic to the group of proper rotations of a cube. By proper rotations, Arnold refers to those that are orientation-preserving, i.e. those

involving no reflections, as reflections are orientation-reversing transformations. He then details how one would go about deconstructing spheres of spectral modes, that is those for given wavenumber k^2, into linear combinations of their irreducible representations. The process will not be described here but this thesis will be taking elements of this approach, specifically reducing the problem into its constituent symmetry classes. The breakdown of the 1:1:1 ABC dynamo into these classes is the subject of chapter 3.

1.4.2 Structure of the ABC flows: chaotic and integrable flow

Dombre *et al.* (1986) provides an in-depth analysis of the structure of different ABC flows and helps to justify why they are so compelling to study. It is in this investigation that the term "ABC flows" is coined, with the A, B and C making reference to Arnold, Beltrami and Childress. As in Arnold (1984), they note the various symmetries, both time-reversing and time-preserving, and use these to restrict the parameters A, B and C so that

$$1 = A \geq B \geq C \geq 0, \tag{1.23}$$

with any choice of parameters being transformable to this situation. By studying where the velocity vanishes, they obtain equations for the stagnation points, finding that eight (four pairs) exist in the periodic cube for the case $A = B = C$ and that for ABC flows with (1.23) and $B^2 + C^2 < 1$, none exist at all. These pairs of stagnation points (in all ABC flows that contain them) are formed of one α-type (two-dimensional stable manifold) and one β-type (two-dimensional unstable manifold) stagnation point with a heteroclinic line, or separatrix, joining the two. Streamlines of the flow are drawn in at the α-type, stretched along the separatrix and are expelled at the β-type. They also note that there are areas of integrable flow, away from the stagnation points and separatrices. They term these the "principal vortices" and are loosely defined as areas of the flow for which motion is predom-

inantly in one direction, that is in the x, y or z directions. The original diagram identifying these is given in figure 1.3. Through numerical integration, Poincaré

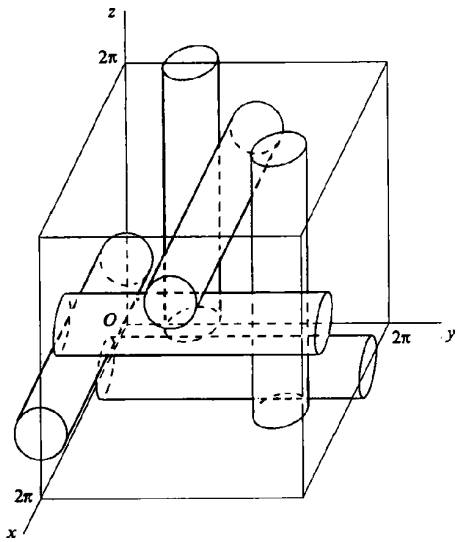

Figure 1.3: Diagram showing the six principal vortices of the ABC flows, taken from Dombre *et al.* (1986). These are areas of flow that have motion in largely one direction (along the length of the tube).

sections were produced for the various flows studied with clear areas of Lagrangian chaos (non-integrable flow). Particularly in the case of $A = B = C = 1$, the areas of chaos are restricted to the stagnation points and along the separatrices. Thin bands of chaos also exist joining the non-paired stagnation points, though these bands are curved and so avoid intersecting the principal vortices.

1.4.3 The 1:1:1 ABC dynamo: the second window

Galloway and Frisch (1986) are the next to address the ABC dynamo problem, in an effort to determine whether the 1:1:1 flow, and others, produce fast or slow dynamos. With access to more computing power, they are able to extend the range of R_m studied from $R_m \approx 20$ as seen in Arnold and Korkina (1983) to $R_m = 550$ in the 1:1:1 flow. For low R_m, they confirm the existence of a window of dynamo action found by Arnold and Korkina (1983) but also discover that a new window of dynamo action appears at $R_m \approx 27$ and persists to the limit of their simulations

($R_m = 550$). The growth rate of the dynamo in the second window is distinctly higher than that of the first, with a maximum of $\gamma = 0.079$ at $R_m = 550$ in comparison to a maximum of $\gamma = 0.0096$ at $R_m = 12.5$. Computational difficulties are encountered, as identified by error bars showing the uncertainty in their results. They indicate that a change from an oscillating mode seen for $R_m < 200$ to a steady mode for $R_m > 350$ is possibly the cause of these difficulties and increasingly long simulations are required to discover if this is the case. Figure 1.4 is taken from

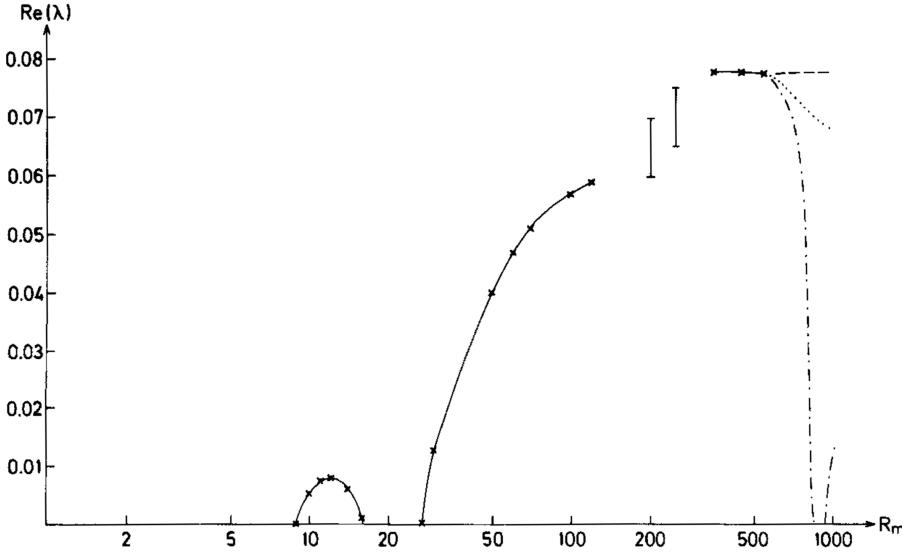

Figure 1.4: Graph of growth rates versus R_m from Galloway and Frisch (1986). Conjectures on three possibilities for future behaviour are shown with dashed, dotted and dash-dotted lines. Error bars indicate uncertain measurements.

Galloway and Frisch (1986) and shows their results for the 1:1:1 dynamo. They go on to provide conjectures for future growth rates, distinguishing three separate cases: a slow dynamo (dotted line) with $\gamma \to 0$ as $R_m \to \infty$, a fast dynamo (dashed line) with $\gamma > 0$ for $R_m \to \infty$ and an "intermediate" dynamo with which the growth rate, for some $R_m > 550$, becomes negative and the dynamo re-emerges at a larger R_m.

Unlike in previous studies, the authors are able to visualise the areas of strong magnetic field and discover that it is restricted to "cigar-like" structures centred on the α-type stagnation points. The strong field appears to be generated at the α-type stagnation points, where the flow draws in the field near to a two-dimensional

stable manifold and then ejects it near to a one-dimensional unstable manifold. On approaching the β-type stagnation point, the magnetic field lines are dispersed by the flow, forming the ends of the cigar structure. The width of these structures is directly linked to R_m, with the authors citing the diffusive scale of $R_m^{-1/2}$. In their investigation of magnetic field structures in the 1:1:1 flow and others, such as the $2^{1/2} : 1 : 1$ flow, it is clear to the authors that the strongest areas of field are those in which the Lagrangian chaos exists.

On further analysis of the field in the first dynamo window, Galloway and Frisch test the field for "symmetry-breaking" structure by choosing random initial conditions and then checking the values of Fourier modes against expected values (this process is similar to the representation "fingerprinting" used in Arnold (1984)). They note that symmetry-breaking does not occur in this window, with the field identified as that in (1.21), which is even under $\mathbf{g}_3$ and odd under $\mathbf{g}_4$ of Arnold (1984). On analysis of the field in the second window, however, the authors argue that symmetry-breaking occurs as none of the symmetries found in the first window are found to be present in the second. This leads the authors to conclude that the magnetic fields in each window are distinctly different modes, with the mode-crossing occurring at $R_m \approx 24$.

As a result of the helical nature of the flow, the α-effect is investigated for these flows. The authors calculate α for values of R_m in both windows of the 1:1:1 flow but note that α saturates in the first window but grows exponentially (with the dynamo mode) in the second. The authors admit that the argument for the importance of the α-effect may not be valid, as the dynamo mode contributes to the mean e.m.f. and so values calculated for the second window may not be sensible.

Galloway and Frisch (1986) outline their numerical procedure and describe that, because of the form of the ABC flows, the calculations can be carried out in spectral space; the numerical method is fully spectral. The resolution scales proportionally with $R_m^{1/2}$ and so the number of Fourier modes scales like $R_m^{3/2}$, a scaling which will be adopted in the numerical work of this thesis.

Part of the investigation by Lau and Finn (1993) explores the ABC dynamo, reproducing the results of the 1:1:1 flow dynamo (to within small error margins) of Galloway and Frisch (1986). They calculate that the change from oscillating mode to steady mode in the second window is between $R_m = 300$ and $R_m = 350$. The authors determine a growth rate $\gamma = 0.076$ for $R_m = 1000$ in the 1:1:1 ABC flow. When considered with the growth rates of Galloway and Frisch (1986), it appears that the growth rate has not changed and may support the fast dynamo conjecture, of Galloway and Frisch, that the growth rate has reached an asymptotic limit. Indeed, Lau and Finn attempt to extrapolate this and previous results to provide the estimate that $\gamma \to \gamma_0 = 0.077$ for $\eta \to 0$ ($R_m \to \infty$), though later studies have shown that $R_m = 1000$ is not sufficiently large to make such a judgement.

1.4.4 Linear hydrodynamic stability of the ABC flows

In a parallel investigation to that of 1986, Galloway and Frisch (1987) discuss the purely hydrodynamical analogue of the dynamo problem. In adding a zero-divergence perturbation $\boldsymbol{v}$ to the ABC flow, applying the Navier–Stokes equation and neglecting non-linear terms, the authors derive the (linear) equation for evolution of the perturbation, given as

$$\partial_t \boldsymbol{v} + \boldsymbol{u} \cdot \nabla \boldsymbol{v} + \boldsymbol{v} \cdot \nabla \boldsymbol{u} = -\nabla p/\rho + R_e^{-1} \nabla^2 \boldsymbol{v}, \tag{1.24}$$

with p as the perturbed pressure. The equation is very similar to that of the magnetic problem, though with an extra pressure term which is eliminated through a projection in Fourier space. The same numerical scheme is adopted as with the magnetic case. Results are reported for three different ABC flows; 1:0:1, 2:1:1 and 1:1:1, with the original graph provided in figure 1.5. At $R_e = 100$, the most unstable (fastest growing) perturbation belongs to that of the 1:0:1 flow, however the growth rate of the 2:1:1 flow perturbation shows signs of overtaking that of the 1:0:1 flow at high R_e. The most stable of the flows is the 1:1:1 flow with an oscillatory mode that first becomes unstable at $R_e \approx 13$. The perturbation grows as R_e increases

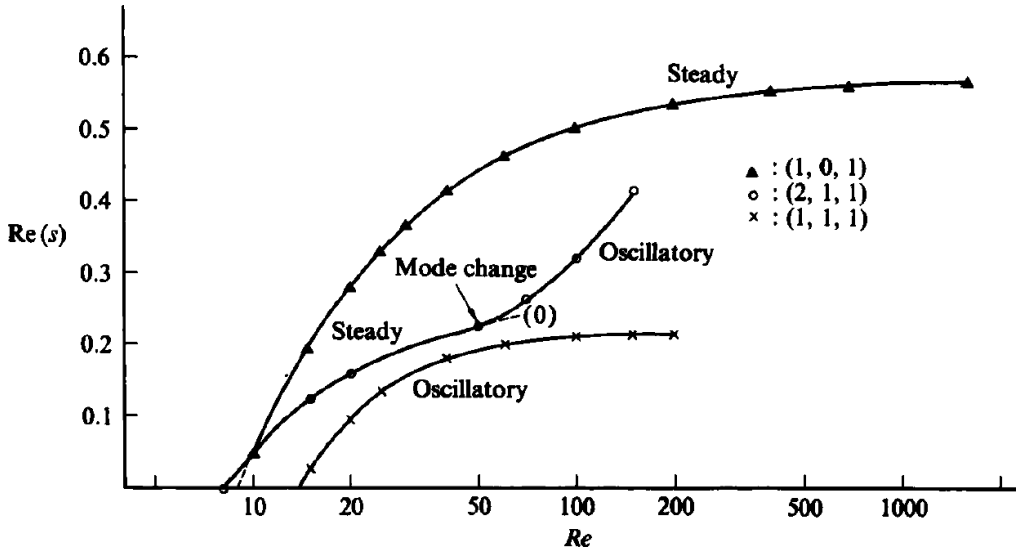

Figure 1.5: Graph of instability growth rates versus R_e for three flows (taken from Galloway and Frisch, 1987).

and by $R_e = 200$, it appears that the growth rate has levelled off at 0.216. The structure of the energy density of the perturbation for the 1:1:1 flow (but not other flows) is visualised at $R_e = 250$ and clearly show cigar-like structures, reminiscent of those in Galloway and Frisch (1986). The authors associate this with the strong advection and straining near to the α- and β-type stagnation points discussed in Dombre *et al.* (1986), as the structures occupy the same locations as with the magnetic problem.

1.4.5 Symmetries and irreducible representations of the ABCS flow

In investigating steady chaotic flows in the absence of diffusion (with $\eta = 0$), two forms of the ABC flows are considered by Gilbert (1992); the 1:1:1 flow and the ABCS ($A = B = 1, C \ll 1$) flow. Of particular relevance to this thesis is that the symmetries of both flows are considered. The ABCS flow contains two generating symmetries: r is a $\pi/2$ rotation (quarter turn) about the axis $x = y = \pi/2$, whereas s is a π rotation (half turn) about the axis $y = z = 0$. These transformations generate a group of eight elements, with this group being isomorphic to the group

D_8, the symmetries of a square. In the method described by Arnold (1984) for the 1:1:1 flow, the five irreducible representations (labelled I–V) of this group are found through use of the group's character table. This allows the author to construct magnetic fields that belong to each of the representations, of which there are five for this flow, and from these fields, initial conditions are generated for numerical calculation. The resulting computations identify large differences in the growth of the field in different representations, suggesting that the symmetries of the field are important in determining how fast a dynamo will grow.

Results are also given from applying the same methods to the 1:1:1 flow. The author provides the complete list of 24 symmetries for this flow and displays the character table, identifying five irreducible representations, again labelled I–V. The initial conditions in each representation lead to roughly similar growth rates, apart from representation I, which shows little or no growth. In drawing a comparison of these results to those of Arnold and Korkina (1983) and Galloway and Frisch (1986), the author identifies that in the first window of dynamo action (in the kinematic 1:1:1 flow dynamo), the field belongs to representation II whereas in the second window, the dominant field belongs to representation V.

1.4.6 Large-scale magnetic fields in the ABC flows

An alternative way to describe the ABC flows is to include a term, $k_0 \in \mathbb{N}$, that controls the periodicity (given by $2\pi/k_0$) of the flows. Galanti *et al.* (1992) define the ABC flows by

$$\boldsymbol{u} = (C \sin k_0 z + B \cos k_0 y, A \sin k_0 x + C \cos k_0 z, B \sin k_0 y + A \cos k_0 x), \quad (1.25)$$

which when $k_0 = 1$ gives the original flows (1.18). In driving the evolution of a magnetic field using the flow (1.25) in a $[0, 2\pi]^3$ box, the magnetic field is allowed to evolve on spatial scales greater than those of the flow, something which is not possible with the original ABC flows.

In addition to the investigations of the kinematic dynamo for $k_0 = 1$ by varying the ratio B/A and of the non-kinematic dynamo problem (also for $k_0 = 1$), the authors compare growth rates for kinematic dynamos driven by ABC flows with $k_0 = 1$ and $k_0 = 2$. The amplitudes are scaled by setting $A = B = C = k_0$, so that R_m remains comparable as k_0 is varied. In addition, the growth rates found with a specific k_0 need to be scaled by $1/k_0^2$. They find that for $k_0 = 2$, the growth rates are considerable larger than those of $k_0 = 1$ and that the distinct windows of dynamo action of the latter do not exist for the former. Although results are limited to R_m up to 45, it is clear that a dynamo driven by (1.25) with $k_0 = 2$ is far more effective. In fact, for $R_m = 12$ and 20, simulations with k_0 from 1 to 10 reveal that for $k_0 \geq 2$ the growth rates are very similar; as the $k_0 = 2$ mode exists in simulations for which $k_0 = 2n$, for $n \in \mathbb{N}$, then the lack of improvement as k_0 is increased suggests that the optimal scale for dynamo action is when the flow is approximately twice the period (i.e. for $k_0 \approx 2$), at least for these low values of R_m.

Notably, in studying the temporal evolution of individual Fourier modes, those with wavenumbers close to k_0 are oscillatory, as opposed to lower wavenumber modes (representing large-scale magnetic fields relative to the flow's scale) which grow monotonically. The authors associate this with the α-effect, determining (from multiple-scale analysis) a monotonic growth rate of αk for a mode with wavenumber k.

As a brief extension to this work, Galanti *et al.* (1993) provide additional results by varying k_0 for $R_m = 35, 60$ and 100. It is seen that the dynamo growth rate is lower at $k_0 = 2$ as R_m is increased. Furthermore, by interpolating the growth rates, it is observed that the optimal scale is actually *decreasing* as R_m increases and goes from $k_0 \approx 2.2$ at $R_m = 20$ to $k_0 \approx 1.8$ at $R_m = 100$. Exploring non-integer values, however, is not possible with this method.

The methods adopted and results achieved in these articles are discussed further in chapter 7, where a comparison is drawn to the method of imposing large-scale fields through the use of Bloch waves, which we use to study the optimal magnetic

field scales for a range of R_m.

1.4.7 Non-linear hydrodynamic stability of the 1:1:1 ABC flow

Podvigina and Pouquet (1994) investigate the stability of the 1:1:1 ABC flow in a non-linear regime, through solving the Navier–Stokes equation (1.5) with initial flow (1.20). From the results it is seen that the onset of instabilities occurs over a very short range of Reynolds number R_e. For instance, at $R_e = 13$, the flow is stable, but at $R_e = 13.05$ the switch to an unstable mode occurs. The nature of these instabilities changes until the flow is again stable in the interval $R_e \in [13.5, 13.9]$ and for $R_e = 13.95$ it again becomes unstable. Interestingly, the authors choose to identify the symmetries of the stable and unstable modes, by comparing spectral mode values to expected symmetric values, and note that when stable, the flow contains all listed symmetries. When unstable, the flow may be asymmetric or contain some of the symmetries and eventually as R_e is increased past 16, no symmetries are present.

In Podvigina (1999), the branches of stable solutions found by Podvigina and Pouquet (1994) are investigated further by applying tangent predictions to steady solutions of lower R_m computations to approximate solutions of the same branch at higher R_m. The Fourier space decompositions of each representation (of the 1:1:1 flow) are also provided and it is here that we see the restrictions on mode values. Similar restrictions will be used for the magnetic field in this thesis, when decomposing a general field into its constituent symmetry classes; see chapter 4.

1.4.8 Growth rate bounds for the 1:1:1 ABC flow

The answer to the question of whether any of the ABC flows contain fast dynamos is still highly sought after. Childress and Gilbert (1995) attempt to place potential bounds on the growth rate of a kinematic dynamo, specifically for the 1:1:1

flow. The flux is calculated through a triangular plane centred on the midpoint of a separatrix (heteroclinic line joining two stagnation points in the flow) and a growth rate is calculated from this. Using the initial condition (1.21), which is the dominant eigenmode of the first window, the authors calculate a figure (for the flux growth rate) of $\Gamma_S \simeq 0.055$, however the growth rate already far exceeds this in calculations made by Galloway and Frisch (1986) and Lau and Finn (1993). Childress and Gilbert acknowledge that the growth rate of the dynamo with initial condition (1.21) is in fact negative for $100 \leq R_m \leq 400$ (see Galloway and O'Brian, 1994).

The rate of line-stretching h_{line} is also calculated for the 1:1:1 flow, giving it as $h_{\text{line}} \simeq 0.09$, which they assert is an upper bound for γ as $R_m \to \infty$. In referring to the results of Galanti *et al.* (1992), however, the growth rates of magnetic fields with twice the period of the ABC flow are given as 0.17 at $R_m = 250$; considerably higher than that of h_{line}, lending credence to the suggestion that the values of R_m reached, in numerical calculations that include diffusion effects, are still not sufficiently large for the growth rate to have become asymptotic.

Also raised by the authors is a question regarding how much the magnetic field in numerical computation can "see" the chaos of the flows. The separation of grid points in the largest simulation undertaken (that of Lau and Finn, 1993, at $R_m = 1000$) is rather large when compared to the size of the bands of chaos and so it may be open to discussion whether the full effect of the chaos could yet be seen. In visualising field evolution for a seed field (1.21) with $\eta = 0$, it is seen that bands of field build up in the chaotic region, with evidence of constructive folding of the field.

1.4.9 Field structure of the 1:1:1 ABC dynamo: investigations into two windows of dynamo action

The ABC dynamos are revisited in Galloway and O'Brian (1994), with simulations carried out for the 1:1:1 flow, the 5:2:2 flow and a third flow, defined by

$\boldsymbol{u} = (\sin z, \sin x, \sin y)$, which is also investigated by Galloway and Proctor (1992). The magnetic field visualisations produced by Galloway and O'Brian are clearer than those of Galloway and Frisch (1986) owing to improvements in visualisation techniques and computer processing power. In the 1:1:1 flow dynamo, these improvements allow the authors to observe (for $R_m = 100$) that the areas of strong field, based around the α-type stagnation points and separatrices, each consist of two distinct cigar-like structures formed on opposing sides of the separatrices. On further investigation, it is found that each cigar has opposing polarity. As the dynamo for $R_m = 100$ is oscillatory, they also resolve the field structure for one quarter of a period later and observe that only one cigar remained where previously there were two, suggesting that the structure evolves significantly over a single cycle.

In taking the initial condition given in (1.21), which is the eigenmode of the first window of dynamo action, Galloway and O'Brian calculate growth rates for $R_m = 100, 200$ and 400, finding that $\gamma \approx -0.03$ for all three values, suggesting that this eigenmode is a "fast non-dynamo". Visualising this eigenmode at $R_m = 100$ reveals a single cigar along the separatrix, with three rotationally symmetric "fins" of field with the opposite polarity to the cigar. Fields of this description can be seen in chapter 5, in the visualisations of fields in symmetry class II.

Archontis and Dorch (1999) discuss the magnetic field structures found in two kinematic ABC dynamos; those of the 1:1:1 and 5:2:2 flows, the latter of which has no stagnation points. Of interest is their discussion of the flow topology, particularly in the region of the stagnation points. The streamlines converge at the α-type stagnation points in three-fold, rotationally-symmetric 'leaves' that lie on the plane perpendicular to the heteroclinic orbit (separatrix) that joints the α-type and β-type stagnation points. Similarly, the streamlines leave the β-type stagnation points in the same three-fold leaf formations.

In their numerical work, growth rates are not explicitly provided, though it is assumed for formerly resolved values of R_m, that the growth rate agrees with those of previous studies. Following prior investigations, the initial condition (1.21) is

used and compared with a uniform seed field. As expected, in the first window ($8 \leq R_m \leq 17$), the results are identical for the two seed fields. In the second window, from a uniform seed, the field settles to the oscillating, exponentially growing energy profile that is expected. Using the initial condition (1.21) at first gives an oscillating, exponentially decaying field that visually resembles that of the first window, though eventually it settles to the dominant field of the second window. Over sufficient time steps, the rounding error becomes large enough to break the symmetry of (1.21) and settles to the fastest growing dynamo mechanism. This makes it difficult to resolve any modes other than the dominant one for a given R_m.

The authors also describe the process of field amplification in the dynamo mechanism of the second window. The two-cigar structure seen in Galloway and O'Brian (1994) is reconfirmed, though they are able to visualise the fields at different times and so observe the field evolution over a cycle. At the start of a cycle (minimum energy), only one cigar exists near the α-type stagnation point. Over time a secondary cigar of opposite polarity forms and grows, while the primary cigar diminishes. By the following minimum, only the secondary cigar remains so that the configuration at this point in the cycle is the same as the start but reversed in polarity. This describes a half-cycle, with the second half-cycle effectively identical so that after a full cycle, the field's original configuration is restored. Essential in the operation of the dynamo mode is the magnetic field reconnection that occurs between the cigars of opposing polarity. According to the authors, this reconnection process is critical for replenishing the magnetic field at the β-type stagnation points and vital in the amplification of the magnetic field. For the non-oscillatory regime, this mechanism remains the same, though the field configuration is now static, with two cigars; one weaker than the other and both with opposing polarity.

A subsequent study, Archontis *et al.* (2003), investigates the magnetic field structure of a dynamo generated by the 1:1:1 ABC flow with $k_0 = 2$, in a similar fashion to Galanti *et al.* (1992). For $k_0 > 1$, the flow can be seen as multiple smaller "copies" of the original ($k_0 = 1$) ABC flow; for $k_0 = 2$, there are eight (2^3) such

copies and so the computational domain contains eight periodic cells. In the fastest growing eigenmode for $k_0 = 2$, the familiar cigar-like structures are seen but these vary in size and strength from cell to cell, which is to be expected as the magnetic field is able to evolve on scales larger than that of the individual cells. It is concluded that the ABC dynamo for $k_0 = 2$ is a fast dynamo, as the growth rate appears to saturate at $\gamma \approx 0.3$.

1.4.10 The most effective ABC dynamo

Through the simulations of Galloway and Frisch (1986), it is seen that the 1:1:1 ABC flow is not the most effective in the generation of magnetic field. Alexakis (2011) further investigates the effectiveness of dynamos in different ABC flows by parameterising A, B and C as spherical coordinates ϕ and ψ and restricting the parameter space through the flows' symmetries. The author's findings show that the effectiveness of the majority of flows varies greatly as R_m is modified, though it is clear that flows with $A = B \simeq 2C/5$ (and permutations), are consistently the most effective up to $R_m = 10^3$. Whether this remains true for $R_m = O(10^4)$ is yet to be seen.

In solving the same problem with the ABC flow defined as in (1.25) with $k_0 = 2$, the author reports that, at least for low and intermediate R_m, the most effective ABC flow has parameters $A = B = C$.

1.4.11 The 1:1:1 ABC dynamo in the high-R_m regime

A recent investigation of the 1:1:1 ABC flow dynamo problem by Bouya and Dormy (2013) extends the results of Galloway and Frisch (1986). The authors report results up to $R_m = 2.5 \times 10^4$, far exceeding the range of previous results achieved (see figure 1.6). The growth rate of the dominant dynamo mode appears to decrease in the range $R_m \in [400, 2000]$ but increases again for $R_m > 2000$. At $R_m = 2.5 \times 10^4$, the growth rate has reached $\gamma \approx 0.1$, which exceeds both the flux growth rate

$\Gamma_S \simeq 0.055$ and the maximum rate of line-stretching $h_{\text{line}} \simeq 0.09$. The former of these terms, Γ_S, is the growth rate of flux in the absence of diffusion (for some initial magnetic field) and Finn and Ott (1988a,b) speculate that Γ_S is equal to the dynamo growth rate in the limit of vanishing diffusion, that is $\Gamma_S = \gamma_0$ for $R_m \to \infty$: this is often referred to as the *flux conjecture*. The latter term, h_{line}, is the maximum possible rate of stretching that any (frozen-in) curve will experience as it is carried by the flow.

That the largest growth rate obtained by Bouya and Dormy exceeds $\Gamma_S \simeq 0.055$ is unsurprising, as this value was calculated for a field of the form (1.21) and so it is unlikely to be realistic limit for the second window. The growth rate exceeding the maximum rate of line-stretching, however, suggests either that the flux conjecture does not hold (for the 1:1:1 ABC flow) or that the values of R_m are not sufficient for the growth rate to have reached an asymptote. The authors' results seem to support the latter suggestion, as the growth rate appears to be increasing in the limit of $R_m = 2.5 \times 10^4$, though this highlights that the question on whether or not this is a fast dynamo, is not easily answered.

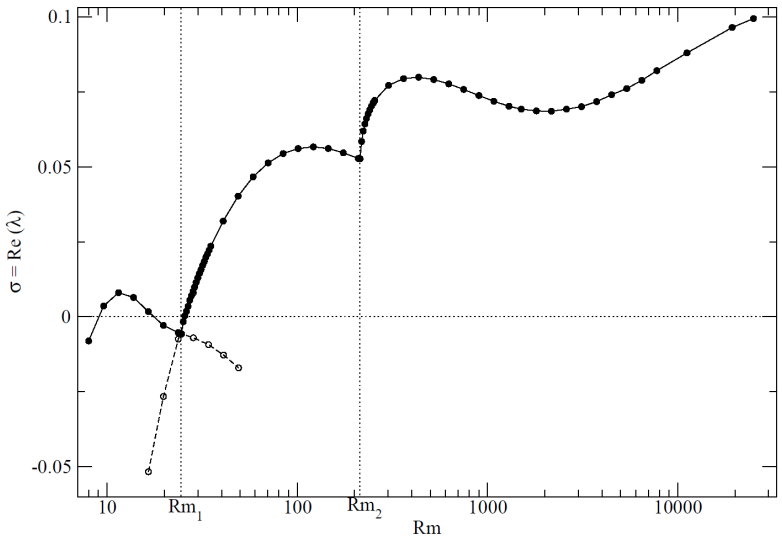

Figure 1.6: Growth rates of the 1:1:1 ABC dynamo for R_m up to 2.5×10^4, taken from Bouya and Dormy (2013).

The paper also follows, numerically, the growth rate of the seed field (1.21) for as high R_m as possible. The limiting factor when resolving this mode is the length

of time that the rounding error remains negligible. Upon reaching a sufficient magnitude (after sufficient time steps), the rounding error leads to the fastest growing mode emerging, meaning that the time evolution of the seed field is seen as transient behaviour. The authors are able to resolve this branch past the mode-crossing at $R_m \approx 24$ but growth rates are given for R_m up to 50, so they have been unable to verify the results of $\gamma \approx -0.03$ for $R_m = 100, 200$ and 400 (Galloway and O'Brian, 1994).

The difficulties encountered by Galloway and Frisch (1986), and later studies, in resolving the growth rate of the fastest growing dynamo mode for $R_m \approx 300$ have been overcome by Bouya and Dormy. As expected, the fastest growing eigenmode switches from an oscillatory mode to a steady mode and the authors have identified that this occurs for $R_m \approx 215$, far lower than the previously accepted range of $R_m \in [300, 350]$. On further analysis, the authors find that the apparent kink in growth rates is not due to another mode crossing (as is suspected by Lau and Finn, 1993), but is instead due to 'coalescence' of eigenvalues in the dominant eigenmode.

1.5 Thesis aims

Now that the context of this work has been described, the purpose of this thesis, and constituent investigations, can be detailed.

This thesis deals with kinematic dynamo action in the 1:1:1 ABC flow. Core to all components of this thesis are the 24 symmetries of the flow. As already discussed, the symmetries are well-known, but as of yet they have not been imposed in any investigations on dynamos or fluid stability. Arnold (1984) details how, for the kinematic problem, a general magnetic field (whose evolution is described by (1.5)), can be described as a linear combination of five irreducible representations. The principal aim of this thesis is to take this decomposition to reduce the full kinematic dynamo problem into five subproblems and solve these independently through time stepping. This has two purposes. Firstly it allows independent investigation of the five possible types of structure than can exist in a general field, two of which (it is

theorised) identify with each of the separate windows of dynamo action. Secondly, solving all five subproblems is time-saving in comparison to solving the full problem, allowing a greater range of R_m to be explored.

In the kinematic dynamo problem, the fluid flow $\boldsymbol{u}$ is prescribed; therefore in solving (1.13), we are solving an eigenproblem with a distinct set of eigenvalues for each R_m. In time stepping, magnetic energy evolution is used to identify the eigenvalue with the largest real part (i.e. the dominant eigenmode) from the growth rate and, if complex, the imaginary part from the oscillation frequency. A more direct method of solving these eigenproblems could be adopted, by applying an eigenvalue solver to the matrix of Fourier mode relations. As the symmetries are applicable here, the aim is to verify the results of time stepping within each representation, and compare the effectiveness of this method to time stepping.

As detailed in Galloway and Frisch (1987), the full linear fluid stability problem can be solved with only minor modifications, which involve projecting the pressure term out and can be incorporated into the numerical scheme. The symmetries can also be applied to this problem, as a result of the linearity in the perturbation term, leading to five independent subproblems; achieving this requires minimal modification to the symmetry-reduced time stepping codes of the magnetic problem. A secondary aim of this thesis is to extend the range of R_e investigated, while resolving all five subproblems.

Following the work of Galanti *et al.* (1992), an investigation into large-scale magnetic fields in the 1:1:1 flow will also be undertaken. Instead of increasing the periodicity of the ABC flow (by using (1.25)), the magnetic field will be rescaled through the introduction of Bloch waves. This is advantageous in that the scale of the magnetic field is not restricted to integer multiples of the scale of the flow, and can be adapted from the full (non large-scale) dynamo problem with relative ease. The symmetries, however, are dependent on the rescaling of the field and so the same decomposition (as in the magnetic and fluid problems) is not used. The aim is to resolve the optimum scale-separation for dynamo action by varying the scale of the magnetic field for as large a range of R_m as possible, with the intent

of exploring and understanding the structures of the resulting magnetic fields.

1.6 Objectives

Summarising the aims as specific objectives allows essential goals to be identified. Each main goal comprises of several smaller goals. The objectives are:

1. Numerically solve the dynamo problem for the five independent symmetry classes
   (a) Reproduce the results of Galloway and Frisch (1986), and identify the symmetries of the two windows of dynamo action
   (b) Resolve the five independent classes for as high R_m as possible
   (c) Parallelise the codes and compare efficiency against serial codes
   (d) Investigate differences in structure between each of the five classes
2. Solve the kinematic ABC dynamo problem using eigenvalue solver
   (a) Reproduce results of time stepping codes for each of the five classes
   (b) Compare accuracy and efficiency of solving through time stepping to eigenvalue solver
   (c) If possible, investigate subdominant branches of solutions in each class
3. Investigate the linear hydrodynamic stability problem by studying the five symmetry classes
   (a) Reproduce the results of Galloway and Frisch (1987) for the 1:1:1 ABC flow
   (b) Solve the hydrodynamic stability problem for highest possible values of R_e
   (c) Investigate the types of structure: compare between each class and compare to magnetic (dynamo) problem

4. Study the effect of large-scale magnetic fields through the use of imposed Bloch waves

   (a) Contrast Bloch wave formulation against Galanti *et al.* (1992) and determine α-effect for this method

   (b) Investigate symmetries and representations of rescaled magnetic fields and reduce system accordingly

   (c) Numerically resolve dynamo over numerous scales for as large a range of R_m as is possible

   (d) Identify most effective dynamos and their corresponding field structures

1.7 Structure of thesis

To summarise, this thesis is divided into eight chapters, of which this introduction is the first. Chapter 2 deals with the symmetries of the problem, with chapter 3 detailing the irreducible representations relevant to these symmetries. Chapter 4 concerns the numerical methods used in this thesis, including exploration of both time stepping and eigenvalue solver routines. Results for the ABC dynamo simulations are given in chapter 5. Chapters 6 and 7 are self-contained and detail the fluid stability and large-scale (magnetic) field problems respectively, with numerical results. The concluding discussion is provided in chapter 8.

The next chapter gives the symmetries of the 1:1:1 ABC flow. First, some background information is given on the three main types of symmetry, with formal definitions of how symmetries apply to vector and scalar fields. Following this, the symmetries of the group D_8 are given and described. The robust process for constructing the symmetries of the ABC flows, by testing permutations of variables, is laid out, with a few specific examples given to identify symmetries and non-symmetries. Finally a brief discussion of the symmetries of other ABC flows is given for completeness.

The subject of chapter 3 is Representation theory and its application to this prob-

lem. The content of this chapter was intentionally kept separate from the symmetries in an effort to avoid presenting too much information at once. With the symmetries established, this chapter first explores irreducible representations in the very simplest case, that of $\mathbb{Z}_2$, to understand the basics of matrix representations. The group D_8, the symmetries of a square, is then explored with a clear demonstration of how the irreducible representations of its five symmetry classes are constructed through knowledge of its conjugacy classes and character tables. This is then applied to the full group G of 24 symmetries of the 1:1:1 ABC flow, using the fact that this group is isomorphic to O_{24}, the group of proper rotations of a cube, and the relevant representation matrices are provided for all five symmetry classes.

Chapter 4 explores the numerical methods adopted, first without any symmetry considerations. The symmetric problem is then depicted within fundamental domains $\mathcal{G}_N$ and $\mathcal{H}_N$, the latter including Hermitian symmetry. Parallelisation of the time stepping problem in $\mathcal{H}_N$ is sketched and results of scaling performance are presented. The method for solving the eigenvalue problem in $\mathcal{G}_N$ with the Implicitly Restarted Arnoldi Method (IRAM) is elaborated, with technical details on reducing the problem to a minimal state. Convergence conditions are also investigated with some results on optimal parameters. The IRAM is compared with the equivalent time stepping problem in $\mathcal{G}_N$, with advantages and disadvantages of both methods discussed.

Results of the 1:1:1 ABC dynamo are given in chapter 5. The five distinct branches of solutions corresponding to each symmetry class are resolved, with growth rates and structures (both oscillating and steady) explored. Differences between the classes are identified and results are placed in the context of previous investigations and literature. Eigenvalue coalescences and separations are investigated, along with mode crossings within the classes. Magnetic energy spectra are also analysed.

In chapter 6, the equivalent fluid stability problem is studied. The chapter begins with the equivalent formulation of the problem and follows with numerical methods to be used. The results for the five symmetry classes are given, with exploration

of the resulting structures. Comparison is drawn with the magnetic problem and a brief discussion of the results is carried out.

Chapter 7 deals with the problem of large-scale magnetic fields in the 1:1:1 ABC flow. The problem is constructed using Bloch wave formulation and this is then related to both the α-effect and the alternative formulation of Galanti *et al.* (1992), with a detailed comparison to the latter. The numerical breakdown of this problem is then detailed, with symmetries and four symmetry classes identified for a particular case; a magnetic field stretched in only one spatial direction. Results of numerical calculations for the full problem (i.e. without symmetry considerations) are given and the optimal scales for dynamo action are found for the range of R_m investigated. The structures of these fields at the optimal scales are also studied and compared to those known for the non-large scale problem.

The final chapter, 8, brings the thesis to its conclusion. This short chapter provides an overview of the results and conclusions of chapters 5, 6 and 7 with the aim of understanding these results in the wider context of the problems. In addition, the main difficulties faced in both time stepping and use of the IRAM are explained, along with each method's shortcomings. Finally, the possible future application of methods developed in this thesis are set out, with the purpose of identifying potentially interesting investigations.

Chapter 2

Symmetries

Wisdom is one of the few things that looks bigger the further away it is.

Terry Pratchett

The aim of this chapter is to provide a basic introduction to the different types of symmetries and how they are relevant to the problems explored in this thesis. Firstly, the three main types of isometry are briefly explored. The isometries, these being reflection, rotation and translation are, by themselves or composed with one another, the basis for symmetries in most systems. Once these are established, clear and intuitive examples are given of geometric constructs (equilateral triangle, square and triangular grid) in which these symmetries are plainly visible. Helical symmetry, being a composition of rotation and translation, is also discussed. The discussion then moves on to how symmetries are formally defined for both scalar and vector functions.

In the next chapter, prior to explaining how symmetries apply to a sufficiently general magnetic field driven by the ABC flow (1.18), a simpler and more easily visualised example needs to be explored. For this purpose, we construct the symmetry group D_8, the symmetries of a square, in this chapter. This allows us to

show how to produce a group of symmetries from their generators and identify each symmetry's physical meaning in relation to the fundamental domain. Next, the symmetries of the ABC flows are discussed, with an explanation on the robust method of testing all possible permutations of variables composed with arbitrary translations, owing to the helical nature of the flow. The methods identify, using a specific example, how orientation-reversing symmetries are not valid symmetries of the flows and show that, dependent on the values of A, B and C, a given permutation of variables can supply a symmetry, given an appropriate translation. The symmetries of the case $A = B = C = 1$ are listed and related to the 24 orientation-preserving symmetries of the cube. The symmetry groups of more general cases are shown to be subsets of the group for $A = B = C = 1$ and the group for each of these other three cases is also listed. The chapter ends with a summary of the main points discussed.

2.1 Definition of symmetries

As explored in the first chapter, a number of sources refer to, and in some cases list, the various symmetries of the ABC flows. Although these symmetries are well-known, it is worth spending some time understanding the concept of symmetries and how they are generated.

One starting point is the definition of the word itself. The Oxford English Dictionary (2013) defines the word symmetry as

> Mutual relation of the parts of something in respect of magnitude and position; relative measurement and arrangement of parts; proportion.

In a scientific context, the word takes on a more precise meaning:

> [A] property by virtue of which something is effectively unchanged by a particular operation; an operation or set of operations that leaves something effectively unchanged; in Physics, a property that is conserved.

Although the word symmetry itself has a few meanings, ranging in technicality, they are based around the same concept; a symmetric object is one in which some property is preserved under a given transformation. We paraphrase the formal mathematical definition from Johnson (2001): Given a metric space X and a subset of X, referred to by F, a symmetry of F is then defined as an isometry of X that fixes F as a set. There may be multiple isometries that meet this requirement and so the set of all symmetries of the subset F can be defined as:

$$\mathrm{Sym}(F) = \{u \in \mathrm{Isom}(X) | Fu = F\}, \tag{2.1}$$

where u is given to be an isometry of X. That is, the symmetries of a particular subset of points are the set of distance-preserving maps (also known as isometries), that when acting on the particular subset F, returns the subset F. This does not imply that the symmetries fix every element of F but does mean that F and Fu are equal as (sub)sets.

In Geometry, distance is a well-defined and easily understood concept and so allows the concept of symmetries to be easily visualised and interpreted. As an example, in n-dimensional Euclidean space, $\mathbb{R}^n$, (Euclidean) distance is simply defined as the shortest distance between two points $\boldsymbol{x}$ and $\boldsymbol{y}$:

$$d = \left(\sum_{i=1}^{n} |x_i - y_i|^2 \right)^{1/2}. \tag{2.2}$$

The next section will discuss the main types of isometries (distance-preserving maps) that form a basis for most types of symmetry.

2.2 Types of symmetry

2.2.1 Reflection

There are three main types of isometries in geometry. Perhaps the most obvious of these is reflection operator, s, this being for us the most intuitive of the isometry types. A reflection is a map that in an n-dimensional space, moves every point to its mirror-image in a fixed $(n-1)$-dimensional hyperplane. Mathematically, this could be written as

$$\mathsf{s} : \mathbb{R}^n \to \mathbb{R}^n, (x_1, x_2, \ldots, x_n) \mapsto (-x_1, x_2, \ldots, x_n)\,, \tag{2.3}$$

for a plane described by $x_1 = 0$. Although the reflection hyperplanes may differ in definition, a coordinate transformation can always be carried out so that a reflection of this type can always be written as in (2.3) (by defining x_1 as perpendicular to the hyperplane).

A reflection is not, however, restricted to mapping points to their mirror-images through a hyperplane of dimension 1 less than the space that we are operating in. For example, reflections are possible through k-dimensional hyperplanes for $k = 0, \ldots, n-1$ in n-dimensional space. Such a transformation would reverse the sign of $j = n - k$ elements of $\boldsymbol{x}$:

$$\mathsf{s} : \mathbb{R}^n \to \mathbb{R}^n, (x_1, x_2, \ldots, x_n) \mapsto (-x_1, \ldots, -x_j, x_{j+1}, \ldots, x_n)\,. \tag{2.4}$$

Reflection through a point ($k = 0, j = n$) is a special case of reflection and is aptly named point reflection. Mathematically it is written as:

$$\mathsf{s} : \mathbb{R}^n \to \mathbb{R}^n, (x_1, x_2, \ldots, x_n) \mapsto (-x_1, -x_2, \ldots, -x_n)\,, \tag{2.5}$$

and can be thought of as a map of every point to its diametrically opposite counterpart through the point of reflection (the origin in this case). In the case of generalised reflections, they can be separated into two categories; those that pre-

serve the orientation of the system and those that do not - we are typically more accustomed to the latter, an example of which would be mirror-symmetry. In relation to (2.4), orientation-preserving reflections are those for which j is even, whereas orientation-reversing reflections require j to be odd. As will be explained, orientation-preserving reflections can always be described through rotations, rendering these reflections redundant. Thus, unless explicitly stated, all references to reflections should be taken as orientation-reversing reflections.

An important feature of the reflection operator is that it is an involution meaning that $\mathsf{s}^2 = \mathsf{i}$, or alternatively $\mathsf{s}^{-1} = \mathsf{s}$.

2.2.2 Rotation

The second type of isometry is rotation, indicated by the symbol r. A rotation is a map that moves every point of a two-dimensional space through a fixed angle about an axis through a fixed point, called the centre of rotation. Similarly, for three-dimensional Euclidean space, a rotation moves every point through a fixed angle about a fixed axis. The centre is often fixed at the origin and likewise, the axis of rotation typically passes through the origin. In k-dimensional Euclidean space, rotations through a fixed angle α about the origin are described by

$$\mathsf{r} : \mathbb{R}^2 \to \mathbb{R}^2, (x, y) \mapsto (x\cos\alpha - y\sin\alpha, x\sin\alpha + y\cos\alpha)\,, \text{ for } k = 2, \quad (2.6)$$

$$\mathsf{r} : \mathbb{R}^3 \to \mathbb{R}^3, (x, y, z) \mapsto (x\cos\alpha - y\sin\alpha, x\sin\alpha + y\cos\alpha, z)\,, \text{ for } k = 3, \quad (2.7)$$

in Cartesian coordinates. Note that a rotation in $\mathbb{R}^3$ involves only one angle, as such a rotation acts as an ordinary two-dimensional rotation in the plane orthogonal to the rotation axis. If the axis of rotation is not defined as one of the coordinate axes, a coordinate transformation can always be carried out so that a rotation can be defined as above. Bear in mind that, in the above definition, for $\alpha > 0$ rotations are anticlockwise, whereas for $\alpha < 0$ rotations are clockwise. Consequently, unless explicitly stated, any rotations defined or used will be anticlockwise.

We can write a rotation as $\mathsf{r} = \mathsf{r}(\boldsymbol{n}, \alpha)$, where $\boldsymbol{n} = \mathbf{0}$ for rotations in $\mathbb{R}^2$ or is the unit vector in the direction of the axis of rotation in $\mathbb{R}^3$. In this thesis, angles will be measured in radians and so angles should be assumed to be in this unit unless stated otherwise. We therefore have that $-\pi < \alpha \leq \pi$, as angles are defined $\mathrm{mod} 2\pi$.

For two rotations $\mathsf{r}_1 = \mathsf{r}(\boldsymbol{n}, \alpha)$ and $\mathsf{r}_2 = \mathsf{r}(\boldsymbol{n}, \beta)$ with the same centre (or axis) of rotation, their composition is simply $\mathsf{r}_1\mathsf{r}_2 = \mathsf{r}_2\mathsf{r}_1 = \mathsf{r}(\boldsymbol{n}, \alpha + \beta)$. If their centres (or axes) of rotation are different, the composition of r_1 and r_2 creates a unique rotation, which itself can be described as a single rotation through a fixed angle around its own centre or axis of rotation. Because angles are defined $\mathrm{mod} 2\pi$ and to compose two rotations, we simply sum the angles, then it can be seen that rotations can be cyclic, with the property $\mathsf{r}^p = \mathsf{i}$, which implies that $p\alpha = 0 \bmod 2\pi$ for some p. Rotations are not implicitly cyclic however, as if the fraction $\frac{\alpha}{2\pi}$ is irrational, then no integer p exists such that $\mathsf{r}^p = \mathsf{i}$.

For the special case of $\alpha = \pi$, the rotation is also an involution, since

$$\mathsf{r}^2 = \mathsf{r}(\boldsymbol{n}, 2\pi) = \mathsf{r}(\boldsymbol{n}, 0) = \mathsf{i}, \tag{2.8}$$

or in other words, applying this rotation twice returns all points to their original position. In our definition of rotations, we define *proper* (orientation-preserving) rotations. Improper (orientation-reversing) rotations are not considered in the subsequent sections and all references to rotations are proper rotations unless specifically stated.

2.2.3 Translation

The third type of isometry is called translation. A translation, t, is a map that moves every point of n-dimensional Euclidean space through a fixed distance in a fixed direction and can be written as

$$\mathsf{t} : \mathbb{R}^n \to \mathbb{R}^n, (x_1, x_2, \ldots, x_n) \mapsto (x_1 + a_1, x_2 + a_2, \ldots, x_n + a_n), \tag{2.9}$$

where $(a_1, a_2, \ldots, a_n) = \boldsymbol{a}$ is a constant vector. Translations can be written as $\mathsf{t} = \mathsf{t}(\boldsymbol{a})$ and the direction of $\boldsymbol{a}$ is known as the axis of t.

Translations are easily composable, as translating all points by $\boldsymbol{a}$, then by $\boldsymbol{b}$ is the same as translating all points by $\boldsymbol{a} + \boldsymbol{b}$, or

$$\mathsf{t}(\boldsymbol{a})\mathsf{t}(\boldsymbol{b}) = \mathsf{t}(\boldsymbol{a} + \boldsymbol{b}). \tag{2.10}$$

More general isometries are composed of different types of isometries. For example, a reflection and then a rotation or a rotation followed by a translation. One composition particularly relevant to this thesis is that of translation composed with rotation, which is required when working with helical structures such as those of the ABC flows.

2.3 Symmetries in Geometry

Symmetries are best understood through the use of geometry. Triangles and squares are examples of two-dimensional shapes that contain both reflectional and rotational symmetry. In geometry, the particular invariant properties are the locations

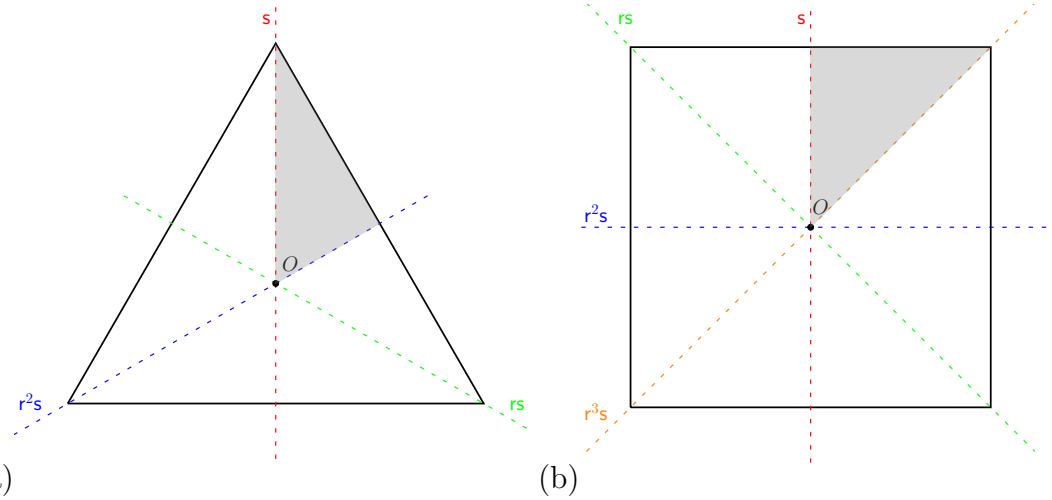

(a) (b)

Figure 2.1: The axes of symmetry of (a) an equilateral triangle and (b) a square. The grey areas mark one possible choice for the fundamental domain.

of vertices and equations of edges. When transforming the coordinates of the ver-

tices and the equations of lines using an isometry, that isometry is a symmetry if the resulting description of the transformed shape, by means of its vertex coordinates and edge equations, is the same as the original description. In a geometric construct, this means that an isometry is a symmetry if an object is visually identical before and after the isometry is applied.

In figure 2.1 are visualisations of the two examples with their axes of reflection marked and centres of rotation located at the origin for ease. The triangle (figure 2.1(a)) has three-fold rotational symmetry, that is, $\mathsf{r} = \mathsf{r}(\mathbf{0}, 2\pi/3)$ and so $\mathsf{r}^3 = \mathsf{i}$. The triangle also has reflectional symmetry about three axes, though a shape is often described by its unique axes of reflection. For example, a reflection in the blue axis can be described by a reflection in the red axis and a $4\pi/3$ rotation about O; only one axis of reflection is therefore required.

The square has four axes of reflection though only one of them is considered unique, owing to the square's four-fold rotational symmetry whose transformation is given by $\mathsf{r} = \mathsf{r}(\mathbf{0}, \pi/2)$. The choice of unique reflection axis is arbitrary. In this case the line $x = 0$ has been set and is identified by the red dashed line. Reflection through the other three (coloured) axes can instead be described in terms of the symmetries rs (green), $\mathsf{r}^2\mathsf{s}$ (blue) and $\mathsf{r}^3\mathsf{s}$ (orange), that is a reflection using s (through the red axis) followed by a $\pi/2$, π or $3\pi/2$ rotation respectively. Combining the rotation and reflection operators means that only one reflection axis is required to construct the whole square from its fundamental domain.

For an n-sided polygon, there is rotational symmetry of order n with $\mathsf{r} = \mathsf{r}(\mathbf{0}, 2\pi/n)$, $\mathsf{r}^n = \mathsf{i}$ and of course a single unique reflection s with an axis intersecting opposing sides, for even n, or intersecting a face and its opposing vertex, for odd n. For finite n, these rotations are referred to as discrete rotations. In the limiting case $n \to \infty$, the regular polygon becomes a circle, though the circle cannot be considered a polygon due to their respective definitions. In this case, the rotations are considered continuous, as opposed to discrete, with its symmetry group, the two-dimensional special orthogonal group, denoted by SO_2, being an example of a continuous group. This idea also extends to three dimensions with the sphere having (continuous)

symmetry group SO_3.

Examples of translational symmetry in geometry require different geometric constructions, as translations move all points and do not have fixed points or axes. Figure 2.2 visualises two triangular grid constructs that could be considered to

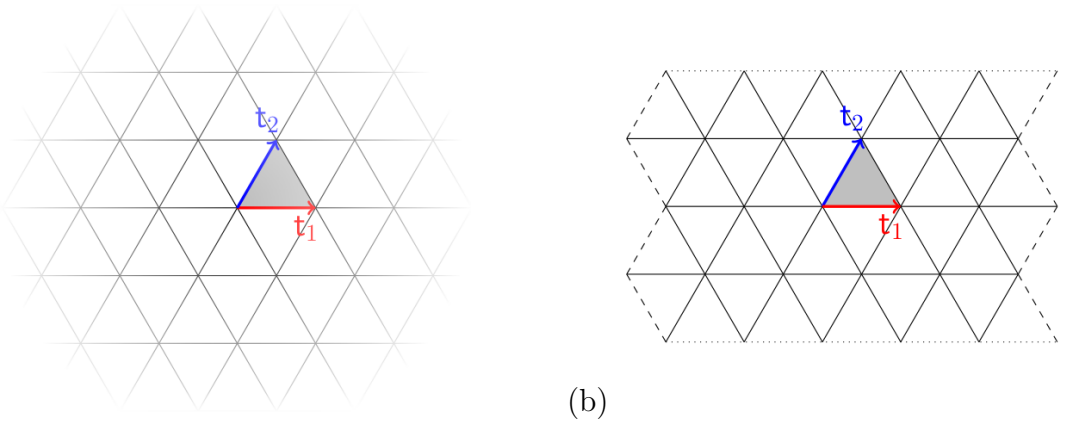

(a) (b)

Figure 2.2: Constructs containing translational symmetry with (a) infinite and (b) periodic domains.

have translational symmetry. In the first case, we have a triangular tiling pattern that we can consider infinite in all directions. The second case is that of the same pattern filling a periodic finite domain; the dotted and dashed lines respectively show connected edges. In the infinite domain there are, in theory, infinitely many symmetries. In the periodic domain, however, we have that $\mathsf{t}_1^p = \mathsf{t}_2^q = \mathsf{i}$ for some p and q. In either case, there are only two unique translations t_1 and t_2; for an n-dimensional domain, there are n unique translations $\mathsf{t}_1, \ldots, \mathsf{t}_n$, with each defined in the direction of the corresponding spatial axis. Rotation and reflection symmetries are also seen in these examples, particularly in the infinite domain.

In some situations, it is the case that certain transformations are not symmetries of the system but their compositions are. One example that is highly relevant to this thesis is helical symmetry. Figure 2.3 identifies a blue helical curve, C_1, and its transformed counterpart, C_2 (in green), with two transformations also indicated: a rotation r about the x-axis and a translation t along the x-axis in the direction of positive x. Although it is easily seen that C_2 is a translation of C_1, it is not so clear that C_2 is also a rotated form of C_1. The position at $x = 0$ of each curve is marked

by a circle of the respective colour to aid in showing this relation by rotation. In

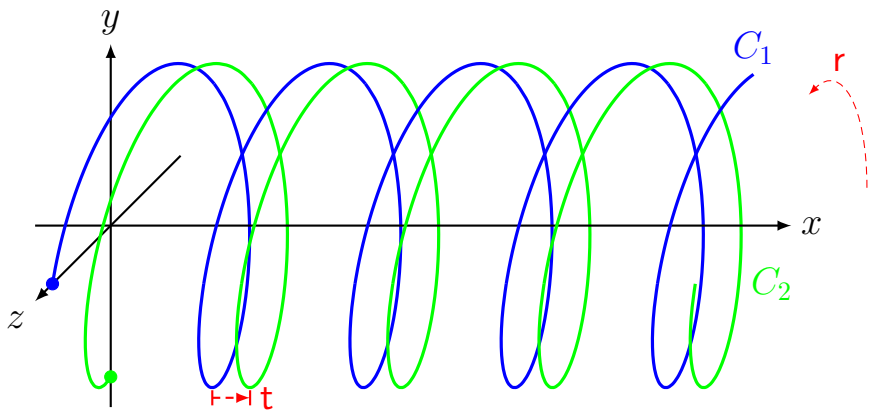

Figure 2.3: Circular helices C_1 and C_2 with relevant isometries r and t. C_2 can be considered the product of C_1 and either r or t.

the figure, C_2 represents C_1 after a rotation by $\pi/2$ about the x-axis but can also be described as a translation by $x = T/4$, where T is the period of the helical curve. It is plain to see that a rotation r by an angle $\alpha > 0$, followed by a translation by $a < 0$ will be a symmetry of this curve if $2\pi/\alpha = -T/a$. As the curves in this example are circular helices they will have continuous (as opposed to discrete) symmetries owing to the composed rotation-translation isometries being continuous by nature. We can note that if this structure is infinite in the x-direction, more symmetries arise as rotations (still defined mod2π) can be composed with an infinite number of translations, each translating first by a and then by an integer multiple of T. Under the assumption that these curves are periodic or even infinite, then a translation by $x = T$ will be equivalent to the action of the identity i.

When discussing symmetries, the fundamental domain is also an important concept. In a geometric sense, the fundamental domain is the smallest piece of a geometric object that is required to be able to reconstruct the object using the symmetries. In figure 2.1, the fundamental domains of the triangle and square are shown by the shaded areas. As any one of the triangle's six or square's eight sections can be chosen, the definition typically restricts the fundamental domain to be a section for which $x, y \geq 0$, or in a more general case, $x_i > 0$ for all i.

In the case of the geometric pattern used to discuss translations, the fundamental

domain for all three cases is a single tile, non-unique examples of which are shaded in Figure 2.2. Of course, in these cases we are neglecting rotations and reflections which certainly exist in all three examples. With these extra symmetries, the fundamental domain could potentially be reduced to an area resembling that in the triangle example of figure 2.1(a).

2.4 Symmetric Functions

In the same way that geometric objects or spaces can be symmetric, so can functions. As an immediate example, the function $\phi(x, y) = x^2 + y^2$ is symmetric about the line $y = x$. That is, if we switch x and y, the function remains unchanged. This is one of the multiple axes of reflectional symmetry inherent in this function.

One way that this reflection could be formally described is how it acts on the function's domain

$$\mathsf{s} : \mathbb{R}^2 \to \mathbb{R}^2, (x, y) \mapsto (y, x), \tag{2.11}$$

but the problem here is that we could easily see that the function ϕ is invariant under this transformation, yet for many functions it may not always be so obvious as to what symmetries are present. We can, instead, define symmetries by how they act on the functions themselves.

2.4.1 Scalar functions

Take ϕ to be a scalar function that maps $\mathbb{R}^n$ onto $\mathbb{R}$. We can define g as a general symmetry that maps $\mathbb{R}^n$ onto $\mathbb{R}^n$, as we have done previously. We can then say that g induces a map on functions such as ϕ, which we can define as

$$\mathsf{g}^* : S \to S,\ (\mathsf{g}^*\phi)(\boldsymbol{x}) = \phi(\mathsf{g}^{-1}\boldsymbol{x}), \tag{2.12}$$

where S is simply the set containing all such functions, ϕ

$$S = \{\phi : \mathbb{R}^n \to \mathbb{R}\}. \tag{2.13}$$

To simplify, we can define $\psi = \mathsf{g}^*\phi$ to get

$$\psi(\boldsymbol{x}) = \phi(\mathsf{g}^{-1}\boldsymbol{x}). \tag{2.14}$$

Referring back to the previous example with $\phi(x, y) = x^2 + y^2$, we can define ψ by

$$\psi(x, y) = \phi(\mathsf{s}^{-1}(x, y)) = \phi(y, x), \tag{2.15}$$

where ψ is the function ϕ reflected through the axis $y = x$.

Of course, g is a symmetry of the function ϕ if

$$\psi(\boldsymbol{x}) = \phi(\mathsf{g}^{-1}\boldsymbol{x}) = \phi(\boldsymbol{x}), \tag{2.16}$$

which is equivalent to saying

$$\phi(\boldsymbol{x}) = \phi(\mathsf{g}\boldsymbol{x}), \tag{2.17}$$

identifying an important feature of symmetries: if a transformation g is a symmetry of a function, the transformation g^{-1} is also a symmetry of that function.

2.4.2 Vector functions

For scalar functions, we have a single map from $\mathbb{R}^n$ to $\mathbb{R}$, that is, we have a single value as an output of such a function. With a vector function, however, we have a vector of scalar functions meaning that we have a map from $\mathbb{R}^n$ to $\mathbb{R}^n$. The function is not assigning values to individual points but vectors, the dimension of these corresponding to the number of variables in the space (n). We define a

function $\boldsymbol{f}$ by

$$\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^n, \ \boldsymbol{f}(\boldsymbol{x}) = (f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), \ldots, f_n(\boldsymbol{x})), \tag{2.18}$$

so we can think of $\boldsymbol{f}$ as a vector (of length n) of scalar functions $(f_1, \ldots, f_n)$. We are operating in a space of dimension n, so our symmetry g will be defined as in the scalar function case. When we use g to induce a map on such functions $\boldsymbol{f}$, we now have

$$\mathsf{g}^* : S \to S, \ (\mathsf{g}^*\boldsymbol{f})(\boldsymbol{x}) = \boldsymbol{J}_{\mathsf{g}}\boldsymbol{f}(\mathsf{g}^{-1}\boldsymbol{x}), \tag{2.19}$$

where S is now the set of all vector functions $\boldsymbol{f}$ and $\boldsymbol{J}_{\mathsf{g}}$ is the Jacobian matrix of g, defined by

$$\boldsymbol{J}_{\mathsf{g}} = \begin{pmatrix} \frac{\partial \mathsf{g}^{(1)}}{\partial x_1} & \cdots & \frac{\partial \mathsf{g}^{(1)}}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathsf{g}^{(n)}}{\partial x_1} & \cdots & \frac{\partial \mathsf{g}^{(n)}}{\partial x_n} \end{pmatrix}, \tag{2.20}$$

with $\mathsf{g} = \left(\mathsf{g}^{(1)}, \ldots, \mathsf{g}^{(n)}\right)$ a function of $\boldsymbol{x} = (x_1, \ldots, x_n)$. As with the scalar function, we can simply define $\mathsf{g}^*\boldsymbol{f}$ as another function $\boldsymbol{F}$ giving us a distinct function that is defined as a rotated form of our original function evaluated at the preimage of the point $\boldsymbol{x}$.

If $\boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x})$, then g (and thus g^{-1}) is a symmetry of the function $\boldsymbol{f}$. By finding all such symmetries g, the group of all symmetries (G) of the function can be constructed. It is often the case that there are only a few truly independent symmetries of a system, with all other transformations being compositions of these. The independent symmetries are referred to as *generators*, as they are used to generate the group of symmetries of a system.

2.5 Symmetries of the square

We will first explore the symmetries of the square, prior to those for the ABC flows, as this example will be used to introduce the ideas and help develop the methods for constructing the irreducible group representations in the next chapter. The

symmetries of a square are easily visualised due to the 2-dimensional geometry, providing a more comprehensible understanding of the irreducible representations and their physical meanings. Furthermore, the group of symmetries of the square, D_8, is well-known and easily appreciated, even by those with little background in group theory.

Referring to figure 2.1, there are four possible reflection axes, though for the purpose of this example the axis we choose is horizontal, passing through the square's centre at the origin O (see figure 2.4). Formally, the reflection operator is defined as

$$\mathsf{s}(x, y) = (x, -y), \tag{2.21}$$

with the assumption that the sides of the square are parallel to the x and y axes. The rotation transformation can be clockwise or anticlockwise, but in keeping with the definitions in section 2.2.2, the rotation in this system will be defined as anticlockwise. This choice isn't necessarily important, as the groups of symmetries constructed through either choice will be isomorphic to one another, as with the axes of reflection. The $\pi/2$ rotation about O is given as

$$\mathsf{r}(x, y) = (-y, x). \tag{2.22}$$

Through compositions of these two symmetries, it is possible to construct the entire group of symmetries for a square. Firstly, we can produce all rotational symmetries by repeatedly rotating the square, as if r is a symmetry, then r^j must also be, though we have that $\mathsf{r}^4 = \mathsf{i}$, leaving r, r^2 and r^3 as valid symmetries. Of course, $\mathsf{r}^0(x, y) = (x, y)$ is unique but is already described by the identity transformation i, which exists as the trivial symmetry in every system. Secondly, by composing each of the rotational symmetries with a reflection we produce a set of rotoreflections. For example, applying a reflection $\mathsf{s}(x, y) = (x, -y)$ and then a single rotation $\mathsf{r}(x, y) = (-y, x)$ will give us $\mathsf{r}(\mathsf{s}(x, y)) = (y, x)$, which can also be seen as a reflection through the axis $x = y$. We will shorten $\mathsf{r}(\mathsf{s})$ to rs for brevity. The full group of eight symmetries is listed in (2.23). Due to the size of the group D_8, it is more

convenient to represent symmetries as compositions of their generators.

$$\begin{aligned} \mathsf{i}(x,y) &= (x,y), & \mathsf{s}(x,y) &= (x,-y), \\ \mathsf{r}(x,y) &= (-y,x), & \mathsf{rs}(x,y) &= (y,x), \\ \mathsf{r}^2(x,y) &= (-x,-y), & \mathsf{r}^2\mathsf{s}(x,y) &= (-x,y), \\ \mathsf{r}^3(x,y) &= (y,-x), & \mathsf{r}^3\mathsf{s}(x,y) &= (-y,-x), \end{aligned} \tag{2.23}$$

Since we have acknowledged that if an isometry is a symmetry, then the isometry's inverse is also a symmetry, then it needs to be pointed out that each symmetry in (2.23) also has its inverse listed. In figure 2.4, each section has been assigned a number so that the equivalence of symmetries can be demonstrated. If we "start"

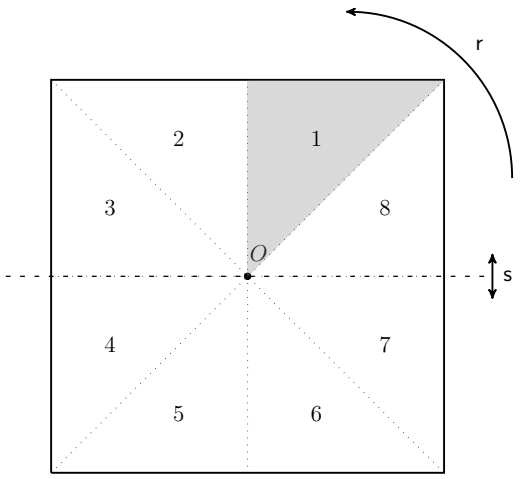

Figure 2.4: Diagram depicting the geometry of a square with symmetry group D_8.

in position 1, the section highlighted in grey, an application of $\mathsf{r}^3\mathsf{s}$ will first reflect through the horizontal, taking us to position 6 and then three rotations by $\pi/2$ will leave us at position 4. To find the inverse of this symmetry, we simply have to identify which in (2.23) will return us to position 1 and it is seen that $\mathsf{r}^3\mathsf{s}$ will indeed achieve this. This means that $\mathsf{r}^3\mathsf{s}$ is an involution, and the same can also be shown for rs and $\mathsf{r}^2\mathsf{s}$. The inverse of r is clearly r^3 as composition of these leads to a full revolution. The transformation r^2 is plainly an involution and so the group of symmetries also contains all inverses, confirming that we indeed have the full group

of symmetries. Any composition of rotations and reflections not listed in (2.23) can always be reduced to one of those that is listed. Due to the lack of commutativity, we have that

$$\mathsf{rs} = \mathsf{s}^{-1}\mathsf{r}^{-1} = \mathsf{sr}^3, \tag{2.24}$$

which implies that regardless of the order that the isometries are composed, they may always be reduced to a form in which the reflection is the first isometry applied, as displayed in (2.23). The result of a system containing D_8 symmetry is that the fundamental domain is 1/8 of the full domain; one possible choice is shaded grey in figure 2.4. Using only the fundamental domain, the full domain can be reconstructed using the seven non-trivial symmetries.

The next chapter will be exploring how a sufficiently general function, which does not contain the same symmetries as the square, can be decomposed into independent functions, each behaving differently but consistently under the symmetries.

2.6 Symmetries of the ABC flows

As discussed in chapter 1, the ABC flows are known to have varying degrees of symmetry, dependent on the choices for the parameters A, B and C. This section concerns constructing the group of symmetries G for the 1:1:1 ABC flow and then discussing which transformations are valid for more general ABC flows. Dombre *et al.* (1986) investigates the different parameter regimes and reduces the different ABC flows to four cases, these being

1. $A = B = C$,
2. $A = B \neq C$,
3. $A \neq B = C$,
4. $A \neq B \neq C \neq A$,

with $1 = A \geq B \geq C \geq 0$, without loss of generality. (Cases 2 and 3 are distinct cases because of the restriction $A \geq B \geq C$.) Symmetries for special cases such as $A = B = 1, C = 0$ and $A = 1, B = C = 0$ are covered by the more general regimes. Although the symmetries of the ABC flows have been known for many years (Arnold, 1984; Gilbert, 1992; Podvigina and Pouquet, 1994), it is important to demonstrate how the group is actually constructed.

2.6.1 Constructing the symmetries of the ABC flows

To construct the symmetries of the ABC flows, a good place to start is with all possible (signed) permutations of the variables, due to the symmetries being bijective maps. We have three variables that are each 2π-periodic, and thus our domain is the cube $[0, 2\pi]^3$. Identifying these symmetries is therefore analogous to constructing the symmetries of a cube, in which axes are mapped to axes through rotations, forming the Octahedral group O_h. There are 48 possible permutations, due to the fact that $x \mapsto -x$ is considered a separate transformation to $x \mapsto x$: there are six choices for the first variable, four choices for the second and two choices for the third. This accounts for the fact that say, once x has been chosen for a transformation, we can not choose x or $-x$ again for that same transformation. Due to the helical nature of the flow, it is assumed that transformations also involve translations, leading to a group of transformations that are of the form:

$$\mathsf{g}_1(x, y, z) = (-x + \alpha_1, y + \beta_1, z + \gamma_1), \tag{2.25}$$

$$\mathsf{g}_2(x, y, z) = (-x + \alpha_2, -y + \beta_2, z + \gamma_2), \tag{2.26}$$

$$\vdots$$

$$\mathsf{g}_{48}(x, y, z) = (-z + \alpha_{48}, -y + \beta_{48}, -x + \gamma_{48}). \tag{2.27}$$

To test whether g_1 through to g_{48} are symmetries of (1.18), it is a matter applying them to the ABC flows as described in (2.19). For the ABC flow and our 48

potential transformations, this becomes

$$\mathsf{g}_j \boldsymbol{u}(x,y,z) = \boldsymbol{J}_{\mathsf{g}_j} \boldsymbol{u}(\mathsf{g}_j^{-1}(x,y,z)). \tag{2.28}$$

Due to the generality of the translations in the transformations, it cannot be assumed that for each g_j there is a unique choice of parameters α_j, β_j and γ_j that will solve (2.28). By testing each transformation and attempting to solve for $(\alpha_j, \beta_j, \gamma_j)$ and also (A, B, C), we can confirm whether they are symmetries of the flows and also how many solutions, of the given form, exist. The ABC flows are naturally 2π-periodic, defined in $\mathbb{T}^3$, due to their definition through trigonometric functions, sine and cosine. This means that translations in the helical symmetries will also be defined modulo 2π and thus will be defined over the interval $(-\pi, \pi]$ to standardise the symmetries.

At this point, we should note that transformations involving time-reversal also exist. Transformations involving both time-reversal and orientation-reversing reflections are actually valid symmetries of the ABC flows; these are discussed by Dombre *et al.* (1986). However, we do not consider these symmetries, as magnetic diffusion is not a reversible process and thus these symmetries would be invalid in the context of evolving magnetic fields described by the induction equation (1.13). We therefore only consider transformations of the three spatial variables x, y and z.

It is possible to test all 48 transformations, but the scale of the problem can be reduced if we take into account the helical nature of the flow, which allows us to discount any transformations that involve (orientation-reversing) reflections. To provide an example of how transformations are tested and also to identify how a reflection allows for no solutions for $(\alpha_j, \beta_j, \gamma_j)$, we have taken the first two transformations g_1 and g_2. Referring to (2.25) and (2.26) we recognise that g_1 represents an orientation-reversing reflection in the y-z plane, whereas g_2 represents

a rotation by π about the origin. Applying g_1 to the ABC flow (1.18) gives us

$$\mathsf{g}_1^* \boldsymbol{u} = \boldsymbol{J}_{\mathsf{g}_1} \boldsymbol{u}(\mathsf{g}_1^{-1}(x, y, z)). \tag{2.29}$$

We can see that $\mathsf{g}_1^{-1} = (-x + \alpha_1, y - \beta_1, z - \gamma_1)$ and the Jacobian of g_1 is easily identified as

$$\boldsymbol{J}_{\mathsf{g}_1} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.30}$$

meaning that (2.29) is resolved as

$$[\mathsf{g}_1^* \boldsymbol{u}]^\intercal = \begin{bmatrix} -C\sin(z - \gamma_1) - B\cos(y - \beta_1) \\ A\sin(-x + \alpha_1) + C\cos(z - \gamma_1) \\ B\sin(y - \beta_1) + A\cos(-x + \alpha_1) \end{bmatrix}. \tag{2.31}$$

If g_1 is a symmetry of $\boldsymbol{u}$, then there will be at least one solution $(\alpha_1, \beta_1, \gamma_1)$ for which $\boldsymbol{u}$ is invariant. Taking the first component of $\mathsf{g}_1^* \boldsymbol{u}$, we require that

$$-C\sin(z - \gamma_1) - B\cos(y - \beta_1) = C\sin z + B\cos y, \tag{2.32}$$

to which the only solution must be $\beta_1 = \gamma_1 = \pi$, as we are only identifying solutions within the interval $(-\pi, \pi]$. With this solution, (2.31) becomes

$$\mathsf{g}_1^* \boldsymbol{u} = \left[C\sin z + B\sin y, A\sin(-x + \alpha_1) - C\cos z, -B\sin y + A\cos(-x + \alpha_1)\right], \tag{2.33}$$

and it is easily seen that we have not returned to the original form of the ABC flows. Even if we set $B = C = 0$, we still require that

$$\sin(-x + \alpha_1) = \sin x, \quad \cos(-x + \alpha_1) = \cos x, \tag{2.34}$$

which has no solution. Since we are unable to show that $\mathsf{g}_1^* \boldsymbol{u} = \boldsymbol{u}$, it is concluded that g_1 is not a symmetry of the ABC flows. The same approach is taken for g_2

and we are looking for solutions in terms of $(\alpha_2, \beta_2, \gamma_2)$ that satisfy

$$\mathrm{g}_2^* \boldsymbol{u} = \boldsymbol{J}_{\mathrm{g}_2} \boldsymbol{u}(\mathrm{g}_2^{-1}(x, y, z)), \tag{2.35}$$

with the Jacobian

$$\boldsymbol{J}_{\mathrm{g}_2} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.36}$$

and $\mathrm{g}_2^{-1} = (-x + \alpha_2, -y + \beta_2, z - \gamma_2)$. We therefore have that

$$[\mathrm{g}_2^* \boldsymbol{u}]^\intercal = \begin{bmatrix} -C\sin(z - \gamma_2) - B\cos(-y + \beta_2) \\ -A\sin(-x + \alpha_2) - C\cos(z - \gamma_2) \\ B\sin(-y + \beta_2) + A\cos(-x + \alpha_2) \end{bmatrix}. \tag{2.37}$$

We want $(\mathrm{g}_2^* \boldsymbol{u})_x = u_x$ and since cosine is an even function, we have the same problem (and thus solution) as for g_1, with $\beta_2 = \gamma_2 = \pi$. This leaves us with

$$\mathrm{g}_2^* \boldsymbol{u} = \left[C\sin z + B\cos y, -A\sin(-x + \alpha_2) + C\cos z, B\sin y + A\cos(-x + \alpha_2)\right], \tag{2.38}$$

and so to return to (1.18), we require that

$$-\sin(-x + \alpha) = \sin x, \quad \cos(-x + \alpha) = \cos x, \tag{2.39}$$

to which the solution is $\alpha = 0$. Thus

$$\mathrm{g}_2 = (-x, -y + \pi, z + \pi), \tag{2.40}$$

is a symmetry of the ABC flows. Using this method, it can be shown that no reflections are symmetries of the ABC flows. All orientation-reversing reflections share the property that $\det(\boldsymbol{J}_{\mathrm{g}_i}) = -1$ and so we can remove all that have this characteristic. This leaves only 24 valid transformations, as for every orientation preserving transformation (of which there are 24) there exists the same transformation with

an extra reflection. Although the derivation of every symmetry is not shown here in the same detail as for g_2, we show how some transformations are only symmetries under specific choices of A, B and C. The symmetry $\mathsf{g}_2 = (-x, -y + \pi, z + \pi)$, for instance, is a symmetry of all four cases, due to its lack of dependence on the values of A, B and C. An example chosen to demonstrate this is another transformation from the remaining list of 24 valid ones, let's call it g_3, defined as

$$\mathsf{g}_3 = (-y + \alpha_3, z + \beta_3, -x + \gamma_3), \tag{2.41}$$

$$\mathsf{g}_3^{-1} = (-z + \gamma_3, -x + \alpha_3, y - \beta_3), \tag{2.42}$$

with Jacobian

$$\boldsymbol{J}_{\mathsf{g}_3} = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & 1 \\ -1 & 0 & 0 \end{pmatrix}. \tag{2.43}$$

Applying g_3 to $\boldsymbol{u}$, we have

$$[\mathsf{g}_3^* \boldsymbol{u}]^\intercal = \begin{bmatrix} -A\sin(-z+\gamma) - C\cos(y-\beta) \\ B\sin(-x+\alpha) + A\cos(-z+\gamma) \\ -C\sin(y-\beta) - B\cos(-x+\alpha) \end{bmatrix}, \tag{2.44}$$

and by comparing the first component of $\mathsf{g}_3^* \boldsymbol{u}$ with that of $\boldsymbol{u}$, we obtain the following relations

$$-A\sin(-z+\gamma) = C\sin z, \tag{2.45}$$

$$-C\cos(y-\beta) = B\cos y, \tag{2.46}$$

of which we only have solutions for $A = B = C = 1$, since we have that $1 = A \geq B \geq C \geq 0$. This must mean that $\gamma = 0$ and $\beta = \pi$, with this solution also satisfying related conditions (for x and y) in the second and third components.

Two relations for α remain, these being

$$\sin(-x+\alpha) = \sin x, \tag{2.47}$$

$$-\cos(-x+\alpha) = \cos x, \tag{2.48}$$

to which the obvious solution (of both) is $\alpha = \pi$. We can say, therefore, that

$$\mathsf{g}_3 = (\pi - y, z + \pi, -x), \tag{2.49}$$

is a symmetry of the ABC flows, when $A = B = C = 1$ but is not a symmetry for more general A, B and C.

Referring to the symmetry g_2, we can see in the derivation that, regardless of the values of A, B and C chosen, this transformation will be a valid symmetry. If we investigate further, we can identify that this is due to the fact that variables are mapped to themselves (or their negations). Taking a counter-example, if we describe a transformation that, for instance, maps x to $\pm y$, then due to the form of the ABC flows (1.18), all terms involving y will have coefficients of A, instead of their pre-transformed coefficient B. This means that unless $A = B$, the transformation will not be a symmetry of the flow and so is only a symmetry for the cases $A = B = C$ and $A = B \neq C$. We can thus identify the general form of symmetries in each of the four regimes:

- $A \neq B \neq C \neq A$ - transformations of the form: $\mathsf{g}_j : (x, y, z) \mapsto (\pm x, \pm y, \pm z)$;
- $A = B \neq C$ - those for $A \neq B \neq C \neq A$, plus transformations of the form $\mathsf{g}_j : (x, y, z) \mapsto (\pm y, \pm x, \pm z)$;
- $A \neq B = C$ - those for $A \neq B \neq C \neq A$, plus transformations of the form $\mathsf{g}_j : (x, y, z) \mapsto (\pm x, \pm z, \pm y)$;
- $A = B = C$ - all valid (signed) permutations of the variables,

with, of course, no reflections allowed in above descriptions. Most importantly, we find that symmetries of the case $A \neq B \neq C \neq A$ will be symmetries of the cases

$A = B \neq C$ and $A = B = C$, or to put it another way, symmetries of more general cases will also be symmetries of the less general cases. Thus we can expect that the symmetry group of the case A, B, C unequal to be a subset of the groups for cases $A \neq B = C$ and $A = B \neq C$, themselves being subsets of the group for $A = B = C$. With this in mind, the full group of symmetries will be discussed for $A = B = C$ before moving onto the subsets belonging to other regimes.

2.6.2 Symmetries of ABC flows for $A = B = C$

With regards to previous discussion on generating the symmetries, each of the 24 orientation-preserving transformations is applied to $\boldsymbol{u}$ to test for invariance. It is found that all 24 were symmetries for $A = B = C$, with only one appropriate choice of $(\alpha_j, \beta_j, \gamma_j)$ for each. This group of symmetries (G) is displayed in (2.50).

$$
\begin{aligned}
\mathsf{i}(\boldsymbol{x}) &= (x, y, z), & \mathsf{a}(\boldsymbol{x}) &= (\tfrac{\pi}{2} - y, x + \tfrac{\pi}{2}, z - \tfrac{\pi}{2}), \\
\mathsf{a}^2(\boldsymbol{x}) &= (-x, \pi - y, z + \pi), & \mathsf{a}^3(\boldsymbol{x}) &= (y - \tfrac{\pi}{2}, \tfrac{\pi}{2} - x, z + \tfrac{\pi}{2}), \\
\mathsf{b}(\boldsymbol{x}) &= (x - \tfrac{\pi}{2}, \tfrac{\pi}{2} - z, y + \tfrac{\pi}{2}), & \mathsf{b}^2(\boldsymbol{x}) &= (x + \pi, -y, \pi - z), \\
\mathsf{b}^3(\boldsymbol{x}) &= (x + \tfrac{\pi}{2}, z - \tfrac{\pi}{2}, \tfrac{\pi}{2} - y), & \mathsf{c}(\boldsymbol{x}) &= (z + \tfrac{\pi}{2}, y - \tfrac{\pi}{2}, \tfrac{\pi}{2} - x), \\
\mathsf{c}^2(\boldsymbol{x}) &= (\pi - x, y + \pi, -z), & \mathsf{c}^3(\boldsymbol{x}) &= (\tfrac{\pi}{2} - z, y + \tfrac{\pi}{2}, x - \tfrac{\pi}{2}), \\
\mathsf{d}(\boldsymbol{x}) &= (z, x, y), & \mathsf{d}^2(\boldsymbol{x}) &= (y, z, x), \\
\mathsf{e}(\boldsymbol{x}) &= (-z, \pi - x, y + \pi), & \mathsf{e}^2(\boldsymbol{x}) &= (\pi - y, z + \pi, -x), \\
\mathsf{f}(\boldsymbol{x}) &= (z + \pi, -x, \pi - y), & \mathsf{f}^2(\boldsymbol{x}) &= (-y, \pi - z, x + \pi), \\
\mathsf{g}(\boldsymbol{x}) &= (\pi - z, x + \pi, -y), & \mathsf{g}^2(\boldsymbol{x}) &= (y + \pi, -z, \pi - x), \\
\mathsf{h}(\boldsymbol{x}) &= (\tfrac{\pi}{2} - x, z + \tfrac{\pi}{2}, y - \tfrac{\pi}{2}), & \mathsf{j}(\boldsymbol{x}) &= (z - \tfrac{\pi}{2}, \tfrac{\pi}{2} - y, x + \tfrac{\pi}{2}), \\
\mathsf{k}(\boldsymbol{x}) &= (y + \tfrac{\pi}{2}, x - \tfrac{\pi}{2}, \tfrac{\pi}{2} - z), & \mathsf{l}(\boldsymbol{x}) &= (-x - \tfrac{\pi}{2}, -z - \tfrac{\pi}{2}, -y - \tfrac{\pi}{2}), \\
\mathsf{m}(\boldsymbol{x}) &= (-z - \tfrac{\pi}{2}, -y - \tfrac{\pi}{2}, -x - \tfrac{\pi}{2}), & \mathsf{n}(\boldsymbol{x}) &= (-y - \tfrac{\pi}{2}, -x - \tfrac{\pi}{2}, -z - \tfrac{\pi}{2}).
\end{aligned}
\tag{2.50}
$$

The approach taken to construct this group is robust and does not require any prior assumptions in order to generate the list of symmetries. Alternative approaches

(see Dombre *et al.* (1986) and Arnold (1984)) involve identifying the generators, of which there are two, and building the group through unique compositions of these generators. The generators, however, are not unique in this problem and displaying the full set of symmetries purely in terms of the generators would be relatively arbitrary. More importantly, since we are making extensive use of the symmetries in later chapters, it would soon become clumsy. We adopt the user-friendly notation (a to n) of Gilbert (1992), also the only study in which they are explicitly listed (albeit in the appendix) and the letters match those of Gilbert for the sake of simplicity. As with the general rotations, reflections and translations earlier in the chapter, we use a sans-serif font to indicate a symmetry, to avoid confusion with other parameters and variables. We note that for the remainder of this thesis, when we apply these transformations to a vector function, we lose the convention of adding an asterisk * (used to indicate that g^* is a map associated with g) to avoiding overcomplicated notation.

The symmetries each represent an orientation-preserving permutation of variables which, as already mentioned, can be linked to the geometry of the cube. It is well-known that the group G is isomorphic to the group O_{24} of orientation-preserving symmetries of the cube (Arnold and Korkina, 1983; Arnold, 1984; Gilbert, 1992). Due to the isomorphism, we can visualise the rotational axes of the symmetries in G simply by visualising the rotational axes of the cube (refer to figure 2.5). All except three symmetries in G involve some form of translation, which are essential due to the helical nature of the flow.

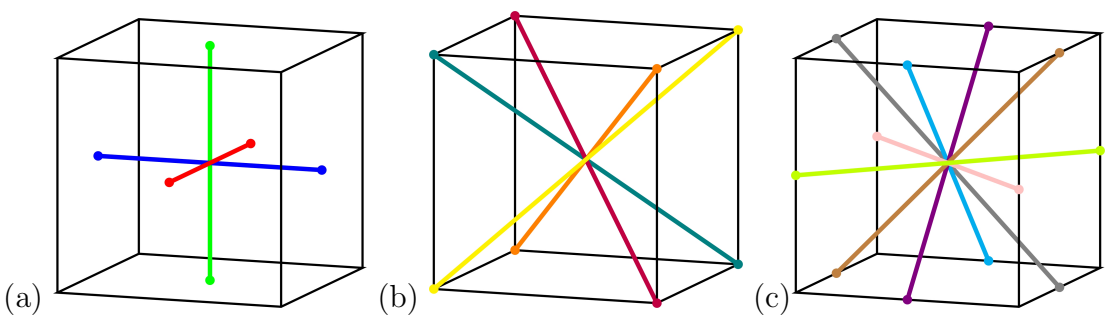

Figure 2.5: Visualisation of the rotation axes of symmetries (a) a to c, (b) d to g and (c) h to n relative to the periodic cube.

The symmetries a, b and c represent $\pi/2$ rotations about the z, x and y axes respectively. The flow is 2π-periodic and so these symmetries are four-fold with

$$\mathsf{a}^4 = \mathsf{b}^4 = \mathsf{c}^4 = \mathsf{i}, \tag{2.51}$$

where i is the identity transformation. The symmetries a^2 and a^3 represent successive applications of a and so are π and $3\pi/2$ rotations, respectively, about the z axis. Equivalently b^2, b^3 and c^2, c^3 are π, $3\pi/2$ rotations about the x and y axes, respectively. These axes are shown in figure 2.5(a), though are not labelled due to arbitrariness.

The symmetries d, e, f and g are $2\pi/3$ rotations about axes that connect two opposite corners of the cube (see figure 2.5(b)), though d is purely a rotation due to the invariance of (1.20) under cyclic permutation of variables. Here d^2, e^2, f^2 and g^2 are $4\pi/3$ rotations about the same respective axes with

$$\mathsf{d}^3 = \mathsf{e}^3 = \mathsf{f}^3 = \mathsf{g}^3 = \mathsf{i}, \tag{2.52}$$

and so these are three-fold symmetries.

The last set of symmetries are equivalent to π rotations about axes that join the midpoints of opposite edges (refer to figure 2.5(c)). The symmetries h, j, k, l, m and n are two-fold and are naturally involutions with

$$\mathsf{h}^2 = \mathsf{j}^2 = \mathsf{k}^2 = \mathsf{l}^2 = \mathsf{m}^2 = \mathsf{n}^2 = \mathsf{i}. \tag{2.53}$$

With regards to generating the group G, we require only two generators; one four-fold symmetry in $\{\mathsf{a}, \mathsf{a}^3, \mathsf{b}, \mathsf{b}^3, \mathsf{c}, \mathsf{c}^3\}$ and one three-fold in $\{\mathsf{d}, \mathsf{d}^2, \mathsf{e}, \mathsf{e}^2, \mathsf{f}, \mathsf{f}^2, \mathsf{g}, \mathsf{g}^2\}$. The symmetries $\{\mathsf{h}, \mathsf{j}, \mathsf{k}, \mathsf{l}, \mathsf{m}, \mathsf{n}\}$ can all be constructed through compositions of four-fold and three-fold symmetries. Relating these symmetries to those of Arnold (1984), we have that symmetries g_3 and g_4 of Arnold correspond to d^2 and b^3 of G. Likewise S_2S_1, S_3S_2 and S_1S_3 of Dombre *et al.* (1986) correspond to $\mathsf{a}^2, \mathsf{b}^2$ and c^2 of G, as noted by Gilbert (1992).

2.6.3 Symmetries of other ABC flows

As has been explained, the symmetries of more general ABC flows are subsets of the group of symmetries G of the case $A = B = C$. In our definition of the ABC flows (1.18), the parameter A is associated with terms involving x, B with terms involving y and C with terms involving z. The validity of the symmetries in G for more general cases are dependent on how they map variables, as demonstrated by testing g_2 and g_3. For the case $A = B \neq C$, we can deduce that only symmetries in which $z \mapsto \pm z$ will be valid. The group of symmetries for this case is

$$\begin{aligned}
\mathsf{i}(\boldsymbol{x}) &= (x, y, z), & \mathsf{a}(\boldsymbol{x}) &= (\tfrac{\pi}{2} - y, x + \tfrac{\pi}{2}, z - \tfrac{\pi}{2}), \\
\mathsf{a}^2(\boldsymbol{x}) &= (-x, \pi - y, z + \pi), & \mathsf{a}^3(\boldsymbol{x}) &= (y - \tfrac{\pi}{2}, \tfrac{\pi}{2} - x, z + \tfrac{\pi}{2}), \\
\mathsf{b}^2(\boldsymbol{x}) &= (x + \pi, -y, \pi - z), & \mathsf{c}^2(\boldsymbol{x}) &= (\pi - x, y + \pi, -z), \\
\mathsf{k}(\boldsymbol{x}) &= (y + \tfrac{\pi}{2}, x - \tfrac{\pi}{2}, \tfrac{\pi}{2} - z), & \mathsf{n}(\boldsymbol{x}) &= (-y - \tfrac{\pi}{2}, -x - \tfrac{\pi}{2}, -z - \tfrac{\pi}{2}),
\end{aligned} \tag{2.54}$$

where the labelling has been retained for ease of comparison. In an identical approach, we can list the symmetries for the case $A \neq B = C$, as the group will only contain symmetries in which $x \mapsto \pm x$:

$$\begin{aligned}
\mathsf{i}(\boldsymbol{x}) &= (x, y, z), & \mathsf{a}^2(\boldsymbol{x}) &= (-x, \pi - y, z + \pi), \\
\mathsf{b}(\boldsymbol{x}) &= (x - \tfrac{\pi}{2}, \tfrac{\pi}{2} - z, y + \tfrac{\pi}{2}), & \mathsf{b}^2(\boldsymbol{x}) &= (x + \pi, -y, \pi - z), \\
\mathsf{b}^3(\boldsymbol{x}) &= (x + \tfrac{\pi}{2}, z - \tfrac{\pi}{2}, \tfrac{\pi}{2} - y), & \mathsf{c}^2(\boldsymbol{x}) &= (\pi - x, y + \pi, -z), \\
\mathsf{h}(\boldsymbol{x}) &= (\tfrac{\pi}{2} - x, z + \tfrac{\pi}{2}, y - \tfrac{\pi}{2}), & \mathsf{l}(\boldsymbol{x}) &= (-x - \tfrac{\pi}{2}, -z - \tfrac{\pi}{2}, -y - \tfrac{\pi}{2}).
\end{aligned} \tag{2.55}$$

These two cases are obviously related through cyclic permutation of A, B and C. This means that any transformations not common to both can be related through cyclic permutation of variables (x, y, z), i.e. applying one of d or d^2. We can easily show that

$$\mathsf{a}^3 = \mathsf{b}\mathsf{d}^2, \qquad \mathsf{a} = \mathsf{d}\mathsf{b}^3 = (\mathsf{b}\mathsf{d}^2)^{-1}, \qquad \mathsf{k} = \mathsf{b}^3\mathsf{d}^2, \qquad \mathsf{n} = \mathsf{l}\mathsf{d}^2. \tag{2.56}$$

As the group G is isomorphic to O_{24} (rotation group of the cube), in both of these cases the group of symmetries is isomorphic to the symmetry group D_8 (the symmetries of a square) which were discussed in section 2.5. The final case with A, B and C unequal is the least symmetric and contains only symmetries that map each variable to themselves or their negations. As discussed, these symmetries will be common to the other cases we have studied and so the group for this case is displayed as

$$\begin{aligned} \mathsf{i}(\boldsymbol{x}) &= (x, y, z), & \mathsf{a}^2(\boldsymbol{x}) &= (-x, \pi - y, z + \pi), \\ \mathsf{b}^2(\boldsymbol{x}) &= (x + \pi, -y, \pi - z), & \mathsf{c}^2(\boldsymbol{x}) &= (\pi - x, y + \pi, -z), \end{aligned} \tag{2.57}$$

which contains the three π rotations about each of the x, y and z axes. This group is isomorphic to D_4, which is also described as the direct group product of two order 2 cyclic groups: $\mathbb{Z}_2 \times \mathbb{Z}_2$.

2.7 Summary

This chapter has provided an overview of the three fundamental types of symmetries: reflection, rotation and translation. Under reflection and rotation, there are always fixed points or lines that are invariant under the action of these symmetries, whereas under the action of a translation, there are no fixed points. Compositions of these different types may also form valid symmetries, for example helical symmetry consists of rotations with translations.

Before constructing the symmetries of the 3-dimensional ABC flows, the symmetry group of a square, labelled D_8, is constructed from its generators r, a $\pi/2$ rotation, and s a reflection through the x-axis. The group has eight elements, meaning that the fundamental domain is 1/8 of the full domain. For a group of this size, it is appropriate to describe the symmetries as compositions of their generators.

To construct the symmetries, the ABC flows were tested for invariance under all possible permutations of variables, which were combined with translations to make

them more general. If a particular permutation was a valid symmetry, there would be at least one unique translation that made it a solution of (2.28). It was found that no (orientation-reversing) reflections could be symmetries of the ABC flows in fitting with the helical (and so chiral) nature of the flows.

It is seen that ABC flows contain varying degrees of symmetry, dependent on the parameters A, B and C. The most symmetric of the ABC flows is the case $A = B = C$, which is invariant under 24 individual transformations, with these symmetries forming the group G. The cases with two of A, B and C equal have symmetry groups which are a subset of G, each contain eight distinct symmetries and are isomorphic to the symmetry group D_8. For all three unequal, the group of symmetries is reduced to four elements. Three symmetries in G contain no translations: the identity i, d and d^2 rotations about the axis $x = y = z$ by $2\pi/3$ and $4\pi/3$ respectively, these being cyclic permutations of the identity. All other symmetries, which are four-fold ($\pi/2$), three-fold ($2\pi/3$) and two-fold (π) rotations, involve translations which can be seen as phase-shifts.

The case $A = B = C = 1$ is the focus of the main thread of this thesis, as it is the most interesting and complex, and thus the best to demonstrate the tools and methods developed in the process. The group G of 24 symmetries is isomorphic to the group of orientation-preserving symmetries (proper rotations) of a cube. The fundamental domain of this flow is 1/24 of the full domain.

A magnetic field that is driven by the ABC flows, however, will not necessarily be invariant under these symmetries. In fact, a magnetic field that is general enough will never be invariant under all of the symmetries. The next chapter will show how, in the linear dynamo problem, a general field can be deconstructed into subfields that each behave differently under the symmetries of the ABC flow through the use of irreducible group representations. A simpler example is first explored with a general magnetic field in a geometry with D_8 symmetry to give an idea of how such a problem is approached.

In chapter 7, the symmetries of large-scale fields will also be explored.

Chapter 3

Representation Theory

Some humans would do anything to see if it was possible to do it. If you put a large switch in some cave somewhere, with a sign on it saying 'End-of-the-World Switch. PLEASE DO NOT TOUCH', the paint wouldn't even have time to dry.

Terry Pratchett

In chapter 2, the symmetry group G of the 1:1:1 ABC flow (1.20) was constructed. As we are studying the evolution of a magnetic field driven by such a flow, it is wise to consider how the symmetries of the flow can affect the types of magnetic field structure that emerge. This chapter therefore concerns the effect that these symmetries have on a magnetic field $\boldsymbol{b}$ with no assumed symmetries, i.e. a general magnetic field, through the use of representation theory.

Through development of an incredibly simple example with one symmetry ($\mathbb{Z}_2$) and then a more informative and complex example with eight symmetries (D_8), the important processes are seen in constructing 'new' fields, from linear combinations of the original field $\boldsymbol{b}$, that are orthogonal to each other and (most importantly)

are independent under the action of all symmetries in their respective groups. This independence, owing to the linear evolution equation, means that these fields can be studied independently. To aid the understanding of this process, the fields in these examples are visualised. Throughout these examples, key concepts such as matrix representations and irreducible representations are introduced.

A more formal discussion on representations and irreducible representations is then provided to clarify ideas introduced in the examples. The process is outlined for a general case, whilst referencing the examples so that the concepts may have physical meaning. The importance of "characters" in relation to irreducible representations is then briefly explored.

Finally, we construct the irreducible representations of the 1:1:1 ABC flow using specific ideas obtained earlier in the chapter and from other investigations of this flow. It is shown that there are five types of structure (symmetry classes) and their irreducible representations are found, without requiring knowledge of the relevant linear combinations of the original field $\boldsymbol{b}$. The findings and ideas of the chapter are then summarised.

3.1 Irreducible representations of $\mathbb{Z}_2$

This first example concerns dipolar and quadrupolar fields, which behave differently under the application of a single reflection operator. The concept is similar to even and odd functions; even functions are those for which $f(-x) = f(x)$ whereas odd functions have $f(-x) = -f(x)$. A general function, for example one that is neither odd nor even, can be described as the sum of an even and an odd function. The same is true of a general magnetic field in a sphere, which contains reflective symmetry in the horizontal plane, and so the field can be described as a sum of dipolar and quadrupolar components. The following example details this in an axisymmetric field, so that the situation is reduced to that of a field in a circle.

The group $\mathbb{Z}_2$ contains two elements: the identity and an element that is an invo-

lution (is its own inverse). Both reflections and π rotations are involutions, though we choose a reflection for the sake of simplicity. For a circle centred at the origin, let a reflection through the line $y = 0$ be defined as

$$\mathsf{s}(x, y) = (x, -y), \tag{3.1}$$

with $\mathsf{s}^2 = \mathsf{i}$. An arbitrary magnetic field $\boldsymbol{b}$ is placed in the lower half of the circle, with the resulting layout visualised in figure 3.1(b). Note that the magnetic field has been chosen to best demonstrate this example, as will be seen, but it is not restricted to any particular form. Applying the reflection s to $\boldsymbol{b}$, the field in the

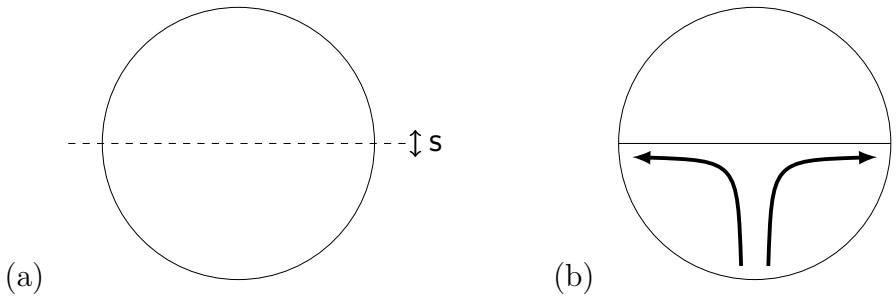

Figure 3.1: Shown are (a) reflection symmetry (horizontal) in a circle, and (b) a magnetic field $\boldsymbol{b}$ confined to the lower-half of the circle.

upper-half of the circle, $\mathsf{s}\boldsymbol{b}$, can be obtained and is linearly independent of $\boldsymbol{b}$. Linear combinations of the fields $\boldsymbol{b} = \mathsf{i}\boldsymbol{b}$ and $\mathsf{s}\boldsymbol{b}$, that transform simply under the reflection, can be found, with these being

$$\boldsymbol{b}^{\mathrm{I}} = \boldsymbol{b} + \mathsf{s}\boldsymbol{b}, \tag{3.2}$$

$$\boldsymbol{b}^{\mathrm{II}} = \boldsymbol{b} - \mathsf{s}\boldsymbol{b}, \tag{3.3}$$

with $2\boldsymbol{b} = \boldsymbol{b}^{\mathrm{I}} + \boldsymbol{b}^{\mathrm{II}}$. The fields are visualised in figure 3.2 and it is clearly seen that $\boldsymbol{b}^{\mathrm{I}}$ is quadrupolar in form, whereas $\boldsymbol{b}^{\mathrm{II}}$ is dipolar (through cancellation of the field near to the axis of symmetry). It can be shown that $\boldsymbol{b}^{\mathrm{I}}$ is even under reflection, as

$$\mathsf{s}\boldsymbol{b}^{\mathrm{I}} = \mathsf{s}\boldsymbol{b} + \mathsf{s}^2\boldsymbol{b} = \boldsymbol{b} + \mathsf{s}\boldsymbol{b} = \boldsymbol{b}^{\mathrm{I}}, \tag{3.4}$$

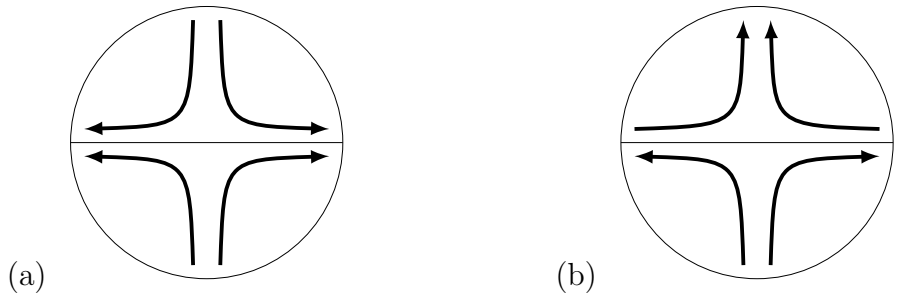

Figure 3.2: Examples of field configurations for symmetry classes (a) I and (b) II.

and that $\boldsymbol{b}^{\mathrm{II}}$ is odd under reflection, as

$$\mathsf{s}\boldsymbol{b}^{\mathrm{II}} = \mathsf{s}\boldsymbol{b} - \mathsf{s}^2\boldsymbol{b} = -\boldsymbol{b} + \mathsf{s}\boldsymbol{b} = -\boldsymbol{b}^{\mathrm{II}}, \tag{3.5}$$

thus satisfying the requirement that the symmetries behave simply under application of s. General fields can then be constructed through linear combinations of $\boldsymbol{b}^{\mathrm{I}}$ and $\boldsymbol{b}^{\mathrm{II}}$. Both (3.4) and (3.5) can be summarised by

$$\mathsf{s}\begin{pmatrix}\boldsymbol{b}^{\mathrm{I}}\\ \boldsymbol{b}^{\mathrm{II}}\end{pmatrix} = \begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix}\begin{pmatrix}\boldsymbol{b}^{\mathrm{I}}\\ \boldsymbol{b}^{\mathrm{II}}\end{pmatrix}. \tag{3.6}$$

The 2×2 matrix is the *matrix representation* of s and describes how the reflection maps these two types of field (even and odd). The other element of the $\mathbb{Z}_2$ group, the identity i, also has a matrix representation describing how it acts on $\boldsymbol{b}^{\mathrm{I}}$ and $\boldsymbol{b}^{\mathrm{II}}$, and is simply the two-dimensional identity matrix $\boldsymbol{I}_2$. As both elements of the group $\mathbb{Z}_2$ maps each field to a linear function of itself only (and not to the other), these fields are clearly independent under the group action and so the matrix representations contain independent *subrepresentations*. The action of the elements of the group $\mathbb{Z}_2$ on the fields can thus be described by

$$\mathsf{g}\boldsymbol{b}^{\alpha} = M^{\alpha}(\mathsf{g})\boldsymbol{b}^{\alpha}, \tag{3.7}$$

for $\mathsf{g}\in\mathbb{Z}_2 = \{\mathsf{i},\mathsf{s}\}$ and α as I or II. The objects labelled as M are called the

irreducible representations of $\mathbb{Z}_2$, with these clearly given by

$$M^{\mathrm{I}}(\mathsf{i}) = M^{\mathrm{II}}(\mathsf{i}) = 1, \qquad M^{\mathrm{I}}(\mathsf{s}) = 1, \qquad M^{\mathrm{II}}(\mathsf{s}) = -1. \tag{3.8}$$

In referring to the 'labels', denoted by I and II in this example, the terminology *symmetry class* will be adopted, in fitting with Metropolis *et al.* (1991). Thus $\boldsymbol{b}^{\mathrm{I}}$ and $\boldsymbol{b}^{\mathrm{II}}$ are identified as magnetic fields belonging to symmetry classes I and II respectively. Roman numeral notation has been adopted for labelling symmetry classes, so as to avoid confusion in later examples where the magnetic fields can have multiple components which are assigned Arabic numerals.

A few points of clarification are required before exploring a system with a greater number of symmetries in the next section. Firstly, in the kinematic problem, the induction equation is linear in $\boldsymbol{b}$ and so the magnetic fields belonging to individual symmetry classes remain independent of fields in other classes and so can be resolved independently, without loss of information. Any kinematic problem can therefore be decomposed into its respective symmetry classes, each of these being resolved separately. In the non-linear regime, however, there can then be symmetry breaking and mixed-mode solutions, but that is not possible in the linear approximation.

Secondly, this example is overly simple so that it is easier to grasp the basics of the application of the use of representations in the context of magnetic fields. The definitions (3.2) and (3.3) may seem rather arbitrary, but this will make more sense with a larger symmetry group, as in the next example.

Thirdly, in this very simple example, it was possible to construct the fields (3.2) and (3.3) so that the matrix representations for i and s were diagonal, meaning that both fields $\boldsymbol{b}^{\mathrm{I}}$ and $\boldsymbol{b}^{\mathrm{II}}$ were independent and so belonged to different symmetry classes. With larger symmetry groups, it is not necessarily possible to construct fields so that the matrix representations are all diagonalised. However it is always possible to make these representations to be block diagonal. This means that certain symmetry classes can contain multiple magnetic fields, with the irreducible

representations M^α being matrices instead of single values.

Finally, a reflection was chosen to demonstrate the construction of fields in each symmetry class but in practice, any involution would have been suitable. For example, a π rotation about the origin, $\mathsf{r}(x, y) = (-x, -y)$, could also have been used, rather than the reflection s. The resulting irreducible representations M^α, however, would have been the same for symmetry classes I and II, as they only depend on the abstract group which is, in this case, $\mathbb{Z}_2$.

3.2 Irreducible representations of D_8

The second example concerns the group D_8: symmetries of the square. This example isn't meant to be particularly physical, as planar fields and flows are not conducive for dynamo action (see chapter 1). However, it demonstrates how irreducible representations are found for a sufficiently complex group but allows the fields in each symmetry class to still be visualised, something which is not possible with the larger group of 1:1:1 ABC flow symmetries, G. We take the square to

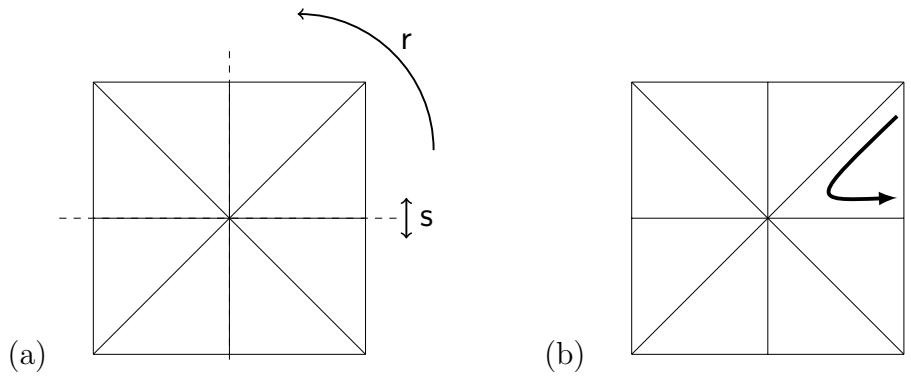

Figure 3.3: Shown are (a) the geometry for a square with symmetry group D_8, and (b) a magnetic field $\boldsymbol{b}$ confined to one eighth of the square.

be centred at the origin with sides parallel to the x and y axes, depicted in figure 3.3(a). D_8 is generated by an anti-clockwise rotation r of $\pi/2$ about the origin and

a reflection s in the x-axis,

$$\mathsf{r}(x,y) = (-y,x), \qquad \mathsf{s}(x,y) = (x,-y), \tag{3.9}$$

and contains eight elements

$$\{\mathsf{i}, \mathsf{r}, \mathsf{r}^2, \mathsf{r}^3, \mathsf{s}, \mathsf{rs}, \mathsf{r}^2\mathsf{s}, \mathsf{r}^3\mathsf{s}\}, \tag{3.10}$$

with $\mathsf{r}^4 = \mathsf{s}^2 = \mathsf{i}$ and $\mathsf{rs} = \mathsf{sr}^{-1}$.

As with the first example, suppose we start with a general magnetic field $\boldsymbol{b}$. Then we may map it with each element $\mathsf{g} \in D_8$ to obtain 8 linearly independent magnetic fields, spanning a vector space W,

$$W = \text{span}\{\boldsymbol{b}, \mathsf{r}\boldsymbol{b}, \mathsf{r}^2\boldsymbol{b}, \mathsf{r}^3\boldsymbol{b}, \mathsf{s}\boldsymbol{b}, \mathsf{rs}\boldsymbol{b}, \mathsf{r}^2\mathsf{s}\boldsymbol{b}, \mathsf{r}^3\mathsf{s}\boldsymbol{b}\}. \tag{3.11}$$

A convenient way to visualise this is to consider a magnetic field depicted schematically in figure 3.3(b) which is confined to one-eighth of the square and has an anti-clockwise, or positive, sense. Then each transformation in D_8 maps the field to a distinct eighth of the square and these fields are all plainly linearly independent. Note that under a reflection s, the resulting field $\mathsf{s}\boldsymbol{b}$ has a clockwise or negative sense. As with the example of $\mathbb{Z}_2$, we can now seek combinations of the eight fields in (3.11) that transform simply under the group elements. There are four of these, which are labelled I–IV,

$$\boldsymbol{b}^{\text{I}} = \boldsymbol{b} + \mathsf{r}\boldsymbol{b} + \mathsf{r}^2\boldsymbol{b} + \mathsf{r}^3\boldsymbol{b} + \mathsf{s}\boldsymbol{b} + \mathsf{rs}\boldsymbol{b} + \mathsf{r}^2\mathsf{s}\boldsymbol{b} + \mathsf{r}^3\mathsf{s}\boldsymbol{b}, \tag{3.12}$$

$$\boldsymbol{b}^{\text{II}} = \boldsymbol{b} + \mathsf{r}\boldsymbol{b} + \mathsf{r}^2\boldsymbol{b} + \mathsf{r}^3\boldsymbol{b} - \mathsf{s}\boldsymbol{b} - \mathsf{rs}\boldsymbol{b} - \mathsf{r}^2\mathsf{s}\boldsymbol{b} - \mathsf{r}^3\mathsf{s}\boldsymbol{b}, \tag{3.13}$$

$$\boldsymbol{b}^{\text{III}} = \boldsymbol{b} - \mathsf{r}\boldsymbol{b} + \mathsf{r}^2\boldsymbol{b} - \mathsf{r}^3\boldsymbol{b} + \mathsf{s}\boldsymbol{b} - \mathsf{rs}\boldsymbol{b} + \mathsf{r}^2\mathsf{s}\boldsymbol{b} - \mathsf{r}^3\mathsf{s}\boldsymbol{b}, \tag{3.14}$$

$$\boldsymbol{b}^{\text{IV}} = \boldsymbol{b} - \mathsf{r}\boldsymbol{b} + \mathsf{r}^2\boldsymbol{b} - \mathsf{r}^3\boldsymbol{b} - \mathsf{s}\boldsymbol{b} + \mathsf{rs}\boldsymbol{b} - \mathsf{r}^2\mathsf{s}\boldsymbol{b} + \mathsf{r}^3\mathsf{s}\boldsymbol{b}. \tag{3.15}$$

It is seen that transforming $\boldsymbol{b}^{\text{I}}$ to $\boldsymbol{b}^{\text{IV}}$ under any group element $\mathsf{g} \in D_8$ simply leads to ± 1 times the original field and so again it is the case that (3.7) is applicable,

where the α now indicates symmetry classes I–IV. As the fields transform simply, the irreducible representations M^α for each of these four symmetry classes is simply a single value, either 1 or -1. It is clear that for class I, $M^{\mathrm{I}}(\mathsf{g}) = 1$ for all $\mathsf{g} \in D_8$. For representation II we have

$$M^{\mathrm{II}}(\mathsf{i}) = M^{\mathrm{II}}(\mathsf{r}) = M^{\mathrm{II}}(\mathsf{r}^2) = M^{\mathrm{II}}(\mathsf{r}^3) = 1, \tag{3.16}$$

$$M^{\mathrm{II}}(\mathsf{s}) = M^{\mathrm{II}}(\mathsf{rs}) = M^{\mathrm{II}}(\mathsf{r}^2\mathsf{s}) = M^{\mathrm{II}}(\mathsf{r}^3\mathsf{s}) = -1. \tag{3.17}$$

Plainly, it is the case that

$$M^\alpha(\mathsf{g})M^\alpha(\mathsf{h}) = M^\alpha(\mathsf{gh}), \tag{3.18}$$

for any elements g and h in the group D_8. This result is very important as to construct the irreducible representations of all elements of a group for a particular symmetry class, only the irreducible representations of the group's generators need to be known for that class. The values of the $M^\alpha(\mathsf{g})$ for $\mathsf{g} \in D_8$ can be read directly from the character table, table 3.1, for representations I–IV. (More discussion on the importance and use of character tables is given in section 3.3.) Given the initial field $\boldsymbol{b}$ shown in figure 3.3(b) the four fields $\boldsymbol{b}^{\mathrm{I}}$ to $\boldsymbol{b}^{\mathrm{IV}}$ are seen schematically in figure 3.4; similar qualitative illustrations of field configurations for D_8 can also be seen in Matthews (1999). Where the field has a negative (clockwise) sense, the segment has been shaded.

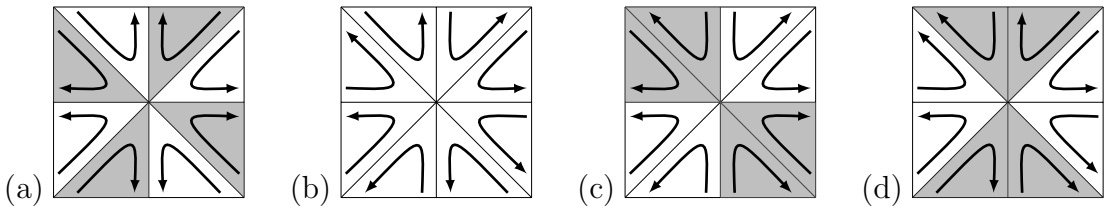

Figure 3.4: Fields for symmetry classes I–IV of D_8: (a) $\boldsymbol{b}^{\mathrm{I}}$, (b) $\boldsymbol{b}^{\mathrm{II}}$, (c) $\boldsymbol{b}^{\mathrm{III}}$, and (d) $\boldsymbol{b}^{\mathrm{IV}}$.

Assuming that field lines, in each of the eight segments, form closed curves and that adjacent field line sections with opposing polarity cancel out, whereas those with

matching polarity form a magnetic 'pole', it is possible to classify these fields in terms of the number of poles, as was done in the example of $\mathbb{Z}_2$. In this way, $\boldsymbol{b}^{\mathrm{I}}$ can be described as being octopolar (eight poles), with $\boldsymbol{b}^{\mathrm{III}}$ and $\boldsymbol{b}^{\mathrm{IV}}$ being quadrupolar (four poles), though the number of poles in $\boldsymbol{b}^{\mathrm{II}}$ is unclear.

Table 3.1: Character table for the group D_8. Entries give the character, that is the trace of the matrix $\boldsymbol{M}^{\alpha}(\mathsf{g})$, for each transformation g and irreducible representation α from I–V.

	$\{\mathsf{i}\}$	$\{\mathsf{r},\mathsf{r}^3\}$	$\{\mathsf{r}^2\}$	$\{\mathsf{s},\mathsf{r}^2\mathsf{s}\}$	$\{\mathsf{rs},\mathsf{r}^3\mathsf{s}\}$
I	1	1	1	1	1
II	1	1	1	−1	−1
III	1	−1	1	1	−1
IV	1	−1	1	−1	1
V	2	0	−2	0	0

Constructing these new fields can be considered as a change of basis within the vector space W. Though these are the only four fields that transform simply (multiplied by ± 1) under the transformations, it is clear that they do not form a full basis and cannot span the full eight-dimensional vector space of fields W in (3.11). We see, for example, that they do not include the possibility of a non-zero mean field in the system, as under the $\pi/2$ rotation r such a field could not transform to ± 1 times itself. It is, in fact, the case that the set of fields (3.12)–(3.15) exhaust the one-dimensional representations of the group D_8, and the next best we can do is to define pairs of fields in the two-dimensional representations,

$$\begin{aligned}\boldsymbol{b}_1^{\mathrm{Va}} &= \boldsymbol{b} + \mathsf{r}\boldsymbol{b} - \mathsf{r}^2\boldsymbol{b} - \mathsf{r}^3\boldsymbol{b} + \mathsf{s}\boldsymbol{b} - \mathsf{rs}\boldsymbol{b} - \mathsf{r}^2\mathsf{s}\boldsymbol{b} + \mathsf{r}^3\mathsf{s}\boldsymbol{b},\\ \boldsymbol{b}_2^{\mathrm{Va}} &= \boldsymbol{b} - \mathsf{r}\boldsymbol{b} - \mathsf{r}^2\boldsymbol{b} + \mathsf{r}^3\boldsymbol{b} - \mathsf{s}\boldsymbol{b} - \mathsf{rs}\boldsymbol{b} + \mathsf{r}^2\mathsf{s}\boldsymbol{b} + \mathsf{r}^3\mathsf{s}\boldsymbol{b},\end{aligned} \tag{3.19}$$

and

$$\begin{aligned}\boldsymbol{b}_1^{\mathrm{Vb}} &= \boldsymbol{b} + \mathsf{r}\boldsymbol{b} - \mathsf{r}^2\boldsymbol{b} - \mathsf{r}^3\boldsymbol{b} - \mathsf{s}\boldsymbol{b} + \mathsf{rs}\boldsymbol{b} + \mathsf{r}^2\mathsf{s}\boldsymbol{b} - \mathsf{r}^3\mathsf{s}\boldsymbol{b},\\ \boldsymbol{b}_2^{\mathrm{Vb}} &= \boldsymbol{b} - \mathsf{r}\boldsymbol{b} - \mathsf{r}^2\boldsymbol{b} + \mathsf{r}^3\boldsymbol{b} + \mathsf{s}\boldsymbol{b} + \mathsf{rs}\boldsymbol{b} - \mathsf{r}^2\mathsf{s}\boldsymbol{b} - \mathsf{r}^3\mathsf{s}\boldsymbol{b}.\end{aligned} \tag{3.20}$$

Note that (Arabic numeral) subscripts label different fields, not their Cartesian

components. These four fields are depicted in figure 3.5. Focusing on the first pair of fields (3.19) shown in figure 3.5(a,b), we have

$$\mathsf{r}\boldsymbol{b}_1^{\mathrm{Va}} = -\boldsymbol{b}_2^{\mathrm{Va}}, \qquad \mathsf{r}\boldsymbol{b}_2^{\mathrm{Va}} = \boldsymbol{b}_1^{\mathrm{Va}}, \qquad \mathsf{s}\boldsymbol{b}_1^{\mathrm{Va}} = \boldsymbol{b}_1^{\mathrm{Va}}, \qquad \mathsf{s}\boldsymbol{b}_2^{\mathrm{Va}} = -\boldsymbol{b}_2^{\mathrm{Va}}, \tag{3.21}$$

and so now the action of the group elements is given by 2×2 matrices defined by

$$\boldsymbol{M}^{\mathrm{Va}}(\mathsf{r}) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{Va}}(\mathsf{s}) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{3.22}$$

with

$$\mathsf{g}\boldsymbol{b}_i^{\alpha} = \sum_j M_{ij}^{\alpha}(\mathsf{g})\, \boldsymbol{b}_j^{\alpha}, \tag{3.23}$$

where the label α corresponds to the label 'Va'. The matrices for the remaining transformations follow by matrix multiplication, using the requirement that (3.18) hold, now as a matrix equation. Note that the fields $\boldsymbol{b}_1^{\mathrm{Va}}$ and $\boldsymbol{b}_2^{\mathrm{Va}}$ allow the possibility of mean fields in the x and y directions respectively. These two fields give a two-dimensional irreducible representation.

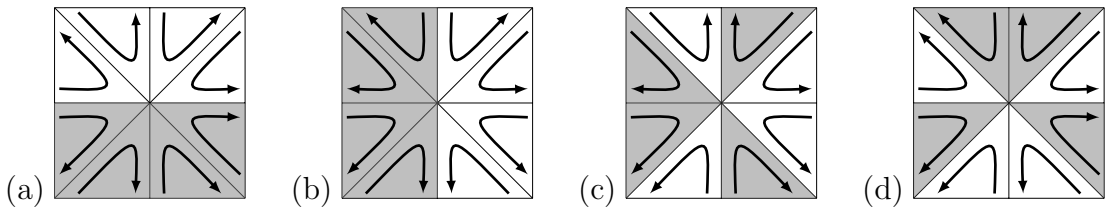

Figure 3.5: Fields in symmetry class V of D_8: (a) $\boldsymbol{b}_1^{\mathrm{V}a}$, (b) $\boldsymbol{b}_2^{\mathrm{V}a}$, (c) $\boldsymbol{b}_1^{\mathrm{V}b}$, and (d) $\boldsymbol{b}_2^{\mathrm{V}b}$.

The second pair of fields, in (3.20), corresponds to another two-dimensional symmetry class 'Vb' isomorphic to Va, and is subtly different. It can be shown that

$$\mathsf{r}\boldsymbol{b}_1^{\mathrm{Vb}} = -\boldsymbol{b}_2^{\mathrm{Vb}}, \qquad \mathsf{r}\boldsymbol{b}_2^{\mathrm{Vb}} = \boldsymbol{b}_1^{\mathrm{Vb}}, \qquad \mathsf{s}\boldsymbol{b}_1^{\mathrm{Vb}} = -\boldsymbol{b}_1^{\mathrm{Vb}}, \qquad \mathsf{s}\boldsymbol{b}_2^{\mathrm{Vb}} = \boldsymbol{b}_2^{\mathrm{Vb}}, \tag{3.24}$$

with the irreducible representations then given by

$$\boldsymbol{M}^{\mathrm{Vb}}(\mathsf{r}) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{Vb}}(\mathsf{s}) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}. \tag{3.25}$$

Both classes Va and Vb have the same irreducible representation for r, however $M^{\mathrm{Vb}}(\mathsf{s}) = -M^{\mathrm{Va}}(\mathsf{s})$. This difference can be seen in the field structures, see figure 3.5(c,d), which resemble hexapolar fields, as opposed to those of class Va that resemble bipolar fields. Despite this, mean fields remain possible in class Vb.

All the symmetry classes, I, II, III, IV and two copies of V, span the vector space W in (3.11) and any field $\boldsymbol{b}$ can be decomposed into components in each of these spaces as above with

$$8\boldsymbol{b} = \boldsymbol{b}^{\mathrm{I}} + \boldsymbol{b}^{\mathrm{II}} + \boldsymbol{b}^{\mathrm{III}} + \boldsymbol{b}^{\mathrm{IV}} + \boldsymbol{b}_1^{\mathrm{Va}} + \boldsymbol{b}_2^{\mathrm{Va}} + \boldsymbol{b}_1^{\mathrm{Vb}} + \boldsymbol{b}_2^{\mathrm{Vb}}. \tag{3.26}$$

Some comments are in order. We have four one-dimensional symmetry classes I–IV, in which the action on the corresponding field is multiplication by 1×1 matrices (see (3.7)). There are two copies of what is essentially the same two-dimensional class V in which the action is on pairs of fields through multiplication by 2×2 matrices. These matrices are not unique under a change of basis, though we have provided a rational choice with $\boldsymbol{M}^{\alpha}(\mathsf{g}) = \boldsymbol{J}_{\mathsf{g}}$. The forms of the fields in (3.19), (3.20) are dependent on this (somewhat arbitrary) choice. Only the traces of the matrices are invariant, and these are given in the character table, table 3.1. As the matrix for the group identity operation i is always the identity matrix, this column gives the dimension of the irreducible representations in each symmetry class. It should be noted that subclasses Va and Vb are in the same symmetry class (and so are isomorphic to one another) because the traces of their irreducible representations are the same.

Going back to the basis of fields of W given in (3.11), the action of an element $\mathsf{g} \in D_8$ is simply to permute these fields and so give rise to a (reducible) representation with 8×8 matrices $\boldsymbol{M}(\mathsf{g})$ having a permutation form (i.e. a single entry of 1 per

column and row). Representation theory can be seen as a way to choose a basis (of magnetic fields here) so as far as possible to diagonalise simultaneously the matrices for all the $\mathsf{g} \in D_8$. Under the basis in (3.12)–(3.15), (3.19) and (3.20), the 8×8 matrix for each g takes a block diagonal form

$$\boldsymbol{M}(\mathsf{g}) = \begin{pmatrix} M^{\mathrm{I}}(\mathsf{g}) & 0 & 0 & 0 & \cdots & \\ 0 & M^{\mathrm{II}}(\mathsf{g}) & 0 & 0 & & \\ 0 & 0 & M^{\mathrm{III}}(\mathsf{g}) & 0 & & \vdots \\ 0 & 0 & 0 & M^{\mathrm{IV}}(\mathsf{g}) & \ddots & \\ \vdots & & & \ddots & \boldsymbol{M}^{\mathrm{Va}}(\mathsf{g}) & \mathbf{0} \\ & & \cdots & & \mathbf{0} & \boldsymbol{M}^{\mathrm{Vb}}(\mathsf{g}) \end{pmatrix}. \tag{3.27}$$

Any linear problem that respects the symmetries can be considered in each of the six subspaces independently, in particular any linear eigenvalue problem. Thus instead of solving the full problem, we have five distinct reduced problems (we need only consider one subclass in class V), each with its own set of eigenvalues.

Finally, we note that it is not necessary to give the fields of each irreducible representation as they are displayed in (3.12)–(3.15), (3.19) and (3.20) for D_8. We have done so simply to demonstrate how these fields relate to those of the vector space W in this example. For larger groups G this becomes an increasingly arduous and uninteresting task. The key information is how each of these fields is transformed by the elements of G, which simply requires knowing matrices $\boldsymbol{M}^{\alpha}(\mathsf{g})$ for each irreducible representation α.

3.3 Representations and characters

So far in the two examples, some of the key concepts behind transforming reducible representations into irreducible representations have been introduced. Before moving on to the irreducible representations of the 1:1:1 ABC flow, a brief clarification of these ideas is necessary. For more detailed information on representation theory and character theory, see Hamermesh (1962).

3.3.1 Reducible and irreducible representations

Representation theory is a field of study that involves describing the action of a group's elements as linear transformations of vector spaces. Commonly, elements of a group can be written as matrices and the group's operation can therefore be reduced to matrix multiplication. This allows a potentially abstract group of objects to be interpreted in a physical sense, through the use of linear algebra. The use of group representations is commonplace in particular areas of physics, for example crystallography where symmetries are important in determining the properties of materials.

Matrix representations can be constructed for any finite group of elements. Through an appropriate similarity transformation, the representation of every group element can be transformed to the same block-diagonal form, with the submatrices along the diagonal termed irreducible representations. As the actions of all group elements can be described by block-diagonal matrices of the same form, the vector space must contain independent subspaces, within which the action of each group element is described by the irreducible representations. The term irreducible representation and its contraction *irrep* are used interchangeably for the remainder of this thesis.

One way to view this is if we have a vector field, say $\boldsymbol{b}$, that is defined in three-dimensional space and a group of n symmetries, $G = \{\mathrm{g}_1, \ldots, \mathrm{g}_n\}$, then a vector space W can be constructed by applying each g_i in G to $\boldsymbol{b}$ so that we obtain a space W of n independent fields, giving

$$W = \mathrm{span}\{\mathrm{g}_1\boldsymbol{b}, \ldots, \mathrm{g}_n\boldsymbol{b}\}. \tag{3.28}$$

The action of each g_i on each of these independent fields is then simply to map it to another field so that $\mathrm{g}_a\mathrm{g}_b\boldsymbol{b} = \mathrm{g}_c\boldsymbol{b}$ where a, b and c are between 1 and n. Constructing a vector of the fields in W allows us to represent each g_i in matrix

form and so if $\boldsymbol{z}^\intercal = (\mathrm{g}_1\boldsymbol{b}, \ldots, \mathrm{g}_n\boldsymbol{b})$, then

$$\mathrm{g}_i\boldsymbol{z} = \boldsymbol{A}(\mathrm{g}_i)\boldsymbol{z}, \tag{3.29}$$

where $\boldsymbol{A}$ are the matrix representations of the g_i and are permutation matrices. If this system is reducible, there exists at least one non-singular matrix $\boldsymbol{P}$ that simultaneously transforms the $\boldsymbol{A}(\mathrm{g}_i)$ for all g_i to the same block diagonal form $\boldsymbol{M}$ so that

$$\boldsymbol{M}(\mathrm{g}_i) = \boldsymbol{P}^{-1}\boldsymbol{A}(\mathrm{g}_i)\boldsymbol{P} = \begin{pmatrix} \boldsymbol{M}^{(1)}(\mathrm{g}_i) & \boldsymbol{0} & \cdots & & \\ \boldsymbol{0} & \boldsymbol{M}^{(2)}(\mathrm{g}_i) & & \ddots & \\ \vdots & & \ddots & & \vdots \\ & \ddots & & \boldsymbol{M}^{(m-1)}(\mathrm{g}_i) & \boldsymbol{0} \\ & & \cdots & \boldsymbol{0} & \boldsymbol{M}^{(m)}(\mathrm{g}_i) \end{pmatrix}, \tag{3.30}$$

where $\boldsymbol{M}^{(j)}$ are the irreducible representations and where $m \leq n$. The representation $\boldsymbol{M}$ can then be written as a direct sum of its irreducible representations, with

$$\boldsymbol{M} = \boldsymbol{M}^{(1)} \oplus \boldsymbol{M}^{(2)} \oplus \ \ldots \ \oplus \boldsymbol{M}^{(m)}. \tag{3.31}$$

The effect of diagonalising the matrices $\boldsymbol{A}$ through the similarity transformation $\boldsymbol{P}$ is essentially a change of basis from fields in W that are related through the symmetries to a new set of orthogonal fields that are independent through application of the symmetries. This is seen in the example of D_8, where the eight fields $\boldsymbol{b}^{\mathrm{I}}$ to $\boldsymbol{b}_2^{\mathrm{Vb}}$ are linear functions of the fields in (3.11). In that example, $\boldsymbol{P}$ can be inferred from the coefficients of the fields $\mathrm{g}\boldsymbol{b}$ given in (3.12)–(3.15), (3.19) and (3.20). It must be noted, however, that $\boldsymbol{P}$ may be non-unique, meaning that the representations $\boldsymbol{M}$ will also be non-unique.

Starting from the n fields in (3.28), n mutually orthogonal fields are produced. However, as $\boldsymbol{M}$ is not necessarily purely diagonal but is block diagonal in form, these fields may not all be independent under the action of the symmetries and

instead may form pairs, triples and so forth, of fields that the symmetries map to each other. This occurs when the $\boldsymbol{M}^{(j)}$ are multi-dimensional and so symmetry classes have the possibility to contain more than one field. In the D_8 example, this is precisely the case with the four fields in symmetry class V, which form two independent pairs, divided into subclasses Va and Vb. Only abelian groups may form purely diagonal representations $\boldsymbol{M}$; all irreducible representations will be one-dimensional and all symmetry classes contain a single field, with this being the case in the example of $\mathbb{Z}_2$. It is then clear that for non-abelian groups the number of irreducible representations m must be less than the number of symmetries n.

3.3.2 Characters and character tables

The character table provides key information for constructing irreducible representations. The table is constructed from characters of the irreducible representations for each element of a group G in each symmetry class. The characters $\chi_j(\mathsf{g}_i)$ are the traces of the irreducible representation matrices $\boldsymbol{M}^{(j)}(\mathsf{g}_i)$ and are fixed for each group G, as all possible irreducible representations of a transformation can be related through similarity transformations. For example, for a non-unique irrep $\boldsymbol{M}^{(j)}(\mathsf{g}_i)$ with character χ_{g_i}, we could identify another irrep as $\boldsymbol{N}^{(j)}(\mathsf{g}_i) = \boldsymbol{Q}\boldsymbol{M}^{(j)}(\mathsf{g}_i)\boldsymbol{Q}^{-1}$ with $\boldsymbol{Q}$ as the corresponding similarity matrix. The character of this new irrep is given by

$$\chi'_{\mathsf{g}_i} = \mathrm{tr}\left[\boldsymbol{N}^{(j)}(\mathsf{g}_i)\right] = \mathrm{tr}\left[\boldsymbol{Q}\boldsymbol{M}^{(j)}(\mathsf{g}_i)\boldsymbol{Q}^{-1}\right], \tag{3.32}$$

and by the property of the trace being invariant under cyclic permutations, we have that $\mathrm{tr}[\boldsymbol{Q}\boldsymbol{M}^{(j)}(\mathsf{g}_i)\boldsymbol{Q}^{-1}] = \mathrm{tr}[\boldsymbol{Q}^{-1}\boldsymbol{Q}\boldsymbol{M}^{(j)}(\mathsf{g}_i)] = \mathrm{tr}[\boldsymbol{M}^{(j)}(\mathsf{g}_i)]$, which is simply χ_{g_i}. Therefore, the character is invariant under a similarity transformation and so all possible irreducible representations of a transformation g_i have the same character. The rows of a character table correspond to the individual symmetry classes and the columns to the conjugacy classes. Two elements of a group G are conjugate if an equivalence relation can be made between the two, using only elements within

G. Formally, elements g_a and g_b are conjugate if there exists some g_g such that

$$\mathsf{g}_g \mathsf{g}_a \mathsf{g}_g^{-1} = \mathsf{g}_b \tag{3.33}$$

and so the conjugacy class of g_a is simply a set of all the elements g_b that fit the requirement (3.33) and, of course, includes g_a. All irreducible representations of elements g_i in the same conjugacy class have the same character and so are grouped in the tables for compactness. This can be seen in the character table for D_8 (table 3.1), where the symmetries are organised into their respective conjugacy classes, and is easily demonstrated. From (3.33), the character for g_b is simply $\chi_{\mathsf{g}_b} = \mathrm{tr}[\boldsymbol{M}^{(j)}(\mathsf{g}_b)] = \mathrm{tr}[\boldsymbol{M}^{(j)}(\mathsf{g}_g)\boldsymbol{M}^{(j)}(\mathsf{g}_a)\boldsymbol{M}(\mathsf{g}_g^{-1})]$, using the result (3.18). On cycling the elements (as the trace is invariant under this operation) and again using (3.18) to show that $\boldsymbol{M}^{(j)}(\mathsf{g}_g^{-1})\boldsymbol{M}^{(j)}(\mathsf{g}_g) = \boldsymbol{M}^{(j)}(\mathsf{i}) = \boldsymbol{I}$, we see that

$$\chi_{\mathsf{g}_b} = \mathrm{tr}\left[\boldsymbol{M}^{(j)}(\mathsf{g}_a)\right] = \chi_{\mathsf{g}_a}, \tag{3.34}$$

for two conjugate elements g_a and g_b. Of course, for abelian groups, no two transformations g_a and g_b satisfy (3.33) and so each element of the group has its own conjugacy class.

It must be noted that the first conjugacy class is always $E = \{\mathsf{i}\}$; no element is conjugate to the identity. The action of the identity is informative, as the irreducible representation for this trivial transformation is simply the identity matrix $\boldsymbol{I}$. The character (trace) of the identity thus reveals the dimensionality, and so the number of fields, in a given symmetry class. Referring to table 3.1, of the D_8 characters, this clearly reveals four single-field symmetry classes (with one-dimensional irreducible representations), in which the fields $\boldsymbol{b}^{\mathrm{I}}$ to $\boldsymbol{b}^{\mathrm{IV}}$ belong. The characters for these symmetry classes are simply the irreducible representations themselves; the trace of a single value (i.e. a 1×1 matrix) is the value itself. The irreducible representations for these classes can therefore be read straight from the character table, as was done in the example of D_8.

Still referring to the previous example, from table 3.1 it is seen that symmetry

class V is two-dimensional, though there are four fields: $\boldsymbol{b}_1^{Va}, \boldsymbol{b}_2^{Va}, \boldsymbol{b}_1^{Vb}$ and $\boldsymbol{b}_2^{Vb}$. This is because a symmetry class of dimension k will contain k^2 fields, as each class contains k subclasses, each with k fields. Class V (of D_8) must therefore have two subclasses, which we have already identified as Va and Vb, each containing two fields. Subclasses within a symmetry class are isomorphic, as their irreducible representations (which may actually be different matrices) have the same character. An important result from representation theory is

$$\sum_{\alpha} \chi_{\alpha}^2 = n, \tag{3.35}$$

where n is the order of G. A further result of representation theory is that if two groups are isomorphic, i.e. it is possible to find a one-to-one map between the elements of both groups, then the groups will share the same character table. This is not surprising, as above we have shown that under similarity transformations, the characters are unchanged; this identifies that the characters are an invariant property of a group of elements. As character tables are well-known for all the main finite symmetry groups, this feature will be used when constructing the irreducible representations of the 1:1:1 ABC flow.

3.3.3 Action of symmetries on fields of a given symmetry class

In chapter 2, the action of a symmetry on a vector function was given in (2.19), with the transformation rotating the function at the preimage of the space $\boldsymbol{x}$. This allowed the symmetry group for the ABC flows to be generated, by finding all distinct transformations g in (2.19). Sufficiently general vector functions defined in a space which has these symmetries will not be invariant under the same symmetries and so the purpose of finding the irreducible representations is to find independent linear combinations that behave as simply as possible under the action of the symmetries. In a given symmetry class α, the action of a symmetry g is now

described by

$$(\mathsf{g}\boldsymbol{b}^{\alpha})(\boldsymbol{x}) = \boldsymbol{M}^{\alpha}(\mathsf{g})\,\boldsymbol{J}_{\mathsf{g}}\boldsymbol{b}^{\alpha}\left(\mathsf{g}^{-1}\boldsymbol{x}\right). \tag{3.36}$$

For a one-dimensional symmetry class, $\boldsymbol{b}^{\alpha}$ will contain one field and $\boldsymbol{M}^{\alpha}(\mathsf{g})$ will be ± 1. Alternatively, in multi-dimensional symmetry classes, $\boldsymbol{b}^{\alpha}$ will contain multiple fields with $\boldsymbol{M}^{\alpha}$ describing how the symmetry g links these fields. Owing to the linearity of $\boldsymbol{b}$ in the kinematic induction equation, fields in different symmetry classes are independent of one another and so can be studied separately. Importantly, $\boldsymbol{b}$ is still only defined in the fundamental domain, as discussed in chapter 2, and to study the evolution of a field in a particular symmetry class, we need only know the irreducible representations for all symmetries in that class. This idea is expanded in chapter 4, when discussing the numerical methods.

3.4 Irreducible representations of G

Having detailed two simple but visualisable examples and touching on the key concepts of characters and irreducible representations, we now return to the problem in hand and discuss the 1:1:1 ABC flow and its representations. Given a sufficiently general magnetic field $\boldsymbol{b}$ we obtain a 24-dimensional vector space W spanned by the images $\mathsf{g}\boldsymbol{b}$ under the various symmetries. We know from Arnold (1984) that the symmetry group G of the 1:1:1 ABC flow is isomorphic to the rotation group of a cube and so rather than writing down a decomposition of this space explicitly, we start by giving the character table for G (i.e. that of the orientation-preserving symmetries of a cube) in table 3.2. Here the elements are grouped into conjugacy classes with

$$\begin{aligned} E &= \{\mathsf{i}\}, & C_4 &= \{\mathsf{a}, \mathsf{a}^3, \mathsf{b}, \mathsf{b}^3, \mathsf{c}, \mathsf{c}^3\}, & C_4^2 &= \{\mathsf{a}^2, \mathsf{b}^2, \mathsf{c}^2\}, \\ & & C_3 &= \{\mathsf{d}, \mathsf{d}^2, \mathsf{e}, \mathsf{e}^2, \mathsf{f}, \mathsf{f}^2, \mathsf{g}, \mathsf{g}^2\}, & C_2 &= \{\mathsf{h}, \mathsf{j}, \mathsf{k}, \mathsf{l}, \mathsf{m}, \mathsf{n}\}. \end{aligned} \tag{3.37}$$

On investigation of the character table, it can be seen immediately, from the first column, that there are two one-dimensional (symmetry) classes, I and II, one two-

dimensional class, III, and two three-dimensional classes IV and V. In classes I and II, the field is simply multiplied by ± 1 under each symmetry operation and, as before, the values may be read from the character table. Class III contains two pairs of field (each pair belonging to a subclass) and classes IV and V each contain three subclasses, with three fields. As the subclasses in classes III, IV and V are isomorphic (to other subclasses within the class), then we need only consider one subclass of each class. Using the isomorphism between G and O_{24} allows us to

Table 3.2: Character table for the group O_{24}. Entries give the character, that is the trace of the matrix $\boldsymbol{M}^{\alpha}(\mathsf{g})$, for transformations g in the conjugacy classes (3.37) and irreducible representations α from I–V.

	E	C_3	C_4^2	C_2	C_4
I	1	1	1	1	1
II	1	1	1	-1	-1
III	2	-1	2	0	0
IV	3	0	-1	1	-1
V	3	0	-1	-1	1

take advantage of irreducible representations that have already been constructed for these groups. As they are not unique, there is an element of arbitrariness in the choice but a good starting point is to review the example of D_8. We note that the irreducible representations of subclass Va are simply the Jacobians of the symmetries. Taking the Jacobians of each transformation $\mathsf{g} \in G$, we notice that these do indeed form a set of irreducible representations for class V by confirming that the characters match those of the character table, table 3.2. As they are to be used in the next chapter, specific matrices for class V are

$$\boldsymbol{M}^{\mathrm{V}}(\mathsf{b}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{V}}(\mathsf{b}^3) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{V}}(\mathsf{d}) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{V}}(\mathsf{d}^2) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \qquad (3.38)$$

with the full set given in Appendix A.

Arnold (1984) describes class IV as a tensor-product of II and V. Having already identified irreducible representations for both of these classes, we can construct the irreducible representation matrices of class IV by

$$\boldsymbol{M}^{\mathrm{IV}}(\mathsf{g}) = \boldsymbol{M}^{\mathrm{II}}(\mathsf{g}) \otimes \boldsymbol{M}^{\mathrm{V}}(\mathsf{g}). \tag{3.39}$$

From table 3.2, all symmetries in classes C_2 and C_4 have character -1 in class II, so the irreducible representations of class IV will simply be those of class V, except with the opposite signs for those in C_2 and C_4. The same matrices as in (3.38) for class IV are given by

$$\boldsymbol{M}^{\mathrm{IV}}(\mathsf{b}) = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{IV}}(\mathsf{b}^3) = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix},$$
$$\boldsymbol{M}^{\mathrm{IV}}(\mathsf{d}) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{IV}}(\mathsf{d}^2) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \tag{3.40}$$

also with the full set provided in Appendix A.

Somewhat less intuitive is the two-dimensional symmetry class III. Arnold (1984) describes this class as a "permutation of the three coordinate axes", which could be considered as a $2\pi/3$ rotation in the plane of perpendicular to the line $x = y = z$, and thus a good starting point is a rotation matrix with $\theta = 2\pi/3$. This, however, is not enough information to construct irreps for all symmetries in G, so we look to predefined tables, which are prevalent in any scientific field that relies on the geometry of molecules, such as crystallography or condensed matter physics. As the irreps of class III are non-unique, there are a variety of different possibilities, though Nussbaum (1971) (see table 5-10 on page 271) supplies a set of irreps that are real-valued, in keeping with the irreps of our other symmetry classes. Naturally,

these irreps satisfy (3.18), with the useful matrices given as

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{b}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{b}^3) = \frac{1}{2}\begin{pmatrix} 1 & -\sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{d}) = \frac{1}{2}\begin{pmatrix} -1 & \sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{III}}(\mathsf{d}^2) = \frac{1}{2}\begin{pmatrix} -1 & -\sqrt{3} \\ \sqrt{3} & -1 \end{pmatrix}, \tag{3.41}$$

with, again, the full set given in Appendix A.

To demonstrate the use of the irreducible representations, we use the symmetry b. We know from table 3.2 that $\boldsymbol{M}^{\mathrm{I}}(\mathsf{b}) = 1$ and $\boldsymbol{M}^{\mathrm{II}}(\mathsf{b}) = -1$ and so

$$\boldsymbol{b}^{\mathrm{I}} = \mathsf{b}\boldsymbol{b}^{\mathrm{I}}, \qquad \boldsymbol{b}^{\mathrm{II}} = -\mathsf{b}\boldsymbol{b}^{\mathrm{II}}. \tag{3.42}$$

The three-dimensional representations for classes IV and V are taken from (3.40) and (3.38) respectively. It is understood that $\boldsymbol{b}^{\alpha}$ for these two classes must contain three fields each and so the action of the irreducible representation is to map these fields to each other, given by

$$\mathsf{b}\boldsymbol{b}_1^{\mathrm{IV}} = -\boldsymbol{b}_1^{\mathrm{IV}}, \qquad \mathsf{b}\boldsymbol{b}_2^{\mathrm{IV}} = \boldsymbol{b}_3^{\mathrm{IV}}, \qquad \mathsf{b}\boldsymbol{b}_3^{\mathrm{IV}} = -\boldsymbol{b}_2^{\mathrm{IV}}, \tag{3.43}$$

and

$$\mathsf{b}\boldsymbol{b}_1^{\mathrm{V}} = \boldsymbol{b}_1^{\mathrm{V}}, \qquad \mathsf{b}\boldsymbol{b}_2^{\mathrm{V}} = -\boldsymbol{b}_3^{\mathrm{V}}, \qquad \mathsf{b}\boldsymbol{b}_3^{\mathrm{V}} = \boldsymbol{b}_2^{\mathrm{V}}. \tag{3.44}$$

For class III, the picture is a little more complicated but we can proceed in the same fashion. There are two fields in this class and from (3.41), we have that

$$2\mathsf{b}\boldsymbol{b}_1^{\mathrm{III}} = \boldsymbol{b}_1^{\mathrm{III}} - \sqrt{3}\boldsymbol{b}_2^{\mathrm{III}}, \qquad 2\mathsf{b}\boldsymbol{b}_2^{\mathrm{III}} = -\sqrt{3}\boldsymbol{b}_1^{\mathrm{III}} - \boldsymbol{b}_2^{\mathrm{III}}, \tag{3.45}$$

so when transformed, $\boldsymbol{b}_1^{\mathrm{III}}$ and $\boldsymbol{b}_2^{\mathrm{III}}$ each form a linear combination of themselves and the other.

The 24-dimensional space W will not be decomposed explicitly, as we did for D_8 in section 3.2, owing to the large size of the group G. This, however, is not strictly important as all the information required to investigate the individual symmetry classes is contained in the irreducible representations, which are known. The decomposition of the previous example was, in part, to aid visualising the types of structure in each class. For the 1:1:1 ABC flow, $\boldsymbol{b}$ is three-dimensional, making it impractical to visualise in the same fashion. It should also be noted that the decomposition of W into symmetry classes satisfies (3.35), as it contains one copy of representations I and II, two copies of III, and three copies of each of IV and V, with these spanning W; $1^2 + 1^2 + 2^2 + 3^2 + 3^2 = 24$, the order of G as required.

3.5 Summary

In these examples, the idea of defining single fields or vector spaces that are independent under the symmetries is key. To summarise the process, we first define a general field $\boldsymbol{b}$ and apply all symmetries in the relevant group, to produce a vector space W of fields. These fields are all related through various symmetries; that is, applying a symmetry g to this set of fields is equivalent to the action of a permutation matrix. This permutation matrix is a matrix representation of the action of g, with a specific representation for each g. The next step is to construct new fields from linear combinations of the fields in W, with the intention that W can be decomposed into as many independent subspaces as possible. In terms of representations, it involves finding an equivalent representation $\boldsymbol{M}(\mathsf{g})$ for each g and with $\boldsymbol{M}$ having exactly the same block diagonal form for every g. The action of any g can then only relate fields of each block to one another, resulting in a series of fields $\boldsymbol{b}^\alpha$ that evolve independently, provided that the evolution equation is linear (as is the kinematic induction equation).

To demonstrate how the symmetries of the 1:1:1 ABC flow apply to a general magnetic field $\boldsymbol{b}$, the process described above was applied to the group G and the irreducible representations were constructed. As the symmetry groups of both $\mathbb{Z}_2$

and D_8 are smaller and defined in a two-dimensional vector space, it is possible to find a suitable equivalence matrix $\boldsymbol{P}$ to identify the explicit decompositions used. Although this is strictly unnecessary, it is advantageous as it allows for visualising the fields belonging to the different symmetry classes. Ultimately, only the irreducible representations are required.

For the 1:1:1 ABC symmetries, it is not practical to visualise all 24 fields and so it is not necessary to know $\boldsymbol{P}$. Using the fact that G is isomorphic to the well-studied cubic rotation group O_{24}, information such as the character table, table 3.2, is revealed. Combined with what is already known of this group (see Arnold, 1984), it is not too laborious to construct a full set of irreducible representations for the five symmetry classes. The first two symmetry classes (I, II) are one-dimensional; they each contain one field and their representations can be read from the character table for G. The third class (III) is two-dimensional and the fourth and fifth (IV, V) are three-dimensional. Classes III, IV and V contain multiple subclasses but only one of each needs to be considered.

In the next chapter, we detail how the symmetries and irreducible representations are actually utilised in the numerical framework.

Chapter 4

Numerical Methods

Build a man a fire, and he'll be warm for a day. Set a man on fire, and he'll be warm for the rest of his life.

Terry Pratchett

Now that the necessary theory of symmetries and representations has been developed and the irreducible representations have been constructed, we may return to the problem at hand; that of solving the induction equation whilst utilising the various symmetries of the flow. We saw from the last chapter that five invariant subspaces exist, which we named symmetry classes, with each symmetry class behaving differently under the transformations of G. The main purpose of this chapter is to show how this affects the numerical procedures, and how only a small subset of Fourier modes need be used.

The chapter begins by describing the spectral numerical scheme for the full 1:1:1 ABC dynamo problem, with a brief discussion of parameters for the numerical analysis. Following this, a reduction of the full domain of Fourier modes to the fundamental domains $\mathcal{G}_N$ and $\mathcal{H}_N$ is detailed. Both domains use all symmetries

in G but the first does not utilise the Hermitian symmetry of the field, whereas the second does. The additional computations required to time step in both $\mathcal{G}_N$ and $\mathcal{H}_N$ are provided in full. The process of parallelising the numerical problem in $\mathcal{H}_N$ is explained, with details on balancing the work load and additional operations involving the order of communication between consecutive cores. Finally, the Implicitly Restarted Arnoldi iteration method (IRAM) is explored as a potential alternative to time stepping, with comparisons drawn between these two methods.

4.1 Time stepping in the ABC dynamo problem

In this section, we identify the numerical methods of time stepping the full ABC dynamo problem, with these methods also used in the symmetry-reduced problems. Our interest is in solving the kinematic induction equation

$$\partial_t \boldsymbol{b} + \boldsymbol{u} \cdot \nabla \boldsymbol{b} = \boldsymbol{b} \cdot \nabla \boldsymbol{u} + \eta \nabla^2 \boldsymbol{b}, \tag{4.1}$$

for the 1:1:1 ABC flow, given by

$$\boldsymbol{u} = (\sin z + \cos y, \sin x + \cos z, \sin y + \cos x). \tag{4.2}$$

We wish to solve this as an eigenvalue problem with

$$\boldsymbol{b} = \tilde{\boldsymbol{b}}(x, y, z)\mathrm{e}^{\lambda t}, \tag{4.3}$$

as a consequence of the linearity of the problem. The eigenvalues of the system are denoted (as typical) by λ. In solving this problem, we aim to obtain growth rates (and frequencies) of the magnetic field from the dominant eigenvalues, as well as the structure of these fields in physical space, along with any other useful analytical tools. We consider magnetic fields with the usual Fourier decomposition

$$\boldsymbol{b} = \sum_{l,m,n} \boldsymbol{b}_{l,m,n}(t)\, \mathrm{e}^{\mathrm{i}lx + \mathrm{i}my + \mathrm{i}nz}, \qquad \boldsymbol{b}_{l,m,n} = (X_{l,m,n}, Y_{l,m,n}, Z_{l,m,n}), \tag{4.4}$$

with wavenumbers l, m and n. Strictly, only two components of the field are necessary, as the third may be reconstructed from the other two via the divergence-free condition $\nabla \cdot \boldsymbol{b} = 0$. In the time stepping codes, this property is used to remove divergence at set intervals with a projection, though all three components $X_{l,m,n}, Y_{l,m,n}$ and $Z_{l,m,n}$ are resolved fully at every time step. In the Arnoldi iteration method, this property is used to remove degrees of freedom and is discussed in section 4.4.3.

As the magnetic field is real, we also have the Hermitian symmetry property,

$$\boldsymbol{b}_{l,m,n} = \boldsymbol{b}^*_{-l,-m,-n}, \tag{4.5}$$

where * denotes the complex conjugate. Numerically the values of l, m and n are limited by the maximum resolution N and the zero mode (which would correspond to a constant mean field) is always set to zero. The full cube of wave-vectors (with truncation at N for each index) is referred to as $\mathcal{F}_N$ and is defined by

$$\mathcal{F}_N = \{(l, m, n) : |l|, |m|, |n| \leq N\}. \tag{4.6}$$

4.1.1 Numerical scheme

In the investigation of Galloway and Frisch (1986) it is noted that the ABC flows are particularly suited for spectral methods (those that operate in Fourier space), as the terms involving $\boldsymbol{u}$ and $\boldsymbol{b}$ result in neighbouring modes being coupled; i.e. to time step a Fourier mode, only its six neighbouring modes are needed. We are therefore adopting a purely spectral code for the time stepping procedure. For procedures that require the magnetic field to be in real space, for example in visualising these fields, a Fast Fourier Transform is adopted; more detail is given later in the chapter.

The advective and stretching terms of (4.1) can be solved through the use of explicit methods. For these terms, a three-step Adams–Bashforth scheme is adopted as a

balance between accuracy and speed. Initially a two-step scheme was used but found to be unstable in certain conditions. The diffusive terms are integrated exactly and give exponential damping of each mode every time step. This scheme can be written analytically, with time $t = j\Delta t$, as

$$\boldsymbol{b}^{j+1}_{l,m,n} = E_{l,m,n}\left[\boldsymbol{b}^{j}_{l,m,n} + \tfrac{1}{12}\Delta t\Big(23\boldsymbol{\Lambda}^{j}_{l,m,n} - 16E_{l,m,n}\boldsymbol{\Lambda}^{j-1}_{l,m,n} + 5E^{2}_{l,m,n}\boldsymbol{\Lambda}^{j-2}_{l,m,n}\Big)\right], \quad (4.7)$$

where

$$E_{l,m,n} = \exp\left[-\eta\left(l^2 + m^2 + n^2\right)\Delta t\right]. \quad (4.8)$$

The advective and stretching terms are represented by $\boldsymbol{\Lambda}^{j}_{l,m,n} = \left[(\Lambda_x)^{j}_{l,m,n}, (\Lambda_y)^{j}_{l,m,n}, (\Lambda_z)^{j}_{l,m,n}\right]$ at time step j, with corresponding x, y and z components. The x-component is given by

$$\begin{aligned}
2(\Lambda_x)^{j}_{l,m,n} &= \mathrm{i}mB\left(Y^{j}_{l,m-1,n} + Y^{j}_{l,m+1,n}\right) - \mathrm{i}mC\left(X^{j}_{l,m,n-1} + X^{j}_{l,m,n+1}\right)\\
&+ mC\left(Y^{j}_{l,m,n-1} - Y^{j}_{l,m,n+1}\right) - mA\left(X^{j}_{l-1,m,n} - X^{j}_{l+1,m,n}\right)\\
&+ \mathrm{i}nB\left(Z^{j}_{l,m-1,n} + Z^{j}_{l,m+1,n}\right) - \mathrm{i}nA\left(X^{j}_{l-1,m,n} + X^{j}_{l+1,m,n}\right)\\
&+ nC\left(Z^{j}_{l,m,n-1} - Z^{j}_{l,m,n+1}\right) - nB\left(X^{j}_{l,m-1,n} - X^{j}_{l,m+1,n}\right), \qquad (4.9)
\end{aligned}$$

with the y-component as

$$\begin{aligned}
2(\Lambda_y)^{j}_{l,m,n} &= \mathrm{i}lC\left(X^{j}_{l,m,n-1} + X^{j}_{l,m,n+1}\right) - \mathrm{i}lB\left(Y^{j}_{l,m-1,n} + Y^{j}_{l,m+1,n}\right)\\
&+ lA\left(X^{j}_{l-1,m,n} - X^{j}_{l+1,m,n}\right) - lC\left(Y^{j}_{l,m,n-1} - Y^{j}_{l,m,n+1}\right)\\
&+ \mathrm{i}nC\left(Z^{j}_{l,m,n-1} + Z^{j}_{l,m,n+1}\right) - \mathrm{i}nA\left(Y^{j}_{l-1,m,n} + Y^{j}_{l+1,m,n}\right)\\
&+ nA\left(Z^{j}_{l-1,m,n} - Z^{j}_{l+1,m,n}\right) - nB\left(Y^{j}_{l,m-1,n} - Y^{j}_{l,m+1,n}\right), \qquad (4.10)
\end{aligned}$$

and the z-component given by

$$\begin{aligned}2(\Lambda_z)^j_{l,m,n} = &\mathrm{i}lA\left(X^j_{l-1,m,n} + X^j_{l+1,m,n}\right) - \mathrm{i}lB\left(Z^j_{l,m-1,n} + Z^j_{l,m+1,n}\right)\\ &+ lB\left(X^j_{l,m-1,n} - X^j_{l,m+1,n}\right) - lC\left(Z^j_{l,m,n-1} - Z^j_{l,m,n+1}\right)\\ &+ \mathrm{i}mA\left(Y^j_{l-1,m,n} + Y^j_{l+1,m,n}\right) - \mathrm{i}mC\left(Z^j_{l,m,n-1} + Z^j_{l,m,n+1}\right)\\ &+ mB\left(Y^j_{l,m-1,n} - Y^j_{l,m+1,n}\right) - mA\left(Z^j_{l-1,m,n} - Z^j_{l+1,m,n}\right),\end{aligned} \tag{4.11}$$

for a general ABC flow; for the 1:1:1 flow, we set $A = B = C = 1$. In the full code, this numerical scheme was used to time step all modes in $\mathcal{F}_N$ by varying l, m and n from their minimum value of $-N$ to their maximum value of N. This code was tested against results of Arnold and Korkina (1983) for the first window and against Galloway and Frisch (1986) for both the first window and up to $R_m = 50$ in the second window, with a (randomly generated) divergence-free seed field. The results were found to agree well (to within three significant figures) with those of Arnold and Korkina (1983) and agreed to within $\sim 2\%$ with those of Galloway and Frisch (1986); this discrepancy has been also found in other investigations (see Lau and Finn, 1993; Bouya and Dormy, 2013).

For testing purposes, Hermitian symmetry was implemented in the full code, as was done by Galloway and Frisch (1986). The numerical scheme is unchanged but only Fourier modes with $l \geq 0$ were time stepped. As the $l = 0$ modes did not have access to all neighbouring modes, specifically those with $l = -1$, an additional "layer" was stored but not time stepped. This layer ($\boldsymbol{b}_{-1,m,n}$) was updated using the Hermitian symmetry (4.5) of the magnetic field so that

$$\boldsymbol{b}_{-1,m,n} = \boldsymbol{b}^*_{1,-m,-n}. \tag{4.12}$$

The Hermitian code was tested against results of the non-Hermitian code and agreed to within rounding error. This method of updating mode values, without stepping them numerically, is core to the reduction of the problem through the symmetries and is discussed in the next section.

4.1.2 Numerical operations and parameters

As the numerical scheme has been detailed, we provide information about codes in which these schemes are realised. It should first be noted that all numerical codes were written in Fortran, unless otherwise specified, as it is the most appropriate choice for managing huge calculations involving arrays, owing to its efficiency (with these types of calculation) and ease of programming, as well as good portability.

One aim of the various investigations in this thesis is to study the evolution of magnetic fields in high-R_m regimes (or instabilities in high-R_e regimes for the fluid stability problem). Increasing R_m means that the diffusive terms need to be resolved for finer and finer resolutions. Galloway and Frisch (1986) identified the relationship between resolution N and R_m: N scales as $R_m^{1/2}$. In practice, the number of Fourier modes in the full domain $\mathcal{F}_N$ is $(2N+1)^3$ and so the number of modes to resolve scales as $R_m^{3/2}$. In the high-R_m regimes, this becomes problematic on account of the size of the numerical problems. Also for consideration is the temporal resolution, i.e. the time step size, though this is more dependent on the numerical scheme adopted. Through testing the full and reduced codes (see section 4.2), 'safe' values of time step size and resolution for ranges of R_m were found; typical parameters are given in table 4.1. In practice, more divisions were made than the ones displayed but it is impractical to note them all here. Safe values of the spatial

Table 4.1: Typical parameters adopted for numerical calculations.

R_m	N	Δt	R_m	N	Δt
5	6	0.01	500	40	0.002
25	10	0.01	1000	54	0.002
120	20	0.005	5000	122	0.001
250	28	0.005	10000	170	0.001

resolution parameter N were found to scale roughly as $1.7R_m^{1/2}$ for all values of R_m investigated. The size of time steps, however, was more difficult to determine and involved trial-and-error to find appropriate values, though it was easier to identify when the time steps were too large, owing to unrealistically large growth rates (i.e. numerical 'blow-up'). In testing, these two parameters were deemed of appropriate

sizes (safe) when growth rates had converged to three significant figures; increasing the resolution and decreasing the time step size further did not affect the growth rate at this precision.

In locating safe parameters, we also tested for conservation of zero-divergence in the magnetic field for long run times. We find that in using divergence-free random initial conditions, the field remains divergence-free for the entire length of the simulation without the need for correction. Nevertheless, all time stepping codes include a subroutine to remove divergence through a projection, with a call to this subroutine placed several times per simulation at equally spaced intervals.

An important point to consider is what occurs at the edge of the truncated domain $\mathcal{F}_N$, i.e. for a mode which has one or more index equal to N or $-N$. Of course, these modes do not have all neighbouring modes and so cannot be time stepped. In the simulations carried out for this thesis, all simulation domains are truncated at index N; modes up to index $N-1$ are time stepped, modes with at least one index equal to N are set to zero. This intentionally damps the effect of neighbouring modes that correspond to smaller scale fields than those we aim to resolve.

The two main outputs of these simulations are the total magnetic energy E_{M} and the magnetic field 'snapshots'. To calculate E_{M}, we simply sum the squares of all modes; for the reduced problems, only the modes in the fundamental domain are used. From (4.3), we see that the magnetic energy grows or decays exponentially and so the natural logarithm is also taken. At every time step, a new value of $\ln(E_{\mathrm{M}})$ is calculated so that at the end of the simulation, a time series is obtained. From this time series, the exponential growth rate of magnetic energy is calculated from the linear trend in $\ln(E_{\mathrm{M}})$. This provides the real part of the eigenvalue λ of this mode. For the imaginary part, the period of the oscillations in $\ln(E_{\mathrm{M}})$ is used. Each period in $\ln(E_{\mathrm{M}})$ corresponds to a half-period in the magnetic field evolution, and we calculate $\mathrm{Im}(\lambda)$ as $2T/2\pi = T/\pi$, where T is the period of a single cycle of $\ln(E_{\mathrm{M}})$. The snapshots are essentially records of the magnetic field at a given point in the simulation. These are used for visualisations and spectral analysis, amongst other things. To visualise, we convert from spectral space to real space using the

FFTW routine supplied by NAG. For 3-D visualisations, a script written in IDL then converts it to a format that can be read by VAPOR (see Clyne *et al.*, 2007). For 2-D visuals (cross-sections), IDL's visualisation tools were used.

Returning to the dynamo code, the order of operations in each time step (for the full code) is as follows:

- If Hermitian symmetry is included, calculate modes $\boldsymbol{b}_{-1,m,n}$ from $\boldsymbol{b}_{1,-m,-n}$ using (4.5),

- loop over all l, m and n ($|l| < N$ for non-Hermitian case, $0 < l < N$ for Hermitian); for each unique (l, m, n) calculate (4.9)–(4.11), then step (4.7),

- copy old arrays to newer arrays, i.e. arrays storing $\boldsymbol{\Lambda}$ for previous steps moved into arrays for one step older, in preparation for next step,

- $\ln(E_{\mathrm{M}})$ calculated for relevant domain; loop sums squares of modes for all mode indices (l, m, n) and stores in predefined array,

- magnetic field snapshot is printed if at required (pre-determined) time step,

- Finally, $\ln(E_{\mathrm{M}})$ is compared against smallest/largest allowed values due to sensitivity to rounding error. If $\ln(E_{\mathrm{M}})$ is too small, all arrays $\boldsymbol{b}^j$, $\boldsymbol{\Lambda}^j, \boldsymbol{\Lambda}^{j-1}$ and $\boldsymbol{\Lambda}^{j-2}$ are scaled up by a predefined value. Similarly if too large, all array values are scaled down.

The process detailed above is valid for full numerical codes (i.e. those without symmetry) and the symmetry-reduced codes. It must be noted, however, that for the reduced codes, the first step is replaced by an equivalent step which impose the symmetries of a particular symmetry class. The procedures are described in the next section.

4.2 Symmetry-reduced problem

The main purpose of this thesis is to take advantage of the intrinsic symmetries of the 1:1:1 ABC flow, with the intent being two fold. Firstly, through identifying the representations of the group of symmetries G, we have found that there are five different types of magnetic field structure that can exist, with every possible magnetic field described by linear combinations of these. It is interesting to investigate not only the dominant mechanisms but also understand the less effective structures that may still be amenable to dynamo action. Secondly, decomposing a field into its constituent symmetry classes allows for reduced computational cost; to numerically solve the 1:1:1 ABC dynamo for a field belonging to a particular symmetry class, we need only know, at any given time, the field contained within the fundamental domain.

There are two reduced problems for consideration. The first is that of solving the dynamo problem using all symmetries of G but neglecting Hermitian symmetry, and this case is explored for use in the eigenvalue solver using the Implicit Restarted Arnoldi method. This method is only applicable to problems in which the evolution is linear and in using Hermitian symmetry, the action of conjugating modes destroys this linearity and so is not appropriate. Solving via the Arnoldi method (IRAM) is discussed in section 4.4. The second is solving the dynamo problem using all symmetries of G and including Hermitian symmetry, with this being the method of choice for time stepping.

In both of these cases, the numerical schemes adopted for the full problem are appropriate. We simply change the range of spectral modes that are time stepped in keeping with the definition of the fundamental domain.

4.2.1 Action of symmetries in Fourier space

First we consider the symmetry-reduced problem in the absence of Hermitian symmetry. From chapters 2 and 3, we know that the action of the symmetries maps

a magnetic field confined to the fundamental domain to the full domain. In spectral space, the effect is similar, though needs some modifications. We recall from chapter 2 that a symmetry acting on a vector function can be described by

$$\mathsf{g}^* : S \to S,\ (\mathsf{g}^* \boldsymbol{f})(\boldsymbol{x}) = \boldsymbol{J}_\mathsf{g} \cdot \boldsymbol{f}(\mathsf{g}^{-1}\boldsymbol{x}), \tag{4.13}$$

where we subsequently drop the * notation for applications of these transformations on the magnetic field. Since we work with the magnetic field written in the Fourier space decomposition (4.4) (with $|l|, |m|, |n| \leq N$), it is useful to understand how the symmetries operate on the Fourier modes themselves. Every symmetry in (2.50), except for d and d^2, is a composition of a rotation and a translation. The rotation can be described by the action of the Jacobian matrix on $\boldsymbol{x}$ and the translation is simply a vector. A general symmetry can therefore be written as

$$\mathsf{g}(\boldsymbol{x}) = \boldsymbol{J}_\mathsf{g}\boldsymbol{x} + \boldsymbol{T}_\mathsf{g}, \tag{4.14}$$

with $\boldsymbol{T}_\mathsf{g}$ being the translation (vector) component of g. We can then apply g to a magnetic field defined by (4.4), giving

$$(\mathsf{g}\boldsymbol{b}) = \sum_{l,m,n} \boldsymbol{J}_\mathsf{g}\boldsymbol{b}_{l,m,n} \exp\big[\mathrm{i}\boldsymbol{k} \cdot (\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{x} + \boldsymbol{T}_{\mathsf{g}^{-1}})\big], \tag{4.15}$$

with $\boldsymbol{k} = (l, m, n)$. The translation and the Jacobian in the exponential term can be separated to get

$$(\mathsf{g}\boldsymbol{b}) = \sum_{l,m,n} \exp\big[\mathrm{i}\boldsymbol{k} \cdot \boldsymbol{T}_{\mathsf{g}^{-1}}\big]\, \boldsymbol{J}_\mathsf{g}\boldsymbol{b}_{l,m,n}\, \exp\big[\mathrm{i}\boldsymbol{k} \cdot (\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{x})\big], \tag{4.16}$$

and now we wish to return the exponential term to the form $\exp[\mathrm{i}\boldsymbol{k} \cdot \boldsymbol{x}]$; we do this by setting $\boldsymbol{k} = \boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k}' = \boldsymbol{J}_{\mathsf{g}^{-1}}(l', m', n')^\intercal$, and then (because all indices l, m and n on the right-hand side are dashed) we can drop the dashes, giving

$$(\mathsf{g}\boldsymbol{b}) = \sum_{l,m,n} \exp\big[\mathrm{i}\,(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k}) \cdot \boldsymbol{T}_{\mathsf{g}^{-1}}\big]\, \boldsymbol{J}_\mathsf{g}\boldsymbol{b}_{(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k})}\, \exp\big[\mathrm{i}\boldsymbol{k} \cdot \boldsymbol{x}\big], \tag{4.17}$$

which is the form we require. Comparing like-for-like modes on both sides, we have

$$(\mathsf{g}\boldsymbol{b})_{l,m,n} = \exp\left[\mathrm{i}\,(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k})\cdot\boldsymbol{T}_{\mathsf{g}^{-1}}\right]\boldsymbol{J}_{\mathsf{g}}\boldsymbol{b}_{\left(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k}\right)}. \tag{4.18}$$

The action of g on a mode with wavenumber $\boldsymbol{k} = (l, m, n)$ is to map it to another mode with wavenumber $\boldsymbol{J}_{\mathsf{g}^{-1}}(l, m, n)^\intercal$ with a phase shift dependent on the translation component of g. This becomes clear when an actual element of G is applied, for example the (l, m, n) modes of $\mathsf{b}\boldsymbol{b}$ and of $\mathsf{d}\boldsymbol{b}$ are given by

$$(\mathsf{b}\boldsymbol{b})_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\boldsymbol{b}_{l,n,-m}, \qquad (\mathsf{d}\boldsymbol{b})_{l,m,n} = \boldsymbol{J}_{\mathsf{d}}\,\boldsymbol{b}_{m,n,l}, \tag{4.19}$$

while for their inverses,

$$(\mathsf{b}^3\boldsymbol{b})_{l,m,n} = \mathrm{e}^{\mathrm{i}(-l+m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}^3}\,\boldsymbol{b}_{l,-n,m}, \qquad (\mathsf{d}^2\boldsymbol{b})_{l,m,n} = \boldsymbol{J}_{\mathsf{d}^2}\,\boldsymbol{b}_{n,l,m}. \tag{4.20}$$

The 24 symmetries in G each map the general mode $\boldsymbol{b}_{l,m,n}$ to a different mode, so through the use of the symmetries, we need only consider a limited range of modes, i.e. those in the fundamental domain. However, the effects of applying g to $\boldsymbol{b}$ still need to be dealt with; evaluating the left-hand sides of (4.19) and (4.20) requires specifying the symmetry class, as the behaviour of the transformation is described by its irreducible representation.

With the trivial remark $\boldsymbol{b} = \mathsf{g}^{-1}(\mathsf{g}\boldsymbol{b})$, it follows that $\boldsymbol{b}^\alpha = \boldsymbol{M}^\alpha(\mathsf{g}^{-1})(\boldsymbol{M}^\alpha(\mathsf{g})\boldsymbol{b}^\alpha)$. We can therefore modify (4.18) to account for the symmetry classes, yielding the result

$$\boldsymbol{b}^\alpha_{l,m,n} = \exp\left[\mathrm{i}\,(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k})\cdot\boldsymbol{T}_{\mathsf{g}^{-1}}\right]\boldsymbol{M}^\alpha(\mathsf{g}^{-1})\,\boldsymbol{J}_{\mathsf{g}}\boldsymbol{b}^\alpha_{\left(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k}\right)}. \tag{4.21}$$

We note how the Jacobian rearranges the components of $\boldsymbol{b}^\alpha_{\left(\boldsymbol{J}_{\mathsf{g}^{-1}}\boldsymbol{k}\right)}$ before the individual fields are mapped around by the irreducible representation of the *inverse* of the transformation g. To demonstrate the result (4.21), we take the example of $\mathsf{b}\boldsymbol{b}$ and obtain

$$\boldsymbol{b}^{\mathrm{I}}_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\boldsymbol{b}^{\mathrm{I}}_{l,n,-m}, \qquad \boldsymbol{b}^{\mathrm{II}}_{l,m,n} = -\mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\boldsymbol{b}^{\mathrm{II}}_{l,n,-m}, \tag{4.22}$$

for classes I and II. Taking $\boldsymbol{M}^{\alpha}(\mathsf{b}^3)$ from (3.40) and (3.38) for classes IV and V respectively, we have

$$\left(\boldsymbol{b}_1^{\mathrm{IV}}\right)_{l,m,n} = -\mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_1^{\mathrm{IV}}\right)_{l,n,-m}, \quad \left(\boldsymbol{b}_1^{\mathrm{V}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_1^{\mathrm{V}}\right)_{l,n,-m}, \tag{4.23}$$

$$\left(\boldsymbol{b}_2^{\mathrm{IV}}\right)_{l,m,n} = -\mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_3^{\mathrm{IV}}\right)_{l,n,-m}, \quad \left(\boldsymbol{b}_2^{\mathrm{V}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_3^{\mathrm{V}}\right)_{l,n,-m}, \tag{4.24}$$

$$\left(\boldsymbol{b}_3^{\mathrm{IV}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_2^{\mathrm{IV}}\right)_{l,n,-m}, \qquad \left(\boldsymbol{b}_3^{\mathrm{V}}\right)_{l,m,n} = -\mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left(\boldsymbol{b}_2^{\mathrm{V}}\right)_{l,n,-m}, \tag{4.25}$$

and for class III,

$$2\left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left[\left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,n,-m} - \sqrt{3}\left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,n,-m}\right], \tag{4.26}$$

$$2\left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\left[-\sqrt{3}\left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,n,-m} - \left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,n,-m}\right]. \tag{4.27}$$

Finally, we should note that these equations have not been fully decomposed; if we were to do so, we would have an equation for each of $(X_i^{\alpha})_{l,m,n}$, $(Y_i^{\alpha})_{l,m,n}$ and $(Z_i^{\alpha})_{l,m,n}$ in every $(\boldsymbol{b}_i^{\alpha})_{l,m,n}$. For example, the equation for $\left(\boldsymbol{b}_2^{\mathrm{V}}\right)_{l,m,n}$ in (4.24) is actually three equations, given by

$$\left(X_2^{\mathrm{V}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\left(X_3^{\mathrm{V}}\right)_{l,n,-m}, \tag{4.28}$$

$$\left(Y_2^{\mathrm{V}}\right)_{l,m,n} = -\mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\left(Z_3^{\mathrm{V}}\right)_{l,n,-m}, \tag{4.29}$$

$$\left(Z_2^{\mathrm{V}}\right)_{l,m,n} = \mathrm{e}^{\mathrm{i}(l-m-n)\pi/2}\,\left(Y_3^{\mathrm{V}}\right)_{l,n,-m}. \tag{4.30}$$

In subsequent use of the symmetries, the full decomposition will not be displayed.

4.2.2 $\mathcal{G}_N$: fundamental domain without Hermitian symmetry

In the absence of Hermitian symmetry, we consider a magnetic field defined in the fundamental domain constructed through the elements of G. As shown above, the

23 non-identity transformations each map a single Fourier mode with wavenumber (l, m, n) to a distinct mode. Knowing the amplitude of a general mode $\boldsymbol{b}_{l,m,n}$ thus tells us the amplitude of 24 Fourier modes in total. The fundamental domain must be 1/24 of the full domain $\mathcal{F}_N$ and so here we define it as

$$\mathcal{G}_N = \{(l, m, n) : 1 \leq l \leq N,\ 0 \leq m \leq l,\ 0 \leq n \leq l\}. \tag{4.31}$$

Although defining the fundamental domain is rather arbitrary, in that any set of points from which $\mathcal{F}_N$ can be reconstructed using the symmetries can be called the fundamental set, the definition (4.31) seems logical and more appropriate in the numerical analysis than others. $\mathcal{G}_N$ is visualised in figure 4.1. To numerically solve

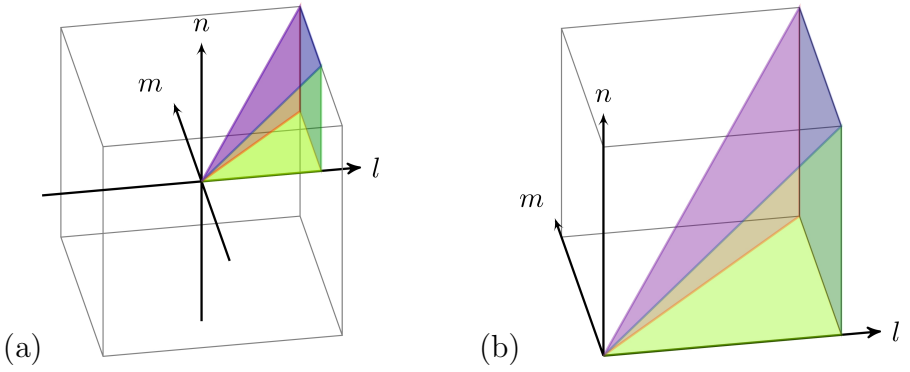

Figure 4.1: Fundamental domain $\mathcal{G}_N$ (a) within the full cube of wave vectors $\mathcal{F}_N$ and (b) within the first octant.

the induction equation in $\mathcal{G}_N$, we need to make sure that every Fourier mode has all six neighbouring modes; to time step $\boldsymbol{b}_{l,m,n}$, we need to know $\boldsymbol{b}_{l\pm1,m,n}, \boldsymbol{b}_{l,m\pm1,n}$ and $\boldsymbol{b}_{l,m,n\pm1}$ as well as $\boldsymbol{b}_{l,m,n}$ itself. We can immediately see that modes on the 'edge' of $\mathcal{G}_N$ will have neighbouring modes that lie outside of $\mathcal{G}_N$. To resolve this, we store an extra layer of modes around $\mathcal{G}_N$ so that all necessary modes for time stepping are available. As only modes within $\mathcal{G}_N$ are updated each time step, the values of these 'outlying' modes will need to be calculated from modes within $\mathcal{G}_N$. To help identify these edge modes, a cross-section of constant l has been visualised in figure 4.2.

There are five 'sets' of modes whose values need to be calculated. These are:

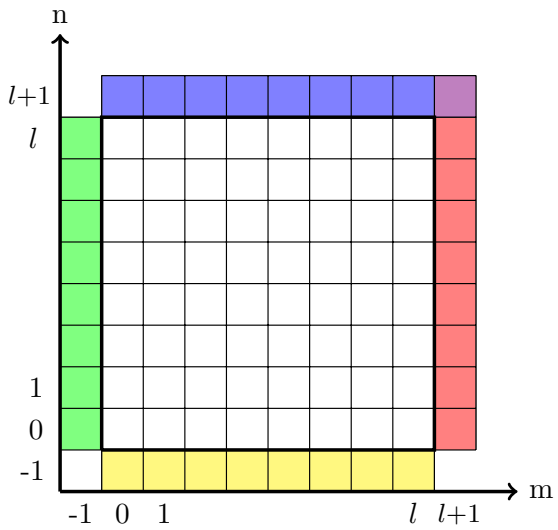

Figure 4.2: A cross-section of $\mathcal{G}_N$ for unspecified l. Unshaded modes are time stepped, coloured modes need to be calculated using symmetries; different colours identify different symmetries used.

- $\boldsymbol{b}_{l,m,-1}$, coloured yellow in figure 4.2, needed to time step $\boldsymbol{b}_{l,m,0}$;
- $\boldsymbol{b}_{l,-1,n}$, coloured green, needed to step $\boldsymbol{b}_{l,0,n}$;
- $\boldsymbol{b}_{l,m,l+1}$, coloured blue, needed to step $\boldsymbol{b}_{l,m,l}$;
- $\boldsymbol{b}_{l,l+1,n}$, coloured red, needed to step $\boldsymbol{b}_{l,l,n}$;
- $\boldsymbol{b}_{l,l+1,l+1}$, coloured purple, needed to step $\boldsymbol{b}_{l+1,l+1,l+1}$.

Note that it is easy to miss the requirement that modes $\boldsymbol{b}_{l,l+1,l+1}$ are filled; these are needed to step modes $\boldsymbol{b}_{l+1,l+1,l+1}$ which cannot be seen in a single cross-section, such as in figure 4.2. Through appropriate choices of (l, m, n), the five sets of modes listed above are mapped, using the relevant symmetries, to modes within $\mathcal{G}_N$. For the second of these, $\boldsymbol{b}_{l,-1,n}$, we have already dealt with the relevant symmetry, namely b, which calculates $\boldsymbol{b}_{l,m,n}$ from $\boldsymbol{b}_{l,n,-m}$ (see (4.19)) and so is used find the value of $\boldsymbol{b}_{l,-1,n}$ from $\boldsymbol{b}_{l,n,1}$. Additionally, we require the symmetries b^3, d and d^2 for the other modes. The transformations involve specific choices of m and n for

(4.19) and (4.20), and are summarised by

$$
\begin{aligned}
(\mathsf{b}\boldsymbol{b})_{l,-1,n} &= \mathrm{e}^{\mathrm{i}(l+1-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\boldsymbol{b}_{l,n,1}, & (\mathsf{d}\boldsymbol{b})_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\,\boldsymbol{b}_{l+1,n,l},\\
\left(\mathsf{b}^3\boldsymbol{b}\right)_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-l+m+1)\pi/2}\,\boldsymbol{J}_{\mathsf{b}^3}\,\boldsymbol{b}_{l,1,m}, & \left(\mathsf{d}^2\boldsymbol{b}\right)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\,\boldsymbol{b}_{l+1,l,m},
\end{aligned}
\tag{4.32}
$$

with the modes on the right-hand sides always belonging to $\mathcal{G}_N$. It should be noted that the purple modes, $\boldsymbol{b}_{l,l+1,l+1}$ can be calculated using either d or d^2.

For classes I and II, (4.32) becomes

$$
\begin{aligned}
\boldsymbol{b}^{\alpha}_{l,-1,n} &= \pm\mathrm{e}^{\mathrm{i}(l+1-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\boldsymbol{b}^{\alpha}_{l,n,1}, & \boldsymbol{b}^{\alpha}_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\,\boldsymbol{b}^{\alpha}_{l+1,n,l},\\
\boldsymbol{b}^{\alpha}_{l,m,-1} &= \pm\mathrm{e}^{\mathrm{i}(-l+m+1)\pi/2}\,\boldsymbol{J}_{\mathsf{b}^3}\,\boldsymbol{b}^{\alpha}_{l,1,m}, & \boldsymbol{b}^{\alpha}_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\,\boldsymbol{b}^{\alpha}_{l+1,l,m},
\end{aligned}
\tag{4.33}
$$

with the upper sign for α as I and the lower sign for α as II. For III, it becomes

$$
\begin{aligned}
2\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,-1,n} &= \mathrm{e}^{\mathrm{i}(l+1-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\left[\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,n,1} - \sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,n,1}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,-1,n} &= \mathrm{e}^{\mathrm{i}(l+1-n)\pi/2}\,\boldsymbol{J}_{\mathsf{b}}\,\left[-\sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,n,1} - \left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,n,1}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-l+m+1)\pi/2}\,\boldsymbol{J}_{\mathsf{b}^3}\,\left[\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,1,m} - \sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,1,m}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-l+m+1)\pi/2}\,\boldsymbol{J}_{\mathsf{b}^3}\,\left[-\sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,1,m} - \left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,1,m}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\,\left[\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l+1,n,l} - \sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l+1,n,l}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\,\left[\sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l+1,n,l} - \left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l+1,n,l}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\,\left[-\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l+1,l,m} + \sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l+1,l,m}\right],\\
2\left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\,\left[-\sqrt{3}\left(\boldsymbol{b}^{\mathrm{III}}_1\right)_{l+1,l,m} - \left(\boldsymbol{b}^{\mathrm{III}}_2\right)_{l+1,l,m}\right],
\end{aligned}
\tag{4.34}
$$

and for classes IV and V, we have

$$
\begin{aligned}
(\boldsymbol{b}_1^\alpha)_{l,-1,n} &= \mp e^{i(l+1-n)\pi/2}\, \boldsymbol{J}_{\mathsf{b}}\ (\boldsymbol{b}_1^\alpha)_{l,n,1}\,, & (\boldsymbol{b}_1^\alpha)_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\ (\boldsymbol{b}_2^\alpha)_{l+1,n,l}\,,\\
(\boldsymbol{b}_2^\alpha)_{l,-1,n} &= \mp e^{i(l+1-n)\pi/2}\, \boldsymbol{J}_{\mathsf{b}}\ (\boldsymbol{b}_3^\alpha)_{l,n,1}\,, & (\boldsymbol{b}_2^\alpha)_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\ (\boldsymbol{b}_3^\alpha)_{l+1,n,l}\,,\\
(\boldsymbol{b}_3^\alpha)_{l,-1,n} &= \pm e^{i(l+1-n)\pi/2}\, \boldsymbol{J}_{\mathsf{b}}\ (\boldsymbol{b}_2^\alpha)_{l,n,1}\,, & (\boldsymbol{b}_3^\alpha)_{l,l+1,n} &= \boldsymbol{J}_{\mathsf{d}}\ (\boldsymbol{b}_1^\alpha)_{l+1,n,l}\,,\\
(\boldsymbol{b}_1^\alpha)_{l,m,-1} &= \mp e^{i(-l+m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{b}^3}\ (\boldsymbol{b}_1^\alpha)_{l,1,m}\,, & (\boldsymbol{b}_1^\alpha)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\ (\boldsymbol{b}_3^\alpha)_{l+1,l,m}\,,\\
(\boldsymbol{b}_2^\alpha)_{l,m,-1} &= \pm e^{i(-l+m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{b}^3}\ (\boldsymbol{b}_3^\alpha)_{l,1,m}\,, & (\boldsymbol{b}_2^\alpha)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\ (\boldsymbol{b}_1^\alpha)_{l+1,l,m}\,,\\
(\boldsymbol{b}_3^\alpha)_{l,m,-1} &= \mp e^{i(-l+m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{b}^3}\ (\boldsymbol{b}_2^\alpha)_{l,1,m}\,, & (\boldsymbol{b}_3^\alpha)_{l,m,l+1} &= \boldsymbol{J}_{\mathsf{d}^2}\ (\boldsymbol{b}_2^\alpha)_{l+1,l,m}\,,
\end{aligned}
\tag{4.35}
$$

with the upper sign for α as IV and the lower sign for α as V.

4.2.3 $\mathcal{H}_N$: fundamental domain with Hermitian symmetry

In previous studies (see Galloway and Frisch, 1986, for instance), it has been natural to exploit the Hermitian symmetry (4.5) to use half the Fourier modes in a traditional truncation of (4.4), for example in place of the full cube $|l|, |m|, |n| \leq N$, to use the half-cube $0 \leq l, |m|, |n| \leq N$, thus reducing the computer time required by a factor of two. With Hermitian symmetry, every symmetry in G maps a general mode $\boldsymbol{b}_{l,m,n}$ to a distinct 1/24 piece of the half-domain, thus the fundamental domain constitutes 1/48 of $\mathcal{F}_N$. We define the fundamental domain with Hermitian symmetry as

$$
\mathcal{H}_N = \{(l,m,n) : 1 \leq l \leq N, 0 \leq m \leq l, 0 \leq n \leq m\}. \tag{4.36}
$$

This is shown in figure 4.3 and can be seen to be half of $\mathcal{G}_N$; the half for which $n \leq m$. If a symmetry maps a general mode $\boldsymbol{b}_{l,m,n}$ to a mode $\boldsymbol{b}_{-l,\pm m,\pm n}$, we can then also use Hermitian symmetry to make $\boldsymbol{b}_{l,m,n}$ map to $\boldsymbol{b}^*_{l,\mp m,\mp n}$, thereby ensuring that every mode in $\mathcal{F}_N$ can be calculated from a mode within $\mathcal{H}_N$. As with the case of solving within $\mathcal{G}_N$, we wish to know which of the modes lie on the edge of the domain, so as to store an extra layer of modes outside of $\mathcal{H}_N$. A slice of constant l is displayed schematically in figure 4.4. The picture is more complicated than with

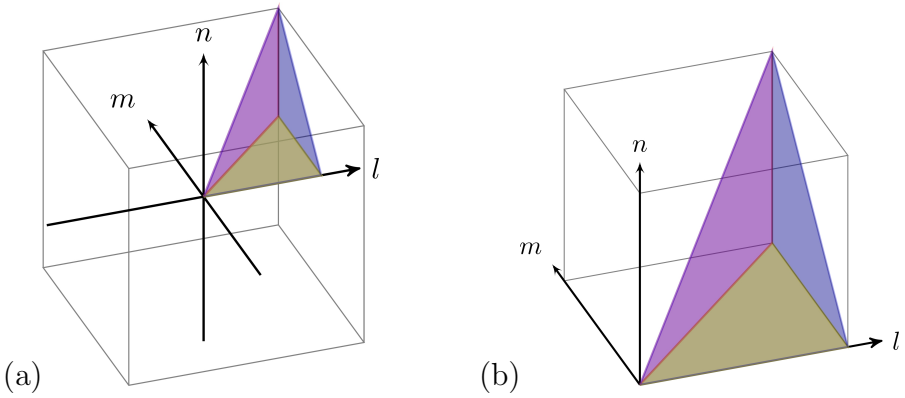

Figure 4.3: Fundamental domain $\mathcal{H}_N$ (a) within the full cube of wave vectors $\mathcal{F}_N$ and (b) within the first octant.

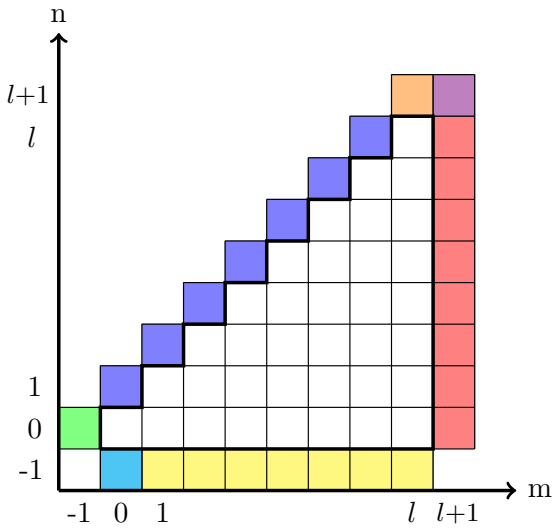

Figure 4.4: A cross-section of $\mathcal{H}_N$ for unspecified l. Unshaded modes are time stepped, coloured modes need to be calculated using symmetries; different colours identify different symmetries used.

$\mathcal{G}_N$ and the same symmetries cannot necessarily be used. There are now seven sets of modes whose values need to be calculated. These are:

- $\boldsymbol{b}_{l,m,-1}$ with $m > 0$, coloured yellow in figure 4.4, needed to time step $\boldsymbol{b}_{l,m,0}$, $m > 0$;
- $\boldsymbol{b}_{l,0,-1}$, coloured turquoise, needed to step $\boldsymbol{b}_{l,0,0}$;
- $\boldsymbol{b}_{l,-1,0}$, coloured green, needed to step $\boldsymbol{b}_{l,0,0}$;
- $\boldsymbol{b}_{l,m,m+1}$ with $m < l$, coloured blue, needed to step $\boldsymbol{b}_{l,m,m}$ and $\boldsymbol{b}_{l,m+1,m+1}$ (for

$m < l$);

- $\boldsymbol{b}_{l,l,l+1}$, coloured orange, needed to step $\boldsymbol{b}_{l,l,l}$;
- $\boldsymbol{b}_{l,l+1,n}$, coloured red, needed to step $\boldsymbol{b}_{l,l,n}$;
- $\boldsymbol{b}_{l,l+1,l+1}$, coloured purple, needed to step $\boldsymbol{b}_{l+1,l+1,l+1}$.

The reason for more sets of modes, even though $\mathcal{H}_N$ has fewer geometrical faces than $\mathcal{G}_N$, will become clear when determining the symmetries used and is due, in part, to the irregular shape of the fundamental domain. We take the example of $\mathsf{I}\boldsymbol{b}$ to demonstrate how Hermitian symmetry is used in conjunction with the symmetries. From (2.50), we have

$$\mathsf{I}(\boldsymbol{x}) = \left(-x - \tfrac{\pi}{2}, -z - \tfrac{\pi}{2}, -y - \tfrac{\pi}{2}\right), \tag{4.37}$$

and using (4.18), we have that

$$(\mathsf{I}\boldsymbol{b})_{l,m,n} = \mathrm{e}^{\mathrm{i}(l+m+n)\pi/2}\, \boldsymbol{J}_{\mathsf{I}} \boldsymbol{b}_{-l,-n,-m}. \tag{4.38}$$

Using the Hermitian symmetry property (4.5) then provides

$$(\mathsf{I}\boldsymbol{b})_{l,m,n} = \mathrm{e}^{\mathrm{i}(l+m+n)\pi/2}\, \boldsymbol{J}_{\mathsf{I}} \boldsymbol{b}^*_{l,n,m}, \tag{4.39}$$

with the resulting transformation exhibiting qualities similar to a reflection; it allows for two indices to be switched without either one changing sign. With the choice $n = m + 1$ for $m < l$, we obtain an expression for calculating the values of modes $\boldsymbol{b}_{l,m,m+1}$ described by

$$(\mathsf{I}\boldsymbol{b})_{l,m,m+1} = \mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{I}} \boldsymbol{b}^*_{l,m+1,m}. \tag{4.40}$$

As was carried out for extra mode layers of $\mathcal{G}_N$, the equations needed to fill each set of modes is provided, though first we avoid repetition by considering the three sets of modes outside $\mathcal{H}_N$ that are transformed in the same way as those outside $\mathcal{G}_N$.

These sets are $\boldsymbol{b}_{l,0,-1}, \boldsymbol{b}_{l,l,l+1}$ and $\boldsymbol{b}_{l,l+1,l+1}$, which are special cases of $\boldsymbol{b}_{l,m,-1}, \boldsymbol{b}_{l,m,l+1}$ and $\boldsymbol{b}_{l,l+1,n}$, respectively. The reason that these specific cases are treated separately to their respective general mode sets is that the transformation used to map these modes back to the fundamental domain, for $\mathcal{G}_N$, remains only valid for the special cases and so alternative transformations are needed for the general cases. With this in mind, the modes $\boldsymbol{b}_{l,0,-1}, \boldsymbol{b}_{l,l,l+1}$ and $\boldsymbol{b}_{l,l+1,l+1}$ adopt the same transformations as in $\mathcal{G}_N$, namely those for $\boldsymbol{b}_{l,m,-1}, \boldsymbol{b}_{l,m,l+1}$ and $\boldsymbol{b}_{l,l+1,n}$ (respectively), described in (4.33)–(4.35), and with the appropriate choices of m and n.

The remaining four sets of modes are $\boldsymbol{b}_{l,m,-1}$ with $m > 0$, $\boldsymbol{b}_{l,-1,0}$, $\boldsymbol{b}_{l,m,m+1}$ for $m < l$ and $\boldsymbol{b}_{l,l+1,n}$ with $n < l + 1$. These are transformed, respectively, by a^2 with Hermitian (4.5), b^2, l with (4.5) and n with (4.5). These transformations are summarised by

$$
\begin{aligned}
(\mathsf{a}^2\boldsymbol{b})_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, \boldsymbol{b}^*_{l,m,1}, & (\mathsf{b}^2\boldsymbol{b})_{l,-1,0} &= \mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, \boldsymbol{b}_{l,1,0},\\
(\mathsf{l}\boldsymbol{b})_{l,m,m+1} &= \mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, \boldsymbol{b}^*_{l,m+1,m}, & (\mathsf{n}\boldsymbol{b})_{l,l+1,n} &= \mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, \boldsymbol{b}^*_{l+1,l,n},
\end{aligned}
\tag{4.41}
$$

with modes on the right-hand side belonging to $\mathcal{H}_N$.

From (4.41), we have for classes I and II that

$$
\begin{aligned}
\boldsymbol{b}^{\alpha}_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, \boldsymbol{b}^{\alpha*}_{l,m,1}, & \boldsymbol{b}^{\alpha}_{l,-1,0} &= \mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, \boldsymbol{b}^{\alpha}_{l,1,0},\\
\boldsymbol{b}^{\alpha}_{l,m,m+1} &= \pm\mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, \boldsymbol{b}^{\alpha*}_{l,m+1,m}, & \boldsymbol{b}^{\alpha}_{l,l+1,n} &= \pm\mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, \boldsymbol{b}^{\alpha*}_{l+1,l,n},
\end{aligned}
\tag{4.42}
$$

with class I being the upper sign and II being the lower. To continue, we will need the relevant irreducible representations for classes III, IV and V; these are given by

$$
\boldsymbol{M}^{\mathrm{III}}(\mathsf{a}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{b}^2) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \boldsymbol{M}^{\mathrm{III}}(\mathsf{l}) = \frac{1}{2}\begin{pmatrix} 1 & -\sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\mathrm{III}}(\mathsf{n}) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},
$$

$$\boldsymbol{M}^{\mathrm{IV}}(\mathsf{a}^2) = \boldsymbol{M}^{\mathrm{V}}(\mathsf{a}^2) = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{M}^{\mathrm{IV}}(\mathsf{b}^2) = \boldsymbol{M}^{\mathrm{V}}(\mathsf{b}^2) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix},$$
$$\boldsymbol{M}^{\mathrm{IV}}(\mathsf{l}) = -\boldsymbol{M}^{\mathrm{V}}(\mathsf{l}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{IV}}(\mathsf{n}) = -\boldsymbol{M}^{\mathrm{V}}(\mathsf{n}) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (4.43)$$

From (4.18) and (4.43), we calculate the extra layers of modes for class III by

$$\begin{aligned} \left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, \left(\boldsymbol{b}_1^{\mathrm{III}}\right)^*_{l,m,1}, \qquad \left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,-1,0} = \mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, \left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,1,0}, \\ \left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, \left(\boldsymbol{b}_2^{\mathrm{III}}\right)^*_{l,m,1}, \qquad \left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,-1,0} = \mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, \left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,1,0}, \end{aligned}$$

$$\begin{aligned} 2\left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,m,m+1} &= \mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, \left[\left(\boldsymbol{b}_1^{\mathrm{III}}\right)^*_{l,m+1,m} - \sqrt{3}\left(\boldsymbol{b}_2^{\mathrm{III}}\right)^*_{l,m+1,m}\right], \\ 2\left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,m,m+1} &= \mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, \left[-\sqrt{3}\left(\boldsymbol{b}_1^{\mathrm{III}}\right)^*_{l,m+1,m} - \left(\boldsymbol{b}_2^{\mathrm{III}}\right)^*_{l,m+1,m}\right], \end{aligned}$$

$$\begin{aligned} \left(\boldsymbol{b}_1^{\mathrm{III}}\right)_{l,l+1,n} &= -\mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, \left(\boldsymbol{b}_1^{III}\right)^*_{l+1,l,n}, \\ \left(\boldsymbol{b}_2^{\mathrm{III}}\right)_{l,l+1,n} &= \mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, \left(\boldsymbol{b}_2^{III}\right)^*_{l+1,l,n}, \end{aligned} \quad (4.44)$$

and for classes IV and V,

$$\begin{aligned} (\boldsymbol{b}_1^{\alpha})_{l,m,-1} &= -\mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, (\boldsymbol{b}_1^{\alpha})^*_{l,m,1}, & (\boldsymbol{b}_1^{\alpha})_{l,-1,0} &= \mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, (\boldsymbol{b}_1^{\alpha})_{l,1,0}, \\ (\boldsymbol{b}_2^{\alpha})_{l,m,-1} &= -\mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, (\boldsymbol{b}_2^{\alpha})^*_{l,m,1}, & (\boldsymbol{b}_2^{\alpha})_{l,-1,0} &= -\mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, (\boldsymbol{b}_2^{\alpha})_{l,1,0}, \\ (\boldsymbol{b}_3^{\alpha})_{l,m,-1} &= \mathrm{e}^{\mathrm{i}(-m-1)\pi}\, \boldsymbol{J}_{\mathsf{a}^2}\, (\boldsymbol{b}_3^{\alpha})^*_{l,m,1}, & (\boldsymbol{b}_3^{\alpha})_{l,-1,0} &= -\mathrm{e}^{\mathrm{i}l\pi}\, \boldsymbol{J}_{\mathsf{b}^2}\, (\boldsymbol{b}_3^{\alpha})_{l,1,0}, \\ (\boldsymbol{b}_1^{\alpha})_{l,m,m+1} &= \pm\mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, (\boldsymbol{b}_1^{\alpha})^*_{l,m+1,m}, & (\boldsymbol{b}_1^{\alpha})_{l,l+1,n} &= \pm\mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, (\boldsymbol{b}_2^{\alpha})^*_{l+1,l,n}, \\ (\boldsymbol{b}_2^{\alpha})_{l,m,m+1} &= \pm\mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, (\boldsymbol{b}_3^{\alpha})^*_{l,m+1,m}, & (\boldsymbol{b}_2^{\alpha})_{l,l+1,n} &= \pm\mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, (\boldsymbol{b}_1^{\alpha})^*_{l+1,l,n}, \\ (\boldsymbol{b}_3^{\alpha})_{l,m,m+1} &= \pm\mathrm{e}^{\mathrm{i}(l+2m+1)\pi/2}\, \boldsymbol{J}_{\mathsf{l}}\, (\boldsymbol{b}_2^{\alpha})^*_{l,m+1,m}, & (\boldsymbol{b}_3^{\alpha})_{l,l+1,n} &= \pm\mathrm{e}^{\mathrm{i}(2l+n+1)\pi/2}\, \boldsymbol{J}_{\mathsf{n}}\, (\boldsymbol{b}_3^{\alpha})^*_{l+1,l,n}, \end{aligned} \quad (4.45)$$

with the upper sign for α as IV and the lower sign for α as V. Updating these modes, as explained for $\mathcal{G}_N$, is carried out prior to usual stepping of the modes within $\mathcal{H}_N$.

4.2.4 Time stepping in the fundamental domain

Now that both fundamental domains, $\mathcal{G}_N$ and $\mathcal{H}_N$ have been defined, we can discuss how this affects the time stepping procedure, as compared to the numerical procedure in $\mathcal{F}_N$. It is important to note, however, that any time stepping carried out in the domain $\mathcal{G}_N$ was done so only for comparison against the results of the IRAM solver (discussed later in this chapter). All time stepping results provided in this section are calculated with a numerical code operating in $\mathcal{H}_N$ and although the process is identical for time stepping in $\mathcal{G}_N$, this discussion primarily deals with $\mathcal{H}_N$.

As updating the values of these extra layers of modes (which will henceforth be referred to as *mode copying*) is carried out so that modes within the fundamental domain can be time stepped, these calculations must be made prior to any time stepping activities. Referring to the order of operations within each time step, the calculations in (4.42), (4.44) and (4.45) replace the updating of the $\boldsymbol{b}_{-1,m,n}$ modes in the first step. Regarding the phase shifts of such mode copying operations, we observe that they are integer multiples of $\pi/2$. In (4.18), instead of explicitly calculating $\exp\left[\mathrm{i}\left(\boldsymbol{J}_{\mathrm{g}^{-1}}\boldsymbol{k}\right)\cdot\boldsymbol{T}_{\mathrm{g}^{-1}}\right]$, we access a predefined array whose elements are $1, \mathrm{i}, -1$ and $-\mathrm{i}$, i.e. the four distinct phase shifts of multiple $\pi/2$. For instance, instead of calculating $\exp[\mathrm{i}(l{+}m{+}n)\pi/2]$, we access index $(l{+}m{+}n)$ of the "phase" array, whose entry will be one of $[1, \mathrm{i}, -1, -\mathrm{i}]$ depending on the outcome of $(l{+}m{+}n)\mathrm{mod}4$. This eliminates a potentially large source of rounding error.

We are simulating in a reduced domain and so only modes within $\mathcal{H}_N$ (i.e. those included in (4.36)) are time stepped. Additionally, the reduced set of modes is used to calculate $\ln(E_{\mathrm{M}})$. In the cases where there are multiple fields (III, IV and V), these fields evolve independently and only interact at the boundaries (i.e. when carrying out mode copying operations). The individual fields are therefore time stepped independently in $\mathcal{H}_N$, with (4.9)–(4.11) and, subsequently, (4.7) solved for each field in turn and for each $(l, m, n) \in \mathcal{H}_N$. The time taken per time step is thus, in part, dependent on the symmetry class of the field: simulating fields in classes

IV or V will take approximately three times longer than for classes I or II. When snapshots of the field are required, the field in $\mathcal{H}_N$ is 'unfolded'; the procedure is briefly sketched. Using (4.39), we calculate modes $\boldsymbol{b}_{l,n,m}$ from $\boldsymbol{b}_{l,m,n}$ to construct the field in $\mathcal{G}_N$. We then use d and d^2 of (4.19) and (4.20) respectively to calculate modes $\boldsymbol{b}_{m,n,l}$ and $\boldsymbol{b}_{n,l,m}$ for all modes in $\mathcal{G}_N$, thus reconstructing the field in the first octant $(l, m, n \geq 0)$. From here, we calculate $\boldsymbol{b}_{m,-l,n}, \boldsymbol{b}_{-l,-m,n}$ and $\boldsymbol{b}_{-m,l,n}$ using a, a^2 and a^3 respectively for all modes with $l, m, n \geq 0$, giving the field for all modes with $n \geq 0$. Finally, the transformation b^2 is used to calculate modes $\boldsymbol{b}_{l,-m,-n}$ from $\boldsymbol{b}_{l,m,n}$ for modes $|l|, |m| \leq N, 0 \leq n \leq N$, giving the field in $\mathcal{F}_N$. For class III, we are 'unfolding' two fields and in classes IV and V, we unfold all three fields; we keep only the first field $\boldsymbol{b}_1^\alpha$, as the additional fields are simply transformed (rotated) versions of the first. We then visualise in the normal fashion.

It must be noted that the parameter values adopted in the full problem are also adopted for the reduced problems, as the core operations remain the same. This works both ways and so the reduced code is useful for identifying minimum yet safe values for parameters. This also makes it easier when verifying that the reduced codes are functioning as required by testing against the well-known growth rates of the full code, though this is, in itself, a difficult task. The first problem lies in finding seed magnetic fields that belong to each symmetry class. This can be achieved by starting with a seed field of the form $\boldsymbol{b} = (\sin z, 0, 0)$, equivalent to $\boldsymbol{b}_{0,0,1} = -\boldsymbol{b}_{0,0,-1} = (-\mathrm{i}/2, 0, 0)$, and testing how the symmetries map this field around for each symmetry class. The same can be done with an initial field $\boldsymbol{b} = (\cos z, 0, 0)$ and by taking linear combinations of these fields and their transformations (by symmetries), initial conditions can be constructed. We construct initial (seed) fields, each belonging to a distinct symmetry class, and use these in the full (Hermitian) code to test whether the respective $\mathcal{H}_N$ codes are functioning

correctly. These seed fields are given by

$$\begin{aligned}
\text{I: } \boldsymbol{b} &= (\sin 3z - \cos 3y, \sin 3x - \cos 3z, \sin 3y - \cos 3x),\\
\text{II: } \boldsymbol{b} &= (\sin z - \cos y, \sin x - \cos z, \sin y - \cos x),\\
\text{III: } \boldsymbol{b} &= (2\sin z, -\sin x, -\sin y),\\
\text{IV: } \boldsymbol{b} &= (\sin 2z, 0, -\sin 2x),\\
\text{V: } \boldsymbol{b} &= (\sin 2z, 0, \sin 2x),
\end{aligned} \tag{4.46}$$

and are tested against the corresponding $\mathcal{H}_N$ codes using *random* seed fields, as the symmetries imposed in each reduced code restrict the resulting magnetic field configurations to the specific class. For low R_m and for small simulation times, the evolution of the fields in (4.46) agrees well with the evolution of fields in the respective reduced $\mathcal{H}_N$ simulations. The second difficulty, however, is that the field evolution seen for a particular symmetry class at a specific R_m in the $\mathcal{H}_N$ code is only seen as transient evolution in the full code, before the field settles on a faster growing, or slower decaying mode. This is to be expected, as the rounding error can build up sufficiently, over large enough simulation times, to allow the mode with the highest growth rate to emerge and take over. Owing to the way that the reduced codes are constructed, however, the symmetry is imposed in the mode copying operation (as the symmetry class determines the values of these modes), making it impossible for a mode belonging to another symmetry class to emerge.

Finally, we mention the computational saving achieved by simulating in the fundamental domain $\mathcal{H}_N$. In the limit of large N, the time spent mode copying is of $O(N^2)$, negligible compared to that spent stepping the modes within, which is of $O(N^3)$. The computational time can naively be seen as reducing to 1/24 of that of the original (Hermitian) code *per field*: for classes I and II, the saving is a factor of 1/24, a factor of 1/12 for III and for IV and V, a factor of 1/8. Resolving fields for all five classes could reduce computational time to 10/24 of the time taken for the full problem. In practice, the mode copying costs are not negligible and so the effective saving is roughly 1/2; still a significant result. This, of course, does not

account for the wealth of additional information that we are able to obtain from following five separate branches and that we are easily able to resolve eigenmode crossings between different symmetry classes.

4.3 Parallelisation in $\mathcal{H}_N$

The purpose of parallelising any operation is to reduce a task to a series of smaller tasks that can be carried out simultaneously by a number of processors or processor cores. This serves to reduce the total time taken to complete the given task. The same features that make the serial ABC dynamo simulations amenable for numerical analysis, make them amenable for parallelisation. In fully parallel numerical operations, we have that parts of the problem are assigned to different computer nodes, each with its own local memory. One difficulty lies in determining the information required by a node that is stored on another node and how to best communicate this information. For the ABC dynamos, the local mode interaction means that if different parts of spectral space are assigned to different nodes, only information on Fourier modes located at the common edges needs communicating. This section discusses how numerical simulations carried out in $\mathcal{H}_N$ are constructed for parallel processing.

4.3.1 Dividing the calculations

First we consider the parallelisation of the full serial problem. Typically, this is done by cutting down the cube of Fourier modes into distinct but equally sized sections. For instance, if $|l| \leq N$, then the first section could contain all modes with $-N \leq l \leq -N + k - 1$, the second section with $-N + k \leq l \leq -N + 2k - 1$ for arbitrary k, with each section assigned to a particular core within a node. Assuming that the range of l can be divided into roughly equal subranges, this allows for fairly balanced breakdown of the problem.

This can be done for $\mathcal{H}_N$ but owing to its fairly irregular shape, as seen in figure 4.3,

the breakdown into discrete sections is not so straightforward. As the maximum values of m and n are determined by l, it still seems sensible for the domain to be divided into ranges of l. Figure 4.5 identifies two possible methods of breaking $\mathcal{H}_N$ down into six sections. It is obvious that the breakdown depicted in figure

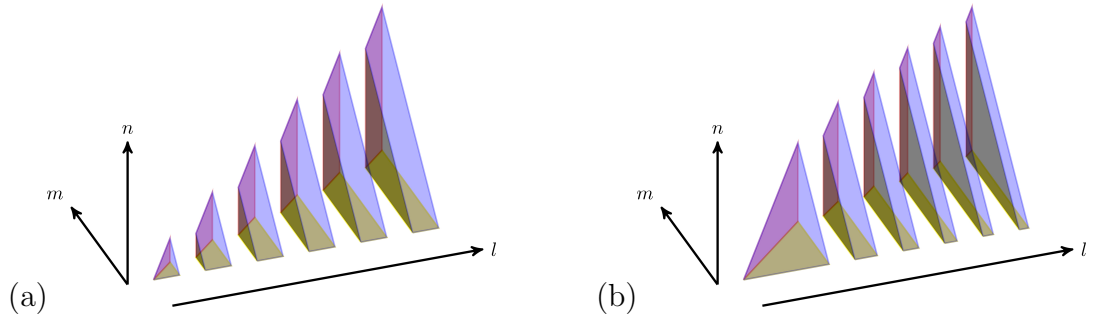

Figure 4.5: Breaking up $\mathcal{H}_N$ into six pieces by (a) dividing l into equal subranges and (b) dividing l into subranges with (approximately) equal numbers of modes.

4.5(a), where l is divided into equal subranges, will be much less efficient than that of 4.5(b), where the number of modes in each section is approximately the same. Programming of method (a), however, is much easier as it simply involves dividing l by the number of sections required.

In any form of parallel processing, it is not common to achieve a breakdown of a task into completely independent parts so that no communication is required. The ABC dynamo problem is no exception; the need to access the value of all neighbouring modes to time step a specific mode means that no matter how the problem is decomposed, there will always be some modes whose neighbouring modes are assigned to another core (potentially on another node). As solving the dynamo problem involves time stepping, communication between nodes will have to occur every time step and importantly, a given core cannot embark on the calculations of the next time step before it has received all information from neighbouring cores. Solving the problem in parallel therefore means that each time step can only be solved as quickly as it takes the core with the largest workload to complete its task, as a time step cannot be completed until all cores have finished their individual tasks. Thus, it is crucial that at least some attempt is made to balance the workload between cores.

In attempting to load-balance, we have to consider that array sizes and the number of nodes to be used has to be known before compiling. This means that the breakdown of the problem is carried out prior to even writing the programs. For this purpose, a bash script was written to calculate a reasonable breakdown, write the program and then compile it for a number of nodes specified by the user. Before going into details of the script, a few important features of the serial codes must be mentioned.

Firstly, in any operation that involves cycling through the mode indices, i.e. time stepping, mode copying, calculating E_{M}, the index that changes the slowest is l, as l is always incremented in the outermost loop. This is why dividing l into subranges makes the most sense, as it simply involves carrying out the particular operation for that core's starting l index to the finishing l index, these are henceforth referred to as l_{min} and l_{max} respectively. Secondly, on investigating all operations that cycle through the l indices, we have to identify any that are affected by the breakdown. Certainly the time stepping of modes is affected, as modes with indices $l = l_{\mathrm{min}}$ and $l = l_{\mathrm{max}}$ will not have all their neighbours. We find that the mode copying operations (4.42), (4.44) and (4.45) require modes with index no less than $l = l_{\mathrm{min}}-1$ and no more than $l = l_{\mathrm{max}}+1$; this is already covered by the 'neighbour requirement' of the time stepped modes. When calculating the magnetic energy E_{M} of the respective sections, no additional information is required. To calculate the total magnetic energy E_{M}, however, the E_{M} from all sections can be communicated to a single core for summing and storing. Thirdly, as information from one core cannot be freely accessed from another core (assuming that they are on different nodes) because they do not share memory, the required information has to be communicated and stored locally. The memory assigned to any core will therefore require modes of $l = l_{\mathrm{min}}-1$ to $l_{\mathrm{max}}+1$ to be stored in its assigned memory. Values of modes in slices $l = l_{\mathrm{min}}-1$ and $l_{\mathrm{max}}+1$ are communicated to the core each time step, though only the slices $l = l_{\mathrm{min}}$ to $l = l_{\mathrm{max}}$ are time stepped on this core; we discuss this communication in more detail in the next subsection.

The script's algorithm determines the l_{min} and l_{max} for all cores, with this informa-

tion written into the program at compile time. In the script's method for balancing, we consider a slice of constant l of $\mathcal{H}_N$; see figure 4.4. The number of modes in a given slice is the triangular number of the index plus one (as the indices m and n start at zero), i.e. $(l+1)(l+2)/2$; this excludes the additional layers that are filled by mode copying and we take this as the number of modes time stepped in any one slice. The last slice, $l = N$, is set to zero, as discussed for the serial problem, and so the last slice in which the modes are time stepped is with $l = N - 1$. In this slice there are $N(N + 1)/2$ modes; we do not wish to split a slice of modes over multiple cores, so we know that $F_{\max} = N(N + 1)/2$ is the smallest possible maximum number of modes that any core will be assigned. As we also have to specify the number of nodes required, we already know the maximum number of cores that are available to use; number of nodes multiplied by number of cores per node. This is identified by $p_{\max}$. The aim of the script is to assign a section of $\mathcal{H}_N$ to every core.

As the process contains nested operations, it is best summarised as a series of numbered instructions. This script effectively starts with a limit of the number of modes $F_{\lim} = F_{\max}$ and tests whether integer multiples of $F_{\lim}$ can accommodate a breakdown into $p_{\max}$ parts. If this isn't possible, $F_{\lim}$ is increased by a small fraction and the process restarts until it is broken down into $p_{\max}$ sections, at which point it ends as it has succeeded. The steps are:

1. Starting from the largest slice at $l = N - 1$, the script tests the number of modes against (at first) $F_{\lim} = F_{\max}$: if the number of modes is less than $F_{\lim}$, it moves to the next slice (i.e. $l = N - 2$) and adds the number of modes in that slice too. It repeats this process until it reaches a slice (e.g. $l = N - k$) where including the next slice ($l = N - k - 1$) would make the total number of modes greater than $F_{\lim}$; at this point it has found the range of l for one core, that is $l_{\max} = N-1$ and $l_{\min} = N-k$. This process is repeated, starting from the next largest unassigned slice, until all slices have been assigned to a core. There are then p cores with work assigned to them.

2. The value of p is tested against $p_{\max}$: if $p > p_{\max}$, the breakdown is not viable and step 1 is repeated but with $tF_{\lim}$ for $t = 2$, as the maximum number of modes. Again, if this fails, t is incremented (by 1) and step 1 is repeated until $p \leq p_{\max}$; the value of p and the breakdown are then stored.

3. Steps 1 and 2 are repeated, starting with $F_{\lim} = (1 + 0.02q)F_{\max}$ for $q = 1, \ldots, 50$. Each repeat produces a value for p with a particular breakdown associated with it. If at any stage, $p = p_{\max}$, the routine ends and the breakdown (values of $l_{\min}, l_{\max}$ for all cores) are stored.

4. Parallel program is written with all $l_{\min}$ and $l_{\max}$ stored in an array. Program is compiled.

Although the process seems rather convoluted, it is simply cycling through many possible arrangements but starting with the smallest $F_{\lim}$ possible ($F_{\max}$). As we are increasingly $F_{\lim}$ slowly (see step 3), when we reach the first breakdown of the problem with $p = p_{\max}$, we have automatically reached the best possible breakdown, whilst keeping $F_{\lim}$ as low as possible and distributing the work in the most equal way over the $p_{\max}$ cores. Of course, owing to the irregular shape of the domain, it is highly unlikely that the cores will have the same workload, yet there is no foolproof way of assigning the work in a balanced way. There may be more effective methods such as giving the same core non-contiguous slices but this presents problems in communication and with data storage in the actual programs.

The value of 0.02 in step 3 may seem arbitrary, though in fact it was chosen so that the script did not take too long to find an optimal breakdown. Bash is not necessarily the most efficient environment for such calculations, but as the script had to be passed variables (maximum number of nodes), had to write a program and finally compile it, other programming environments would not have been suitable. The values of q are also intentional as step 3 was intended to try $F_{\lim} = F_{\max}$ to $F_{\lim} = 2F_{\max}$ in as small steps as possible, thus incrementing by $0.02F_{\max}$ on every repeat. Setting the increment to $0.01F_{\max}$ was tested but took much longer and did not achieve better domain breakdowns than 0.02.

The use of this script is limited to high R_m simulations, where the resolution N is large. As the parallel codes are intended for use in this regime, then this is less of an issue, though there are still situations where no solution with $p = p_{\max}$ can be found. In these cases, the breakdown with the highest p is chosen, as this is generally the best solution. Typically this occurs when N is not large enough compared to $p_{\max}$, and we rerun the script but with a lower number of nodes (i.e. lower $p_{\max}$) until a solution with $p = p_{\max}$ is found.

4.3.2 Communicating between nodes

Now that the method for optimal breakdown of $\mathcal{H}_N$ has been discussed, we can move on to what information is communicated and how. We wish only to communicate necessary information, as passing extraneous information may serve to slow down the simulations. An important feature of parallel processing to consider is that at the initialisation of a (parallel) simulation, all cores involved are assigned a *rank*, simply the number of the core. If we have $p_{\max}$ cores, they are numbered 0 to $p_{\max} - 1$. For ease of operation, we assign neighbouring sections of $\mathcal{H}_N$ to neighbouring cores. That is to say, if a core of rank p has $l_{\max} = k$ for some k, then the core with rank $p + 1$ will have $l_{\min} = k + 1$. As parallel codes use the same set of instructions, the only way to have them perform different operations is to use locally defined variables; the rank p, $l_{\min}$ and $l_{\max}$ are all locally defined (i.e. the variable has the same name on every core but not the same value). This is the key to communication in parallel processing environments, as the same communication command can be used for all cores.

We first consider communicating neighbouring modes so that time stepping and mode copying are possible. In the discussion above, we know that each core has a distinct $l_{\min}$ and $l_{\max}$, though the arrays are defined for $l = l_{\min} - 1$ to $l = l_{\max} + 1$ to store the extra layers of modes. We conceptualise the cores as being arranged in a linear fashion and with ranks of ascending order. Thus for a core of rank p, we consider rank $p - 1$ to be to the *left* and rank $p + 1$ to be to the *right*; this allows

the process to be more easily described.

It is easier to deal with communication in one direction at a time and so passing left will be discussed first. As the communication call has to deal with the multiple cores, a 'send' followed by a 'receive' is required, though both are covered by the 'send and receive' command, which avoids the blocking effect of the individual commands; a core can place a receive call without waiting for confirmation (i.e. without being 'blocked' from carrying out other operations until receiving confirmation) that the data has been successfully received by another core. To summarise, a core p sends slice $l = l_{\min}$ to core $p-1$. The core then receives data from core $p+1$ and stores it in slice $l = l_{\max}+1$. It must be noted that, as core 0 has no core to its left, it does not pass any data during this step, but only receives from core 1. The process is shown in figure 4.6: the green arrows represent passing data to the left, with the solid green modes being transmitted (sent) and the hashed green modes being received.

Passing to the right occurs in same fashion. Each core p places a send and receive call, sending slice $l = l_{\max}$ to core $p+1$ and then receiving data from core $p-1$ and storing it in slice $l = l_{\min}-1$. This process is summarised by the blue arrows, with solid blue modes the data sent by a core and the hashed blue modes the data received by the core. Clearly, core $p_{\max}-1$ does not have a neighbour to the right and so only places a receive call (this is managed through an 'if' condition, testing for rank value). In figure 4.6, we also see the dashed outline, which indicates the full extent of the storage arrays. The only modes to be time stepped in each core are identified by those contained within the thick black border, though can also be identified as those modes with solid colour (no hashing).

The process described above outlines the 'mode communication', which occurs once per time step. The process, however, is made more complicated by the irregular shape of $\mathcal{H}_N$. Fortran is not particularly efficient at handling irregular arrays and so we are limited to defining regular (cuboid) arrays with $l_{\min}-1 \leq l \leq l_{\max}+1$, $0 \leq m \leq l_{\max}+1$, $0 \leq n \leq l_{\max}+1$. Approximately half the space in the arrays is unused, as $n \leq m$ in $\mathcal{H}_N$ and so to communicate the whole slice would

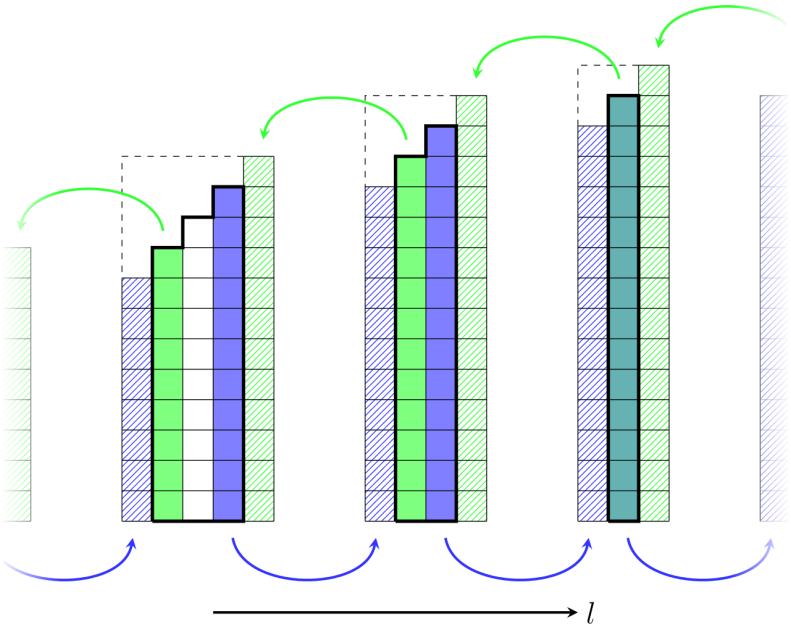

Figure 4.6: Visualisation of the communication process. First all cores pass $\boldsymbol{b}_{l_{\min},m,n}$ to cores one rank lower (left), then receive the same data from the core with one rank higher and store it in $\boldsymbol{b}_{l_{\max}+1,m,n}$. All cores then pass $\boldsymbol{b}_{l_{\max},m,n}$ to cores one rank higher (right), then receive same data from the core with one rank lower and store it in $\boldsymbol{b}_{l_{\min}-1,m,n}$. Only solid coloured and white modes are time stepped.

be incredibly wasteful. A method involving pointers was considered but could not be implemented due to its ever increasing complexity, as it would involve pointers pointing to pointers and so forth. The low-tech method implemented involved copying only the mode values to be sent from the slice into a vector of the correct length. This vector is sent and then copied (or unpacked) into the correct slice by the receiving core. In testing this method against communicating the original slice (with extraneous zero modes), it was found to be faster, despite the additional copying operations that it entails.

As the mode copying operation, for a given core, requires modes with index l no less than $l_{\min} - 1$ and no greater than $l_{\max} + 1$, then no further mode communication is required, other than that detailed above. To carry out mode copying, therefore, the communication described above must already have occurred, placing this communication as the first operation in a time step, immediately followed by the mode copying. The order of operations is discussed in more detail along with

the specifics of the codes used (see next subsection).

The other communication that is carried out every time step is that of the individually calculated E_{M}. Unlike the mode communication, which involves each core communicating data with its neighbours, reconstructing E_{M} simply involves all individual E_{M} (once calculated) being sent to one core. Once the script, described in the last subsection, finds the optimal breakdown, it assigns the section of $\mathcal{H}_N$ with $l_{\max} = N - 1$ to the core with the highest rank, and then works in descending order, finally assigning the last section to core rank 0. Because of this, core 0 is always (as far as testing has identified) assigned less work than the core with the maximum workload. For this reason, core 0 is assigned the task of putting together the E_{M} with all other cores placing a 'send' call to core 0 and with core 0 receiving $p_{\max} - 1$ values. Once summed, this is then redistributed to all cores; this is important as the value of $\ln(E_{\mathrm{M}})$ is tested against minimum and maximum values; if it is too small or large, all arrays values are resized by a fixed factor and this has to be carried out for all modes stored on all cores, thus requiring every core to know the total E_{M}. This 'E_{M} communication' has to occur after the respective E_{M} are calculated.

4.3.3 Fortran MPI codes: additional information

We briefly provide some additional details of the parallel code. The parallel codes were adapted from the serial $\mathcal{H}_N$ Fortran codes to use the MPI 1.3 implementation on the University of Exeter supercomputer, Zen. Zen has 12 cores per node (two hex-core processors) with these 12 cores sharing random access memory. The implementation of only MPI in these codes is efficient at communication between nodes and the cores within them, though it does not take advantage of the shared memory in a given node and instead assigns distinct portions of the memory to each core. This means that cores within a node still have to communicate information. OpenMP, however, is designed to allow all cores, within a node, to access the whole of the shared memory, removing the need for communication. The weakness

of OpenMP is that it does not support inter-node communication and so can only be implemented in simulations on a single node. A hybrid of the two systems can be used, so that OpenMP avoids communication within a node and MPI allows communication between nodes; this would have been ideal for use the parallel $\mathcal{H}_N$ codes, though the actual speed-up may have arguably been marginal. The reason behind this is that a time step can only begin once all cores have the required data and there would still be a minimum of 1 or 2 cores per node that required inter-node communication to access this data. As the time step is no faster than the slowest core, every (non-communicating) core would still have to wait for these communication operations to complete.

In terms of the one-off operations such as snapshots of the field, these were all handled by core 0. Rather than having each core write the respective sections to a single file, all cores send their respective sections of $\mathcal{H}_N$ back to core 0, which then reconstructs $\mathcal{H}_N$ in a dynamically allocated array and finally writes the data to a file. As Zen has a maximum limit for simulation time, it is necessary to store all relevant information at the end of a given run, should the magnetic field need more time to evolve. This requires saving not only $\boldsymbol{b}$ (as with the snapshot) but also $\boldsymbol{\Lambda}^j, \boldsymbol{\Lambda}^{j-1}$ and $\boldsymbol{\Lambda}^{j-2}$, with each of these stored in a separate file. As the fundamental domain for large R_m can take up large amounts of memory, each of these variables is dealt with one at a time: data is received from all cores and stored in a dynamically allocated array, the data is then written to file and the array is deallocated before allocating a new array for the next variable. In restarting the simulation with the stored data, the process is reversed with core 0 sending the sections of each variable to the correct core, in turn.

For completeness, the order of a single time step in the parallel code is given below:

- All cores (except rank 0) pack modes in slice $l_{\min}$ into vector arrays. All cores except $p_{\max} - 1$ (highest rank) pack modes in slice $l_{\max}$ into vector arrays.

- All cores (except 0) send $l_{\min}$ vector to the core on their left (i.e. to core with one rank lower). All cores (except $p_{\max} - 1$) receive their respective vectors

and unpack into position of slice $l_{\max}+1$. The vector unpacking procedure mirrors the packing procedure.

- All cores (except $p_{\max}-1$) send $l_{\max}$ vector to their right (one rank higher). All cores (except 0) receive their respective vectors and unpack into position of slice $l_{\min}-1$.
- Each core carries out the mode copying operation for $l=l_{\min}$ to $l=l_{\max}$: external mode layers are calculated by (4.42), (4.44) or (4.45) depending on the symmetry class.
- Each core time steps modes from $l=l_{\min}$ to $l=l_{\max}$
- Each core calculates a local E_{M}^2 by taking the sum of the magnitudes of modes in slices $l_{\min}$ to $l_{\max}$.
- All cores apart from 0 send their local E_{M}^2 to core 0; core 0 receives the local E_{M}^2 in turn. Once all have been received, it adds them all to its own local E_{M}^2, takes the square root and sends the global (i.e. for $\mathcal{H}_N$) E_{M} to all other cores in turn; all other cores place a receive call and store this value.
- The global $\ln(E_{\mathrm{M}})$ is tested against minimum and maximum cut-off values on all cores. If $\ln(E_{\mathrm{M}})$ is smaller than the lower cut-off or larger than the upper cut-off, all array values in $\boldsymbol{b}$, $\boldsymbol{\Lambda}^j$, $\boldsymbol{\Lambda}^{j-1}$ and $\boldsymbol{\Lambda}^{j-2}$ are rescaled by a fixed quantity.

The numerical scheme adopted for time stepping the modes remains the same as in the serial codes. Thus the memory allocated to each core stores its respective piece of $\boldsymbol{b}$, $\boldsymbol{\Lambda}^j$, $\boldsymbol{\Lambda}^{j-1}$ and $\boldsymbol{\Lambda}^{j-2}$. In the initialisation of the serial simulations in $\mathcal{H}_N$, random initial conditions are used, as the evolution of the field is determined by the given symmetry class. The same is true of the parallel codes, with each piece being filled with random noise by its respective core. As typical in the serial codes, the field snapshots in the parallel code are taken after the last time step.

A brief discussion on the performance of the parallel codes is necessary. These codes were tested against the serial code results for $R_m=1000$ and 2000. For all

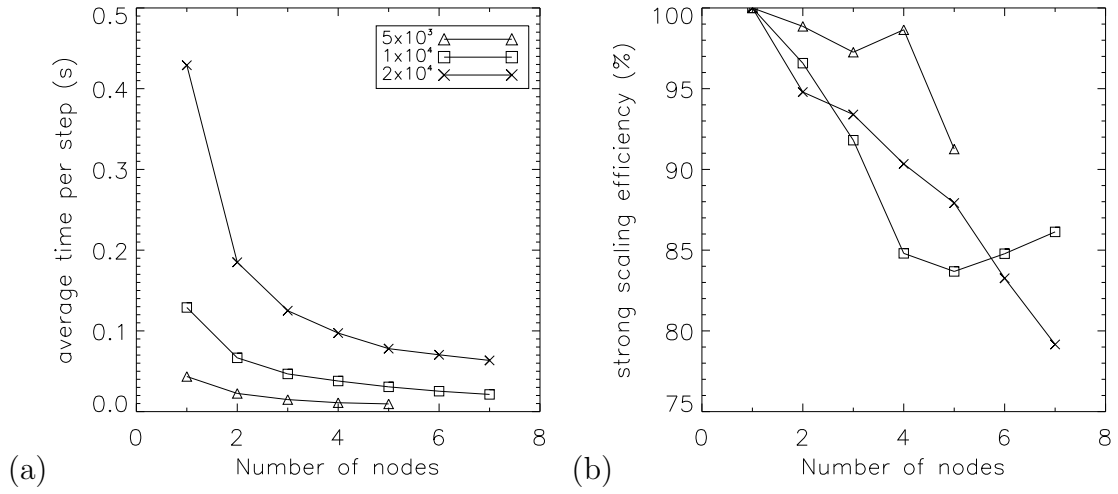

Figure 4.7: Test runs of parallel code for $R_m = 5{\times}10^3, 10^4$ and $2{\times}10^4$ identifying (a) average time per step and (b) strong scaling efficiency (eff_s) as a percentage.

five symmetry classes, results were found to agree to within the required precision (three significant figures) at both values of R_m, even when using different random initial conditions, though long run times facilitated the comparison.

Although code optimisation isn't necessarily a priority, a brief analysis of scalability has been undertaken with results given in figure 4.7. With strong scaling efficiency, the assumption is that the problem size (in our case total number of modes) stays the same while the number of processing units is increased; it is calculated by

$$\text{eff}_\text{s} = t_1/(pt_p), \tag{4.47}$$

where t_1 is the total time taken for one processing unit and t_p is the time taken for p processing units. Alternatively, weak scaling efficiency, $\text{eff}_\text{w} = (t_1/t_p)$, assumes that the workload per processing unit is constant but the problem size increases; this is not appropriate in our case. Average time per step and strong scaling efficiency (as a percentage), are both found for three values of R_m: $5000, 10000$ and 20000. Although t_1 and t_p typically represent total time using 1 and p processing units respectively, we use average time per step, and define the number of processing units as the number of nodes.

The expected p^{-1} decay profile of scaling is seen (figure 4.7(a)) but the scaling

efficiency reveals more information. With $R_m = 5\times10^3$, we see that the efficiency does not drop significantly until the problem is spread over 5 nodes. Contrariwise, with $R_m = 10^4$, efficiency drops significantly for each extra node used but then slows down and actually increases for greater than 5 nodes. This suggests that there can be significant differences in efficiency of particular breakdowns of $\mathcal{H}_N$, owing to the way that the script produces the decompositions. The increase for $R_m = 10^4$ at 6 and 7 nodes is interesting, though without further analysis is not yet understood. With $R_m = 2\times10^4$, however, the scaling efficiency decrease is approximately linear, which is more expected in these types of parallel problems. It may be the case, especially with breakdowns over many cores, that the CPU cache hit-rate is improved; the numerical operations are reduced in size and so can be carried out in a CPU's own memory cache more often, reducing the frequency of time-consuming memory-access requests. It is also possible that these results also reflect varying performance of the various nodes, despite being apparently physically identical.

Ideally, a fine balance would be struck between spreading the workload and the amount of communication carried out, as with parallelising any such problem. Clearly, communication is costly but spreading the work over more cores counteracts this cost. The object of this thesis was never to optimise the problem for such supercomputers, but given more time and in the effort to obtain more results, a more detailed analysis would be appropriate.

Finally, it should be noted that the parallel code was used in determining growth rates for $R_m \geq 3000$ for all symmetry classes. These results are given, in combination with serial results, in chapter 5.

4.4 The Arnoldi eigensolver

In the introduction, it was made clear that the kinematic dynamo problem has well-defined eigenvalues. Through time stepping, these eigenvalues are manifested in the exponential growth or decay of the magnetic field and it was explained,

earlier in the chapter, how the eigenvalues are produced using a time-series of the total magnetic energy E_{M}. It is also possible to solve this problem through the use of an eigensolver, that is a routine designed to calculate the eigenvalues and corresponding eigenvectors. This section describes how we construct the reduced 1:1:1 ABC dynamo problem for use with a particular eigensolver.

4.4.1 Formulation of eigenproblem

We consider the dynamo problem as a matrix eigenvalue system in the form

$$\lambda\tilde{\boldsymbol{b}} = -\boldsymbol{u}\cdot\nabla\tilde{\boldsymbol{b}} + \tilde{\boldsymbol{b}}\cdot\nabla\boldsymbol{u} + \eta\nabla^2\tilde{\boldsymbol{b}} \equiv \boldsymbol{A}\tilde{\boldsymbol{b}}, \tag{4.48}$$

where we have replaced $\boldsymbol{b}$ using (4.3) in (1.13). If it is possible to describe the action of the terms on the right-hand side of (1.13) as the action of a matrix $\boldsymbol{A}$ on $\tilde{\boldsymbol{b}}$, then the eigenvalues of the system are eigenvalues of $\boldsymbol{A}$, with distinct eigenvalues for each value of $R_m = \eta^{-1}$.

We aim to use the symmetries and so the reduced domains for this problem. As will be explained when the particular method is introduced, Hermitian symmetry cannot be used, thus requiring the problem to be solved in $\mathcal{G}_N$. Dealing with this problem in Fourier space, $\tilde{\boldsymbol{b}}$ is the vector containing all Fourier modes, and we know that $\mathcal{G}_N$ contains approximately $N^3/24$ modes. In addition, when constructing $\tilde{\boldsymbol{b}}$ as a vector containing all modes, the components of each mode $(\tilde{X}_{l,m,n}, \tilde{Y}_{l,m,n}, \tilde{Z}_{l,m,n})$ must be given individually, and so the vector $\tilde{\boldsymbol{b}}$ for a given problem is approximately $M = N^3/8 \sim O(N^3)$ in length. The matrix $\boldsymbol{A}$ is of size $M\times M$, therefore making it difficult to construct explicitly, even for small N. We are already aware, however, that a mode couples to only itself and six neighbours from the flow $\boldsymbol{u}$; $\boldsymbol{A}$ must only contain $7M$ non-zero entries. This makes $\boldsymbol{A}$ very sparse and incredibly impractical to construct and store, especially for large R_m. The use of an Arnoldi iterative method is therefore appropriate as $\boldsymbol{A}$ need never be explicitly given. As an added advantage, the method is also most effective at finding eigenvalues of sparse matrices such as ours, though it can still be used for densely-populated matrices.

We choose to implement the Arnoldi Package (ARPACK), available within the NAG scientific library, which solves problems of the form

$$\boldsymbol{A}\boldsymbol{w} = \lambda \boldsymbol{w}, \tag{4.49}$$

using an Arnoldi (iterative) method. The user code is required to return $\boldsymbol{A}\boldsymbol{w}$ when the ARPACK subroutine supplies any vector $\boldsymbol{w}$. Ideally the user should also be able to supply $(\boldsymbol{A} - \sigma)^{-1}\boldsymbol{w}$ when the routine supplies $\boldsymbol{w}$ and a shift σ, but in our problem it is not realistic to invert the resulting linear system because of its size. Shifts can improve the rate of convergence and are discussed in the next subsection. The routine can determine eigenvalues λ of greatest or least magnitude, real part, or imaginary part; we require greatest real part.

4.4.2 Implicitly Restarted Arnoldi iteration method

The process by which the ARPACK solver estimates the eigenvalues is through the Implicitly Restarted Arnoldi iteration method; we will summarise the process by which it locates eigenvalues of $\boldsymbol{A}$. Information in this subsection is an amalgamation of material from different sources, but primarily from Trefethen and Bau (1997). The ARPACK user's guide (Lehoucq *et al.*, 1998) also provides substantial background on Arnoldi and other methods, with additional information on Implicitly Restarted schemes taken from Saad (2011).

One of the most well-known methods for finding the eigenvalue of a matrix $\boldsymbol{A}$ with largest magnitude is the power iteration method. Given a vector $\boldsymbol{w}_0$, which can be randomly generated, a single iteration of the power method is described by

$$\boldsymbol{w}_{k+1} = \frac{\boldsymbol{A}\boldsymbol{w}_k}{\|\boldsymbol{A}\boldsymbol{w}_k\|}, \tag{4.50}$$

that is by applying the matrix $\boldsymbol{A}$ to $\boldsymbol{w}$ and then normalising. Assuming that $\boldsymbol{A}$ has an eigenvalue with magnitude greater than all others and that at least one element of $\boldsymbol{w}_0$ in the direction of the dominant eigenvector is non-zero, then the method will

typically allow $\boldsymbol{w}$ to converge to the dominant eigenvector $\boldsymbol{w}$ with corresponding dominant eigenvalue λ_1. There are additional conditions for convergence, however they are not relevant in this discussion. The power iteration algorithm is incredibly simple though finds only the most dominant eigenvalue. It is also slow to converge, especially if the two eigenvalues of greatest magnitude are close in size, and can be computationally costly for finding just one eigenvalue. An extension of this method is the *Krylov subspace* method, which takes advantage of the fact that the power iteration method generates all successive vectors $\boldsymbol{A}^j\boldsymbol{w}$ in the process, without ever specifying $\boldsymbol{A}$ explicitly. Rather than discarding these vectors, which often retain useful information about other eigenvector directions, the Krylov method considers a linear subspace of these vectors, given by

$$\mathcal{K}_k = \text{span}\{\boldsymbol{w}, \boldsymbol{A}\boldsymbol{w}, \boldsymbol{A}^2\boldsymbol{w}, \ldots, \boldsymbol{A}^{k-1}\boldsymbol{w}\}. \tag{4.51}$$

A *Krylov matrix* $\boldsymbol{K}_k$ is defined as a $k \times n$ matrix with the vectors of the Krylov subspace as its columns. These vectors, and thus matrix columns, are not orthogonal in their definition. As the power method often converges on the dominant eigenvector (i.e. for successively larger k, the last column of the Krylov matrix converges to the eigenvector), the vectors in $\mathcal{K}_k$ may tend closer to linear dependence and so a scheme that orthogonalises all vectors (or as many as possible) in $\mathcal{K}_k$ can be utilised, if multiple eigenvectors and eigenvalues are sought.

Returning to the matrix $\boldsymbol{A}$; a QR decomposition is always possible such that $\boldsymbol{A} = \boldsymbol{Q}\boldsymbol{R}$, where $\boldsymbol{Q}$ is an orthogonal matrix (whose columns form an orthonormal basis for the space of column vectors of $\boldsymbol{A}$) and $\boldsymbol{R}$ is an upper triangular matrix whose diagonal elements are eigenvalues of $\boldsymbol{A}$. Stable algorithms exist to find such decompositions, but such a breakdown relies on $\boldsymbol{A}$ being real, that is $\boldsymbol{A} = \boldsymbol{A}^*$, where * denotes the conjugate. For $\boldsymbol{A} \in \mathbb{C}^{n\times n}$, $\boldsymbol{Q}$ is instead unitary, satisfying $\boldsymbol{Q}^{\text{H}}\boldsymbol{Q} = \boldsymbol{I}$, where $^{\text{H}}$ denotes the conjugate transpose. One method is described by the Gram-Schmidt process, which uses projections on successive vectors in a vector space; this is akin to transforming $\boldsymbol{A}$ one column at a time until every column is

orthonormal. The projection is defined by

$$\mathrm{proj}_{\boldsymbol{c}}(\boldsymbol{d}) = \frac{\langle \boldsymbol{d}, \boldsymbol{c} \rangle}{\langle \boldsymbol{c}, \boldsymbol{c} \rangle} \boldsymbol{c}, \tag{4.52}$$

where $\langle \cdot, \cdot \rangle$ is an inner product. The initial step is setting $\boldsymbol{c}_1 = \boldsymbol{d}_1$, then subsequently, $\boldsymbol{c}_2 = \boldsymbol{d}_2 - \mathrm{proj}_{\boldsymbol{c}_1}(\boldsymbol{d}_2)$ and so forth, with the k^{th} term as

$$\boldsymbol{c}_k = \boldsymbol{d}_k - \sum_{i=1}^{k-1} \mathrm{proj}_{\boldsymbol{c}_i}(\boldsymbol{d}_k). \tag{4.53}$$

It is fairly evident that this produces a QR decomposition, with $\boldsymbol{R}$ being upper triangular. The advantage is that at step k, there are k orthogonalised vectors, $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_k$, thus the partially complete process can yield eigenvalue and eigenvector pairs. This is a necessary feature, as $\boldsymbol{A}$ is $n \times n$, with $n \gg k$. The process, however, is numerically unstable owing to the build up of rounding error, which is unavoidable. An adaptation of this method is *Arnoldi iteration*, which uses a stabilised (modified) form of the Gram-Schmidt algorithm though does not produce a QR decomposition but a Hessenberg reduction (see below). The algorithm is given below.

- take arbitrary $\boldsymbol{w}$, calculate $\boldsymbol{q}_1 = \boldsymbol{w}/\|\boldsymbol{w}\|$
- loop over i from 1 to k
  - + calculate $\boldsymbol{c}_i = \boldsymbol{A}\boldsymbol{q}_i$
  - + loop over j from 1 to i
    - – calculate $h_{j,i} = \boldsymbol{q}_j^* \boldsymbol{c}_i$
    - – calculate $\boldsymbol{c}_i = \boldsymbol{c}_i - h_{j,i}\boldsymbol{q}_j$
  - + calculate $h_{i+1,i} = \|\boldsymbol{c}_i\|$
  - + calculate $\boldsymbol{q}_{i+1} = \boldsymbol{c}_i / h_{i+1,i}$

After k steps, the algorithm has produced $\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots, \boldsymbol{q}_k$. It is clear that this is similar to the Gram-Schmidt process, as each vector $\boldsymbol{q}_j$ is dealt with in turn and

all previous vectors $\boldsymbol{q}_1$ to $\boldsymbol{q}_{j-1}$ are projected out through $\boldsymbol{c}_i = \boldsymbol{c}_i - h_{j,i}\boldsymbol{q}_j$. The process differs from the Gram-Schmidt method, however, as each successive vector is formed from applying $\boldsymbol{A}$ to the previous one, and so the vectors $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_k$ form an orthonormal basis of the Krylov subspace $\mathcal{K}_k$, instead of $\boldsymbol{A}$. As this is based on the Gram-Schmidt algorithm, the process can be stopped at any point to provide partial estimates of the eigenvectors and their corresponding eigenvalues. In systems where the matrix $\boldsymbol{A}$ is very large, the full eigenvalue decomposition is not necessarily possible and it is important to be able to estimate a subset of eigenvalues with a particular property and thus partial orthogonalisation becomes very valuable for such cases. Indeed this is the case for our dynamo eigenvalue problem for which we wish to approximate at most 10 eigenvalues.

From the algorithm above, the terms $h_{i,j}$ have not yet been discussed. These terms form elements of the *Hessenberg matrix*, which describe the projections of $\boldsymbol{A}$ onto $\mathcal{K}_k$. We have

$$\boldsymbol{H}_k = \boldsymbol{Q}_k^{\mathrm{H}} \boldsymbol{A} \boldsymbol{Q}_k, \tag{4.54}$$

with $\boldsymbol{Q}_k$ as the matrix whose columns are $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_k$. The Hessenberg matrix is similar to a triangular matrix, except entries in the super-diagonal (for lower Hessenberg) or sub-diagonal (for upper Hessenberg) are non-zero. As the Gram-Schmidt process reduces a matrix $\boldsymbol{A}$ to a QR form, the Arnoldi iteration reduces a matrix to the form (4.54). The $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_k$ generated by this process are the orthonormal bases of the Krylov subspace $\mathcal{K}_k$ and so in a QR reduction of $\boldsymbol{K}_k$, we would find that

$$\boldsymbol{K}_k = \boldsymbol{Q}_k \boldsymbol{R}_k. \tag{4.55}$$

In the Arnoldi iteration, $\boldsymbol{K}_k$ and $\boldsymbol{R}_k$ never have to be explicitly constructed. It is clear from the algorithm that $\boldsymbol{Q}_{k+1}$ can be generated by finding $\boldsymbol{q}_{k+1}$ and adding it as a column to the end of $\boldsymbol{Q}_k$. The value of $\boldsymbol{q}_{k+1}$ is dependent on the values of all previous $\boldsymbol{q}_i$, making it a $(k+1)$-step recurrence relation with

$$\boldsymbol{A}\boldsymbol{Q}_k = \boldsymbol{Q}_{k+1} \tilde{\boldsymbol{H}}_k, \tag{4.56}$$

where $\tilde{\boldsymbol{H}}_k$ is the $(k+1)\times k$ Hessenberg matrix. As $\boldsymbol{H}$ is a projection of the matrix $\boldsymbol{A}$ onto $\mathcal{K}_k$, the required eigenvector information is contained within $\boldsymbol{H}_k$.

The algorithm above describes a single (k-step) Arnoldi iteration (i runs from 1 to k), after which we have an Arnoldi factorisation given by

$$\boldsymbol{A}\boldsymbol{Q}_k = \boldsymbol{Q}_k\boldsymbol{H}_k + \boldsymbol{c}_k\boldsymbol{e}_k^{\mathsf{T}}, \tag{4.57}$$

with $\boldsymbol{Q}_k^{\mathrm{H}}\boldsymbol{c}_k = 0$ and $\boldsymbol{e}_k$ as the k^{th} basis vector. The parameter k is chosen by the user dependent on the application and conditions of convergence of the particular problem being solved. The eigenvalues of $\boldsymbol{H}_k$ are referred to as the *Ritz values* and are estimates of the first k eigenvalues of $\boldsymbol{A}$. As the Krylov subspace is formed from successive applications of $\boldsymbol{A}$ to $\boldsymbol{w}$ and $\boldsymbol{H}_k$ is the orthogonal projection of $\boldsymbol{A}$ onto $\mathcal{K}_k$, it is not unreasonable to expect that eigenvalues found are 'extreme' in some aspect (owing to powers of $\boldsymbol{A}$), i.e. those of largest magnitude or largest real part. The Ritz values can be calculated through typical methods such as the QR algorithm, which computes a QR decomposition of $\boldsymbol{H}_k$.

As the Ritz values converge to a subset of eigenvalues of $\boldsymbol{A}$, the quantity $\|\boldsymbol{c}_k\|$ becomes increasingly small; we therefore want k to be large enough to allow for good approximations of the eigenvalues of $\boldsymbol{A}$. For large k, this quickly becomes memory-intensive, as all $\boldsymbol{q}_i$ need to be stored. Furthermore, as each Arnoldi vector $\boldsymbol{q}_i$ depends on all prior vectors $\boldsymbol{q}_1$ to $\boldsymbol{q}_{i-1}$, there are $O(nk)$ calculations to be carried out, becoming increasingly costly for larger k. The solution lies in *restarting* the whole process, though keeping sufficient Arnoldi vectors from the previous iteration, and thus retaining valuable eigenvalue information. On restarting, an *implicit shift* is also introduced to improve the convergence of the procedure on the desired eigenvalues by filtering out 'unwanted' eigenvalues. In ARPACK, this is achieved through the implementation of the truncated implicitly shifted QR algorithm or Francis algorithm, the details of which we will not discuss.

Say we want to find r eigenvalues of greatest magnitude $\lambda_1, \ldots, \lambda_r$ of $\boldsymbol{A}$, with $r < k$. After an Arnoldi iteration, the k Ritz values of $\boldsymbol{H}_k$ are calculated and can

be ordered by the required property, in this case descending order of magnitude. The first r Ritz values are those that we want to converge and the remaining $p = k - r$ values are used to calculate the implicit shifts in the shifted QR process. For a sample of how the shifted QR process operates, the single shift QR process for $\boldsymbol{H}_k \in \mathbb{R}^{n\times n}$ is outlined in the algorithm below, which is taken from Golub and Van Loan (1996).

- set $\boldsymbol{H}^{(1)} = \boldsymbol{H}_k$
- loop over i from 1 to p
  - + calculate QR reduction of shifted Hessenberg: $\boldsymbol{H}^{(i)} - \mu_i \boldsymbol{I} = \boldsymbol{V}_i \boldsymbol{R}_i$
  - + calculate $\boldsymbol{H}^{(i+1)} = \boldsymbol{R}_i \boldsymbol{V}_i + \mu_i \boldsymbol{I}$
- set $\boldsymbol{H}^{+} = \boldsymbol{H}^{(p+1)}$

To improve convergence on the wanted eigenvalues, we set the shifts μ_i as the unwanted Ritz values, which slows convergence on the undesired eigenvalues of $\boldsymbol{A}$. Of course, for our particular problem, $\boldsymbol{H}_k$ is complex, and so a double shift is used, which acts as two consecutive single shifts. Instead of the Ritz values, the two eigenvalues used for the double shift are calculated from the trailing 2×2 submatrix of $\boldsymbol{H}^{(i)}$. The above algorithm is computationally expensive and so implicit shifting algorithms, such as the Francis algorithm, are adopted instead. Regardless of the method used, we obtain $\boldsymbol{H}^{+}$ and $\boldsymbol{V} = \boldsymbol{V}_1 \boldsymbol{V}_2 \cdots \boldsymbol{V}_p$, which are related by $\boldsymbol{H}_k^{+} = \boldsymbol{V}^{\mathrm{H}} \boldsymbol{H}_k \boldsymbol{V}$; we use the information in these to restart the Arnoldi process. Thus after the shifts, our Arnoldi factorisation (4.57) becomes

$$\boldsymbol{A}\boldsymbol{Q}_k^{+} = \boldsymbol{Q}_k^{+} \boldsymbol{H}_k^{+} + \boldsymbol{c}_k \boldsymbol{e}_k^{\mathsf{T}} \boldsymbol{V}, \tag{4.58}$$

with $\boldsymbol{Q}_k^{+} = \boldsymbol{Q}_k \boldsymbol{V}$. Owing to the shifts, it is the case that $(\boldsymbol{e}_k^{\mathsf{T}} \boldsymbol{V})_i$ is zero-valued for $i = 1, \ldots, r-1$ and therefore when the first r columns of both sides of (4.58) are equated, we have a new r-step Arnoldi factorisation, given by

$$\boldsymbol{A}\boldsymbol{Q}_r^{+} = \boldsymbol{Q}_r^{+} \boldsymbol{H}_r^{+} + \boldsymbol{c}_r^{+} \boldsymbol{e}_r^{\mathsf{T}}. \tag{4.59}$$

Having kept the information of the r useful Ritz values, the next iteration is carried out by extending this r-step factorisation to a k-step using the Arnoldi iteration algorithm for $i = 1, \ldots, p$, with the process of implicit shifting and restarting repeated until the desired number of eigenvalues with the required property have converged. The obvious advantage of the restarts is that no more than k Arnoldi vectors need to be stored, vastly reducing the memory constraints of the Arnoldi process. The implicit shifts, although computationally costly (for instance the Francis algorithm requires $O(10n^2)$ operations; see Golub and Van Loan, 1996), can drastically improve convergence.

4.4.3 Calculating eigenvalues of the kinematic 1:1:1 ABC dynamo

With an idea of how the ARPACK routines calculate eigenvalues, we now return to the problem at hand. Before any calculations actually take place, the parameters need to be set up in order to solve the problem correctly. The first routine defines the size of the problem, as the length of the vector $\boldsymbol{w}$ is defined, and is calculated from the number of components in $\mathcal{G}_N$. In initial tests, the vector $\boldsymbol{w}$ contained all Fourier mode components in $\mathcal{G}_N$ and so the length corresponded to the number of modes in $\mathcal{G}_N$ multiplied by three. However, the zero-divergence condition (1.3), in its spectral form, can be rearranged to

$$\tilde{X}_{l,m,n} = -l^{-1}(m\tilde{Y}_{l,m,n} + n\tilde{Z}_{l,m,n}), \tag{4.60}$$

and as there are two independent components, only these two components per mode need to be stored (the third can always be reconstructed given the other two). Additional mode restrictions allow $\boldsymbol{w}$ to be further shortened and are discussed later in this section. The size of $\tilde{\boldsymbol{b}}$, along with the number of Arnoldi basis vectors to use in the Arnoldi iteration and the number of eigenvalues to be found, are entered into the first routine "F12ANF". Next, options are set to indicate the type of problem being solved. The first option is the type of eigenvalues found, which

for this problem is those with largest real part. Secondly, the problem being solved is set to 'regular', thus solving a problem of the type $\boldsymbol{A}\lambda = \lambda\boldsymbol{w}$. Thirdly, as we wish to extract the eigenvectors also, we enable the option for calculation of Ritz vectors. These choices are entered into routine "F12ARF". A further option is the number of Arnoldi iterations that should be carried out – if the routine fails to converge on eigenvalues, then after a predefined number of Arnoldi iterations the routine exits; the default number is 300.

Once the problem is set up, the routine "F12APF", which carries out the Arnoldi iterations, can be run. On first use, it provides us with the vector $\boldsymbol{w}$ that is to store the components of modes in $\mathcal{G}_N$, though initially it is filled with random values; this is similar to adopting random initial conditions in the time stepping routines. The term $\boldsymbol{A}\boldsymbol{w}$, the right hand side of (4.48), then needs to be evaluated.

After ARPACK has provided the vector $\boldsymbol{w}$, the following steps are implemented:

- A routine 'unpacks' $\boldsymbol{w}$ by placing successive elements of the vector as magnetic field mode components $\tilde{\boldsymbol{b}}_{l,m,n}$ for modes in the fundamental domain $\mathcal{G}_N$ (4.31). For representations III, IV and V the various fields $(\tilde{\boldsymbol{b}}_j)_{l,m,n}$ are unpacked in turn. The process is detailed later in this section.

- Modes that neighbour $\mathcal{G}_N$ are filled in as detailed in section 4.2.2 according to the representation in use. This employs the same mode copying routines as in the ($\mathcal{G}_N$) time stepping codes.

- $\boldsymbol{A}\tilde{\boldsymbol{b}}$, the right-hand side of (4.48), is now evaluated. Note that this is almost the same as time stepping the time evolution problem (1.13) and only the diffusion terms are treated differently, otherwise the routine used is from the time stepping code. The results for the magnetic field $\tilde{\boldsymbol{b}}$ for I and II, or fields $\tilde{\boldsymbol{b}}_j$ for III, IV and V are stored.

- These fields are now 'packed', mode by mode, back into a vector $\boldsymbol{w}$, precisely the inverse of the first step, and this is returned to the ARPACK routine.

The order that modes are unpacked from $\boldsymbol{w}$ to $\tilde{\boldsymbol{b}}$ is fairly arbitrary, though must be

consistent within each simulation. We choose to unpack all 'central modes' first, followed consecutively by sets of modes that lie on the edge of $\mathcal{G}_N$ (those that do not have all neighbours). Mode components are unpacked in turn, followed by individual fields for the same mode. For instance in class III, $\tilde{Y}^{(1)}_{l,m,n} = w_j, \tilde{Z}^{(1)}_{l,m,n} = w_{j+1}, \tilde{Y}^{(2)}_{l,m,n} = w_{j+2}$ and $\tilde{Z}^{(2)}_{l,m,n} = w_{j+3}$, where $\left(\tilde{\boldsymbol{b}}^{\mathrm{III}}_1\right)_{l,m,n} = \left(\tilde{X}^{(1)}_{l,m,n}, \tilde{Y}^{(1)}_{l,m,n}, \tilde{Z}^{(1)}_{l,m,n}\right)$ and $\left(\tilde{\boldsymbol{b}}^{\mathrm{III}}_2\right)_{l,m,n} = \left(\tilde{X}^{(2)}_{l,m,n}, \tilde{Y}^{(2)}_{l,m,n}, \tilde{Z}^{(2)}_{l,m,n}\right)$. Regardless of how values are unpacked from $\boldsymbol{w}$, the important step is to pack them back into $\boldsymbol{w}$ in exactly the same order that they are unpacked; this ensures that we return $\boldsymbol{Aw}$ when given $\boldsymbol{w}$.

At the beginning of the section, the issue of being unable to exploit Hermitian symmetry in the eigenproblem was raised; we now discuss this in more detail. It is the case that any real function is Hermitian symmetric; the Fourier transform of a real function is a Hermitian function. As our magnetic field is real, it has the Hermitian symmetry property given in (4.5), which allows half of the frequency domain $\mathcal{F}_N$ to be used to simulate the full field, as has been done in previous studies. As already explained in section 4.1, this requires an additional mode copying operation, given by (4.12), involving complex conjugation, which is a non-linear operation.

In numerically resolving magnetic fields driven by the 1:1:1 ABC flow, we also have the flow's 24 symmetries to consider. As described earlier in the chapter, the full frequency domain $\mathcal{F}_N$ can be reduced to $\mathcal{G}_N$. As $\mathcal{G}_N$ is a reduction from the full domain $\mathcal{F}_N$, and not the half-domain, there is no assumption that the field is Hermitian symmetric and so the mode copying operations (4.33)–(4.35) do not involve complex conjugation; they are all linear operations. This means that the map $\boldsymbol{w} \to \boldsymbol{Aw}$ in $\mathcal{G}_N$ is linear and so suitable for the Implicitly Restarted Arnoldi method, as linearity is a prerequisite for the use of such methods. The fundamental domain $\mathcal{H}_N$, however, is found by using all 24 symmetries to reduce the frequency half-domain; each symmetry maps a general mode $\boldsymbol{b}_{l,m,n}$ to two distinct modes in $\mathcal{F}_N$, one map involving the symmetry and the other map involving the symmetry composed with (4.5). This results in mode copying operations (4.42), (4.44) and (4.45), some of which involve a complex conjugation operation. The map $\boldsymbol{w} \to \boldsymbol{Aw}$

in $\mathcal{H}_N$ is therefore a non-linear map and so we are unable to use the Arnoldi method in this reduced domain. We therefore calculate eigenvalues in $\mathcal{G}_N$, the fundamental domain without Hermitian symmetry.

An important difference concerns the degrees of freedom of the magnetic field in Fourier space. Focusing on $\mathcal{G}_N$ for the sake of argument, the symmetries for a given representation may require certain modes on the boundary of $\mathcal{G}_N$ to be zero, or to be related to other modes. For the purposes of time stepping this does not seem to be important, as simulated fields very rarely experience symmetry-breaking (i.e. loss of imposed symmetry) and so these conditions are not imposed explicitly on the field in our simulations. It is suspected, however, that spurious symmetry-breaking fields found in certain heavily decaying regimes may be the result of excess degrees of freedom in the simulations (see discussion based around figure 4.10(a) in section 4.5). Initially these conditions were also not imposed on magnetic fields for the Arnoldi solver, but it was found that the routine often converged to a spurious eigenvalue and corresponding eigenvector (in which the symmetries are broken) or, more frequently, was unable to converge at all. This serious problem led to the formulation of restrictions on Fourier mode values, of which the calculations are now briefly outlined.

The vector $\boldsymbol{w}$ that holds all the magnetic field information is therefore required to contain precisely the correct number of degrees of freedom, without duplication or zeros. To do this we start by exploiting $\nabla \cdot \tilde{\boldsymbol{b}} = 0$ to reduce the number of components per mode to two instead of three through (4.60), as already described, and so only store $\tilde{Y}_{l,m,n}$ and $\tilde{Z}_{l,m,n}$. The next step is to consider modes on the boundary of $\mathcal{G}_N$ which are mapped to themselves or to other boundary modes under symmetries, leading to either duplicated information, or modes forced to be zero. As an example, we take the mode $\tilde{\boldsymbol{b}}_{l,0,0}$ in classes I and II. The relevant transformations here are b, b^2 and b^3, as they all map $x \to x$, leaving the wave number $(l, 0, 0)$ unchanged. Thus these three transformations map $\tilde{\boldsymbol{b}}_{l,0,0}$ to itself.

If we recall that $\mathsf{b}(\boldsymbol{x}) = (x - \pi/2, \pi/2 - z, y + \pi/2)$, then we have that

$$\tilde{\boldsymbol{b}}_{l,0,0} = \mathrm{e}^{\mathrm{i}l\pi/2}\, M^{\alpha}(\mathsf{b}^3)\, \boldsymbol{J}_{\mathsf{b}} \tilde{\boldsymbol{b}}_{l,0,0}. \tag{4.61}$$

After breaking it down to components of the field, we then have

$$\tilde{X}^{\alpha}_{l,0,0} = \mathrm{e}^{\mathrm{i}l\pi/2}\, M^{\alpha}(\mathsf{b}^3)\tilde{X}^{\alpha}_{l,0,0}, \quad \tilde{Y}^{\alpha}_{l,0,0} = -\mathrm{e}^{\mathrm{i}l\pi/2}\, M^{\alpha}(\mathsf{b}^3)\tilde{Z}^{\alpha}_{l,0,0}, \quad \tilde{Z}^{\alpha}_{l,0,0} = \mathrm{e}^{\mathrm{i}l\pi/2}\, M^{\alpha}(\mathsf{b}^3)\tilde{Y}^{\alpha}_{l,0,0}. \tag{4.62}$$

The divergence-free condition implies that $\mathrm{i}l\tilde{X}_{l,0,0} + 0\tilde{Y}_{l,0,0} + 0\tilde{Z}_{l,0,0} = 0$ and thus $\tilde{X}^{\mathrm{I}}_{l,0,0} = \tilde{X}^{\mathrm{II}}_{l,0,0} = 0$. As $\tilde{Y}_{l,0,0}$ maps to $\tilde{Z}_{l,0,0}$ and vice-versa, we find that

$$\tilde{Y}^{\alpha}_{l,0,0} = -\mathrm{e}^{\mathrm{i}l\pi}\, M^{\alpha}(\mathsf{b}^2)\tilde{Y}^{\alpha}_{l,0,0}, \qquad \tilde{Z}^{\alpha}_{l,0,0} = -\mathrm{e}^{\mathrm{i}l\pi}\, M^{\alpha}(\mathsf{b}^2)\tilde{Z}^{\alpha}_{l,0,0}, \tag{4.63}$$

as $M^{\alpha}(\mathsf{b}^3)M^{\alpha}(\mathsf{b}^3) = M^{\alpha}(\mathsf{b}^2)$. Since $M^{\alpha}(\mathsf{b}^2) = 1$ for I and II, we are left with the conclusion that either $\mathrm{e}^{\mathrm{i}l\pi} = -1$, or both $\tilde{Y}^{\alpha}_{l,0,0} = \tilde{Z}^{\alpha}_{l,0,0} = 0$. We know that $\mathrm{e}^{(2k-1)\mathrm{i}\pi/2} = -1$ for $k \in \mathbb{Z}$ and so

$$\tilde{Y}^{\mathrm{I}}_{2k,0,0} = \tilde{Y}^{\mathrm{II}}_{2k,0,0} = \tilde{Z}^{\mathrm{I}}_{2k,0,0} = \tilde{Z}^{\mathrm{II}}_{2k,0,0} = 0. \tag{4.64}$$

Finally for $\tilde{Y}^{\alpha}_{2k-1,0,0}$ and $\tilde{Z}^{\alpha}_{2k-1,0,0}$, we recall the relationships in (4.62) and use the fact that $\mathrm{e}^{(2k-1)\mathrm{i}\pi/2} = (-1)^k\mathrm{i}$ to produce

$$\tilde{Z}^{\mathrm{I}}_{2k-1,0,0} = (-1)^{k+1}\mathrm{i}\tilde{Y}^{\mathrm{I}}_{2k-1,0,0}, \qquad \tilde{Z}^{\mathrm{II}}_{2k-1,0,0} = (-1)^{k}\mathrm{i}\tilde{Y}^{\mathrm{II}}_{2k-1,0,0}. \tag{4.65}$$

This provides a taste of how mode restrictions are calculated. The picture is a little more complicated for classes III, IV and V, as there are multiple fields being mapped to each other through the irreducible representations. This often, but not always, allows more freedom for components (i.e. fewer forced to be zero) of edge modes in these classes. Similar calculations were made for modes $\tilde{\boldsymbol{b}}_{l,l,l}$ which map to themselves through d and d^2. Components that are forced to be zero by the symmetries are not stored in $\boldsymbol{w}$, reducing the number of components needed.

Additionally, pairs of Fourier modes within $\mathcal{G}_N$ that are related by the symmetries

present surplus degrees of freedom and allow for symmetry breaking (and thus spurious eigenvalues) over sufficiently long calculations due to rounding error. We aim to make it so that every component of every mode in $\boldsymbol{w}$ cannot be mapped to another component through the symmetries, ergo removing all degrees of freedom. An example would be modes $\tilde{\boldsymbol{b}}_{l,l,n}$ and $\tilde{\boldsymbol{b}}_{l,m,l}$ with $m = n$, which are related through d and d^2. The solution is to simply remove one of these sets from the modes stored in $\boldsymbol{w}$; $\tilde{\boldsymbol{b}}_{l,m,l}$ is chosen for removal. The same applies to $\tilde{\boldsymbol{b}}_{l,0,n}$ and $\tilde{\boldsymbol{b}}_{l,m,0}$ with $m = n$, and so we remove $\tilde{\boldsymbol{b}}_{l,0,n}$. Figure 4.8 visualises the different types of modes for unspecified l; removed modes are hashed, coloured modes have imposed symmetry restrictions and white modes are unrestricted by the symmetries. After the removal

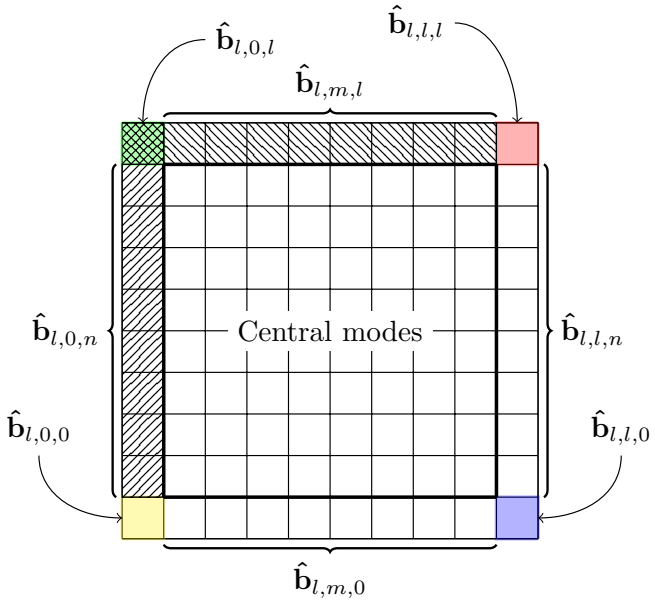

Figure 4.8: Cross-section of $\mathcal{G}_N$, for general l, displaying different mode types. Hashed modes are removed, coloured modes are restricted and white modes are unrestricted.

of all mode components that are either zero or that can be calculated from other components within $\mathcal{G}_N$, $\boldsymbol{w}$ is left with only the minimal set of components to solve the problem, and so zero degrees of freedom. On unpacking $\boldsymbol{w}$ to $\tilde{\boldsymbol{b}}$, the removed modes and components are reconstructed from the remaining modes, in order to have all values required to carry out the operation $\boldsymbol{Aw}$. Obviously, the length of $\boldsymbol{w}$ differs in each symmetry class, owing to different numbers of fields and now, additionally, differing numbers of restricted mode components. The full summary

of all restrictions of mode values is given in Appendix B, along with the length of the vector $\boldsymbol{w}$ for a given N.

In defining the length of $\boldsymbol{w}$, which is a different function of N for each class I–V, the 'safe' values of N were taken from the time stepping procedures. The number of Arnoldi (basis) vectors (n_{bv}), however, was determined at simulation time. The documentation for this suite of routines sets a lower bound for n_{bv} as $2n_{\mathrm{ev}} + 1$, where n_{ev} is the number of eigenvectors sought. Clearly, it is in our interest to keep n_{bv} as low as possible, in order to keep computing time low for the Arnoldi iterations. However, finding a minimum value that is still large enough for the routine to converge on all required eigenvalues is a significant difficulty and will be discussed in the next section.

4.5 Comparison of time stepping and Arnoldi solver

In this section, we draw a comparison between the time stepping methods and the use of the Arnoldi solver in locating the eigenvalues of the respective symmetry classes. For this purpose, $R_m = 10$ and 100 were chosen to demonstrate the major differences between these methods, as these values are relatively small. To facilitate the comparison between these very different methods, the only parameter common to both, the resolution, has been standardised as $N = 2R_m^{1/2}$, as the resolution required scales as $R_m^{1/2}$ (Moffatt and Proctor, 1985; Galloway and Frisch, 1986). These values of N have been tested for convergence of growth rates to three significant figures. We note, however, that for $R_m = 100$, the value of $N = 22$ is used; this is was unintentional but is not significant as both methods use this value. To keep the comparison accurate, the time stepping was carried out in $\mathcal{G}_N$. The results are given in table 4.2 (time stepping) and table 4.3 (Arnoldi iteration).

In the time stepping runs, the typical parameters (N and Δt) from table 4.1 are used. The parameter T represents the total time that simulations needed to be run for: in time stepping, the simulation needs to be run long enough for the time

Table 4.2: Typical values for time stepping codes

R_m	rep	N	Δt	T	memory	comp. time
10	I	8	0.01	200	1196KB	1s
	II	8	0.01	200	1196KB	1s
	III	8	0.01	200	1460KB	1.5s
	IV	8	0.01	200	1720KB	2s
	V	8	0.01	200	1720KB	2s
100	I	22	0.005	800	4892KB	2m 44s
	II	22	0.005	200	3836KB	41s
	III	22	0.005	800	7268KB	3m 47s
	IV	22	0.005	200	8588KB	1m 11s
	V	22	0.005	300	8852KB	1m 48s

Table 4.3: Typical values for Arnoldi eigenvalue solver codes

R_m	rep	N	No. of solns.	basis vectors	memory	comp. time
10	I	8	1	7	1424KB	$\ll$ 1s
	II	8	1	7	1452KB	$\ll$ 1s
	III	8	1	11	1656KB	$\ll$ 1s
	IV	8	1	15	1872KB	$\ll$ 1s
	V	8	1	11	1860KB	$\ll$ 1s
100	I	22	2	110	14MB	2m 24s
	II	22	1	80	11MB	1m 34s
	III	22	5	184	44MB	11m 35s
	IV	22	4	250	86MB	17m 59s
	V	22	2	230	79MB	10m 30s

series of $\ln(E_{\mathrm{M}})$ to have settled sufficiently so that the growth rate γ and frequency ω (if oscillating) can be determined. The values in this column represent the minimum run time required, rounded to the nearest 100, for γ to have converged to three significant figures. Figure 4.9 shows $\ln(E_{\mathrm{M}})$ versus time for three different simulations where calculating λ is relatively easy due to dominant modes that emerge quickly and with periods of the oscillatory modes being manageably short; thus meaning that long simulation times are unnecessary.

For the Arnoldi iteration, the numbers of solutions found are noted in the 4^{th} column of table 4.3 as it identifies one of the difficulties with the Arnoldi solver, namely that it does not always locate the eigenvalue with the largest real part unless it is set to compute multiple eigenvalues. This seems to occur when there are several eigenvalues whose real parts are of the same order of magnitude. To

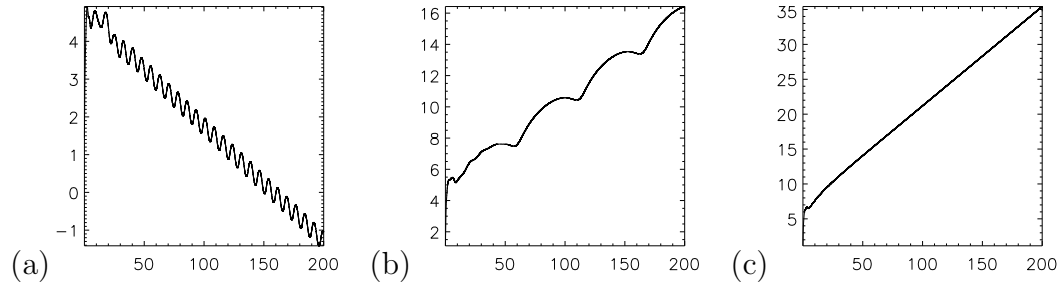

Figure 4.9: Magnetic energy profiles for (a) class II, $R_m = 100$, (b) class V, $R_m = 100$ and (c) class V, $R_m = 250$.

demonstrate this, we take representations II and III at $R_m = 100$: representation II has one overridingly dominant eigenvalue whereas representation III has several modes which are decaying at similar rates, with this being reflected in the number of solutions required. This is an inherent weakness with the Arnoldi method and will be discussed later in this section. To find the minimum number of basis vectors n_{bv} required for convergence, a script was constructed to restart the simulation if the maximum number of iterations (which we set as the default 300) was reached without successful convergence. On each restart, n_{bv} was incremented by 1 thus on successful convergence, the minimum n_{bv} is obtained. For larger simulations, this increment was increased to 5 and 10, to speed up the search. Again, we are interested in keeping this figure as low as possible since higher n_{bv} means more computation required per Arnoldi factorisation.

In general, even for low R_m simulations, the Arnoldi routines require more memory (RAM). For $R_m = 10$, the difference is only slight but at $R_m = 100$, it is substantially larger. For example, with class II the Arnoldi iteration requires approximately 3 times the memory than its time stepping counterpart and with class IV, more than 10 times the memory is required for the Arnoldi method. In terms of the computing time, at low R_m, the Arnoldi method is significantly faster but at $R_m = 100$, only faster for the trivial case of class I. In fact, the Arnoldi solver takes considerably longer for classes II–V. The larger memory usage and increased computing time of the Arnoldi solver may be related to its difficulty in resolving individual eigenvalues that are close in magnitude.

One advantage of testing two methods that solve the same problem is that the eigenvalues obtained can be tested against each other. Other than the first window of dynamo action for class II and the second window for class V, there are no results from prior studies to identify the accuracy (or inaccuracy) of the values calculated. At least for R_m up to 100, the Arnoldi method's results agreed with the dominant eigenvalue obtained for each symmetry class, so we can be sure that the different methods are at least solving the same problem, providing more confidence that the results are reliable and reproducible.

A major difficulty of the Arnoldi iteration scheme was, as already hinted, setting the values of n_{bv}. The manual for the ARPACK routines states that the choices of n_{bv} are problem specific and must be determined empirically. In general, we find that convergence improves for increased n_{bv} but so does the workload and thus computing time. The greatest issue is that n_{bv} can be very sensitive to the value of R_m, as the minimum number of basis vectors required can increase by an order of magnitude with only a small change in R_m. The lack of a clear relationship between R_m (through N) and n_{bv} greatly increases the difficulty of finding eigenvalues for even small ranges of R_m, as the required n_{bv} cannot be extrapolated from those used for lower R_m. However, n_{bv} was not found to decrease for increasing R_m and thus must be in a monotonic (increasing) relationship with R_m. The sensitivity of n_{bv} to R_m could, however, be indicative of more interesting structure, within each representation, that is not revealed through time stepping, such as multiple eigenvalues that are very close in the complex plane. In our experience, the only systematic means of finding the minimum number of basis vectors for convergence is by brute force, rerunning simulations for fixed R_m and incrementing the number of basis vectors until the routine succeeds in finding the desired number of eigenvalues. For low and intermediate R_m this is not too problematic, as simulations are relatively fast (on the order of minutes), but since increasing the number of basis vectors will also increase the computer time for a given simulation, this becomes extremely difficult with $R_m = O(10^3)$. The fact that we are unable to determine in advance whether the routine will succeed (converge) until it has ended, is at the

core of the problem, although a similar problem arises with time stepping, as it can be difficult to determine whether the field has settled to a final state or is in a long transient (an example being representation V at $R_m = 215$). There are many cases for which it can be very difficult to resolve the growth rate and so require very long run times. Three such cases are identified in figure 4.10.

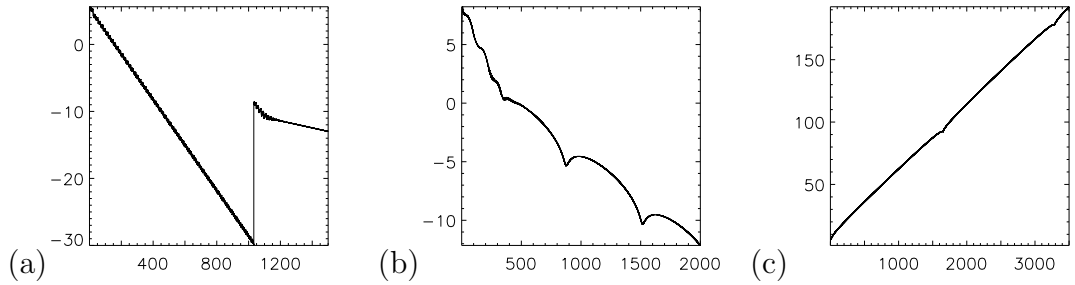

Figure 4.10: Magnetic energy profiles for (a) class II, $R_m = 190$, (b) class III, $R_m = 400$ and (c) class V, $R_m = 215$.

The first, and most difficult to resolve, is class II at $R_m = 190$, which is on the verge of a mode-change, where a subdominant mode overtakes the dominant mode. This results in two eigenmodes with very similar growth rates and requires long simulation times. In figure 4.10(a), a jump in the magnetic energy indicates where the stored arrays are rescaled to move mode values away from the machine precision limit. After the rescaling, however, a spurious steadily decaying mode is seen in which the symmetries of class II are no longer found; this should not be possible in the symmetry-reduced codes. It was initially assumed that as the mode values decayed and approached machine precision, the rounding error would become large enough to destroy the imposed symmetry. In rescaling the Fourier modes at higher values, so that their magnitude cannot approach machine precision, the spurious mode was still seen. In avoiding the rescaling altogether, we still see the spurious mode suggesting that rescaling the fields is not the cause of the issue. We speculate that the mode restrictions, which are employed in the Arnoldi method but not in time stepping, may actually be important in these heavily-decaying regimes: over sufficiently long simulations, sufficient rounding error may build up in one or more pairs of modes in $\mathcal{H}_N$ (or $\mathcal{G}_N$) that are linked by the symmetries, thus destroying the symmetry of the field. So close to a mode-change, there is insufficient time for the

field to settle to a final state before the spurious mode appears and so the growth rate for $R_m = 190$ is estimated from the oscillatory part of the energy profile, immediately prior to the rescaling (energy jump). Although it isn't apparent in the figure, the oscillations are damped, decreasing in magnitude extremely slowly and so appear to be settling towards a steady decay.

The second example is shown by class III, $R_m = 400$. In this case, the field settles to a low frequency oscillatory decay but with additional high-frequency oscillations. To resolve the growth rate properly, simulations required long run times to allow time for the high-frequency oscillations to die away. This difficulty increases with R_m as the high-frequency oscillations remain for longer. The third example is identified in figure 4.10(c) at $R_m = 215$ for class V, where $\mathrm{Im}(\lambda)$ is very small, leading to very low frequency oscillations. To resolve the λ correctly, the simulations must have very long run times. As the ranges of R_m for which the field evolution contains long transients or very low frequency modes, are not known in advance, the only reliable way to achieve this is to set default run times and then to rerun those that have not converged, repeating this process until reliable results have been obtained. This approach means that, often, a large amount of time is spent resolving λ for small ranges of R_m. Regardless of the run times required, only by visual inspection can it be known when $\ln(E_\mathrm{M})$ has converged to the dominant eigenmode, thus λ cannot be automatically calculated with reliability. Thus, for every simulation in every class, λ is manually calculated; a reliable and accurate method but incredibly time consuming.

A distinct advantage that the Arnoldi iteration method has over time stepping is that multiple eigenvalues, and thus eigenfunctions of magnetic field, can be found. In the time stepping procedure, the total magnetic energy is only a function of the fastest growing (or slowest decaying) magnetic mode and so the eigenvalue is that of the dominant mode. In the Arnoldi solver, we are able to set the number of eigenvalues to be found (i.e. n_ev), though a higher number means slower convergence. To demonstrate this ability, we set $n_\mathrm{ev} = 10$ and locate the eigenvalues in the interval $R_m \in [10, 100]$ for class V; growth rates are given in figure 4.11.

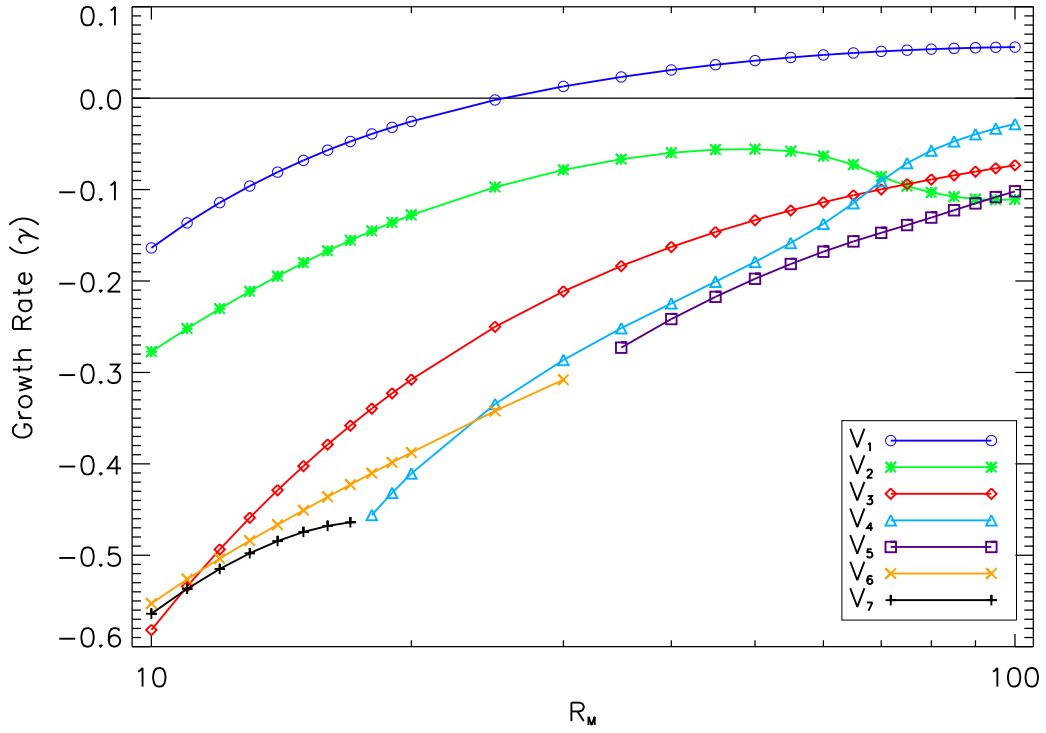

Figure 4.11: Growth rates for seven eigenmodes of class V, $R_m = 10$ to $R_m = 100$.

If the Arnoldi routine finds a complex eigenvalue, it nearly always finds its conjugate, as the eigenvalues occur in conjugate pairs. If both eigenvalues of a complex pair are found, one is discarded. Additionally, eigenmodes with fewer than four eigenvalues located are also discarded. This leaves only seven distinct eigenmodes in figure 4.11, despite locating 10 for every value of R_m. We see interesting behaviour, particularly for $R_m \geq 60$ where three subdominant branches of eigenvalues are very close in their decay rates. This possibly explains why n_{bv} is high for class V at $R_m = 100$ in table 4.3 and, indeed, we find similar circumstances in class IV; see figure 5.20(d). We also see the familiar dominant branch, labelled V_1, which forms the second window of dynamo action in Galloway and Frisch (1986).

To summarise, the Arnoldi method is generally less predictable than the time stepping procedure, owing to the fact that the relationship between R_m and the parameters that determine convergence is not clear. The Arnoldi routine is capable of producing more results and could be useful in studying mode crossings within

symmetry classes and subdominant dynamos. Unfortunately, it presently seems impractical for investigating problems with $R_m > O(10^2)$, owing to the memory and time resources needed for such simulations. Optimising the problem, so that the minimum set of mode components was present for each symmetry class, seemed to improve convergence but without further investigations into the difficulties faced, the Arnoldi method cannot supersede the time stepping for generating results at high R_m. Rather the range of the results it provides should supplement those of the time stepping method by describing subdominant mechanism and structures, and confirming previously known results, which it has done for R_m up to 100. Room for improvement lies in taking advantage of known eigenvalues and Ritz vectors for a specific R_m, and using these as a 'first guess' to improve convergence for nearby R_m runs. Additionally, taking advantage of the in-built diagnostic routine "F12AQF" may reveal more information on the convergence conditions.

In the next chapter, further results of the Arnoldi method are presented, in addition to those of the $\mathcal{H}_N$ time stepping codes.

Chapter 5

Results for the 1:1:1 ABC dynamo

Wisdom comes from experience. Experience is often a result of lack of wisdom.

Terry Pratchett

This chapter is concerned with results of the 1:1:1 ABC dynamo problem. All results presented in this chapter, unless explicitly stated, are found through the time stepping routine in fundamental domain $\mathcal{H}_N$. We present eigenvalues obtained in the form of growth rates $\gamma = \mathrm{Re}(\lambda)$ and frequencies $\mathrm{Im}(\lambda)$. Also provided are visualisations of magnetic field structure, depicting areas of strong field through 3-dimensional snapshots and 2-dimensional cross-sections to help determine finer structural details, in addition to field polarity. The chapter aims to give an overview of the types of structures that can exist within each class and describes the field evolution for oscillatory eigenmodes.

The chapter starts by summarising the eigenvalues of all five classes in two graphs; one for growth rates $\mathrm{Re}(\lambda)$ and another for $\mathrm{Im}(\lambda)$, the frequencies of the field oscillations. We discuss why kinks appear in the growth rates and classify these kinks using the behaviour of $\mathrm{Im}(\lambda)$ in the neighbourhood of the kink. The second section then goes on to investigate these kinks in more detail and looks at an

example of each of eigenvalue coalescence, decoalescence and mode crossing. For eigenvalue (de)coalescence, we show that the eigenmode does not change after a growth rate kink and that in a mode crossing, the eigenmode does change. The third section is concerned with visualising the fields and describing the structures that emerge. Each symmetry class is discussed in turn, in order of the most effective dynamos. A brief review is given of the magnetic fields in the chaotic regions at high R_m and the energy spectra are then discussed. Finally, we explore the results produced using the Arnoldi iteration method at low R_m and reanalyse the mode crossing example previously studied in the time stepping results. We also explore the magnetic field configurations produced using this method and compare these to those of the time stepping approach. The last section summarises the findings of this chapter.

5.1 Growth rates and frequencies

This section of the results focuses on the real and imaginary parts of the dominant eigenvalues within each class. The growth rates for magnetic fields in all of our simulations are summarised in figure 5.1, showing growth rate $\gamma = \mathrm{Re}\,(\lambda)$ versus R_m, with typical time traces for magnetic energy depicted in figure 4.9. Results were obtained with the serial $\mathcal{H}_N$ time stepping codes for $R_m < 3000$ and the parallel codes for $R_m \geq 3000$. The imaginary parts of the eigenvalues, corresponding to (half of) the frequencies of the magnetic energy oscillations, are shown in figure 5.2. Since the eigenvalues occur in complex conjugate pairs, the absolute value of the imaginary part is plotted. It is apparent that some of the modes, within each symmetry class, are non-oscillatory (i.e. have a real eigenvalue), whereas others are oscillatory (complex eigenvalue). Transitions between non-oscillatory and oscillatory modes for a given symmetry class may occur because of a mode crossing (giving a discontinuity in the frequency) or because of eigenvalue coalescence (in which the frequency tends to zero with a square root dependence). Additionally, we see transitions from one oscillatory mode to another, identified by jumps in

$\mathrm{Im}(\lambda)$ from one non-zero value to another non-zero value.

For the 1:1:1 ABC flow, the growth rates identify that fields in only two of the symmetry classes provide the dominant dynamo mechanism for the range of R_m investigated (up to 10^4), namely classes II and V. As found by Galloway and Frisch (1986), and reported in subsequent studies, there are two distinct windows of dynamo action, separated by a brief interval of purely decaying fields over the range of $R_m \approx 17.5$ to $R_m \approx 27$. The fields in the second window were referred to as 'symmetry-breaking', as they have fewer apparent symmetries than fields in the first window. It can now be identified that fields in the second window belong to symmetry class V, and this remains dominant in our simulations up to $R_m = 10^4$, and up to 2.5×10^4 in Bouya and Dormy (2013) (where the growth rate begins to level at approximately 0.1). The fields in classes II, IV and V give growth rates that show no sign of tending to zero in the large R_m limit, suggestive of fast dynamo action. However the limiting values for large R_m cannot be assessed from the present results and further investigation is required. The dominant field of class V emerges as a dynamo in the interval $R_m \in [25, 26]$ and the growth rate increases, reaching a local maximum of $\gamma = 0.0566$ at $R_m \approx 121$, after which the growth rate begins to fall. In the interval $R_m \in [214, 216]$, an apparent jump in the growth rate is observed. Bouya and Dormy (2013) narrow this interval down to $R_m \in [215, 215.4]$, though a more precise value is incredibly difficult to pin down as the frequency of the oscillations, seen prior to the jump, tends to zero as the critical value is approached (see discussion on difficulties of time stepping in section 4.5). The sudden growth rate increase for $R_m > R_m^{\mathrm{crit}}$ is caused by the coalescence of two complex eigenvalues to a real eigenvalue and can be seen in figure 5.2 as $\mathrm{Im}(\lambda) \to 0$; this is explored in more detail in section 5.2. The growth rate reaches another maximum at $R_m \approx 419$, with $\gamma \approx 0.799$ before falling to a minimum at $R_m \approx 2000$, with $\gamma \approx 0.0686$. From here the growth rate increases at a moderate pace and in the limit of R_m explored ($R_m = 10^4$), $\gamma = 0.0863$. There is no evidence of a mode-change within class V in this high-R_m regime and so we assume that the dominant eigenmode remains the same for $R_m > R_m^{\mathrm{crit}}$.

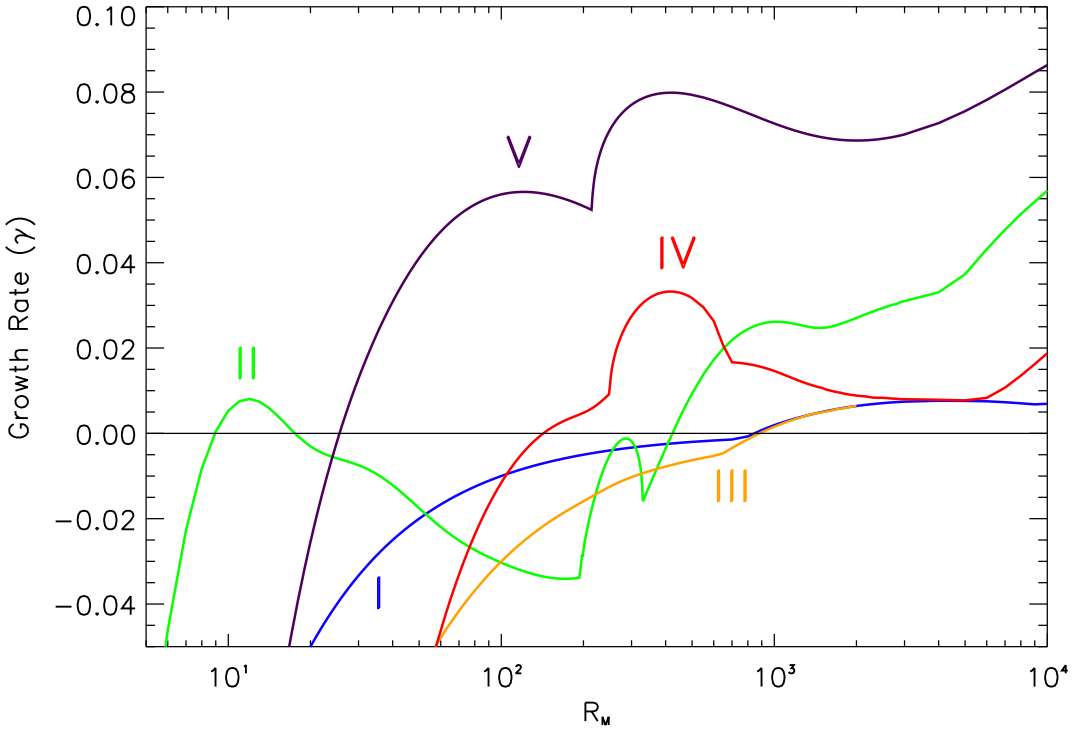

Figure 5.1: Growth rates of symmetry classes I–V in the 1:1:1 ABC dynamo for R_m up to 10^4.

In the low R_m regime, we find that the eigenmode of the first window belongs to class II and confirm the results of Arnold and Korkina (1983), Galloway and Frisch (1986) and later studies. We locate the crossing between classes II and V in the interval $R_m \in [24, 25]$, though we observe that it is at the lower end of this interval. Bouya and Dormy (2013) reduce this interval to $R_m \in [24.05, 24.10]$ and follow the class II eigenmode past the crossing with mode V, using the initial condition (1.21). The growth rates that we obtain are in agreement with those of Bouya and Dormy for as far as they manage to resolve them. The growth rate of the dominant eigenmode in class II continues to fall and reaches and apparent minimum at $R_m \approx 191$, and then experiences a sudden jump in growth rates in the interval $R_m \in [192, 194]$ in a similar fashion to that of class V. From figure 5.2, however, we can deduce that this is a change of mode within class II, and is identified by a discontinuity in $\mathrm{Im}(\lambda)$. We also see an eigenvalue decoalescence in the interval $R_m \in [328, 330]$ and the dominant eigenmode again corresponds to

one of two complex eigenvalues. After the decoalescence, the growth rate increases rapidly and the eigenmode becomes a dynamo at $R_m \approx 423$. The growth rate continues to increase and shows a similar pattern (of slowing down, decreasing and then increasing again) to that of class V. In the interval $R_m \in [4000, 5000]$, a mode-change occurs and is marked by a rapid increase in both $\mathrm{Re}(\lambda)$ and $\mathrm{Im}(\lambda)$. At $R_m = 10^4$, the class II eigenmode appears to be an effective dynamo, showing no sign of slowing in its growth rate increase. Before moving on to the other classes, we note that Galloway and O'Brian (1994) calculated a value of $\gamma \approx -0.03$ for $R_m = 100, 200$ and 400 (refer to chapter 1). We can confirm the first two of these as we have $\gamma = -0.0303$ and -0.0275 for $R_m = 100$ and 200 respectively. For $R_m = 400$, our result is substantially different, at $\gamma = -0.00355$. It is not unreasonable to expect that this mode was not well-resolved in the calculations of Galloway and O'Brian, as they would have seen it only through transient magnetic energy evolution prior to the field settling on the fastest growing mode of class V.

Magnetic fields from classes I, III and IV have not been identified in previous investigations, as they are subdominant eigenmodes and so cannot be found through conventional (full domain) time stepping. As discussed in chapter 3, class IV of the group G is a tensor product of II and V and so it is unsurprising that its growth rate profile is visually similar to that of classes II and V. The dominant eigenmode of class IV becomes a dynamo at roughly the middle of the interval $R_m \in [142, 143]$ and is an oscillating mode, with low frequency oscillations. As with class V, there is an eigenvalue coalescence, which occurs at $R_m \approx 249$ but, unlike class V, this eigenmode experiences decoalescence of its eigenvalues, which again become complex valued in the interval $R_m \in [650, 700]$ (the resolution in R_m is coarser as R_m increases due to increased computational demand). As R_m is further increased, the dominant eigenmode of class IV sees decreasing growth and appears to reach a minimum near to $R_m = 6000$. By $R_m = 7000$, it is clear that a mode change has occurred, as the dynamo effectiveness appears to improve fairly rapidly, and this is also identified in figure 5.2 by the discontinuity in $\mathrm{Im}(\lambda)$. Fields in classes I and III are very different from those of II, IV and V and the first indication of this

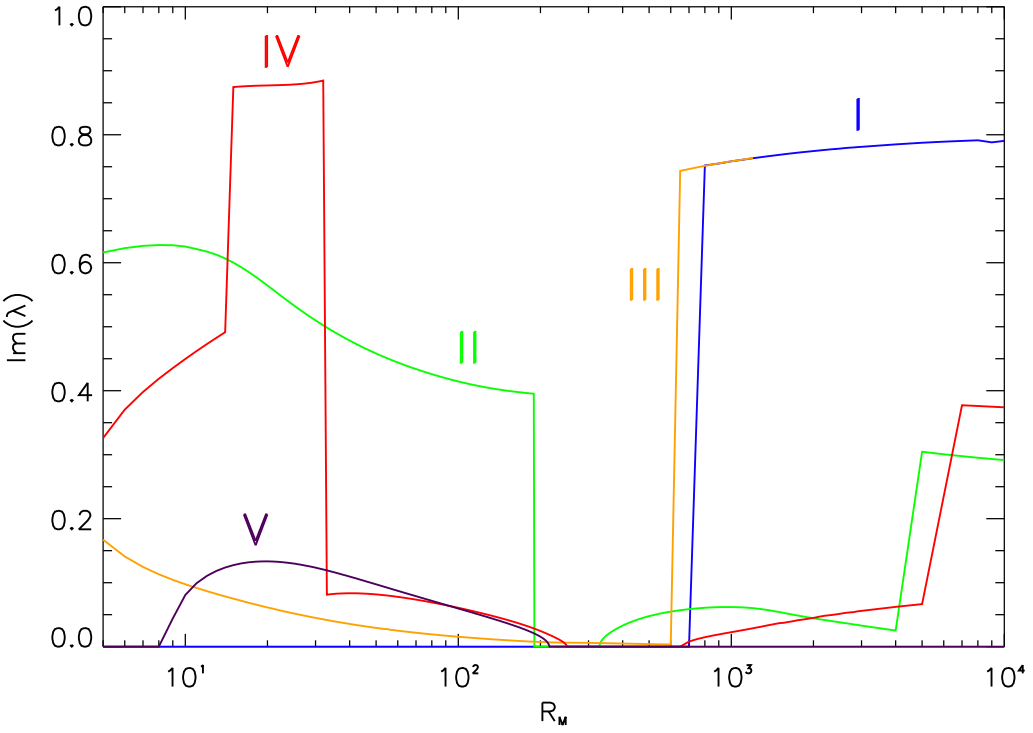

Figure 5.2: Frequencies of symmetry classes I–V for R_m up to 10^4.

is through their growth rate profiles. These have also been established to contain dynamos, but show the lowest growth rates and last onset as R_m is increased. In both cases, difficulties were encountered with calculating the leading growth rate, as fields in these two classes show long, slowly evolving transients, particularly in class III. Limitations of computer time have not allowed us to explore values of R_m for class III beyond those shown in figure 5.1. In class I, until the mode change at $R_m \approx 800$, the field with the largest growth rate is the trivial solution, with $\boldsymbol{b} = \boldsymbol{u}$, $\gamma = -R_m^{-1}$. Other magnetic fields exist within class I, though they are only seen as long transients, in which decay occurs at a greater rate than the trivial solution. In class III, the frequency appears to be decreasing rapidly as R_m increases, prior to the mode change at around 620. For larger R_m, an increase in frequency is seen, closely matching the frequency profile for class I, which may be related to the similarity in the structure of their magnetic fields; this will be discussed in section 5.3.

It is important to note that without enforcing symmetries, the dynamo growth rate will simply be the envelope of the individual growth rates of the symmetry classes, meaning that the growth rate of only the fastest growing dynamo for a given R_m will be calculated. Likewise, when imposing the symmetries, the growth rate of only the fastest growing (or slowest decaying) dynamo eigenmode within each symmetry class is found.

Overall, we are able to show that all symmetry classes of the 1:1:1 ABC dynamo are capable of dynamo action. Three symmetry classes appear to contain fast dynamo mechanisms, these being II, IV and V, as their growth rates are increasing in the limit of our simulations. Although classes I and III also contain viable dynamos at higher R_m, they display profiles that resemble those of slow dynamos, in that they have falling growth rates for increasing R_m. It is possible that the behaviour of any class may change at higher R_m, but we cannot speculate further based on current results.

5.2 Eigenvalue coalescence and mode changes

In the jump in growth rate of the dominant mode in class V (at $R_m \approx 215$), we also see that the frequency falls to zero. The lack of discontinuity suggests that this jump is not a mode change but is caused by a *coalescence* of the complex eigenvalue pair which, as R_m increases past $R_m^{\text{crit}} \approx 215$, become real-valued. Bouya and Dormy (2013) provide an explanation as to why the growth rate would suddenly appear to increase. They describe the eigenvalues in terms of two real differentiable functions α and β such that

$$\lambda = \alpha(R_m) \pm \sqrt{\beta(R_m)}, \tag{5.1}$$

and so a negative β allows for (complex) conjugate eigenvalues. The value of R_m at the mode-change must be a root of the function β such that $\beta\left(R_m^{\text{crit}}\right) = 0$. For $R_m > R_m^{\text{crit}}$, $\beta(R_m) > 0$ and so starts contributing to $\text{Re}(\lambda)$, and can explain the sudden increase in growth rate. This also suggests that a second real eigenvalue

exists for which the growth rate drops for $R_m > R_m^{\text{crit}}$. As a subdominant eigenvalue, it would be only be possible to find the corresponding mode using a more direct method such as Arnoldi iteration.

To confirm that the dominant eigenvalues for $R_m < 215$ and $R_m > 215$ belong to the same eigenmode, simulations were carried out to produce snapshots for $R_m = 200$ and $R_m = 250$. The resulting field snapshot at $R_m = 200$ was used as a seed field for simulations with $R_m = 205$ to 300 in an attempt to follow the eigenmode branch past the critical R_m. The same was carried out for the snapshot from $R_m = 250$ with simulations for $R_m = 120$ to 230. If the growth rates below and above $R_m = R_m^{\text{crit}}$ belong to different modes, this would be identified through different growth rate profiles obtained from the different seed fields. The results are summarised in figure 5.3, with both Re(λ) (solid line) and Im(λ) (dotted line) shown; the two symbols used represent the different initial conditions used. It is

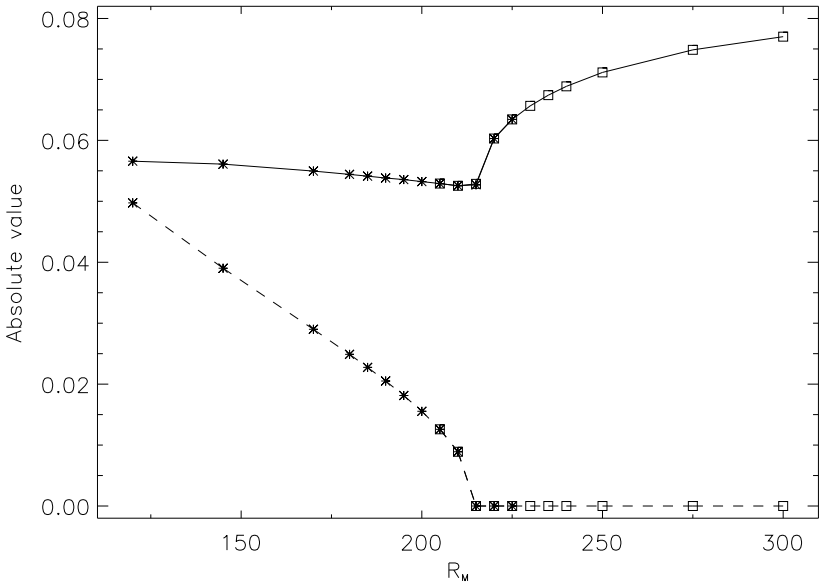

Figure 5.3: Re(λ) (solid line) and Im(λ) (dotted line) for two initial conditions; asterisks identify results using a field snapshot for $R_m > R_m^{\text{crit}}$, squares are results using a field snapshot for $R_m < R_m^{\text{crit}}$.

clear that the eigenvalues of the eigenmode from after the jump are the same as those of the eigenmode before the jump, identifying that these belong to the same eigenmode. If they were from different eigenmodes, we would expect to be able

to follow $\gamma = \mathrm{Re}(\lambda)$ past R_m^{crit} and produce two distinctly different branches of solutions that cross at $R_m = R_m^{\mathrm{crit}}$. We note that the kink in the dominant class V curve around R_m^{crit} was a source of difficulty for Galloway and Frisch (1986) as very slow oscillations led to uncertainty in the growth rate. It was not possible at the time to determine whether this corresponds to a real eigenvalue or two complex eigenvalues with small imaginary parts. It is now clear that $\mathrm{Im}(\lambda)$ goes to zero below this point. Interestingly, our interval and the interval of Bouya and Dormy (2013) are much lower than prior estimates (e.g. of Lau and Finn, 1993) of $R_m \in [300, 350]$; by this time the field evolution is non-oscillatory.

In figure 5.1, we see two kinks in the growth rate of class II for intermediate R_m. On initial inspection, the first kink (where growth rate rapidly increases), might seem to be an eigenvalue coalescence but figure 5.2 indicates that a mode change has occurred, due to the discontinuity in $\mathrm{Im}(\lambda)$. The second kink, however, is an eigenvalue *decoalescence* as $\mathrm{Im}(\lambda)$ increases (from 0) with a square-root dependence. Although classes I and III have mode changes for $R_m < 1000$, the difficulty in resolving λ in the time stepping codes makes investigation of these cases impractical; the changes in class II are easier to investigate as they occur at lower R_m. First we address the mode change at $R_m \approx 190$ by taking a field snapshot from $R_m = 250$ (after the change) and using it to seed simulations over the range $R_m \in [170, 215]$, in an effort to follow the non-oscillatory eigenmode prior to it becoming dominant. As the eigenmode we are attempting to follow would be decaying at a greater rate, we see it as transient field evolution before settling on to the fastest growing (slowest decaying) eigenmode. We then take a field snapshot at $R_m = 170$ and use this as a seed field for simulations in the same range. Again, this field is seen as transient evolution prior to the field evolution settling to that of the dominant mode. Results are seen in figure 5.4(a), with the $\square$ symbol identifying a non-oscillatory field ($R_m = 250$ seed) and $\times$ showing results of the oscillatory ($R_m = 170$ seed) field. As we are able to follow both eigenmodes past the crossing, we have shown that these two eigenmodes are distinct; otherwise we would have seen the same situation as the coalescence in class V, where no transient evolution

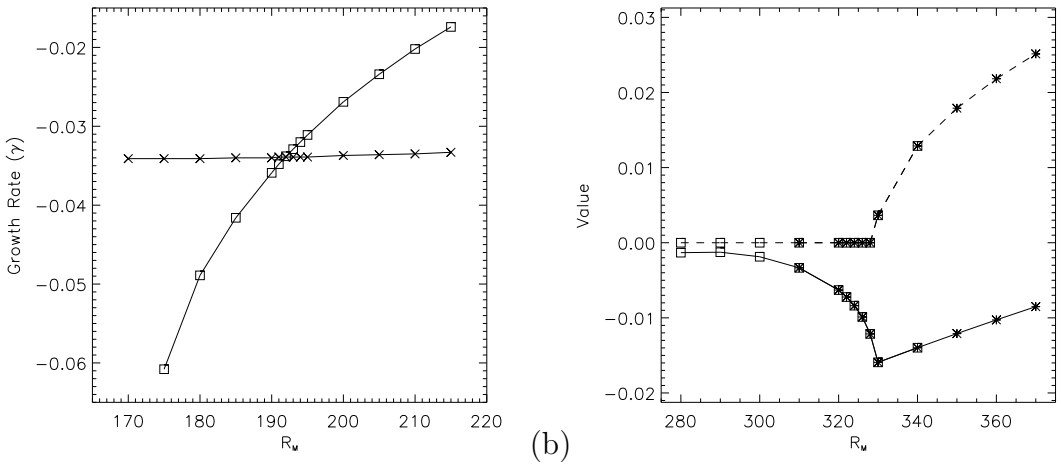

Figure 5.4: Numerical results at (a) the mode crossing and (b) eigenvalue decoalescence of class II. For the mode crossing, only the growth rate is displayed, whereas in (b) both real (solid) and imaginary (dashed) parts are shown. As with previous figure, squares indicate a post-kink seed and asterisks denote pre-kink seed.

occurs, as both seed fields correspond to the same eigenmode. We note that the non-oscillatory field is very heavily decaying for $R_m = 170$ and cannot be identified as the magnetic field is quick to settle to the (dominant) oscillating eigenmode. The non-oscillatory branch in this range of R_m is also investigated as a test case for the Arnoldi method, revealing more information (see section 5.4).

On investigating the decoalescence, two seed fields are again used to rule out the growth rate kink being caused by a mode crossing. Figure 5.4(b) shows the results of two sets of simulations with $\mathrm{Re}(\lambda)$ (solid line) and $|\,\mathrm{Im}(\lambda)|$ (dashed). As with the eigenvalue coalescence of class V, we see that solutions using both initial conditions follow the same branch with no indication of transient evolution in the magnetic energy profiles, implying that the initial conditions belong to the same eigenmode.

As eigenvalue decoalescences occur in both classes II and IV, then it is possible that one may occur in class V. For R_m up to 10^4, we see no decoalescence of the dominant eigenvalue into a complex conjugate pair and Bouya and Dormy (2013) extend this up to $R_m = 2.5{\times}10^4$ but the possibility remains that at even higher R_m, we may see an oscillatory eigenmode re-emerge. Further investigations into the 1:1:1 ABC dynamo will hopefully reveal whether this is (or is not) the case.

5.3 Field structures

In this section, we investigate the structures of the dominant eigenmodes in each symmetry class, through the use of 3-dimensional visualisations of magnetic field strength and 2-dimensional cross-sections. To visualise the fields, the Fourier modes are first 'unfolded' from $\mathcal{H}_N$ through the symmetries and mapped to the whole of the domain $\mathcal{F}_N$; this is detailed in section 4.2.4. After applying a fast Fourier transform routine to obtain the field in real space, the 3-dimensional visualisations are produced using VAPOR, a piece of software designed for viewing and analysing fluid flows or, more generally, any 2-d or 3-d vector field. The visualisations can be chosen to display field strength $|\boldsymbol{b}|$ or magnitude of certain components (i.e. b_x, b_y or b_z); we have opted for field strength. Unless explicitly stated, all 3-d visualisations are oriented so that the origin is located at the bottom back corner and the point $\boldsymbol{x} = (2\pi, 2\pi, 2\pi)$ is located at the top front corner.

For the cross-sections, IDL's contour plotting capabilities are used to visualise the magnitude of b_z at a specified z = constant. The combination of 2-d and 3-d visualisations allows for easier analysis of field structures and comparison between the fields of each class.

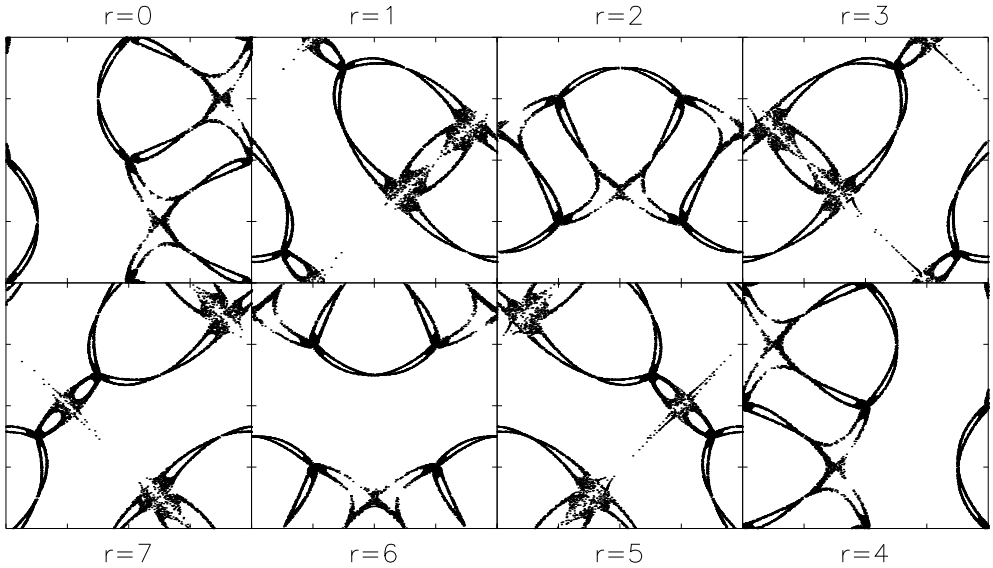

Figure 5.5: Poincaré sections of particle trajectories in the 1:1:1 ABC flow, crossing the planes of $z = r\pi/4$ for $r = 0, 1, ..., 7$, reading across then down.

We begin by giving the Poincaré sections of the 1:1:1 ABC flow for $z = 0, \pi/4, \ldots, 7\pi/4$ in figure 5.5 (after Dombre *et al.*, 1986). There is a network of thin bands of chaos visible, and in between (the white regions), there is integrable motion, referred to as the 'principal vortices' of the flow (discussed in chapter 1). At $z = \pi/4$ (corresponding to the second panel, $r = 1$), there are stagnation points $\boldsymbol{u} = \mathbf{0}$ lying at $\boldsymbol{r}_1 = (5, 3, 1)\pi/4$ (with a 2-d stable manifold, 1-d unstable manifold, type 'α') and $\boldsymbol{r}_2 = (7, 5, 1)\pi/4$ (with a 2-d unstable manifold, 1-d stable manifold, type 'β'). At $z = \pi/4$, we also see midpoints of two different heteroclinic lines that connect α- and β-type stagnation point pairs. One is seen at $\boldsymbol{r}_3 = (1, 1, 1)\pi/4$ and the other at $\boldsymbol{r}_4 = (3, 7, 1)\pi/4$. We will be relating the field structures to these flow features by taking cross-sections at $z = \pi/4$ also.

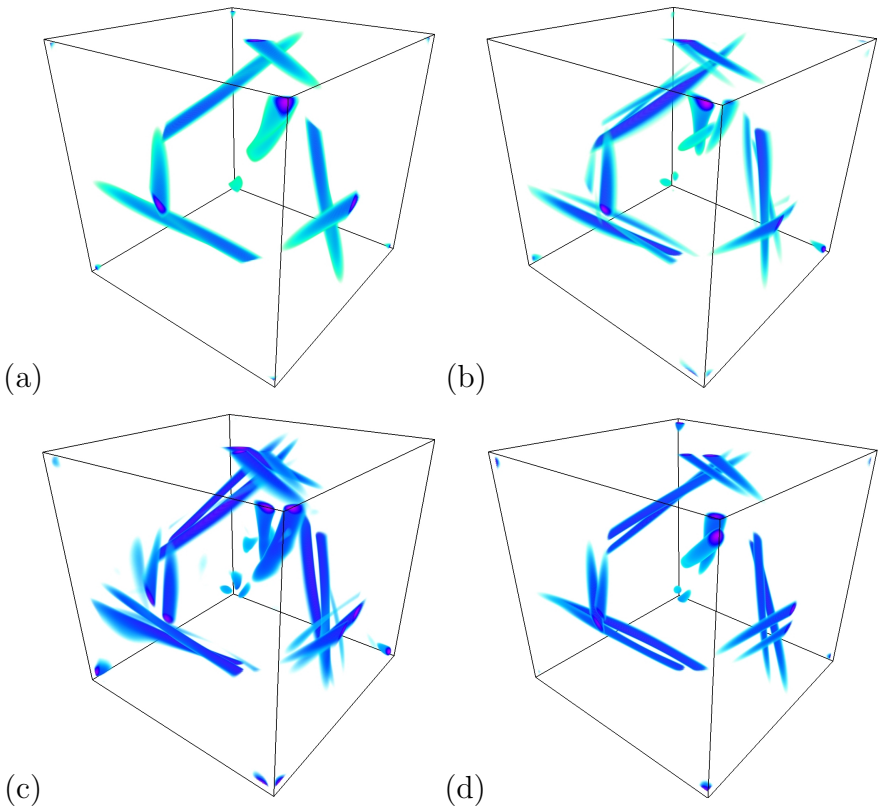

Figure 5.6: Visualisations of magnetic field strength $|\boldsymbol{b}|$ for the dominant eigenmodes of classes (a) II to (d) V for $R_m = 100$.

Figure 4.9 (in the previous chapter) gives examples of plots of total magnetic energy versus time for two oscillatory cases (a,b) and one non-oscillatory case (c). In

the oscillatory cases, the structure of the field is time dependent and we have visualised fields over times during a cycle. The symmetries in each class force the magnetic field to take certain arrangements, leading to the individual structures that are seen in figure 5.6, which shows three-dimensional magnetic field snapshots of the dominant eigenmodes of each class at $R_m = 100$. We see the familiar cigar-like structures (which we simply refer to as 'cigars' in what follows), some single and some multiple, concentrated on the straight-line heteroclinic connections, or separatrices, determined by Dombre *et al.* (1986). Note that we have chosen a moderate value of R_m as visualisations at larger values simply give similar, but more concentrated structures and are in practice harder to interpret. We do not show the field for class I, at this value of R_m, which is just $\boldsymbol{b} = \boldsymbol{u}$; we would simply see an isosurface of fluid speed $|\boldsymbol{u}|$ in the 1:1:1 flow and no cigars.

The remaining classes II to V are all oscillatory at this value of R_m and have similarities in structure but differences in detail. To give another view we show in figure 5.7 the field component b_z perpendicular to a slice $z = \pi/4$ through the fluid domain. In each symmetry class, from II to V, in figure 5.6 the field has

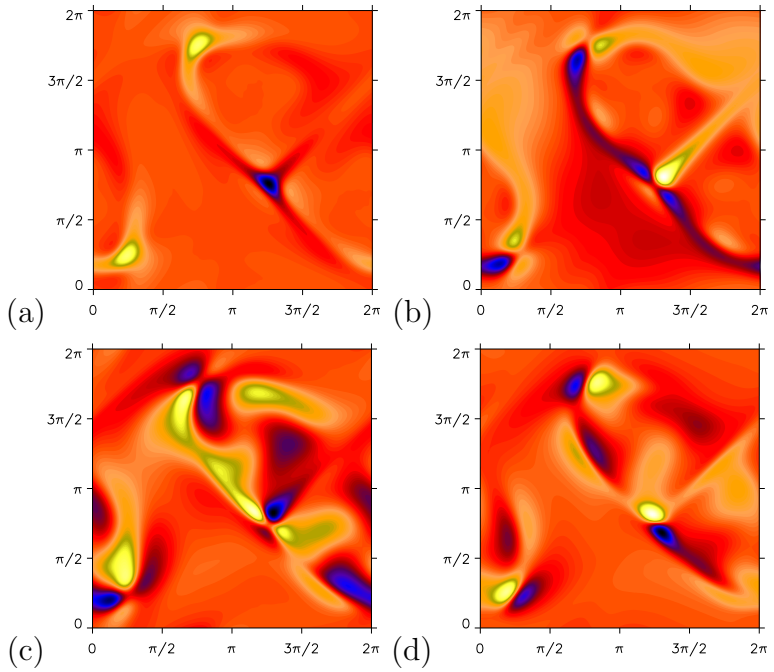

Figure 5.7: Visualisations of magnetic field component b_z on a cross section $z = \pi/4$ for (a) II to (d) V with $R_m = 100$.

concentrations along the separatrices joining stagnation points, and this is evident in the panels in figure 5.7, where the values of b_z are large near to $\boldsymbol{r}_1$, $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$. At the stagnation point $\boldsymbol{r}_1$ the 2-d stable and 1-d unstable manifold acts to concentrate magnetic fields into flux ropes and these are then visible along separatrices, e.g. at $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$. In each case, the field is not significant at $\boldsymbol{r}_2$, unsurprisingly since here the 2-d unstable manifold acts to repel field. As all the modes are oscillatory for this value of R_m, the fields in the cigars change sign during a single cycle. However the structures and the way this occurs differs in the different representations, and we now discuss this.

5.3.1 Symmetry class V: the dominant dynamo

As the dominant dynamo mechanism for the majority of the range of R_m investigated, it is natural that we should look to class V first. In the visualisations above, it is clear that the dominant eigenmode in class V contains two flux tubes, centred around the α-type stagnation point and stretched roughly equally in both directions towards different sides of the same β-type stagnation point. (Depending on which side of the 2-d stable manifold is being approached, a particle in the flow would be carried towards one side or another of the corresponding β-type point.) The field evolution in class V at $R_m = 100$ is oscillatory and so the snapshots in figures 5.6 and 5.7 represent just one state in its cyclical evolution. A full cycle of field evolution is displayed in figure 5.8 at eight different times for $R_m = 100$. The snapshots were taken so that the evolution of the field could be seen as clearly as possible, though some explanation is required. The magnetic energy, for this visualisation, is at a minimum at $t = 0$, with the snapshots normalised to account for the growing field. At the start of the cycle, we see lone flux tubes at the stagnation point $\boldsymbol{r}_1$ and at the midpoints $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$. These flux tubes weaken as secondary flux tubes of opposing polarity develop alongside them until by $0.25T$, at a magnetic energy maximum, the flux tubes are of equal strength. The original flux tubes continue to weaken, so that by $t = 0.4T$ (frame 4), the secondary flux tubes are stronger. At the magnetic energy minimum ($t = 0.5T$), only the sec-

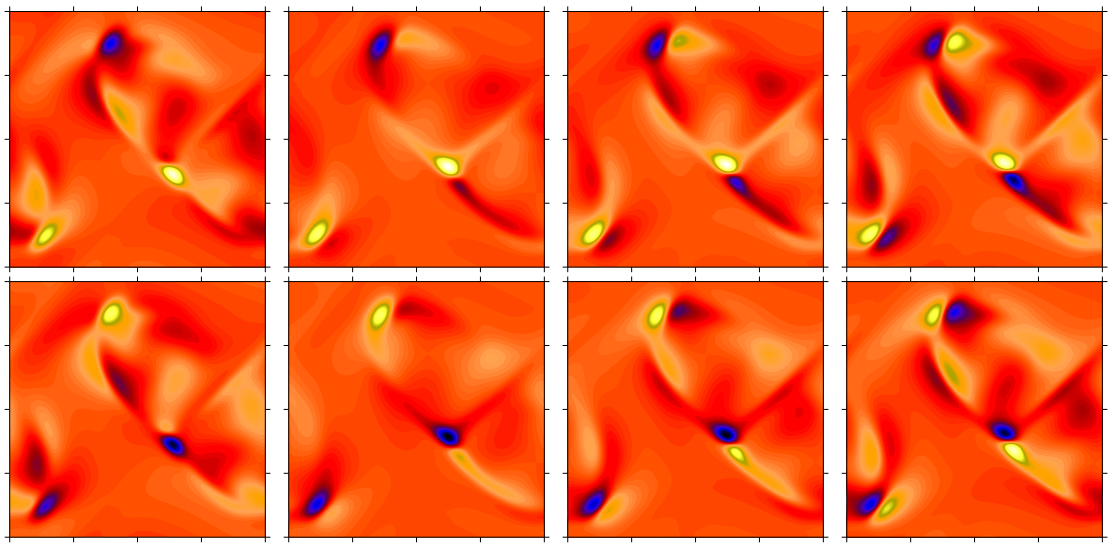

Figure 5.8: Magnetic field component b_z at $z = \pi/4$ for times $t = 0$, $0.1T$, $0.25T$, $0.4T$, $0.5T$, $0.6T$, $0.75T$ and $0.9T$ for class V, $R_m = 100$. Blue (dark) represents negative field strength and yellow (light) represents positive.

ondary flux tubes remain at $\boldsymbol{r}_1$, $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$ and the field is in the same state as at $t = 0$, except with the polarity reversed. The second half of the cycle mirrors the first half (except with inverted polarity) and another energy maximum occurs at $t = 0.75T$ (frame 7), with the flux tubes again of equal strength. At $t = 0.9T$, we see that the newest flux tube at each $\boldsymbol{r}_1$, $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$ is dominant and continues to grow, feeding off the other, so that by $t = T$, which is identical to $t = 0$, only this flux tube remains.

We note that the single flux tubes that exists at $t = 0$ and $t = 0.5T$ appear to be centred on the α-type stagnation points and the separatrices that pass through them. As the secondary tubes form, the original flux tubes drift away from the stagnation point and by the energy maximum, the tubes in each pair are situated on opposite sides of the separatrix.

The above process describes the dominant dynamo mechanism for $R_m > 24$, with the period of this cyclic behaviour increasing until at $R_m \approx 215$, the period becomes infinite and the structure of the dynamo is no longer time dependent, though the field amplification process remains the same (see Dorch, 2000). The final field configuration can be seen in figure 5.9 at $R_m = 250$. It appears that the structure

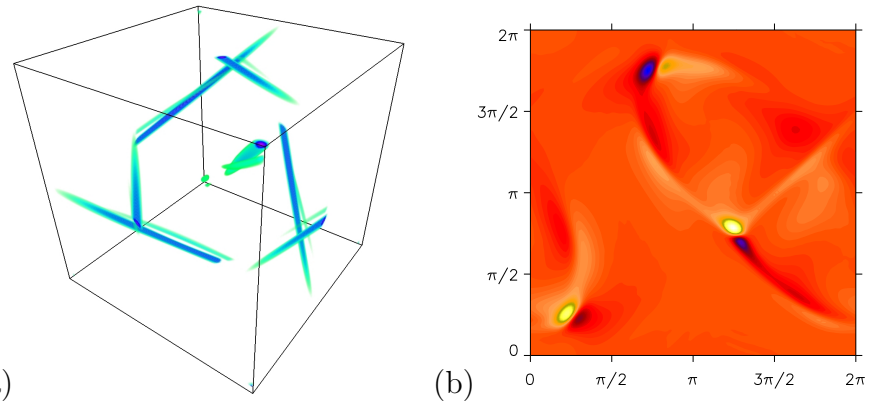

Figure 5.9: Visualisations of (a) magnetic field strength $|\boldsymbol{b}|$ and (b) field component b_z at $z = \pi/4$ for class V, $R_m = 250$.

of the dominant eigenmode consists of one major and one minor flux tube with the opposite polarity, which are fixed at their respective positions either side of the α-type stagnation points and the corresponding separatrices. These results are in agreement with Bouya and Dormy (2013), who also visualise the structure of the dominant dynamo mode. As R_m is increased, the structure of this dominant mode does not change, although the flux tubes are confined to an ever smaller region, making it increasingly difficult to visualise at higher R_m. For this reason, we are limiting visualisations to low and intermediate R_m where possible.

5.3.2 Symmetry class II

Class II contains the only other eigenmode that is dominant in the range of R_m investigated and has the second largest growth rate at $R_m = 10^4$. From prior studies, it is known that the structure of the first window is that of cigar-like structures centred on the α-type stagnation points. The first eigenmode seen, which the first window belongs to, persists until the mode-change at $R_m \approx 190$. This eigenmode is oscillatory and so the cross-section in figure 5.7(a) again only presents one of its many states. Figure 5.10 displays eight stages in a single cycle with period T. As well as the cigars, we see three 'fins' of field that form in a rotationally-symmetric configuration around each cigar and are stretched to form sheets of field; these are not necessarily seen in figure 5.6(a), as the snapshot was

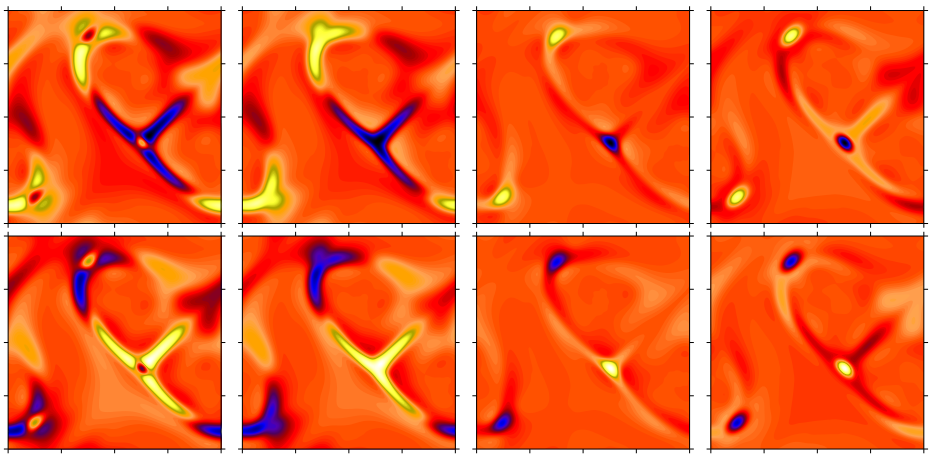

Figure 5.10: Magnetic field component b_z at $z = \frac{\pi}{4}$ for times $t = 0$, $0.07T$, $0.25T$, $0.42T$, $0.5T$, $0.57T$, $0.75T$ and $0.92T$ for class II, $R_m = 100$. Blue (dark) represents negative field strength and yellow (light) represents positive.

taken when these fins were relatively weak. At the point of lowest energy ($t = 0$), there is a weak flux tube with stronger fins of opposing polarity. The weak cigar (flux tube) weakens further and the fins contract inwards. By $t = 0.25T$ (maximum magnetic energy), the fins have joined to form a concentrated cigar. At $t = 0.42T$ (frame 4), we see new fins with the opposite polarity begin to form, strengthening until they are stronger than the weakening cigar ($t = 0.5T$) and the structure has returned to that of $t = 0$, except with the opposite polarity. The second half of the cycle, again, mirrors the first half but with reversed polarity. By $t = 0.9T$, we again see weak rotationally-symmetric fins forming around the cigar, which strengthen so that by $t = T$, the field is in its starting configuration (see frame 1). At the

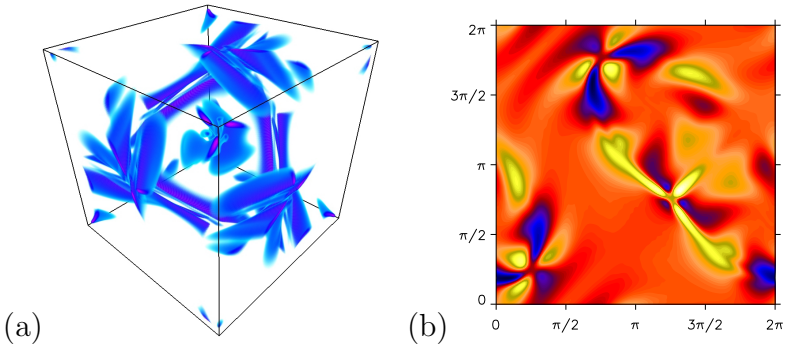

Figure 5.11: Visualisations of (a) magnetic field strength $|\boldsymbol{b}|$ and (b) field component b_z at $z = \pi/4$ for class II, $R_m = 250$.

mode-change, however, the field structure of the new eigenmode is significantly

different; see figure 5.11. In place of the central cigars, are three thin flux tubes centred on the α-type stagnation points with three-fold rotational symmetry. In between each of these narrow tubes are the 'fins' of field of opposing polarity to the tubes. Prior to the eigenvalue decoalescence, this arrangement is static as the field evolution is non-oscillatory but for $R_m > 330$, the field is oscillatory. The dominant

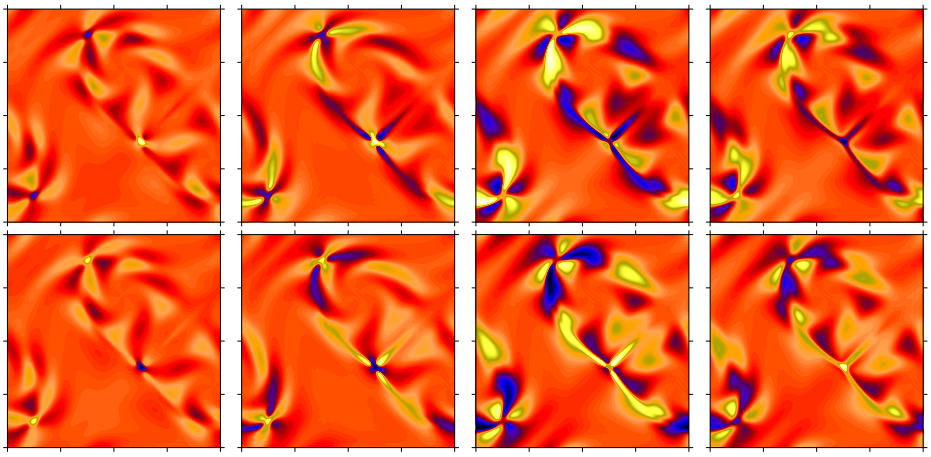

Figure 5.12: Magnetic field component b_z at $z = \frac{\pi}{4}$ for times $t = 0$, $0.1T$, $0.3T$, $0.4T$, $0.5T$, $0.6T$, $0.8T$ and $0.9T$ for class II, $R_m = 450$. Blue (dark) represents negative field strength and yellow (light) represents positive.

eigenmode after the decoalescence is visualised through cross sections at $z = \pi/4$ in figure 5.12. In this sequence, $t = 0$ is at the local energy minimum, with the maximum corresponding to $t = 0.25T$, as is the case for all cycle visualisations. As with the $R_m = 100$ eigenmode, the same three-fold structures typical of symmetry class II are visible.

At the start of the cycle, single cigar-like flux tubes are seen at $\boldsymbol{r}_1$, $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$. The cigar grows fin-like projections and by $t = 0.1T$ (frame 2), flux tubes of the opposite polarity to the cigar have formed between the cigars fins. These flux tubes grow in strength, whilst the original cigar decays and by $t = 0.3T$, these flux tubes meet at the stagnation point $\boldsymbol{r}_1$ (and separatrix midpoints $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$) with only the fins remaining of the original cigar. These fins decay, along with the flux tubes of opposing polarity, leaving a cigar-like concentration of field at $\boldsymbol{r}_1$, $\boldsymbol{r}_3$ and $\boldsymbol{r}_4$ by $t = 0.4T$. At $t = 0.5T$, all that remains is the compact cigar at these points and, of course, the configuration is identical to that of $t = 0$ but with reversed polarity.

As is to be expected, field evolution is identical in the second half of the cycle, with frames 5–8 matching those of 1–4, except with the opposite sign. This loosely describes the evolution of the structure of class II's dominant eigenmode up to $R_m \approx 5000$. For the new oscillatory eigenmode that emerges at $R_m \approx 5000$, we are unable to carry out such a detailed investigation. Unfortunately the visualisation tools are incapable of handling multiple snapshots in the largest simulations, in part owing to the memory limitations of available computing resources. Thus we are unable to visualise the structure of the eigenmode at multiple times during the cycle for $R_m > 5000$ and, as such, it is not discussed in this thesis.

5.3.3 Symmetry class IV

Class IV is interesting in that a single eigenmode dominates in the interval $R_m \in [33, 5000]$ and undergoes both eigenvalue coalescence and decoalescence in this range. This means that the structure in this range does not change significantly, though becomes finer, allowing us to understand the structure of the dynamo at high R_m through visualisations at lower R_m, as with class V. In figure 5.6(c), we see that at $R_m = 100$, there are four cigars twisted around the separatrices and centred on the α-type stagnation point. The field evolution over a single cycle is visualised through cross-sections in figure 5.13 at $R_m = 200$, where this eigenmode acts as a dynamo. The cycle is somewhat similar to class V in that there are cigars of opposing polarity which weaken and strengthen in turn. The difference is that class IV contains four individual flux tubes which form two pairs, with these pairs of tubes weakening and strengthening over a single period. The two flux tubes, within each pair, have opposing polarity. At the beginning of the cycle (point of lowest energy, $t = 0$), one pair is strong whilst the other is weak. The weaker pair increase in strength, appearing to feed on the stronger pair, and by $t = 0.25T$ (point of maximal energy) they are of equal strength. The growing pair continue to strengthen whilst the other pair weaken and are eventually dissipated. The remaining pair of flux tubes then shift to the position of the now dissipated pair by $t = 0.5T$ (frame 5). The field structure is the same as at $t = 0$, except with

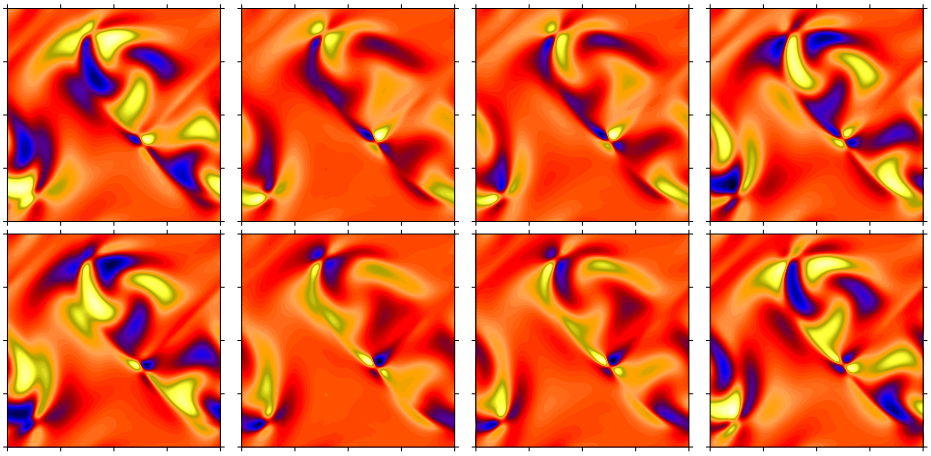

Figure 5.13: Magnetic field component b_z at $z = \frac{\pi}{4}$ for times $t = 0$, $0.1T$, $0.25T$, $0.45T$, $0.5T$, $0.6T$, $0.75T$ and $0.95T$ for class IV, $R_m = 200$.

opposing polarity; thus describing a half-cycle. The flux tubes repeat the process of strengthening, weakening and shifting to return to the starting arrangement. After

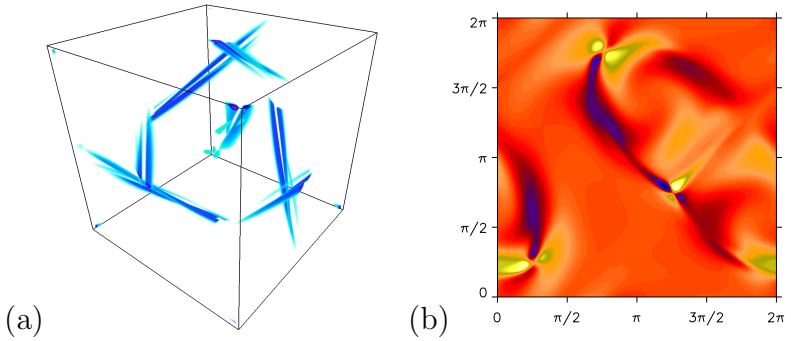

Figure 5.14: Visualisations of (a) magnetic field strength $|\boldsymbol{b}|$ and (b) field component b_z at $z = \pi/4$ for class IV, $R_m = 400$.

the eigenvalue coalescence, we have a steadily evolving field. The field structure is shown in figure 5.14 and most closely resembles the field at $t = 0.1T$ (second panel in figure 5.13). We recall that the pair of flux tubes in class V's dominant eigenmode strengthen and weaken in turn and that after the eigenvalue coalescence, the structure is such that one flux tube is consistently weaker than the other. The same occurs in class IV but now one pair of flux tubes remains dominant and the other subdominant. However, this field arrangement does not persist for a large range of R_m, and returns to an oscillatory state at a higher R_m, in contrast with class V.

We find that symmetry class II appears to be the only one that permits flux tube formation on the separatrices. In classes IV and V we see that the various flux tubes are twisted around the heteroclinic lines but never lie on them, at least for any extended period of time.

5.3.4 Symmetry classes I and III

Classes I and III are discussed together as their field structures are very similar. Prior to the emergence of the growing mode in class I, the slowest decaying mode is simply the trivial solution $\boldsymbol{b} = \boldsymbol{u}$; the field is concentrated on the principal vortices of the flow and we omit visualisations of this eigenmode. The structure of class III's eigenmode prior to its mode-change appears, visually, to be a hybrid of the eigenmode of class I and an eigenmode that resembles those of the other classes, where cigars form near to the α-type stagnation points. In the dominant

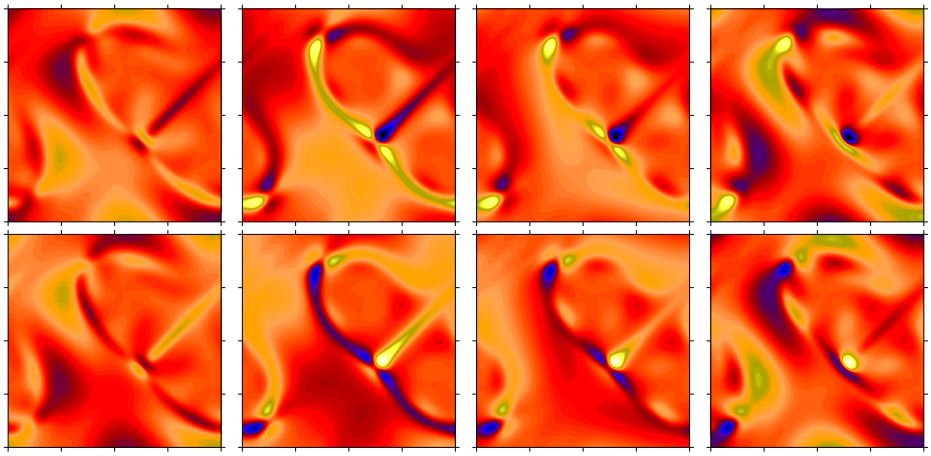

Figure 5.15: Magnetic field component b_z at $z = \frac{\pi}{4}$ for times $t = 0$, $0.05T$, $0.2T$, $0.47T$, $0.5T$, $0.55T$, $0.7T$ and $0.97T$ for class III, $R_m = 100$.

eigenmode for $R_m < 650$ in class III, the field switches between two states. In the first of these states, seen at $t = 0$ in figure 5.15, the strongest areas of field lie in four of the six principal vortices, with the magnetic energy at a local minimum. In the second state, the strongest field is located in three flux tubes, one major tube and two minor with opposite polarity to the major, which are situated around the separatrices and centred on the α stagnation points. This second state can be seen

at $t = 0.2T$ (frame 3). In each half-period, the field spends very little time in the first state and quickly evolves to the second state. By $t = 0.5T$, the magnetic field has the same structure but opposite polarity as at $t = 0$, with evolution in the second half-period mirroring that of the first.

In the growing mode of class I that emerges at $R_m \approx 800$ we see something intriguing: the magnetic field is formed as 'flux ropes', centred around the six principal vortices. Each flux rope consists of six individual flux tubes, each tube with varying strength along its (periodic) length. This structure can be clearly seen in figure 5.16 (a,b) and the cross section of b_z at $z = \pi/4$ reveals the alternating polarity of these tubes. The field in this eigenmode is oscillatory and each flux rope undergoes a reversal (all flux tubes reverse polarity) each half-cycle. In the growing

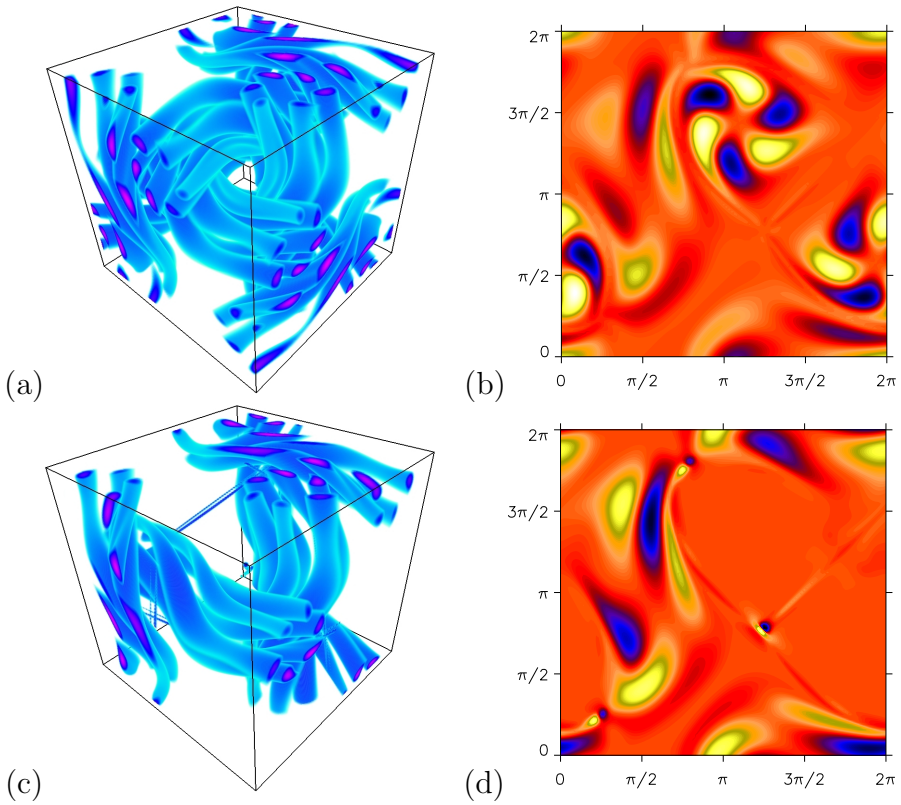

Figure 5.16: Magnetic field strength $|\boldsymbol{b}|$ for (a) class I and (c) class III and field component b_z at $z = \pi/4$ for (b) class I and (d) class III, for $R_m = 1000$.

mode of class III, however, which emerges at $R_m \approx 650$, we again have a hybrid of structures, with the flux ropes, seen in class I, concentrated around four of the six principal vortices and only two apparent flux tubes (of opposite polarities) twisted

around the separatrices. This structure can be seen in figure 5.16 (c,d). Note that the fact that we see field in four of the six principal vortices may seem surprising, but is ultimately a result of our choice of matrices employed for III; with this our visualisations (for III, IV and V) have an element of arbitrariness although the eigenvalues do not. Despite the differences in structure, this eigenmode displays the same evolution cycle as the first mode, switching between the flux ropes and the tightly confined double-cigar structures.

The flux ropes in classes I and III (above) closely resemble the structure of fields in smooth Ponomarenko dynamos (Ponomarenko, 1973). This is well known as a slow dynamo mechanism, relying on diffusion in a curved geometry and shear to give a dynamo cycle (Gilbert, 1988; Ruzmaikin *et al.*, 1988). Our current results indeed show falling growth rates for these branches as $R_m \to \infty$, in keeping with this interpretation.

5.3.5 Symmetries and field configurations

We briefly examine how the individual symmetry classes I to V impact on the presence or otherwise of cigars in the magnetic field problem. We consider the separatrix whose centre is at the α-type stagnation point at $\boldsymbol{r} = -(1,1,1)\pi/4$ and introduce local cylindrical polar coordinates (ρ, θ, ζ) there (see Gilbert *et al.*, 2011), given by

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = - \begin{pmatrix} \pi/4 \\ \pi/4 \\ \pi/4 \end{pmatrix} + \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{6} & 1/\sqrt{3} \\ -1/\sqrt{2} & 1/\sqrt{6} & 1/\sqrt{3} \\ 0 & -2\sqrt{6} & 1/\sqrt{3} \end{pmatrix} \begin{pmatrix} \rho\cos\theta \\ \rho\sin\theta \\ \zeta \end{pmatrix}. \tag{5.2}$$

We choose to test the action of two transformations, d and n, on the potential for single cigars to develop. We identify these particular transformations as they simply rotate about the stagnation point $\boldsymbol{r} = -(1,1,1)\pi/4$; other transformations will act as rotations about other stagnation points or will map one stagnation point to another. In the new coordinate system, the action of d is to take $(\rho, \theta, \zeta) \to$

$(\rho, \theta + 2\pi/3, \zeta)$ and the action of n is to take $(\rho, \theta, \zeta) \to (\rho, -\theta, -\zeta)$. We may consider potential cigar structures of the general form

$$\boldsymbol{b} = \hat{\boldsymbol{\zeta}} \rho^m \mathrm{e}^{\mathrm{i}m\theta}, \tag{5.3}$$

and we assume that a single cigar manifests as a local, approximately axisymmetric (i.e. $m = 0$), ζ-independent structure. Under the action of d, we have that $\mathsf{d}\boldsymbol{b}^\alpha = \boldsymbol{b}^\alpha$ for both classes I and II, as $M^{\mathrm{I}}(\mathsf{d}) = M^{\mathrm{II}}(\mathsf{d}) = 1$. We therefore arrive at the relation $\mathrm{e}^{\mathrm{i}(2\pi m/3)} = 1$, for which the solution is $m = 0 \bmod 3$. However, $M^{\mathrm{I}}(\mathsf{n}) = -M^{\mathrm{II}}(\mathsf{n}) = 1$ and so $\mathsf{n}\boldsymbol{b}^{\mathrm{I}} = \boldsymbol{b}^{\mathrm{I}}$, whereas $\mathsf{n}\boldsymbol{b}^{\mathrm{II}} = -\boldsymbol{b}^{\mathrm{II}}$. For class I, we find no value of m satisfies this relation for the transformation n suggesting that no single axisymmetric cigars can exist in this class. In class II, a single cigar is indeed possible, with $m = 0$ as the only valid solution.

In class III, this becomes more complicated, owing to the presence of multiple fields. We set

$$\boldsymbol{b}_1^{\mathrm{III}} = c_1 \, \hat{\boldsymbol{\zeta}} \rho^m \mathrm{e}^{\mathrm{i}m\theta}, \qquad \boldsymbol{b}_2^{\mathrm{III}} = c_2 \, \hat{\boldsymbol{\zeta}} \rho^m \mathrm{e}^{\mathrm{i}m\theta}, \tag{5.4}$$

where c_1 and c_2 are arbitrary constants and we recall that

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{d}) = \frac{1}{2} \begin{pmatrix} -1 & \sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix}, \qquad \boldsymbol{M}^{\mathrm{III}}(\mathsf{n}) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}. \tag{5.5}$$

Under the action of d, we find that

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{d}) \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \mathrm{e}^{-\mathrm{i}(2\pi m/3)} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \tag{5.6}$$

suggesting that the restrictions on m are determined from the eigenvalues of $\boldsymbol{M}^{\mathrm{III}}(\mathsf{d})$. These eigenvalues are found to be $\lambda_1 = \overline{\lambda_2} = -\frac{1}{2} + \mathrm{i}\frac{\sqrt{3}}{2}$, which can be rewritten as $\lambda_1 = \mathrm{e}^{\mathrm{i}(2\pi/3)}$ and $\lambda_2 = \mathrm{e}^{\mathrm{i}(4\pi/3)}$. Clearly for class III, we restrict $m \neq 0 \bmod 3$ and this is further confirmed by the action of n on these fields, ruling out the possibility of a single cigar in this class.

We carry out the same analysis for classes IV and V, which have fields $\boldsymbol{b}_j^\alpha = c_j\ \hat{\boldsymbol{\zeta}}\rho^m \mathrm{e}^{\mathrm{i}m\theta}$ for $j = 1, 2, 3$ and irreducible representations given by

$$\boldsymbol{M}^\alpha(\mathsf{d}) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^\alpha(\mathsf{n}) = \pm \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{5.7}$$

with α as IV for the upper sign and α as V for the lower. If $\boldsymbol{c} = [c_1, c_2, c_3]^\intercal$, then under the action of d for both fields, we have

$$\boldsymbol{M}^\alpha(\mathsf{d})\boldsymbol{c} = \mathrm{e}^{\mathrm{i}(2\pi m/3)}\boldsymbol{c}, \tag{5.8}$$

and again the restrictions of m are determined by the eigenvalues of $\boldsymbol{M}^\alpha(\mathsf{d})$. These eigenvalues are $\lambda_j = \mathrm{e}^{\mathrm{i}[2\pi(j-1)/3]}$, resulting in $m = 0, 1, 2 \bmod 3$ and thus placing no restriction on m. However, under the action of n, we have that $\mathsf{n}\boldsymbol{b}_3^{\mathrm{IV}} = \boldsymbol{b}_3$, which results in no solution for $m = 0$ suggesting that single cigars are not possible in class IV. For class V we have $\mathsf{n}\boldsymbol{b}_3^{\mathrm{V}} = -\boldsymbol{b}_3$, resulting in the solution $m = 0$ and so single cigars are possible in this symmetry class.

This characterisation of the structures that can occur within a given symmetry class has points of contact with figures showing magnetic fields, earlier in this section, and the growth rates at the largest values of R_m in figure 5.1. We observe that the two modes showing the fastest growth rates are II and V, in which single cigars can exist. On the other hand I, III and IV show rather smaller growth rates to the right of this figure. Although these cases allow multiple cigars (that is $m \neq 0$) or more complex structures, the enhanced dissipation is likely to suppress growth rates at these R_m.

5.3.6 Field structures and regions of chaos

The ABC flows' Lagrangian chaos is believed to play an important role in constructive stretching and folding of the magnetic field, leading to amplification over

successive cycles. We know that for classes II, IV and V (and to some extent III), the magnetic energy builds up at the α-type stagnation points and is stretched in two different directions towards opposite sides of the same β-type point, forming tubes of flux. In the Poincaré sections (figure 5.5), the α- and β-type points are identifiable but indistinguishable from one another, though the mid-points of the separatrices can also be clearly recognised where three narrow bands of chaos meet from different directions.

As pointed out by Childress and Gilbert (1995), the resolution of the numerical models in Lau and Finn (1993) at $R_m = 1000$ is relatively low, meaning that the grid point spacing is large compared to the size of the areas of chaos. From our numerical models, the fields at the highest R_m are visualised near to the α-type stagnation point at $\boldsymbol{r} = (5, 3, 1)\pi/4$ in figure 5.17 and near to the separatrix mid-point at $\boldsymbol{r} = (4, 4, 0)\pi/4$ in figure 5.18. The values of R_m chosen were 5000 for classes IV and V, and 4000 for class II – it was decided to investigate the structure before the high-R_m mode change in class II. At the α-type stagnation point, we see for classes IV and V that the most concentrated areas of magnetic field are located very close to the actual stagnation points. For class II, the cigar actually lies on the stagnation point, as is the case at lower R_m. Additionally, the fin-like features of class II are seen to correspond with the bands of chaos at the α-type stagnation point. The areas of strong field in all classes appear to lie well within the broad chaotic regions associated with the stagnation point, in these high R_m regimes. The observation could be made that there is no fine-scale structure associated with the extremely narrow bands of chaos, though we note that the resolution still may not be small enough to sufficiently resolve the smallest chaotic regions. The distance between two grid points is indicated by the length of the black line in the lower-left of the cross-sections; it is clear that the strongest areas of field may only be spread over the equivalent of a 3×3 or 4×4 grid points. However, on increasing the resolution it is known that dynamo growth rates will not change significantly, having already converged to within desired accuracy (three significant figures). We can thus assume that the cigars themselves, are sufficiently well-resolved. The

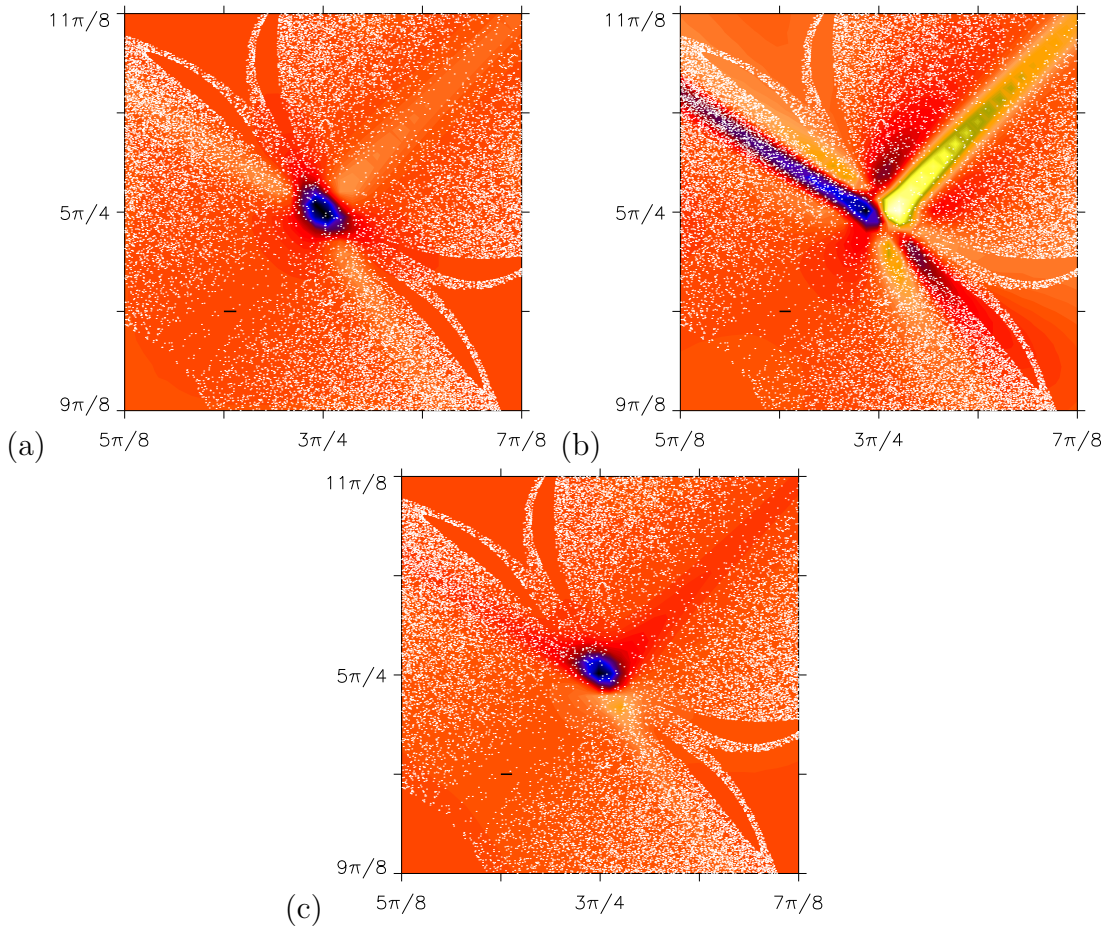

Figure 5.17: Magnetic field component b_z at $\boldsymbol{r} = (5, 3, 1)\pi/4$ for (a) class II, $R_m = 4000$, (b) class IV, $R_m = 5000$ and (c) class V, $R_m = 5000$. Overlaid in white is the Poincaré section and the length of the black bar signifies the distance between two grid points.

second set of overlaid cross-sections (figure 5.18) centres on a separatrix mid-point at $\boldsymbol{r} = (4, 4, 0)\pi/4$ and is where the Lagrangian chaos is most tightly confined. We see that the areas of strongest field, in these cross-sections, coincide with the mid-point itself $(\pi, \pi, 0)$ in all three classes. The fin-like projections of class II correspond well with the thin bands of chaos, with a clear three-fold rotational symmetry. Class IV's four cigar-like structures are visible with one pair much stronger than the other at the particular time the snapshot was taken. The two cigars of the strong pair seem to form in relation to two of the three bands of chaos that meet at the separatrix, whereas the weaker pair are located within the third band. Smaller features are noticeable in this section, with some field structure located in the extremely thin outer chaotic regions. Class V's two cigars appear

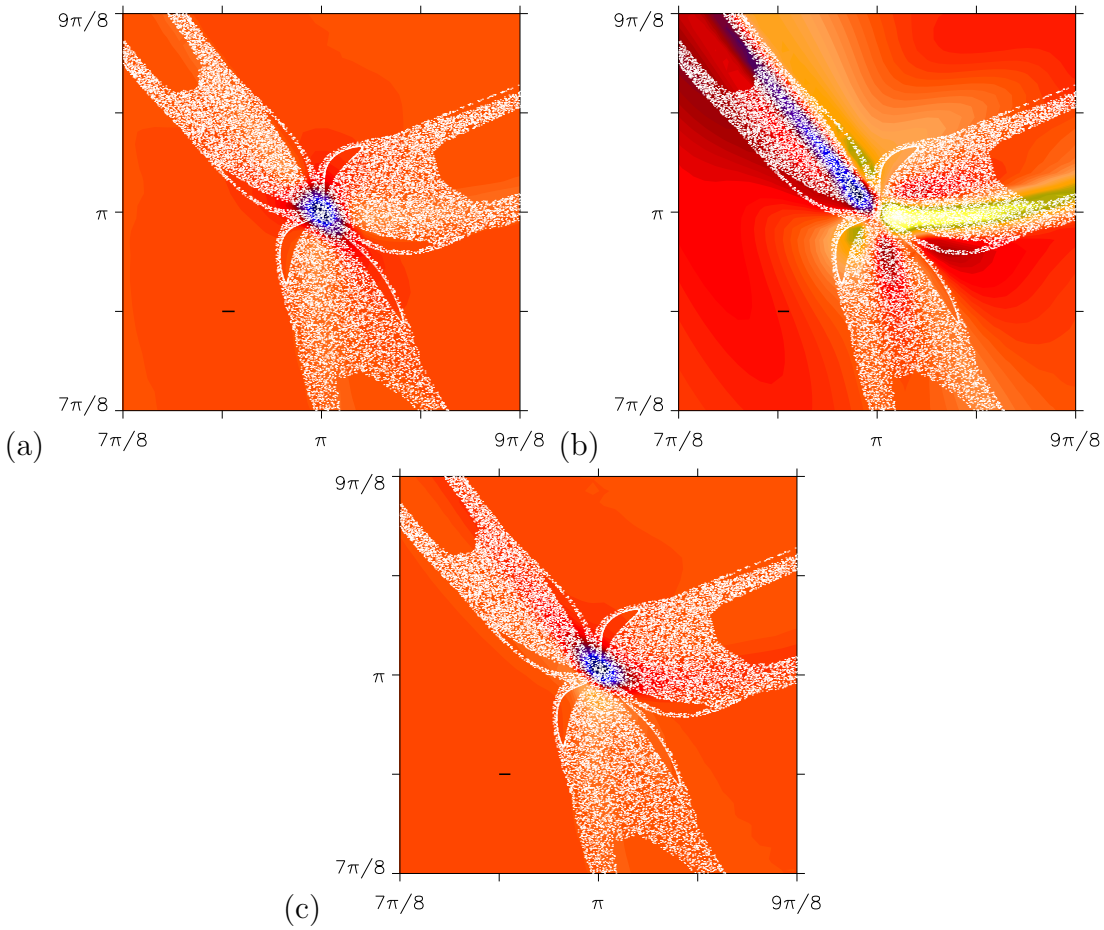

Figure 5.18: Magnetic field component b_z at $\boldsymbol{r} = (4, 4, 0)\pi/4$ for (a) class II, $R_m = 4000$, (b) class IV, $R_m = 5000$ and (c) class V, $R_m = 5000$, with overlaid Poincaré section.

to be located well within the chaotic regions at the separatrix mid-point. Some weaker field features associated with the strong cigar are located in two of the three chaotic bands, with the third band (bottom half of figure 5.18(c)) containing a weak extension of the subdominant flux tube.

Cross-sections were not taken for higher R_m, in part, due to the size of the storage array involved. The full process of unfolding to the full domain and carrying out a fast Fourier transform is incredibly memory-intensive at such high R_m. The difficulty of handling such large files with Fortran and IDL (used to produce cross-sections and convert the data for use in VAPOR), along with the ever reducing scales of the structure involved, means that visualising these fields is complex and time-consuming, and out of the scope of this work.

5.3.7 Magnetic energy spectra

Finally, we briefly study the magnetic energy spectra of simulations in each symmetry class. The magnetic energy of the Fourier modes is calculated and placed into shells, depending upon the wavenumber of the mode. Modes of wavenumber (l, m, n) are placed into shell j such that if $k = (l^2 + m^2 + n^2)^{1/2}$ then $j = \lfloor k \rfloor$. For those classes and R_m regimes in which the dominant eigenmode is oscillatory, the simulation was rerun and the energy of each shell was averaged over an entire cycle. The spectra reveal some differences between the representations that mirror

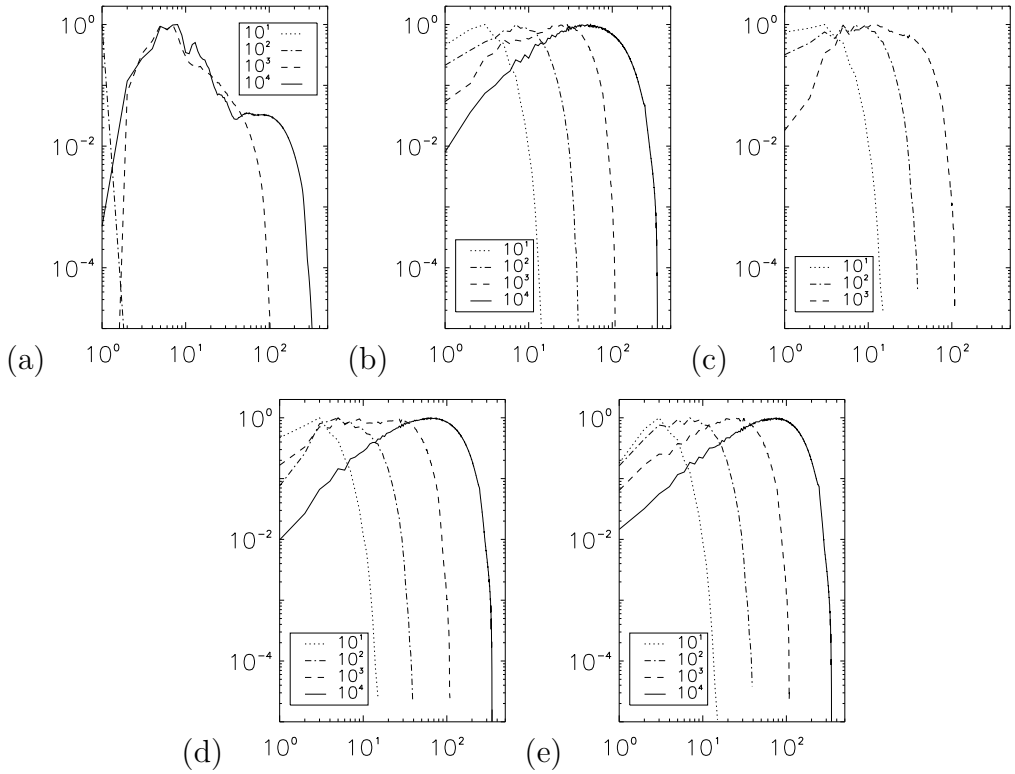

Figure 5.19: Energy spectra displaying normalised magnetic energy against wavenumber k (log-log) for (a) I to (e) V, at $R_m = 10^1, 10^2, 10^3, 10^4$.

the structure of the fields in real space. It is immediately visible, in figure 5.19, that the spectra of class I are very different to those of the other classes, with this reflecting the structural differences identified earlier in the section. Prior to the mode change, nearly all the energy is located in the mode $\boldsymbol{b}_{1,0,0}$ in $\mathcal{H}_N$, with trace amounts of energy in higher wavenumber modes. After the mode change, we see that the majority of magnetic energy is located on fairly large scales, $k \sim O(1)$,

with roughly equal distribution over the range $k \in [4, 8]$ at both $R_m = 10^3$ and 10^4. It appears that some energy remains at the smaller scales with approximately equal amounts of energy over the range $k \in [40, 100]$ at $R_m = 10^4$. The considerable energy at large scales is consistent with the broader spatial scales of a slow dynamo maintained by weak stretching in the principal vortices.

By contrast, the spectra of the other four classes are very similar, with minor differences in class III. Despite areas of strong field forming in the principal vortices, in class III, the energy spectra prior to the mode-change correspond visually to those of the other classes. This may be because the field in the principal vortices is a transient feature, only forming at the point of lowest energy in the cycle and when the field is undergoing a reversal, and so for the majority of the cycle the energy does not remain at the largest scales. For $R_m = 10^3$ (after the mode crossing), the spectrum still roughly resembles those of the other classes, though has features of class I such as the steep slope for low wavenumber k. Class IV also differs for $R_m = 10^3$ to those of the other classes, with the magnetic energy distributed over a broad range of scales. On comparison between figures 5.10 and 5.13, we see that the same eigenmode at $R_m = 200$ (as at $R_m = 10^3$) has structures that are much less confined than in class II at $R_m = 100$.

Most importantly, in figure 5.19, it is seen that for representations II, IV and V a spectrum with a power law dependence of approximately k emerges at large R_m: this is consistent with the concentration of magnetic field in isolated flux tubes (A. Gilbert, private communication). Here we draw a parallel between the magnetic induction equation (1.12) and the vorticity equation for an incompressible flow $\boldsymbol{u}$, given by

$$\partial_t \boldsymbol{\omega} = \nabla \times (\boldsymbol{u} \times \boldsymbol{\omega}) + \nu \nabla^2 \boldsymbol{\omega}, \tag{5.9}$$

where $\boldsymbol{\omega}$ is the fluid *vorticity*, given by $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ and $\nu = \mu/\rho$ is the *kinematic viscosity*. A well-known solution to (5.9) is the Burgers vortex, in which vorticity has a Gaussian distribution about an axis; a vortex tube not unlike the magnetic flux tubes we find in classes II, IV and V. Such tubes of vorticity would correspond

to an energy spectrum of k^{-1} and so a vorticity spectrum of k. Though the magnetic field is not directly linked to the flow as is the vorticity, it may be possible to interpret the formation of magnetic flux tubes, or cigars, through well-studied vorticity transport mechanisms.

Overall, all spectra clearly show a high-k falloff in magnetic energy, well before the largest shell which has $k_{\max} = \sqrt{3}(N-1)^{1/2}$. This identifies that the magnetic fields are well-resolved at the smallest scales.

5.4 Arnoldi method results

This section concerns the results obtained through the use of the ARPACK routine in the NAG library. As discussed in the previous chapter, this method uses an (implicitly restarted) Arnoldi iteration method to identify the eigenvalues of interest, in this case those with largest real part.

As the resolution N scales as $R_m^{3/2}$ and n_{bv} scales with N (though this relationship is not well-understood), the size of the problem quickly becomes unmanageable for $R_m > O(10^2)$. For this reason, we only provide a detailed investigation for R_m up to 100 in all classes, though we carry out two case studies at higher R_m. To obtain the results for $R_m \leq 100$, the number of eigenvalues to be found (n_{ev}) is set to 10. As with the eigenvalues of class V in section 4.5, we discard eigenmodes found at fewer than four values of R_m, in addition to one eigenvalue from a complex pair. This leaves at most 6 unique eigenvalues at any one R_m investigated.

On successful completion, the routine outputs the eigenvalues found and prints off the corresponding eigenvectors to file. The eigenvalues then have to be manually assigned to branches: this task is made more difficult as the NAG ARPACK routine does not output eigenvalues in any particular order and branch crossings are frequent, meaning that the first eigenvalue outputted for a specific R_m may not belong to the same branch of solutions as the first eigenvalue of the previous or next R_m studied. To accurately match the eigenvalues to the correct branches, it

is assumed that their real parts do not change drastically over a short range of R_m. The easiest method of matching involves identifying eigenmodes through $\mathrm{Im}(\lambda)$; the growth rates of different modes are often very similar but their frequencies have, so far, always been distinctly different. The process of matching by hand is time-consuming but rewarding, as the varied behaviour of the various subdominant eigenmodes is revealed; the process could have been automated but undoubtedly would not have accounted for all exceptions.

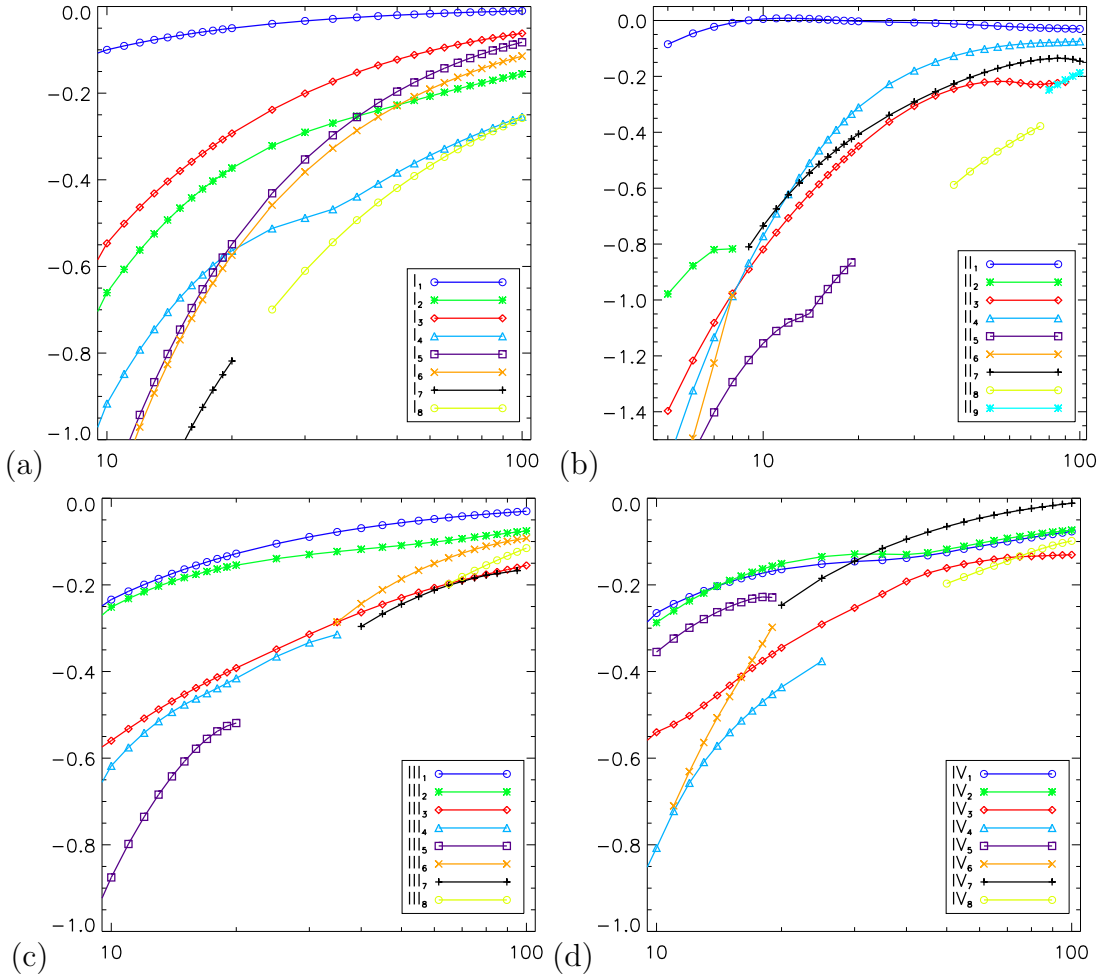

Figure 5.20: Growth rates of eigenmodes with largest real part for classes (a) I to (d) IV, for $R_m \leq 100$. Results for class V are given in figure 4.11.

Figure 5.20 displays the growth rates of the most dominant eigenmodes of classes I–IV for $R_m \leq 100$, in parallel with the results given for class V in figure 4.11. Results for eigenmodes with less than four identifiable eigenvalues were not included. The Arnoldi method reveals an array of previously unknown eigenmodes with their own

kinks and crossings with other subdominant branches. From the imaginary part of the eigenvalues, we know that for classes I, II and III, there are no crossings involving the dominant mode due the lack of discontinuities seen in figure 5.2. However, it is obvious that in class IV, there are two crossings; one in the interval $R_m \in [14, 15]$ and another in $R_m \in [32, 33]$. The Arnoldi method identifies these crossings, as it allows us to track the eigenmodes prior to the crossings. In the time stepping problem, great difficulty is found in resolving two competing modes near to a mode crossing, particularly in heavily decaying regimes, and so despite the additional resources needed to calculate two similar eigenvalues, the Arnoldi method is seen to be far superior for such situations. Furthermore, when time stepping at low R_m, the magnetic fields are typically decaying very quickly, often with the total magnetic energy reaching values on the level of machine precision before being able to settle on a particular eigenmode. The Arnoldi solver avoids this complication, making it the preferred method for low-R_m regimes.

We turn our attention to the mode crossing in class II, discussed in section 5.2. Applying the Arnoldi iteration method to the range $R_m \in [170, 215]$ has two purposes. Firstly, as the growth rates of subdominant eigenmodes were taken from transient evolution profiles, we wish to confirm the results through another method. Secondly, we were not able to locate the non-oscillatory mode at $R_m = 170$ and wish to test the Arnoldi method's ability to find this particular eigenvalue. The results are summarised in figure 5.21, with only the two modes involved in the crossing included (though others were indeed found). In the eigenvalues found for the range $R_m \in [170, 174]$, we see an oscillatory eigenmode with $\mathrm{Im}(\lambda)$ displaying a square root dependence on R_m and appearing to decay to zero; the frequencies are not visualised. It is immediately clear that a non-oscillatory mode could not be found at $R_m = 170$ because the eigenmode is oscillatory at this value. By increasing R_m, the change was found to occur in the interval $R_m \in [174.2, 174.25]$. In section 5.2, we see that the eigenvalue coalescence leads to two real-valued eigenvalues corresponding to $\lambda_+ = \alpha + \beta$ and $\lambda_- = \alpha - \beta$ respectively. In the range $R_m \in [174.25, 215]$, the number of eigenvalues to be found (n_{ev}) was set to a suffi-

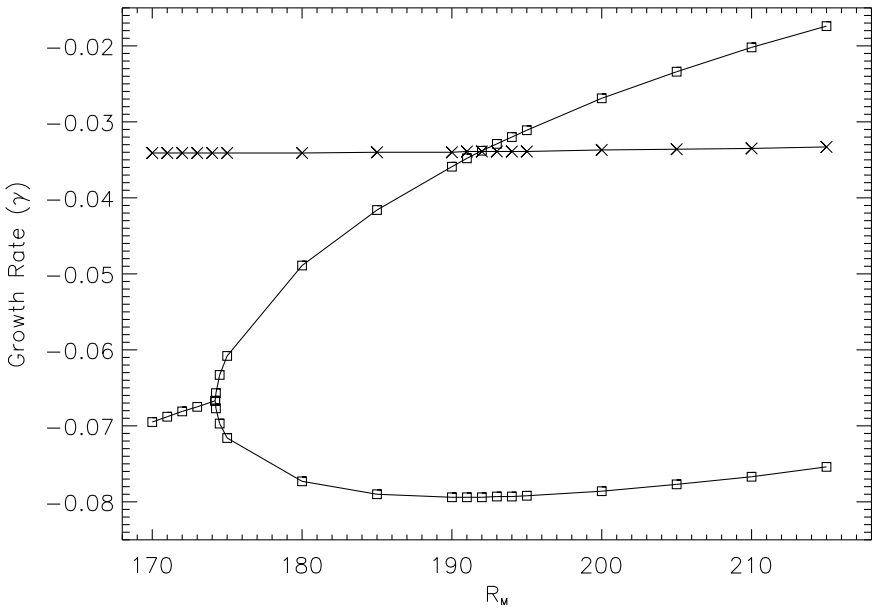

Figure 5.21: Coalescence of class II's subdominant oscillating eigenmode into a pair of non-oscillatory eigenmodes (squares) in the neighbourhood of the mode crossing with the eigenmode of the first window (crosses).

ciently large value so that two real eigenvalues were found, leading to the results in figure 5.21. This identifies another strength of the Arnoldi method; it can be used to follow both branches of solutions after such a bifurcation. Such analysis was not carried out for the coalescence in class V ($R_m \approx 215$), as the Arnoldi method is much more computationally demanding for multi-dimensional representations, owing to the relative size of the matrix $\boldsymbol{A}$, and is currently beyond the realms of this investigation.

Despite the interesting eigenmodes revealed, it must be noted that the Arnoldi method has not found a subdominant dynamo eigenmode, that is one with $\mathrm{Re}(\lambda) > 0$, for $R_m \leq 100$, within any of the five symmetry classes. It is certain that at least one such subdominant mode exists at larger R_m, at least within class V, as the coalescence of eigenvalues leads to a secondary mode corresponding to $\lambda_- = \alpha - \beta$ for $R_m > R_m^{\mathrm{crit}}$.

On successful convergence, the Arnoldi routine has identified the requested number of Ritz vectors, which it then converts to eigenvector estimates. These estimates can be unpacked to $\mathcal{H}_N$ in the same fashion that the vector $\boldsymbol{w}$ is unpacked on every

iteration of the routine. We can then reconstruct the magnetic field for $\mathcal{F}_N$ through the unfolding routines of the time stepping codes and produce field visualisations in the same procedure used in the time stepping routines. The visualisations produced are varied, with some corresponding well to the fields seen in the time stepping and others only vaguely resembling their time stepping counterparts. Figure 5.22 displays the cross-sections and 3-d visualisations of the dominant eigenmodes in classes II to V, constructed from the Ritz eigenvectors. We note that the cor-

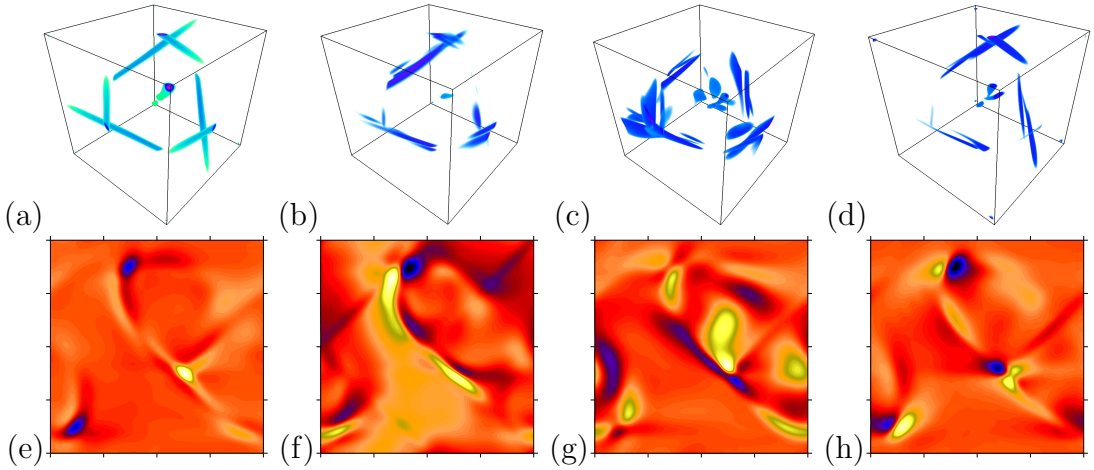

Figure 5.22: Visualisations produced using the IRAM routine. Shown are field magnitude $|\boldsymbol{b}|$ for (a) II to (d) V and field component b_z at $z = \pi/4$ for (e) II to (h) V, at $R_m = 100$.

responding eigenvalues of the above visualisations are identical to those obtained through time stepping to three significant figures. By comparing with figures 5.6 and 5.7, it is clear that only class II closely resembles the structure found in the time stepping. From investigating the evolution of the periodic fields of classes III to V (these visualisations are not included), it is found that such configurations as those seen in figures 5.22(b)–(d) and (f)–(h) do not occur in the time stepping problem. For classes IV (c,g) and V (d,h), the field loosely resembles those of the time stepping problem, inasmuch as they have the required number of flux tubes (4 and 2, respectively), though for class III, there appears to be no such similarity.

One of the difficulties in comparing the fields of the two methods is that the Arnoldi method produces one field configuration for an oscillatory mode, as opposed to a range of structures in time stepping. At first, it may not seem clear how the Arnoldi

fields relate to those produced through time stepping, though we recall that each method is solving a slightly different but related problem. We have that

$$\boldsymbol{b} = \tilde{\boldsymbol{b}}(x, y, z)\mathrm{e}^{\lambda t} + \tilde{\boldsymbol{b}}^{*}(x, y, z)\mathrm{e}^{\lambda^{*} t}, \tag{5.10}$$

once Hermitian symmetry has been considered, as $\boldsymbol{b}$ is real-valued. In the Arnoldi routine, we do not consider Hermitian symmetry and so it calculates λ and λ^{*} as distinct eigenvalues which, indeed, they are. On reconstructing the Arnoldi field, the routine simply returns the time-independent field $\tilde{\boldsymbol{b}}(x, y, z)$ in $\mathcal{G}_N$. Because Hermitian symmetry is imposed, the time stepping routines are resolving $\boldsymbol{b}$ as described in (5.10), which is time-dependent. On decomposing the magnetic fields and eigenvalues into their real and imaginary components, so that $\tilde{\boldsymbol{b}} = \tilde{\boldsymbol{b}}_{\mathrm{r}} + \mathrm{i}\tilde{\boldsymbol{b}}_{\mathrm{i}}$ and $\lambda = \lambda_{\mathrm{r}} + \mathrm{i}\lambda_{\mathrm{i}}$, (5.10) is reduced to

$$\boldsymbol{b} = 2\mathrm{e}^{\lambda_{\mathrm{r}} t}\left[\tilde{\boldsymbol{b}}_{\mathrm{r}}\cos\left(\lambda_{\mathrm{i}} t\right) - \tilde{\boldsymbol{b}}_{\mathrm{i}}\sin\left(\lambda_{\mathrm{i}} t\right)\right], \tag{5.11}$$

and so it is possible to compare the fields directly by constructing $\boldsymbol{b}$ in (5.11) from the Arnoldi field components $\tilde{\boldsymbol{b}}_{\mathrm{r}}$ and $\tilde{\boldsymbol{b}}_{\mathrm{i}}$. To do this, we must know the phase of the time stepping snapshot, i.e. how far through a cycle the snapshot was taken, however it is not necessarily clear whether a time of $\lambda_{\mathrm{i}} t = 0 \bmod 2\pi$ would correspond to the minimum energy state of a cycle. Furthermore, the Arnoldi solver will actually be returning $\mathrm{e}^{\mathrm{i}\phi}\tilde{\boldsymbol{b}}(x, y, z)$, where $\mathrm{e}^{\mathrm{i}\phi}$ is an arbitrary phase factor. For complex eigenvalues, this phase factor can be absorbed into the field's periodic behaviour, though this then makes comparing the resulting fields of both methods more complicated. It is important to note that for $\lambda \in \mathbb{R}$, (5.11) reduces further to $\boldsymbol{b} = 2\mathrm{e}^{\lambda_{\mathrm{r}}}\tilde{\boldsymbol{b}}_{\mathrm{r}}$, in theory making direct field comparisons much simpler. Again, to account for the arbitrary phase factor that is introduced by the Arnoldi solver, we simply take the real part of $\mathrm{e}^{\mathrm{i}\phi}\tilde{\boldsymbol{b}}_{\mathrm{r}}$, which amounts to $\cos(\phi)\tilde{\boldsymbol{b}}_{\mathrm{r}}$ – a scalar multiple of $\tilde{\boldsymbol{b}}_{\mathrm{r}}$. This should not affect the comparison, as the fields are normalised by the visualisation routines.

Additionally, we have to remember that these fields are produced from eigenvector

estimates and as such, may not be accurate depictions of the actual eigenvectors. It is also possible that the Arnoldi fields show valid field configurations that cannot be reached through any time stepping methods. For the oscillatory eigenmodes, the Arnoldi method typically finds the pair of eigenvalues and their corresponding eigenvectors. It would be expected that these eigenmode pairs have identical field configurations but we find that this is not the case. Whilst bearing some resemblance, it seems that these field pairs share only minor details. For the above figure, the eigenmode field (from each pair) with the greatest visual similarity to the time stepping field was chosen which for classes III and V, correspond to negative $\mathrm{Im}(\lambda)$. The difference in the fields of these eigenmode pairs may be explained by the inaccuracy of the eigenvector estimates.

For a final case-study, we look at visualising the fields of non-oscillatory eigenmodes; those of class V at $R_m = 250$. At this R_m, it is difficult to calculate the eigenvectors due to the large size of the problem, and so the routine is only requested to find three eigenvalues. The eigenvalues are numbered 1 to 3, in descending order of $\mathrm{Re}(\lambda)$, with their respective configurations displayed in figure 5.23. The eigenvalues are $\lambda_1 = 0.0711$, $\lambda_2 = 0.0292$ and $\lambda_3 = 0.0167 - 0.488\mathrm{i}$. Importantly, we have identified the existence of two subdominant dynamo eigenmodes.

In the figure, panels (a) and (d) are equivalent, respectively, to (a) and (b) of figure 5.9. They are not substantially different in that two oppositely polarised flux tubes are present near to the α-type stagnation point. However, one cigar is forked, so that in the cross-section, it appears as two tubes of the same polarity; see panel (d). As this is a non-oscillatory eigenmode, it was hoped that the Arnoldi method's field would not deviate too far from that revealed through time stepping. Again, it has to be assumed that the differences are introduced by the eigenvector estimation process. Ironically, the structure of eigenmode 3, which is oscillatory, is (visually) closer to the actual time stepping structure, with two cigar-like flux tubes of opposing polarity.

Finally, we note that there is a distinct possibility that eigenmode 2 – panels (b) and (e) – could be the real partner to eigenmode 1, in that if $\lambda_1 = \lambda_+ = \alpha + \beta$,

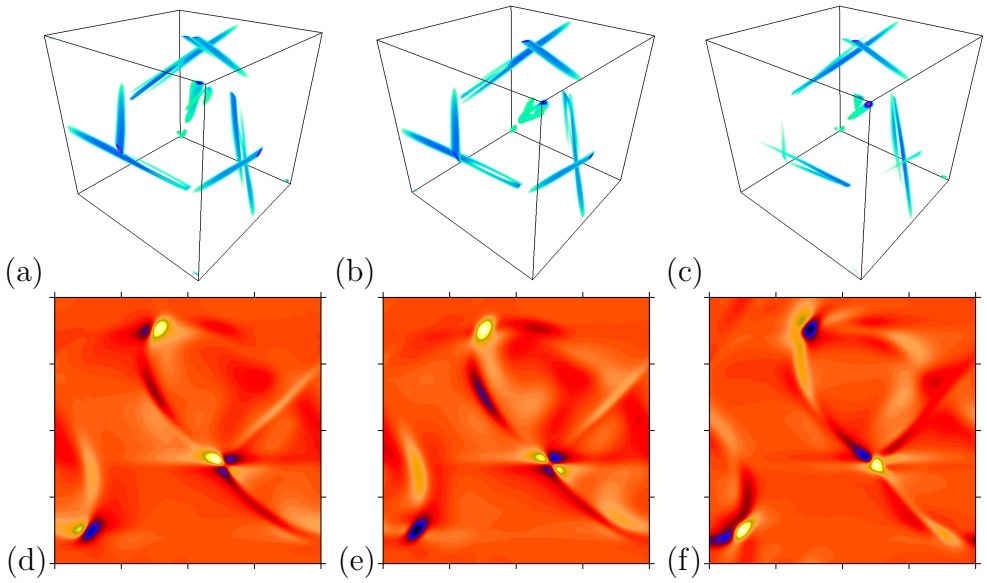

Figure 5.23: Visualisations produced using the IRAM routine. Shown are field magnitude $|\boldsymbol{b}|$ for eigenmodes (a) 1 to (c) 3 and field component b_z at $z = \pi/4$ for eigenmodes (d) 1 to (f) 3, in class V at $R_m = 250$.

then $\lambda_2 = \lambda_- = \alpha - \beta$, with α and β being real differentiable functions of R_m. Some structural similarity exists: eigenmode 1 has only one 'forked' cigar in each pair, whereas in mode 2, both cigars of the pair are forked; see panel (e).

5.5 Summary

In this final section, we provide a brief summary of the important findings and results. Firstly, the breakdown of the problem into five subproblems has been successful, in that we have been able to accurately reproduce (and improve upon) results of previous studies using only two of the five symmetry classes. We now know that the first window of dynamo action is provided by an eigenmode in class II, whereas the eigenmode of the second window is in class V, giving a greater understanding into the physical differences of the fields in each window. The eigenmode in class V remains dominant in the limit of our simulations and at $R_m = 10^4$, it reaches $\gamma = 0.0863$. The dominant eigenmodes in classes I, III and IV have not previously been identified, as they are always subdominant to either class II or class V (or both) at any given R_m. However, we have shown that fields in these

classes are also capable of dynamo action, thus dynamos can exist in all symmetry classes. Classes II, IV and V show no sign of slowing growth by $R_m = 10^4$, making these the best candidates for fast dynamos. It is even possible that another mode crossing occurs between class II and V, at some $R_m > 2.5{\times}10^4$. Classes I and III appear to resemble slow dynamos, as we see that $\gamma \to 0$ as $R_m \to \infty$.

Secondly, a closer look at the growth rate kinks confirm their nature as eigenvalue coalescences, decoalescences or mode crossings. By taking field snapshots from either side of the kinks, and using these as seed fields for a range of R_m, we either produce the same growth rate profile suggesting both seed fields belong to the same eigenmode or produce a distinct profile for each initial condition. The first case indicates coalescence or decoalescence, where coalescence describes a complex eigenvalue pair merging to form a real eigenvalue pair, with the reverse for decoalescence. The second case shows a mode crossing. We find an eigenvalue coalescence (or decoalescence) in classes II, IV and V (i.e. the fast dynamos), and mode crossings in all classes (particularly at low R_m). An example of each of coalescence (class V, $R_m \approx 215$), decoalescence (class II, $R_m \approx 330$) and mode crossing (class II, $R_m \approx 190$) are investigated in greater detail, confirming the assumptions made, about the causes of growth rate kinks, using the $\mathrm{Im}(\lambda)$ profiles.

Thirdly, the different symmetry restrictions of each class lead to different magnetic field configurations. There are essentially only two types of structure, the first type being flux ropes. The flux ropes are formed from six flux tubes in the principal vortices (regions where $|\boldsymbol{u}|$ is large and predominantly in one direction) and are seen in the dynamos of classes I and III at high R_m. For lower R_m, the individual tubes of the rope are merged so that only a single flux tube exists in each principal vortex. The second type of structure seen is that of the cigars, which are centred on the α-type stagnation points and extend towards either side of the same β-type point. These form in all classes except I. The dynamos of classes II, III, IV and V see varying amounts of cigar-like formations with one, three, four and two cigars respectively.

In class II, the first window's eigenmode contains only one cigar (per α-β pair),

centred on the separatrix. This is often accompanied by fin-like field projections of opposite polarity to the cigar. In class II's second dynamo mode, the structure is more complicated, with fins and narrow flux tubes surrounding the cigar, each alternatively joined to the central cigar, dependent on its polarity. Class V's two cigars are of opposite polarity and lie either side of the separatrix. In the oscillatory regime, these cigars strengthen and weaken in turn, alternately feeding off one another. The post-coalescent structure is non-oscillatory and consists of the same two cigars, however one is now permanently weaker than the other. The field-generating mechanisms that form this configuration are seen to be the most effective in the 1:1:1 ABC flow. Class IV is very similar to V, except that each cigar is replaced with a cigar-pair; the cigars within each pair have opposing polarity.

Classes I and III are very similar in structure and growth rates. The dynamo of class I consists of flux ropes in all six principal vortices, though this dynamo mechanism appears to be slow. Class III's dynamo is a hybrid of flux ropes and three cigars, with the field's structure oscillating between the two states over a single cycle. The flux ropes of class III lie in only four of the principal vortices, but this is owing to the non-uniqueness of the irreducible representations used for this class. The flux ropes are visually similar to those seen in a Ponomarenko dynamo, which is also a slow dynamo.

A brief analysis through a transformation to cylindrical coordinates allows us to show (analytically) that single cigars lying on the separatrix and centred at the α-type point can only exist in classes II and V, which have the two highest growth rates by $R_m = 10^4$. Despite single cigars not existing in the other classes, multiple cigars are permissible.

An investigation into the field structures of classes II, IV and V at high R_m allows us to see the scale of the cigars in comparison to the chaotic regions. It is noted that the structures are far smaller than the chaotic regions at the α-type stagnation point and relatively small compared to the chaotic regions of the separatrix mid-point. Class II's cigar (and associated fin-like projections) are three-fold rotationally symmetric, and are clearly related to the three-fold symmetry of the

chaotic regions. With multiple-cigar configurations, the cigars do not lie on the separatrix and are oriented with respect to each of the three connected chaotic regions. In class IV, the two major cigars are each related to one of the three chaotic bands, whereas the minor cigar pair both share one band. The same is seen in class V, with the major cigar formed, seemingly, in relation to two of the three chaotic bands, while the minor cigar forms in relation to the third. The configuration in class IV is unstable, in that it is oscillatory, whereas it is stable in class V. We also find in class IV that very small-scale features exist, these being associated with the very narrow bands of chaos. Since the eigenmode of class II is oscillating at $R_m = 4000$, it is possible that such structures, not seen in the supplied snapshot, exist at other times in the cycle.

The spectra reveal a little more detail on the field structures that exist in each class. Class I and, to some extent, class III show energy distributed over a large range of scales, which is a feature of slow dynamos, in fitting with interpretations of the growth rate in the high-R_m limit and structural similarity to the Ponomarenko dynamo. For $R_m \leq O(10^3)$, this is also seen for class IV, also fitting an apparent slow dynamo profile, with decreasing γ for increasing R_m. After the mode change at $R_m \approx 6000$, the spectrum more closely resembles those of classes II and V, which show a power law dependence of approximately k which is consistent with the magnetic field being concentrated into individual flux tubes (cigars). The spectra also show that the choice of N at each R_m is sufficiently large and so the field is well-resolved on the smallest scales.

Finally, interesting results of the Arnoldi iteration method are provided. The growth rates of dominant and subdominant eigenmodes are given for all five symmetry classes. They reveal a complex array of subdominant modes with their own mode crossings and growth rate kinks. This method verifies the growth rates calculated for the dominant eigenmodes of each class, at least for $R_m \leq 100$. It also provides further insight into the mode crossings of class IV, previously only seen through the discontinuities of $\mathrm{Im}(\lambda)$. In classes II and V, it also shows that there are no subdominant dynamo modes for $R_m \leq 100$. Through the case study

of class II's mode crossing, the Arnoldi method allows both branches to be followed past the crossing and identifies that the non-oscillatory mode emerges from an eigenvalue coalescence in the interval $R_m \in [174.2, 174.25]$, explaining why no such non-oscillatory mode could be found at $R_m = 170$ in the time stepping problem. From the discussion of coalescences in section 5.2, we know that two real eigenmodes exist with corresponding eigenvalues λ_+ and λ_-, and with the Arnoldi method, we are able to not only confirm the existence of the subdominant λ_- eigenmode, but also follow its growth rate for the range of R_m investigated in this particular case study.

The eigenvalues calculated have associated eigenvectors, which the Arnoldi method estimates from the Ritz vectors. We are able to reconstruct the magnetic field of an eigenmode through the same unpacking and mode copying procedures used in the routine itself. Visualisations produced for classes II to V at $R_m = 100$ (oscillatory regime), identify some structural similarities with the time stepping problem, particularly in the case of class II, in which the structure greatly resembles that of the time stepping. For oscillatory fields, the Arnoldi method does not produce the same field $\boldsymbol{b}$ as in the time stepping, though one can, in theory, be reconstructed from the other. In the steadily evolving fields, we would expect the Arnoldi solver's snapshots to resemble those of the time stepping, even despite the additional phase factor introduced by the Arnoldi method. However, in another case study of class V at $R_m = 250$, we find physical differences, such as the cigars being 'forked', as opposed to the standard cigars seen earlier in the chapter. We put these differences down to the fact that there are errors when estimating the eigenvectors; more analysis is needed to determine whether this is truly the case. Furthermore, this case study allows us to find subdominant dynamos: not only do we find the eigenvalue λ_- which corresponds to the dominant mode λ_+ but a third eigenvalue that is oscillating at this value of R_m and has a positive growth rate, thus a subdominant dynamo mode. These additional eigenmodes are visualised and it is found, perhaps coincidentally, that the third (oscillating) eigenvalue most closely resembles the field of the time stepping problem.

Chapter 6

Fluid stability of the 1:1:1 ABC flow

Why bother with a cunning plan when a simple one will do?

Terry Pratchett

In this short chapter, the problem of linear hydrodynamic stability of the 1:1:1 ABC flow is investigated, as a function of the Reynolds number R_e. We explore this problem as a natural extension to our investigation into the 1:1:1 ABC dynamo, due to the similarity between the two. In studying this problem, we adopt a similar numerical scheme with only minor changes made and, again, take advantage of the symmetries of the flow to reduce the full problem to a set of five sub-problems, the sum of which is computationally less intensive than the original. Through the symmetries, we are able to see the different possible independent perturbation configurations that may occur and make direct comparisons to their equivalent magnetic field structures.

The chapter begins by formulating the problem in terms of $\boldsymbol{U} = \boldsymbol{u} + \boldsymbol{v}$ (the ABC flow plus a perturbation) in the Navier-Stokes equation and obtaining an equation

for the evolution of $\boldsymbol{v}$, which is then linearised. The majority of the technical details regarding symmetries remain the same, and so we only give the numerical scheme used to time step each mode. We then show the results of simulations carried out for R_e up to 10^3, providing eigenvalue profiles and visualisations of the structure of the perturbation $\boldsymbol{v}$, supplemented by energy spectra. The chapter concludes with a summary of the main findings.

6.1 Formulating the problem

The problem of linear fluid stability in the ABC flows is, mathematically, very close to the kinematic dynamo problem; the link between the two problems is examined in detail by Moffatt (1985, 1986). The investigation of Galloway and Frisch (1987) into the ABC flow's linear stability is a natural follow-up to that of the dynamo problem (Galloway and Frisch, 1986), as it can be solved with only minor modifications. For the same reason, we investigate the linear fluid stability problem but with the additional symmetry reductions developed for the ABC dynamo problem.

We give the 1:1:1 ABC flow again, so as to avoid referring to previous chapters, as

$$\boldsymbol{u} = (\sin z + \cos y, \sin x + \cos z, \sin y + \cos x), \tag{6.1}$$

In the first chapter, it was mentioned that the ABC flows have the Beltrami property $\nabla \times \boldsymbol{u} = k\boldsymbol{u}$, and this can also be written as

$$0 = \boldsymbol{u} \times (\nabla \times \boldsymbol{u}) = \tfrac{1}{2}\nabla|\boldsymbol{u}|^2 - \boldsymbol{u} \cdot \nabla \boldsymbol{u}. \tag{6.2}$$

From the first chapter, we also recall the non-dimensionalised Navier-Stokes equation as

$$\partial_t \boldsymbol{U} + \boldsymbol{U} \cdot \nabla \boldsymbol{U} = -\nabla P + \varepsilon \nabla^2 \boldsymbol{U} + \boldsymbol{f}, \tag{6.3}$$

for an incompressible flow $\boldsymbol{U}$, with ε as the kinematic viscosity and with the force $\boldsymbol{f}$

set to $\varepsilon\boldsymbol{u}$ to counter viscous dissipation, as $\nabla^2\boldsymbol{u} = -\boldsymbol{u}$. To investigate the stability of the 1:1:1 ABC flow, we introduce a perturbation velocity $\boldsymbol{v}$ and set $\boldsymbol{U} = \boldsymbol{u} + \boldsymbol{v}$ and $P = -\frac{1}{2}\boldsymbol{u}^2 + p$ in view of (6.2) where p is the perturbation pressure. The field $\boldsymbol{v}$ obeys, exactly,

$$\partial_t\boldsymbol{v} + \boldsymbol{u}\cdot\nabla\boldsymbol{v} + \boldsymbol{v}\cdot\nabla\boldsymbol{u} + \boldsymbol{v}\cdot\nabla\boldsymbol{v} = -\nabla p + \varepsilon\nabla^2\boldsymbol{v}, \tag{6.4}$$

which in the linear approximation becomes

$$\partial_t\boldsymbol{v} + \boldsymbol{u}\cdot\nabla\boldsymbol{v} + \boldsymbol{v}\cdot\nabla\boldsymbol{u} = -\nabla p + \varepsilon\nabla^2\boldsymbol{v}, \tag{6.5}$$

with $\nabla\cdot\boldsymbol{v} = 0$. This has a very similar structure to the magnetic field problem, with $\varepsilon = R_e^{-1}$ now as an inverse Reynolds number. In parallel to the dynamo growth rates and their dependence on R_m, we are interested in the evolution of $\boldsymbol{v}$ as R_e is increased, and the resulting perturbation structures. Unlike the magnetic problem, we now have $\boldsymbol{u}\cdot\nabla\boldsymbol{v} + \boldsymbol{v}\cdot\nabla\boldsymbol{u}$ as opposed to $\boldsymbol{u}\cdot\nabla\boldsymbol{b} - \boldsymbol{b}\cdot\nabla\boldsymbol{u}$ of the magnetic problem, as well as an additional pressure gradient ∇p. The pressure is an unknown quantity but we use the fact that the field $\boldsymbol{v}$ is divergence-free to project the pressure out; this is seen in the next section.

6.2 Numerical methods

As with the dynamo problem, we are solving (6.5) which is linear and thus an eigenvalue problem. We therefore are looking for solutions of the form

$$\boldsymbol{v} = \tilde{\boldsymbol{v}}(x, y, z)\mathrm{e}^{\lambda t}, \tag{6.6}$$

with eigenvalues λ. In Fourier space, the perturbation takes the form

$$\boldsymbol{v} = \sum_{l,m,n} \boldsymbol{v}_{l,m,n}(t)\,\mathrm{e}^{\mathrm{i}lx+\mathrm{i}my+\mathrm{i}nz}\,, \qquad \boldsymbol{v}_{l,m,n} = (X_{l,m,n}, Y_{l,m,n}, Z_{l,m,n}), \tag{6.7}$$

and for the sake of simplicity, we keep the notation X, Y and Z as the respective components of $\boldsymbol{v}$, as in the magnetic problem. The perturbation field $\boldsymbol{v}$ is real-valued and so the field has Hermitian symmetry, given by

$$\boldsymbol{v}_{l,m,n} = \boldsymbol{v}^*_{-l,-m,-n}. \tag{6.8}$$

We continue to use a three-step Adams–Bashforth numerical scheme for the advective terms, with the diffusive terms integrated exactly, and solve

$$\boldsymbol{v}^{j+1}_{l,m,n} = E_{l,m,n}\left[\boldsymbol{v}^{j}_{l,m,n} + \tfrac{1}{12}\Delta t\Big(23\boldsymbol{\Gamma}^{j}_{l,m,n} - 16E_{l,m,n}\boldsymbol{\Gamma}^{j-1}_{l,m,n} + 5E^2_{l,m,n}\boldsymbol{\Gamma}^{j-2}_{l,m,n}\Big)\right], \tag{6.9}$$

with

$$E_{l,m,n} = \exp\left[-\varepsilon\left(l^2+m^2+n^2\right)\Delta t\right]. \tag{6.10}$$

The term $\boldsymbol{\Gamma}$ in (6.9) includes the projection of the pressure gradient from the convective acceleration terms. We have that

$$\boldsymbol{\Gamma}^{j}_{l,m,n} = \left[\boldsymbol{k}\left(\boldsymbol{k}\cdot\boldsymbol{\Lambda}^{j}_{l,m,n}\right)/k^2\right] - \boldsymbol{\Lambda}^{j}_{l,m,n}, \tag{6.11}$$

where $\boldsymbol{k} = (l,m,n)$ and $k^2 = |\boldsymbol{k}|^2 = l^2+m^2+n^2$. The term $\boldsymbol{\Lambda}^{j}_{l,m,n}$ differs only slightly to its magnetic counterpart (4.9)–(4.11), with the x-component as

$$\begin{aligned}
2(\Lambda_x)^{j}_{l,m,n} &= \mathrm{i}mB\left(Y^{j}_{l,m-1,n} + Y^{j}_{l,m+1,n}\right) - \mathrm{i}mC\left(X^{j}_{l,m,n-1} + X^{j}_{l,m,n+1}\right)\\
&+ mC\left(Y^{j}_{l,m,n-1} - Y^{j}_{l,m,n+1}\right) - mA\left(X^{j}_{l-1,m,n} - X^{j}_{l+1,m,n}\right)\\
&+ \mathrm{i}nB\left(Z^{j}_{l,m-1,n} + Z^{j}_{l,m+1,n}\right) - \mathrm{i}nA\left(X^{j}_{l-1,m,n} + X^{j}_{l+1,m,n}\right)\\
&+ nC\left(Z^{j}_{l,m,n-1} - Z^{j}_{l,m,n+1}\right) - nB\left(X^{j}_{l,m-1,n} - X^{j}_{l,m+1,n}\right)\\
&- 2C\left(Z^{j}_{l,m,n-1} + Z^{j}_{l,m,n+1}\right) - 2\mathrm{i}B\left(Y^{j}_{l,m-1,n} - Y^{j}_{l,m+1,n}\right).
\end{aligned} \tag{6.12}$$

The y-component is given by

$$\begin{aligned}2(\Lambda_y)^j_{l,m,n} = &\,\mathrm{i}lC\left(X^j_{l,m,n-1} + X^j_{l,m,n+1}\right) - \mathrm{i}lB\left(Y^j_{l,m-1,n} + Y^j_{l,m+1,n}\right)\\ &+ lA\left(X^j_{l-1,m,n} - X^j_{l+1,m,n}\right) - lC\left(Y^j_{l,m,n-1} - Y^j_{l,m,n+1}\right)\\ &+ \mathrm{i}nC\left(Z^j_{l,m,n-1} + Z^j_{l,m,n+1}\right) - \mathrm{i}nA\left(Y^j_{l-1,m,n} + Y^j_{l+1,m,n}\right)\\ &+ nA\left(Z^j_{l-1,m,n} - Z^j_{l+1,m,n}\right) - nB\left(Y^j_{l,m-1,n} - Y^j_{l,m+1,n}\right)\\ &- 2A\left(X^j_{l-1,m,n} + X^j_{l+1,m,n}\right) - 2\mathrm{i}C\left(Z^j_{l,m,n-1} - Z^j_{l,m,n+1}\right), \end{aligned} \tag{6.13}$$

and the z-component as

$$\begin{aligned}2(\Lambda_z)^j_{l,m,n} = &\,\mathrm{i}lA\left(X^j_{l-1,m,n} + X^j_{l+1,m,n}\right) - \mathrm{i}lB\left(Z^j_{l,m-1,n} + Z^j_{l,m+1,n}\right)\\ &+ lB\left(X^j_{l,m-1,n} - X^j_{l,m+1,n}\right) - lC\left(Z^j_{l,m,n-1} - Z^j_{l,m,n+1}\right)\\ &+ \mathrm{i}mA\left(Y^j_{l-1,m,n} + Y^j_{l+1,m,n}\right) - \mathrm{i}mC\left(Z^j_{l,m,n-1} + Z^j_{l,m,n+1}\right)\\ &+ mB\left(Y^j_{l,m-1,n} - Y^j_{l,m+1,n}\right) - mA\left(Z^j_{l-1,m,n} - Z^j_{l+1,m,n}\right)\\ &- 2B\left(Y^j_{l,m-1,n} + Y^j_{l,m+1,n}\right) - 2\mathrm{i}A\left(X^j_{l-1,m,n} - X^j_{l+1,m,n}\right), \end{aligned} \tag{6.14}$$

with $A = B = C = 1$ for the 1:1:1 ABC flow (6.1).

The same symmetry considerations apply to fluid stability as to the magnetic problem, as they are driven by the same flow. We are solving this problem in the fundamental domain $\mathcal{H}_N$, which means that once (6.9) has been solved for all modes within $\mathcal{H}_N$, then the modes immediately outside of $\mathcal{H}_N$ need to be filled. This is carried out in an identical fashion to the magnetic problem, with mode copying carried out as described in (4.42)–(4.45) at the beginning of each time step. As with the magnetic routines, each class is numerically solved using a distinct program and thus there are five branches of results.

To investigate the behaviour of the perturbation $\boldsymbol{v}$, we calculate its kinetic energy E_K at every time step, in the same fashion that E_M is calculated in the magnetic problem, building up a time series of $\ln(E_\mathrm{K})$ allowing us to calculate the growth rate $\gamma = \mathrm{Re}(\lambda)$ from the linear trend and $\mathrm{Im}(\lambda)$ from the oscillation period (if applicable). The same parameters are used as in the magnetic problem, namely

those summarised in table 4.1 of section 4.1.2. Random seed fields are used unless stated otherwise.

Finally, we note that the Arnoldi iteration method was not used in solving the fluid stability problem, though it is possible to do so. On modifying the core operation within each time step to account for differences in $\mathbf{\Lambda}$ and to then calculate $\mathbf{\Gamma}$, no further adjustments are required due to the similarity of these problems. This also means that the parallel codes could easily be adapted to solve this problem, as the communication and mode copying steps are unaffected. We have not parallelised these routines, however, as the values of R_e investigated did not warrant doing so.

6.3 Results

This section provides the results of the linear hydrodynamic stability simulations carried out in all five symmetry classes. To verify that the routines were solving the correct problem, we draw on the results of Galloway and Frisch (1987), who investigate up to $R_e = 200$ for the 1:1:1 ABC flow. We use a full domain ($\mathcal{F}_N$) code with no assumption of symmetry for a comparison with their results. They see an unstable mode emerging at $R_e \approx 13$ and the growth rate appears to become asymptotic by $R_e = 200$, with $\gamma = 0.216$. We are unable to follow this mode up to $R_e = 200$, but for $R_e \leq 160$ our results agree well with growth rates inferred from the graph; see figure 1.5.

6.3.1 Instability growth rates and frequencies

The numerical results are summarised in figures 6.1 and 6.2, the first of which identifies the growth rate of $\boldsymbol{v}$ and the second, the imaginary part of the dominant eigenvalue. It is clear that the results are very different to that of the magnetic case, with many growth rate kinks. In parallel with the magnetic case, all five symmetry classes contain flows in which the growth rate of the instability is positive. The first flow to become unstable is in class V and this occurs somewhere in the interval

$R_e \in [13, 14]$. As R_e is increased, classes III, IV and II follow in the intervals $R_e \in [23, 24], [23, 24]$ and $[26, 30]$ respectively. Class I is the last to have an unstable mode, emerging in the interval $R_e \in [49, 50]$. It must be noted that in the non-linear hydrodynamic problem, investigated by Podvigina and Pouquet (1994), the flow is seen to become unstable in the interval $R_e \in [13, 13.05]$, which our results are in agreement with. In contrast to the magnetic case, we see modes from all but class IV taking the role of dominant eigenmode.

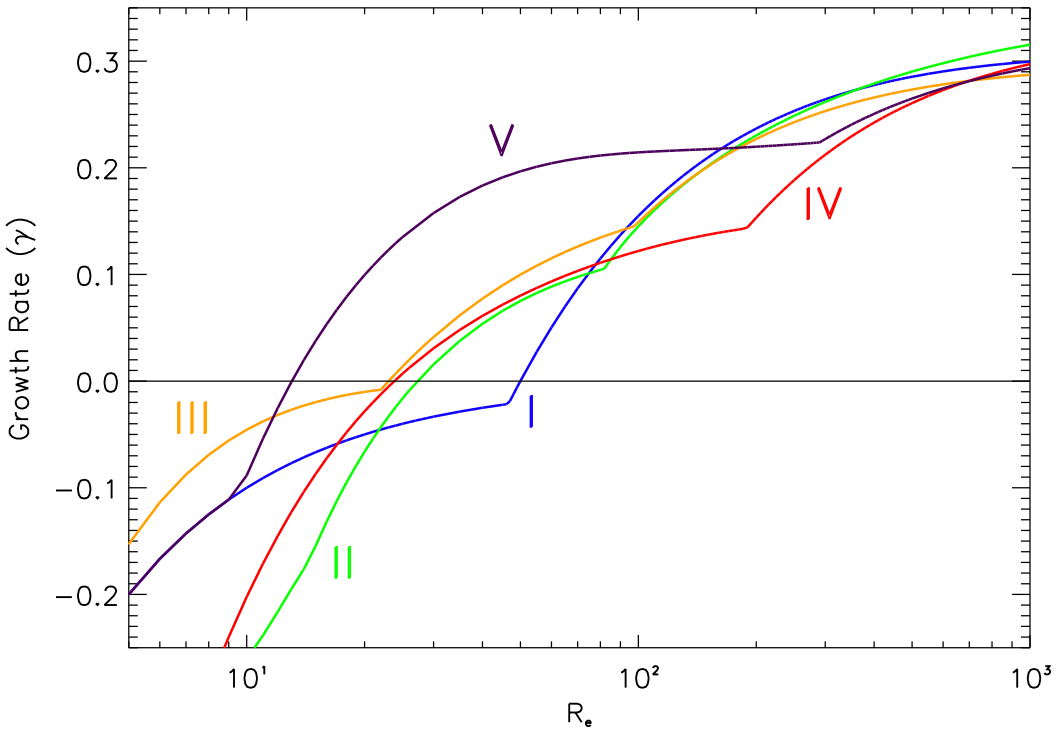

Figure 6.1: Growth rates for fluid instabilities of classes I to V in the 1:1:1 ABC flow.

Class V remains the most unstable flow until $R_e \approx 160$, at which point a flow in class I becomes the dominant mode. It is then superseded by a steady mode in class II at $R_e \approx 370$, with this mode remaining dominant up to the largest R_e simulated and has a growth rate of $\gamma = 0.293$ at $R_e = 10^3$. In contrast to the magnetic problem, all classes have remarkably similar growth rates at large R_e and the instabilities have significantly larger growth rates than their magnetic field counterparts. Regarding the results of Galloway and Frisch (1987), we note that

they do not appear to identify the crossing between eigenmodes of classes V and I for the 1:1:1 ABC flow. It is possible that the spacing (in R_e) of their simulations was not sufficiently small to identify the change, though we would expect to see some evidence of a kink in their final result. At $R_e = 200$, class V's dominant eigenmode has $\gamma = 0.220$, not significantly different from their value of 0.216. They also find that this mode's oscillation period is 15.4 at this R_e, as compared to our result of 14.8. It is therefore possible that they have followed the eigenmode of class V past the crossing, although this would require using a settled field at a lower R_e (e.g. 150) as a seed field at $R_e = 200$, which they may well have done. As a check, we compared the growth rates of the full code (without symmetries) to the envelope of growth rates from all five classes up to $R_e = 250$ and found that the growth rate kink also exists in the full code too, suggesting that it is not, somehow, an artefact of the reduced codes.

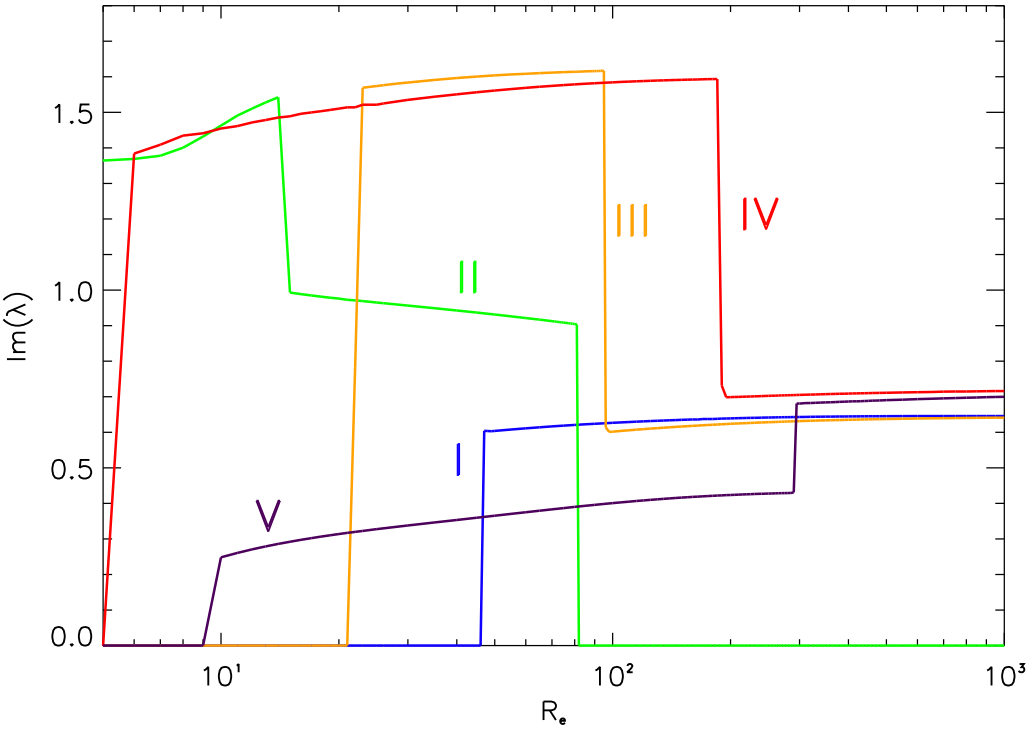

Figure 6.2: Frequencies for fluid instabilities of classes I to V in the 1:1:1 ABC flow.

All growth rate kinks identified, up to the largest R_e explored, are mode crossings.

From figure 6.2, we observe only discontinuous changes in frequency indicative of completely different modes emerging within each class. Furthermore, the growth rate profile of each branch appears to be approaching an asymptotic value prior to the mode changes; for no R_e can we find a falling instability growth rate. We see no cases of eigenvalue coalescence or decoalescence.

For $R_e < 5$, the dominant eigenmodes of all classes are steady, except for class II. As R_e increases, class II is the only mode for which the dominant modes have successively lower frequencies. Subsequently, it is also the only class whose eigenmode at high R_e is steady. The modes of classes I, III, IV and V at $R_e = 10^3$ are oscillatory with surprisingly little difference in their frequencies. Notably, classes I and III are very close in frequency, as are classes IV and V, a potential indicator that their instabilities are similar in nature. It is also important to be aware that the strongest instability at $R_e = 10^3$ is steady as with the dominant mode of the magnetic problem.

With regards to projections for behaviour at even higher R_e, it is entirely possible that an eigenmode from another class becomes dominant. We see at least one mode change for every class, and so it is not unreasonable to expect this behaviour to continue outside of the range of R_e investigated.

6.3.2 Fluid perturbation structures

In this section we provide discussion on visualisations of the magnitude of the fluid perturbation in the form of 3-dimensional snapshots of $|\boldsymbol{v}|$ and cross-sections of v_z at $z = \pi/4$. Unlike the magnetic problem, we choose to visualise only the growing instabilities, as all classes have growing instabilities for moderate values of R_e. We first visualise instabilities at $R_e \sim O(10^2)$; at the specific R_e chosen for each class, the respective eigenmode is the same one to first become unstable in that class. These fields are seen in figures 6.3 (3-d) and 6.4 (2-d) for classes I, IV and V at $R_e = 100$, class II at $R_e = 80$ and class III at $R_e = 90$. It is immediately noticeable that the configurations are, for the most part, very different to those seen in the

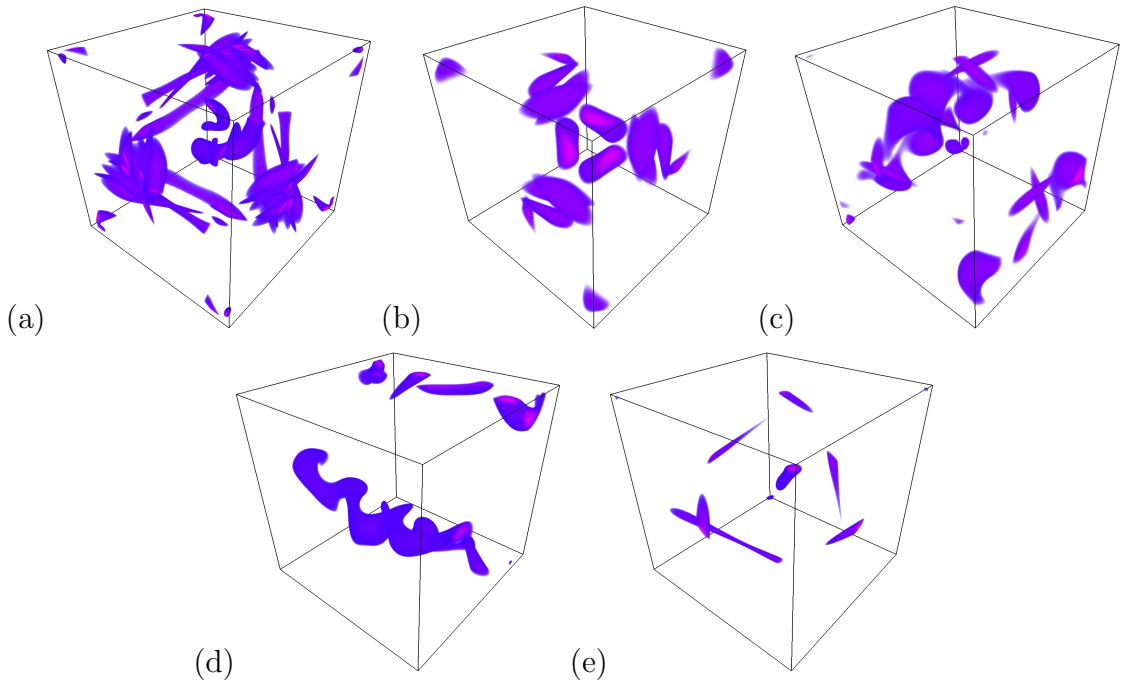

Figure 6.3: Visualisations of magnitude of fluid instability $|\boldsymbol{v}|$ for classes (a) I, $R_e = 100$; (b) II, $R_e = 80$; (c) III, $R_e = 90$; (d) IV, $R_e = 100$ and (e) V, $R_e = 100$.

magnetic problem and in combination with the growth rates in figure 6.1, suggest that the symmetries play a less crucial role in the fluid instability problem. Firstly, the areas of strongest instability in class I's mode are located near to the α-type stagnation point. From snapshots taken at other times during the cycle, it appears that these instabilities are also linked to the principal vortices, particularly those parts that lie close to the separatrices; this is perhaps more clear in figure 6.4(a), where structures are visible roughly vertically along the line $x = \pi/2$. A clear three-fold rotational symmetry is seen in figure 6.3(a), as is to be expected, owing to the action of d and d^2 on class I. The strongest instabilities of class II appear as flattened cigars around the α-type stagnation point. Again, the broader scale of these features means that they appear to be simultaneously related to both the principal vortices and the stagnation points. The perturbation direction in each of these 'cigars' is not entirely clear from the cross-section, as only two clearly separate features are seen. The instabilities of classes III and IV are concentrated primarily in the principal vortices and appear as twisted coil-like structures. It is important to note that only two of the six principal vortices, in each class,

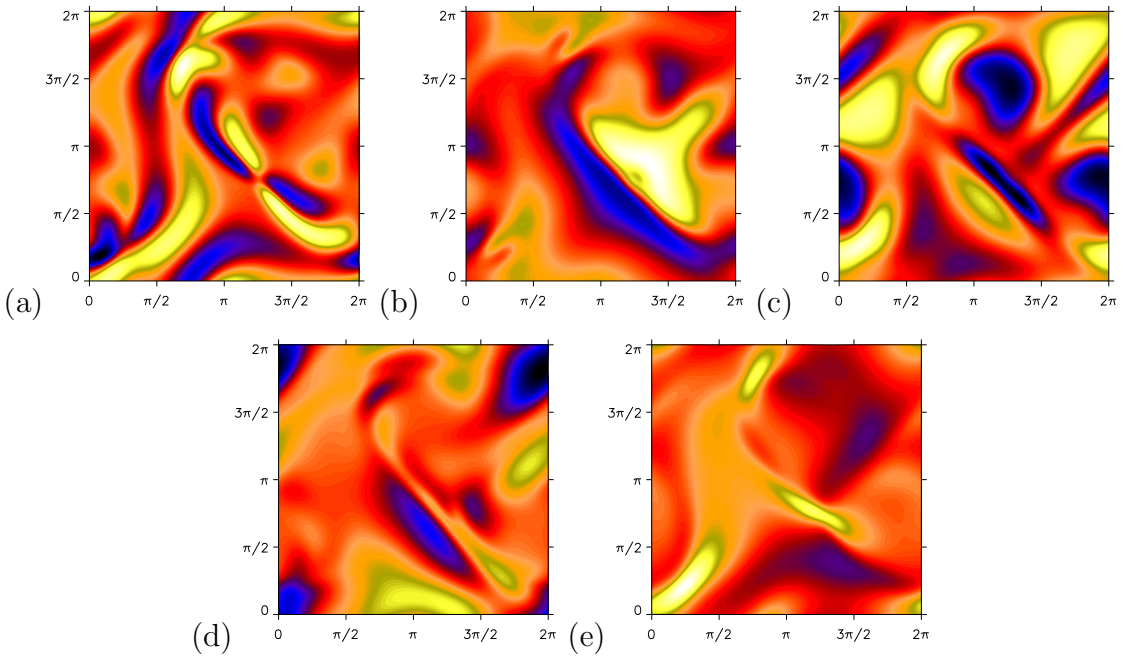

Figure 6.4: Fluid instability component v_z at $z = \pi/4$ for classes (a) I, $R_e = 100$; (b) II, $R_e = 80$; (c) III, $R_e = 90$; (d) IV, $R_e = 100$ and (e) V, $R_e = 100$.

seem to generate the instabilities. The six principal vortices can be paired off by the direction of strong flow, each pair will have large $|u_x|, |u_y|$ or $|u_z|$ with small magnitudes in the other directions. The instabilities in each of these classes seem to form in a single principal vortex pair, suggesting that one direction is more amenable to instability formation than the others. This is not surprising, as classes III, IV and V lack the three-fold rotational symmetry of classes I and II. It may be asked why a certain direction appears to be favoured above others but, as in the magnetic problem, the direction is arbitrary and is related to the choice of irreducible representations. Owing to the multi-dimensional representations being non-unique, alternative choices of irreps could lead to different directions being favoured. Class III appears also to have a single cigar-shaped perturbation near to the α-type stagnation point, though a weaker (secondary) perturbation acting in the opposite direction is revealed by the cross-section. In class V, a flattened cigar-like perturbation is clearly visible and through additional visuals (not given here), it is found to be twisted as it is stretched along the separatrix. The cigar lies on the separatrix and is centred at the α-type stagnation point, in a similar fashion to class II's mode of the magnetic problem.

The second set of visualisations – figures 6.5 and 6.6 – are from snapshots taken at $R_e = 300$ for all classes. This value was chosen as the dominant modes are, by this point, on the same branches as the dominant modes at $R_e = 10^3$, and so the flows have a similar structure but are not so tightly confined to the separatrices. The instabilities of these new eigenmodes now appear to more closely resemble the structures seen in the magnetic problem, with all but class I exhibiting strong perturbations in cigar-like features. However, we still see much more structure in the cross-sections than in the magnetic case, indicating that the instabilities are only partly linked to the chaotic regions near the separatrices.

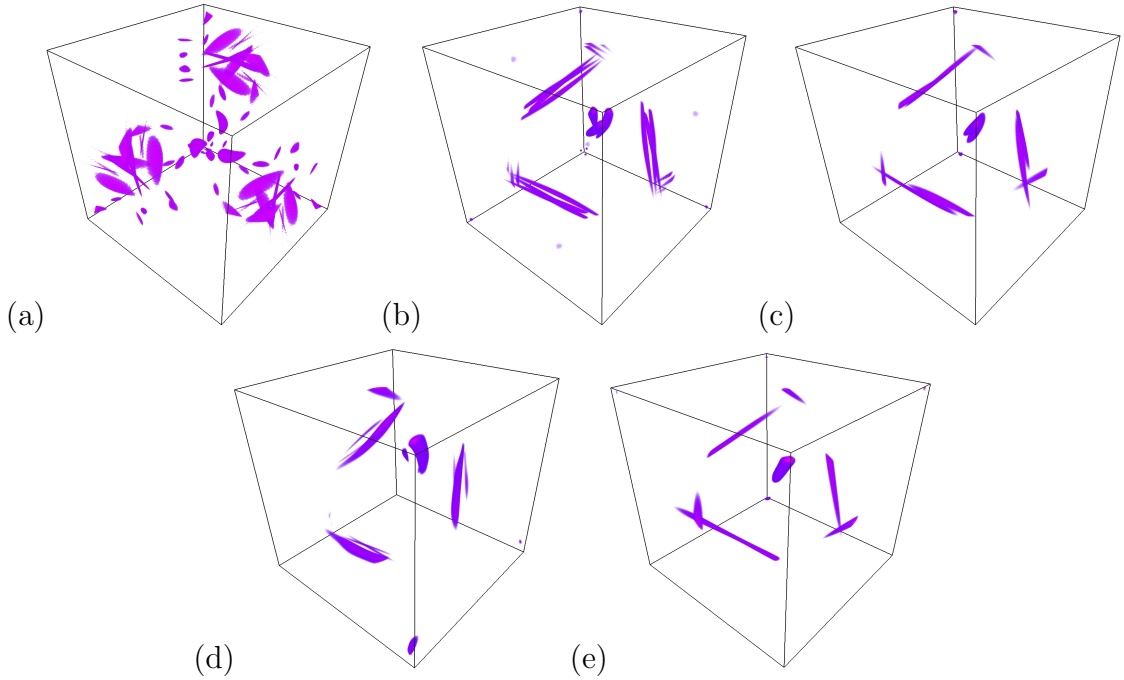

Figure 6.5: Visualisations of magnitude of fluid instability $|\boldsymbol{v}|$ for (a) I to (e) V for $R_e = 300$

In class I, the eigenmode seen at $R_e = 300$ is the same as that seen at $R_e = 100$. In figures 6.5(a) and 6.6(a), we are simply seeing another part of the periodic instability. The perturbations appear to be strongest in the principal vortices but unstable flow is also situated around the separatrix mid-points connecting the α- and β-type stagnation points, rather than centred on the α-type stagnation points. Instabilities also build up at each of the three bands of chaos lying on the two-dimensional stable manifolds of the α-type stagnation points. In the flow cross-

section, the well-developed structure associated with the principal vortices seems to have alternating direction, much like the flux ropes seen in class I's magnetic field. Overall, the structure is complicated and the instabilities of this class appear to be driven by both areas of large $|\boldsymbol{u}|$ and the chaotic regions. The dominant flows in

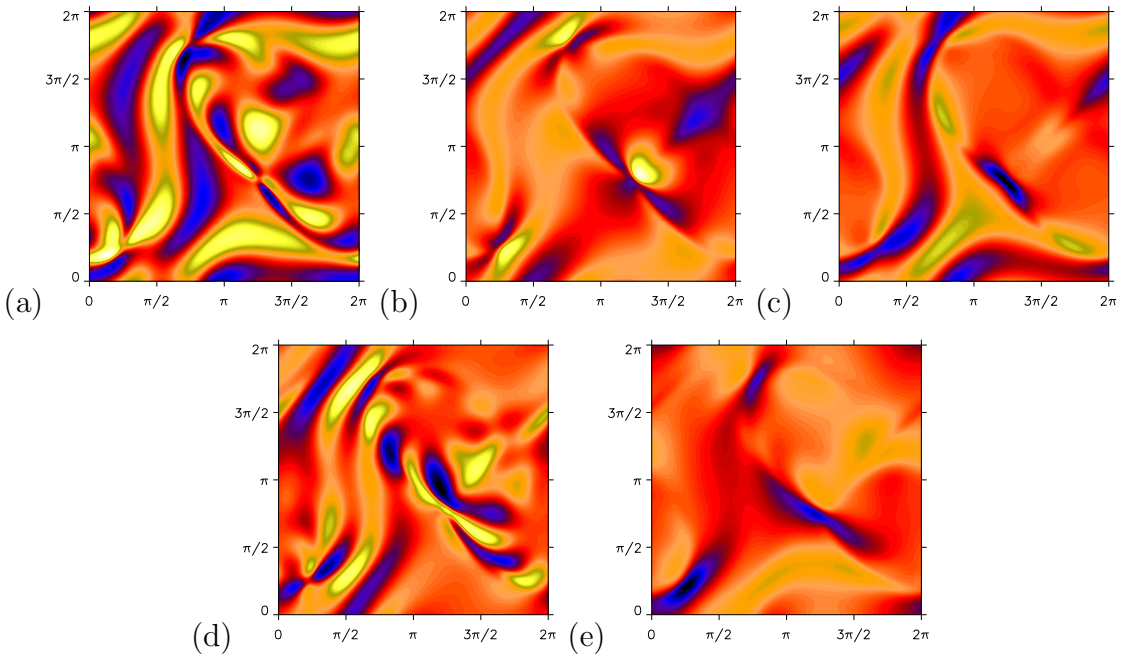

Figure 6.6: Fluid instability component v_z for (a) I to (e) V for $R_e = 300$ at $z = \pi/4$.

classes II–V are reminiscent to those of classes II, IV and V of the magnetic problem, with the familiar multiple cigar configurations. Again, imposing the symmetries forces the flows into certain configurations. Class II includes three cigars of strong flow concentrated around the α-type stagnation points that are stretched equally, in both directions, towards the β-type stagnation points. Referring to the cross section of v_z at $z = \pi/4$, the flow in two of the tubes is of the same sign, or direction, with the third of the opposite sign. Note that in this class, strong perturbations also form in proximity to the β-type stagnation points; these are not seen in figure 6.5(b) but are in figure 6.6(b). Class III has two cigars of flow centred on the α-type stagnation point. These two cigars are joined at the stagnation point, however, with each tube stretched toward a different side of the same β-type stagnation point. Interestingly, these connected tubes are of the same sign, seen as a flat area of strong flow at $\boldsymbol{r}_1 = (5, 3, 1)\pi/4$ in figure 6.6(c).

The most unstable flow in class IV includes a single major and multiple minor tubes centred on the α-type stagnation point. The cross-section is less clear at $\boldsymbol{r}_1 = (5, 3, 1)\pi/4$ than others, though we can identify stronger instabilities near $\boldsymbol{r}_2$ and relatively strong flows in the principal vortex that runs vertically on the left-hand side of the section. Not quite so clear, in the cross-section, are medium-strength instabilities formed in the two principal vortices running perpendicular to the section, further suggesting that the principal vortices are contributing to the instabilities seen in this class. The strongest perturbations in class V are, again, centred on the α-type stagnation points and are stretched into the classic cigars, though appearing more flattened than their counterparts in the magnetic case. These features resemble those of figure 2 in Galloway and Frisch (1987), which appears to belong to the same branch of solutions as those in class V, suggesting that they are indeed following this subdominant branch past the crossing that we identify at $R_e \approx 160$. Drawing on the brief analysis on the possibility of single cigars, from the previous chapter, we note for the fluid instability problem, that only class V is seen to have a single cigar in contrast to classes II and V of the magnetic problem. As the symmetries are unchanged for this problem, the possibility of a single tube of $\boldsymbol{v}$ for class II remains and this configuration is most likely present as a subdominant mode.

Overall, in comparing visualisations of both the magnetic and fluid problems at equivalent parameter regimes, the structures of the fluid problem appear to be less ordered and much more loosely linked to the thin bands of Lagrangian chaos, with substantial instabilities also found in the principal vortices.

6.3.3 Energy spectra

Finally, we compare the structures of the fluid instabilities, seen above, with energy spectra produced at $R_e = 10, 100$ and 1000. To obtain the spectra of oscillating perturbations, the average energy is calculated over a single cycle. The first point we make is that the energy, in all five symmetry classes, is mainly located in the

low-wavenumber modes, suggesting that the instabilities appear as less confined features, even at high R_e. Certainly, this agrees with what we are able to see from the visualisations, as although the strongest instabilities occur in cigar-like features (these being fairly small-scale structures) in classes II-V, we also see moderately unstable flow in other regions such as the principle vortices, which are spatially broad in comparison to the cigars. We see that the majority of the energy is located around $k \sim O(10)$, even for the largest R_e investigated, indicating that all unstable eigenmodes of this problem may share this property. The similarity

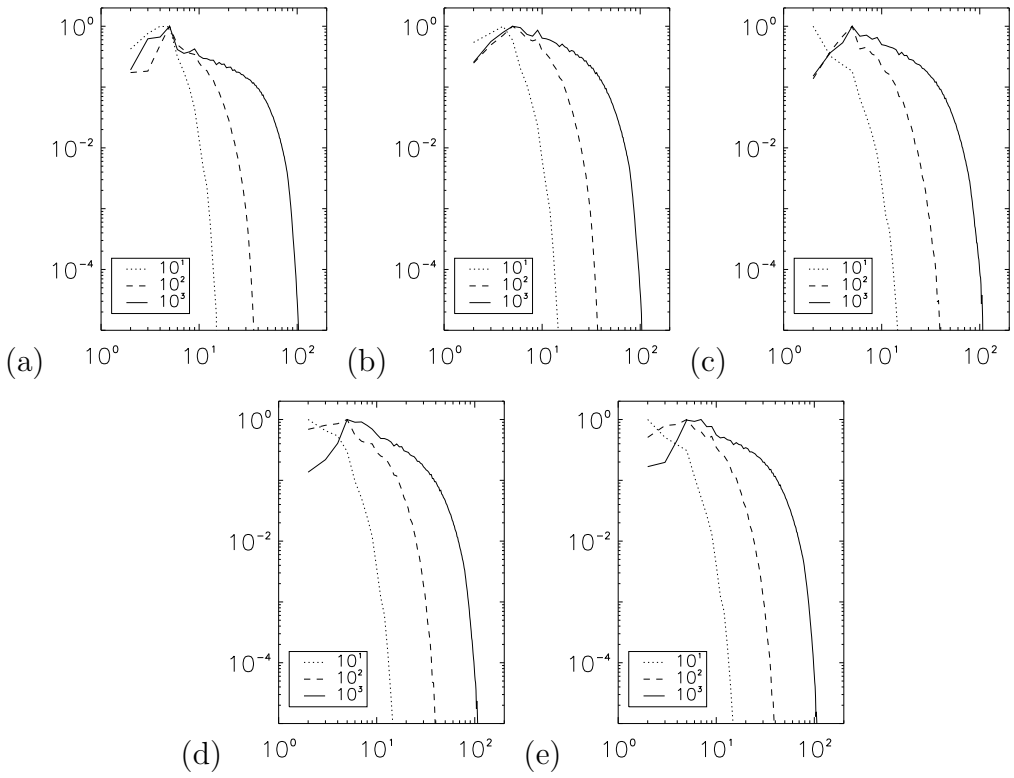

Figure 6.7: Energy spectra displaying normalised kinetic energy against wavenumber k (log-log) for classes (a) I to (e) V, at $R_e = 10^1, 10^2$ and 10^3.

of perturbation structures and eigenvalues of certain classes is mirrored by the spectra. Classes I and III show a clear peak at $k = 5$ for $R_e = 1000$, whereas classes IV and V both show a roughly equal energy distribution over $k = 5$ to $k = 8$. On the whole, as R_e is increased in all classes, there is little change in where the strongest fields are found, though the magnetic energy is clearly distributed over a wider range of scales. This contrasts with the magnetic problem for which

the peak magnetic energy is located at ever-smaller scales as R_m is increased and clearly scales as a function of R_m.

6.4 Summary

This brief investigation into the stability of the 1:1:1 ABC flow is made as a natural extension to the 1:1:1 ABC dynamo problem. The same symmetry classes and their respective irreducible representations, originally constructed for the magnetic problem, are also used to solve the problem of fluid stability in the (1:1:1) ABC flow. The results of Galloway and Frisch (1987) are reproduced using the symmetries up to $R_e = 160$, after which another mode, not indicated by Galloway and Frisch, becomes dominant. We confirm the emergence of a new dominant mode through a kink in the instability growth rates at around $R_e = 160$ with no symmetry assumptions, i.e. using a 'full' numerical code.

Despite the similarities of the two problems, the eigenvalues found are very different. In the magnetic problem, we have eigenmodes which persist over a large range of R_m; for example in class V, the dominant eigenmode does not change for $R_m \in [5, 10^4]$. Conversely, in the fluid stability problem, no eigenmode persists over such a large range of R_e and there are numerous mode crossings within each symmetry class. Also unlike the magnetic problem, the eigenmodes do not seem to exhibit falling growth rates and instead appear to have growth rates that tend towards asymptotic values.

Four of the five symmetry classes provide the dominant instability for $R_e \geq 10$, with these being III, V, I and II (in order of R_e first seen), and with crossings between them occurring at $R_e \approx 12, 160$ and 370. We note, however, that when class III's eigenmode is dominant, the flow is still stable. All five symmetry classes contain unstable flows with class V emerging first ($R_e \in [13, 14]$), followed by classes III ($[23, 24]$), IV ($[23, 24]$), II ($[26, 30]$) and finally class I ($R_e \in [49, 50]$). The dominant eigenmode at $R_e = 10^3$ is that of class II, with $\gamma = 0.293$, though unlike the magnetic problem, all symmetry classes appear to have very similar growth rates,

with the lowest at $R_e = 10^3$ being that of class III with $\gamma = 0.287$. The analysis in section 5.3.5 remains applicable to the fluid instability problem. Owing to the remarkably similar growth rates, however, we cannot link the possibility of a single axisymmetric cigar-like structure (at α-type stagnation points) with considerably higher instability growth rates, as we could in the magnetic problem. There is no reason to believe that mode crossings will not occur for $R_e > 10^3$, allowing the possibility for an eigenmode of another class to become dominant.

Through the profiles of $\mathrm{Im}(\lambda)$ obtained, we identify that no eigenvalue coalescences or decoalescences are seen for any of the five classes, as all growth rate kinks correspond to frequency discontinuities. In terms of frequencies, all classes and their respective eigenmodes are distinctly different. At $R_e = 10^3$, classes I and III have very similar frequencies, as do classes IV and V. Class II is distinctly different to the other classes, in that it sees falling frequencies as R_e increases. For $R_e = 10^3$, class II's eigenmode is steady in contrast to those of the other classes.

The instability structures vary in each class but again are related to either the chaotic regions associated with the stagnation points or the flow's principal vortices. In the magnetic problem, the field structures could easily be distinguished between those formed through the Lagrangian chaos (II, III, IV and V) and those formed in the principal vortices (I and III). This distinction is much less clear-cut for the fluid instabilities, with structures associated with both the chaotic regions and the principal vortices, in all classes except V. In the eigenmodes of classes II–V that persist until $R_e = 10^3$, we see the classic cigar-like structures of the magnetic problem, with varying numbers for each class. Class V sees a single cigar (per α-β pair), lying on the separatrix, whereas class II contains three cigar-like instabilities, organised (rotationally) symmetrically about the separatrix. Cigars of classes III and IV are accompanied by moderate instabilities in, or near to, the principal vortices. With the differences seen between structures of the fluid and magnetic problems (considering these problems are very similar in form), it may be interesting to carry out artificial calculations to isolate the importance of the pressure term as well as the the sign-change in the stretching term.

Overall, the symmetries do not seem to play as important a role in determining the stability of the 1:1:1 ABC flow. Of course, the instabilities of some symmetry classes initially grow faster than others but their respective unstable eigenmodes emerge within the same order of magnitude of R_e, unlike in the magnetic problem, and there is little difference, between the classes, in instability growth rate at $R_e = 10^3$. The symmetries remain highly apparent in the resulting instability configurations.

Chapter 7

Large-Scale fields in the ABC dynamo

The trouble with having an open mind, of course, is that people will insist on coming along and trying to put things in it.

Terry Pratchett

This chapter concerns the effects of increasing the period of the magnetic field relative to that of the flow. We study the growth rates of such large-scale fields and the structures of the resultant dominant fields. Rescaling the magnetic field is achieved through the introduction of Bloch waves that enforce a specific periodicity independently in each spatial direction. This method is advantageous as it requires only minor modification of the non-large-scale dynamo codes. We formulate the problem for a general Bloch wave, though numerically solve for a Bloch wave imposed in a single spatial direction. We draw comparison, both in methodology and results, to the investigation by Galanti *et al.* (1992) who propose an alternative definition of the ABC flow in order to achieve scale separation. Furthermore, using

multiple scale analysis, we derive alpha for our (Bloch wave) large-scale magnetic field driven by the 1:1:1 ABC flow.

The chapter begins by constructing the large-scale field problem through the use of Bloch waves and, using the symmetries of the 1:1:1 ABC flow, provides restrictions on the possible distinct Bloch waves. We then carry out a multiple scale analysis, which results in the expected contribution from the α-effect as a function of R_m and the Bloch wave scale. Subsequently, the Bloch wave formulation is compared to that of Galanti *et al.* (1992) to allow comparison of results and to demonstrate the effectiveness of our methods. The chapter moves on to numerical methods specific to this problem and then develops the theory of symmetries and representations for large-scale fields. We present a summary of the growth rates obtained from a multitude of simulations and provide some analysis of the resulting magnetic field structures. Briefly, we compare our results with those of Galanti *et al.* and with the expected contribution from the α-effect. Finally, the chapter is drawn to a close with a short summary of the main findings.

7.1 Formulation of the problem using Bloch waves

For this investigation, we adopt a method that modifies the scale of the magnetic field, leaving the scale of the flow unchanged. To achieve this, we define a large-scale magnetic field as

$$\boldsymbol{B} = \boldsymbol{b}\,\mathrm{e}^{\mathrm{i}\boldsymbol{K}\cdot\boldsymbol{x}}, \tag{7.1}$$

where $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$ and $\boldsymbol{b}(x, y, z, t)$ is periodic in x, y and z, each with period 2π. Here $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$ are Bloch wavenumbers, determining the scale of the magnetic field in each of the x, y and z directions respectively. We also have that $0 \leq |\mathcal{L}|, |\mathcal{M}|, |\mathcal{N}| \leq 0.5$ so that the scale of the magnetic field can be chosen arbitrarily. We are again solving the eigenvalue problem by setting

$$\boldsymbol{B} = \tilde{\boldsymbol{b}}(x, y, z)\,\mathrm{e}^{\mathrm{i}\boldsymbol{K}\cdot\boldsymbol{x}+\lambda t}. \tag{7.2}$$

Bloch waves are typically used in solid-state physics to describe wavefunctions of particles in periodic potentials, though they can be adopted for use in any situation in which a wave's periodicity needs to be strictly controlled. In MHD, their use is not particularly widespread, though studies including Childress (1970) and Roberts (1970) have adopted these approaches for studies on α-effect dynamos; the low frequency waves, however, are not introduced as Bloch waves.

On setting $\boldsymbol{k} = (l, m, n)$, the Fourier decomposition of the large-scale field (7.2) is given as

$$\begin{aligned}\boldsymbol{B} &= \mathrm{e}^{\lambda t} \sum_{l,m,n} \boldsymbol{b}_{l,m,n}\, \mathrm{e}^{\mathrm{i}(\boldsymbol{K}+\boldsymbol{k})\cdot\boldsymbol{x}} \\ &= \mathrm{e}^{\lambda t} \sum_{l,m,n} \boldsymbol{b}_{l,m,n}\, \mathrm{e}^{\mathrm{i}(\mathcal{L}+l)x+\mathrm{i}(\mathcal{M}+m)y+\mathrm{i}(\mathcal{N}+n)z}, \end{aligned} \tag{7.3}$$

which we can interpret as modifying the harmonics so that, instead of having wavenumber l, m and n in the x, y and z-directions respectively, they have wavenumbers $\mathcal{L}+l$, $\mathcal{M}+m$ and $\mathcal{N}+n$. Due to our choice of each $0 \leq |\mathcal{L}|, |\mathcal{M}|, |\mathcal{N}| \leq 0.5$, we are ensuring that the harmonics are no longer periodic over the interval $[0, 2\pi]^3$ but instead are periodic on the interval $[0, 2\pi/K_j]$ in each direction, assuming that $K_j \in \mathbb{Q}$. This allows certain periodicities to be imposed, so that only fields with the prescribed periodicity can exist in the problem for each choice of $\boldsymbol{K}$. However, we must be aware that the above definition means that the field $\boldsymbol{B}$ is complex-valued and thus the field's complex conjugate does not coincide with the field itself; we can no longer assume that Hermitian symmetry holds for general $\boldsymbol{K}$. As a quick verification, taking the complex conjugate of $\boldsymbol{B}$ in (7.3) results in

$$\boldsymbol{B}^* = \mathrm{e}^{\lambda^* t} \sum_{l,m,n} \boldsymbol{b}^*_{l,m,n}\, \mathrm{e}^{-\mathrm{i}(\boldsymbol{K}+\boldsymbol{k})\cdot\boldsymbol{x}} = \mathrm{e}^{\lambda^* t} \sum_{l,m,n} \boldsymbol{b}^*_{-l,-m,-n}\, \mathrm{e}^{\mathrm{i}(\boldsymbol{k}-\boldsymbol{K})\cdot\boldsymbol{x}}, \tag{7.4}$$

and shows that $\boldsymbol{B}^*$ corresponds to an eigenmode with Bloch wave $\mathrm{e}^{-\mathrm{i}\boldsymbol{K}\cdot\boldsymbol{x}}$ and not the original wave with wavenumber $\boldsymbol{K}$. The lack of Hermitian symmetry comes into play when considering the symmetries of such fields (see section 7.6).

We are solving the kinematic (linear) induction equation (1.13) with the large-scale field defined in (7.2) but we actually resolve the behaviour of $\tilde{\boldsymbol{b}}$ within $\boldsymbol{B}$. We obtain the evolution equation for $\tilde{\boldsymbol{b}}$ as

$$\lambda\tilde{\boldsymbol{b}} + [\boldsymbol{u}\cdot(\nabla + \mathrm{i}\boldsymbol{K})]\,\tilde{\boldsymbol{b}} = \tilde{\boldsymbol{b}}\cdot\nabla\boldsymbol{u} + \eta\,[(\nabla + \mathrm{i}\boldsymbol{K})\cdot(\nabla + \mathrm{i}\boldsymbol{K})]\,\tilde{\boldsymbol{b}}, \tag{7.5}$$

again with $\eta = R_m^{-1}$. We require that the field $\boldsymbol{B}$ satisfies $\nabla\cdot\boldsymbol{B} = 0$, and so the divergence-free condition becomes

$$(\nabla + \mathrm{i}\boldsymbol{K})\cdot\boldsymbol{b} = 0, \tag{7.6}$$

with $\nabla\cdot\boldsymbol{b} \neq 0$ in general. The induction equation (7.5) is solved through time stepping in an identical fashion to the original problem and is discussed in section 7.5. The intention is to calculate the eigenvalue, λ, of the fastest growing (or slowest decaying) eigenmode.

7.2 Restrictions on $\boldsymbol{K}$

There are three parameters to explore in this problem ($\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$) corresponding to the wavevector of the large-scale fields in each direction. To fully explore this parameter space, only a subset need actually be considered. This is addressed in a similar fashion to the parameters A, B and C of the ABC flows in Dombre *et al.* (1986), in which the symmetries of the system allow them to restrict the parameters to

$$1 = A \geq B \geq C \geq 0. \tag{7.7}$$

Before this is done for $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$, it is useful to note that the growth rate is an even function of $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$. One way to demonstrate this is by conjugating $\boldsymbol{B}$, as is done in (7.4) which we expand to give

$$\boldsymbol{B}^* = \mathrm{e}^{\lambda^* t}\sum_{l,m,n}\boldsymbol{b}^*_{l,m,n}\,\mathrm{e}^{-\mathrm{i}(l+\mathcal{L})x - \mathrm{i}(m+\mathcal{M}) - \mathrm{i}(n+\mathcal{N})}, \tag{7.8}$$

which means that we can identify that

$$\lambda(-\mathcal{L},-\mathcal{M},-\mathcal{N}) = \lambda^*(\mathcal{L},\mathcal{M},\mathcal{N}). \tag{7.9}$$

We remark that if $\boldsymbol{B}$ is a solution to the induction equation (7.5), then $\mathsf{g}\boldsymbol{B}$ will also be a solution for any $\mathsf{g} \in G$. This is important in further restricting the values of $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$.

It can be shown through use of the symmetries of the ABC flow that reversing one or all three of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$ will give us $\lambda^*(\mathcal{L},\mathcal{M},\mathcal{N})$ and reversing the sign of two of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$ will give us $\lambda(\mathcal{L},\mathcal{M},\mathcal{N})$. As an example, we take the symmetry b^2 and apply it to $\boldsymbol{B}$, giving

$$\mathsf{b}^2\boldsymbol{B} = \mathrm{e}^{\lambda t}\mathrm{e}^{\mathrm{i}(\mathcal{L}x-\mathcal{M}y-\mathcal{N}z)}\,\mathrm{e}^{\mathrm{i}\pi(\mathcal{L}+\mathcal{N})}\,\mathsf{b}^2\tilde{\boldsymbol{b}}. \tag{7.10}$$

The eigenvalue λ remains the same but we now have changed the sign of $\mathcal{M}$ and $\mathcal{N}$ in the Bloch wave, and so

$$\lambda(\mathcal{L},\mathcal{M},\mathcal{N}) = \lambda(\mathcal{L},-\mathcal{M},-\mathcal{N}). \tag{7.11}$$

This is possible for all transformations where two signs are reversed, without also switching positions of $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$. It is also possible to transform the conjugated field to obtain

$$\mathsf{b}^2\boldsymbol{B}^* = \mathrm{e}^{\lambda^* t}\mathrm{e}^{\mathrm{i}(-\mathcal{L}x+\mathcal{M}y+\mathcal{N}z)}\,\mathrm{e}^{-\mathrm{i}\pi(\mathcal{L}+\mathcal{N})}\,\mathsf{b}^2\tilde{\boldsymbol{b}}^*, \tag{7.12}$$

and so

$$\lambda(-\mathcal{L},-\mathcal{M},-\mathcal{N}) = \lambda(-\mathcal{L},\mathcal{M},\mathcal{N}). \tag{7.13}$$

In combination with (7.9), we have that

$$\begin{aligned}\lambda(\mathcal{L},\mathcal{M},\mathcal{N}) = {}& \lambda(\mathcal{L},-\mathcal{M},-\mathcal{N}) = \lambda(-\mathcal{L},\mathcal{M},-\mathcal{N}) = \lambda(-\mathcal{L},-\mathcal{M},\mathcal{N}) \\ = {}& \lambda^*(-\mathcal{L},-\mathcal{M},-\mathcal{N}) = \lambda^*(-\mathcal{L},\mathcal{M},\mathcal{N}) = \lambda^*(\mathcal{L},-\mathcal{M},\mathcal{N}) = \lambda^*(\mathcal{L},\mathcal{M},-\mathcal{N}).\end{aligned} \tag{7.14}$$

Of course, eigenvalues occur in complex conjugate pairs and so if λ is an eigenvalue of the problem for $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$, then so is λ^*. In practice, therefore, we only need to study one case of the eight above, this being $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$ for $0 \leq \mathcal{L}, \mathcal{M}, \mathcal{N}$. The upper bound on values of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$ of 0.5 is also necessary as for each of these parameters, we have a symmetry such that $\lambda(K_j) = \lambda^*(1-K_j)$. We show that this is the case by taking $\boldsymbol{K} = (1-\mathcal{L}, \mathcal{M}, \mathcal{N})$, which gives

$$\boldsymbol{B} = \mathrm{e}^{\lambda t} \sum_{l,m,n} \boldsymbol{b}_{l,m,n}\, \mathrm{e}^{\mathrm{i}(l+1-\mathcal{L})x + \mathrm{i}(m+\mathcal{M})y + \mathrm{i}(n+\mathcal{N})z} \tag{7.15}$$

$$= \mathrm{e}^{\lambda t} \sum_{l,m,n} \boldsymbol{b}_{l+1,m,n}\, \mathrm{e}^{\mathrm{i}(l-\mathcal{L})x + \mathrm{i}(m+\mathcal{M})y + \mathrm{i}(n+\mathcal{N})z}, \tag{7.16}$$

and so we end with the same situation at $\boldsymbol{K} = (-\mathcal{L}, \mathcal{M}, \mathcal{N})$ with an eigenvalue $\lambda(1-\mathcal{L}, \mathcal{M}, \mathcal{N}) = \lambda(-\mathcal{L}, \mathcal{M}, \mathcal{N}) = \lambda^*(\mathcal{L}, \mathcal{M}, \mathcal{N})$. The upshot of this is that for $\mathcal{L} > 0.5$, we obtain the same eigenvalues as for $\mathcal{L} < 0.5$, with reflective symmetry for $\mathrm{Re}(\lambda)$ through $\mathcal{L} = 0.5$ and reflective anti-symmetry for $\mathrm{Im}(\lambda)$. The same can be shown for each $\mathcal{M}$ and $\mathcal{N}$ and of course, if we have $\boldsymbol{K} = (1-\mathcal{L}, 1-\mathcal{M}, \mathcal{N})$, then this is equivalent to finding the eigenvalue $\lambda(-\mathcal{L}, -\mathcal{M}, \mathcal{N}) = \lambda(\mathcal{L}, \mathcal{M}, \mathcal{N})$. In the same way, it can be shown that $\lambda(1-\mathcal{L}, 1-\mathcal{M}, 1-\mathcal{N}) = \lambda^*(\mathcal{L}, \mathcal{M}, \mathcal{N})$. This provides the restrictions of $0 \leq \mathcal{L}, \mathcal{M}, \mathcal{N} \leq 0.5$ and also restricts the eigenvalue for $\mathcal{L} = 0.5$, and any permutation, to be real as λ and λ^* coincide.

If we consider a magnetic field for which a Bloch wave is imposed in only one direction, then we can choose any one of $\mathcal{L} \neq 0$, $\mathcal{M} \neq 0$ or $\mathcal{N} \neq 0$. Due to the threefold rotational symmetry around the $x = y = z$ axis (corresponding to transformations d and d^2), we can always rotate the flow so that the large-scale field is in the required direction. For example, if we choose $\boldsymbol{K} = (0, \mathcal{M}, 0)$ or $\boldsymbol{K} = (0, 0, \mathcal{N})$, we can rotate the field so that the Bloch wave is acting in the x-direction:

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{M}y}\boldsymbol{b}, \quad \mathsf{d}^2\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{M}x}\mathsf{d}^2\boldsymbol{b}, \tag{7.17}$$

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{N}z}\boldsymbol{b}, \quad \mathsf{d}\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{N}x}\mathsf{d}\boldsymbol{b}. \tag{7.18}$$

It is therefore only necessary to vary one of the three parameters without loss of generality and so the choice for $\boldsymbol{K} = (\mathcal{L}, 0, 0)$ seems suitable. With two non-zero parameters, the same applies and we can always rotate the flow using d and d^2 to obtain a large-scale field in the required direction. The choices $\boldsymbol{K} = (0, \mathcal{M}, \mathcal{N})$ and $\boldsymbol{K} = (\mathcal{L}, 0, \mathcal{N})$ can be rotated so that the Bloch wavenumbers are associated with only x and y:

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{M}y+\mathrm{i}\mathcal{N}z}\boldsymbol{b}, \quad \mathsf{d}^2\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{M}x+\mathrm{i}\mathcal{N}y}\boldsymbol{J}_{\mathsf{d}^2}\boldsymbol{b}, \tag{7.19}$$

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{L}x+\mathrm{i}\mathcal{N}z}\boldsymbol{b}, \quad \mathsf{d}\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{N}x+\mathrm{i}\mathcal{L}y}\boldsymbol{J}_{\mathsf{d}}\boldsymbol{b}, \tag{7.20}$$

thus we can choose $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, 0)$ without loss of generality. Further to this we are able to restrict $L \geq M$, by the transformation k:

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\mathcal{L}x+\mathrm{i}\mathcal{M}y}\,\boldsymbol{b} \quad \mathsf{k}\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\pi/2(\mathcal{L}-\mathcal{M})}\,\mathrm{e}^{\mathrm{i}\mathcal{M}x+\mathrm{i}\mathcal{L}y}\,\mathsf{k}\boldsymbol{b}, \tag{7.21}$$

which allows us to switch the Bloch wavenumbers $\mathcal{L}$ and $\mathcal{M}$ if $\mathcal{M} > \mathcal{L}$, without affecting the eigenvalue. A final additional restriction can be achieved, as through the symmetries l, m and n in combination with (7.9), $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$ can always be rearranged so that the indices are in descending order of size. For example, for l we have

$$\mathsf{l}\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\pi/2(-\mathcal{L}-\mathcal{M}-\mathcal{N})}\,\mathrm{e}^{\mathrm{i}(-\mathcal{L}x-\mathcal{N}y-\mathcal{M}z)}\,\mathsf{l}\boldsymbol{b}, \tag{7.22}$$

and thus determine that $\lambda(-\mathcal{L}, -\mathcal{N}, -\mathcal{M}) = \lambda(\mathcal{L}, \mathcal{M}, \mathcal{N}) = \lambda^*(\mathcal{L}, \mathcal{N}, \mathcal{M})$. The set of unique $\boldsymbol{K}$ is therefore defined as

$$0.5 \geq \mathcal{L} \geq \mathcal{M} \geq \mathcal{N} \geq 0, \tag{7.23}$$

with this reduced parameter-space not unlike the reduced set of modes in $\mathcal{H}_N$, defined in (4.36). We therefore only need to investigate a subset of all the possible values to find the optimal magnetic field scale for effective dynamo action.

Even after reducing the parameter-space to that given in (7.23), the number of simulations required becomes unmanageable for lower frequency Bloch waves, es-

pecially considering that these parameters are varied for each R_m studied. In choosing a case to study, we have to be realistic in regards to limiting factors such as time and computational resources. We make the choice to study the simplest case (in terms of large-scale field configuration), which is that of $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, being also the most interesting, as it is the most symmetric of all cases (see section 7.6).

7.3 Multiple-scale analysis and Bloch waves

In chapter 1, we briefly introduce the concept of electromotive force (e.m.f.) and note that it provides a method for small-scale flows to generate large-scale fields, referred to as the α-effect. In this section, we use scale separation to identify how α manifests in the growth rates obtained when solving (7.5). In a very similar fashion to Gilbert (2003), in which the G.O. Roberts flow (a specific case of the ABC flows, in which $A = B = 1$ and $C = 0$) is the subject of multiple-scale (asymptotic) analysis, we carry out this same analysis with a large-scale magnetic field imposed, in the form of a Bloch wave with wavenumber $\boldsymbol{K}$.

Let us define our magnetic field without Bloch waves as

$$\boldsymbol{b} = \left(\boldsymbol{b}_0 + \varepsilon\boldsymbol{b}_1 + \varepsilon^2\boldsymbol{b}_2 + \ldots\right) = \sum_{n=0}^{\infty} \varepsilon^n \boldsymbol{b}_n, \tag{7.24}$$

though, of course, we are solving this (as with the original problem) as an eigenvalue problem, which we do by defining the magnetic field as

$$\boldsymbol{b} = \mathrm{e}^{\lambda t}\tilde{\boldsymbol{b}}(x, y, z) = \mathrm{e}^{\nu(\lambda_0 + \varepsilon\lambda_1 + \varepsilon^2\lambda_2 + \ldots)t}\left(\tilde{\boldsymbol{b}}_0 + \varepsilon\tilde{\boldsymbol{b}}_1 + \varepsilon^2\tilde{\boldsymbol{b}}_2 + \ldots\right). \tag{7.25}$$

This notation, however, becomes problematic when linking these ideas to those of Moffatt (1978). The tilde identifying that the fields are only spatially-dependent will be dropped, and so $\tilde{\boldsymbol{b}}_i$ will be identified instead as $\boldsymbol{b}_i$. Therefore, each $\boldsymbol{b}_i$ considered from here onwards is purely a function of $\boldsymbol{x}$, with the time-dependence

extracted in the form of exponential growth or decay represented by the eigenvalues λ_i.

We define the Bloch waves as

$$\mathrm{e}^{\mathrm{i}\kappa \boldsymbol{K}\cdot\boldsymbol{x}} = \mathrm{e}^{\mathrm{i}\kappa\mathcal{L}x+\mathrm{i}\kappa\mathcal{M}y+\mathrm{i}\kappa\mathcal{N}z}, \tag{7.26}$$

so that

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\kappa\mathcal{L}x+\mathrm{i}\kappa\mathcal{M}y+\mathrm{i}\kappa\mathcal{N}z+\nu\left(\lambda_0+\varepsilon\lambda_1+\varepsilon^2\lambda_2+\ldots\right)t}\left(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\varepsilon^2\boldsymbol{b}_2+\ldots\right). \tag{7.27}$$

The extra terms κ and ν are in place to control the scales of the associated parameters ($\boldsymbol{K}$ and λ_i respectively) by appropriate choices. We consider a field that evolves on the scale of $|\boldsymbol{K}| = \mathcal{O}(R_m^2)$ over short time scales $t = \mathcal{O}(R_m^{-1})$. This approximation is suitable for very small scales of R_m, say $R_m = \varepsilon \ll 1$ and so appropriate choices would be for $\kappa = \varepsilon^2$ and $\nu = \varepsilon^{-1}$. With these choices, we have

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\varepsilon^2\mathcal{L}x+\mathrm{i}\varepsilon^2\mathcal{M}y+\mathrm{i}\varepsilon^2\mathcal{N}z+\left(\lambda_0+\varepsilon\lambda_1+\varepsilon^2\lambda_2+\ldots\right)t/\varepsilon}\left(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\varepsilon^2\boldsymbol{b}_2+\ldots\right). \tag{7.28}$$

We can derive expressions for the growth rates λ_i, though it is a fairly lengthy affair and so is summarised below. Splitting the field $\boldsymbol{B}$ into scalar and vector parts, we let

$$\varphi = \mathrm{e}^{\mathrm{i}\varepsilon^2\mathcal{L}x+\mathrm{i}\varepsilon^2\mathcal{M}y+\mathrm{i}\varepsilon^2\mathcal{N}z+\left(\lambda_0+\varepsilon\lambda_1+\varepsilon^2\lambda_2+\ldots\right)t/\varepsilon}, \tag{7.29}$$

and

$$\boldsymbol{\Psi} = \boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\varepsilon^2\boldsymbol{b}_2+\ldots. \tag{7.30}$$

From the definitions above, it follows that

$$\partial_t\varphi = \varepsilon^{-1}(\lambda_0+\varepsilon\lambda_1+\varepsilon^2\lambda_2+\ldots)\varphi, \tag{7.31}$$

$$\nabla\varphi = \mathrm{i}\varepsilon^2\boldsymbol{K}\varphi, \tag{7.32}$$

$$\nabla^2\varphi = -\varepsilon^4K^2\varphi, \tag{7.33}$$

with $K = |\boldsymbol{K}|$. Now substituting (7.29) and (7.30) into the induction equation (7.5) gives us

$$\boldsymbol{\Psi}\partial_t\varphi = \varphi\nabla\times(\boldsymbol{u}\times\boldsymbol{\Psi}) + \nabla\varphi\times(\boldsymbol{u}\times\boldsymbol{\Psi}) + \varepsilon^{-1}\left(\boldsymbol{\Psi}\nabla^2\varphi + 2(\nabla\boldsymbol{\Psi})\cdot(\nabla\varphi) + \varphi\nabla^2\boldsymbol{\Psi}\right), \tag{7.34}$$

and using the results in (7.31) to (7.33) we have

$$\begin{aligned}
&\varepsilon^{-1}(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots)(\lambda_0+\varepsilon\lambda_1+\varepsilon^2\lambda_2+\ldots)\varphi = \\
&\quad\varphi\nabla\times(\boldsymbol{u}\times(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots)) + \mathrm{i}\varepsilon^2\boldsymbol{K}\times(\boldsymbol{u}\times(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots))\varphi + \\
&\quad\varepsilon^{-1}\left[-\varepsilon^4K^2(\boldsymbol{u}\times(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots)) + 2\mathrm{i}\varepsilon^2(\boldsymbol{K}\cdot\nabla)(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots) + \right. \\
&\quad\left.\nabla^2(\boldsymbol{b}_0+\varepsilon\boldsymbol{b}_1+\ldots)\right]\varphi.
\end{aligned} \tag{7.35}$$

We identify different scales, firstly by multiplying through by ε so that the lowest power is zero and then arranging by powers of ε to get

$$\varepsilon^0: \quad \lambda_0\boldsymbol{b}_0 = \nabla^2\boldsymbol{b}_0, \tag{7.36}$$

$$\varepsilon^1: \quad \lambda_1\boldsymbol{b}_0 + \lambda_0\boldsymbol{b}_1 = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_0) + \nabla^2\boldsymbol{b}_1, \tag{7.37}$$

$$\varepsilon^2: \quad \lambda_2\boldsymbol{b}_0 + \lambda_1\boldsymbol{b}_1 + \lambda_0\boldsymbol{b}_2 = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_1) + 2\mathrm{i}(\boldsymbol{K}\cdot\nabla)\boldsymbol{b}_0 + \nabla^2\boldsymbol{b}_2, \tag{7.38}$$

$$\varepsilon^3: \quad \lambda_3\boldsymbol{b}_0 + \lambda_2\boldsymbol{b}_1 + \lambda_1\boldsymbol{b}_2 + \lambda_0\boldsymbol{b}_3 = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_2) + \mathrm{i}\boldsymbol{K}\times(\boldsymbol{u}\times\boldsymbol{b}_0) + 2\mathrm{i}(\boldsymbol{K}\cdot\nabla)\boldsymbol{b}_1 + \nabla^2\boldsymbol{b}_3, \tag{7.39}$$

$$\varepsilon^4: \quad \lambda_4\boldsymbol{b}_0 + \lambda_3\boldsymbol{b}_1 + \lambda_2\boldsymbol{b}_2 + \lambda_1\boldsymbol{b}_3 + \lambda_0\boldsymbol{b}_4 = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_3) + \mathrm{i}\boldsymbol{K}\times(\boldsymbol{u}\times\boldsymbol{b}_1) - K^2\boldsymbol{b}_0 + 2\mathrm{i}(\boldsymbol{K}\cdot\nabla)\boldsymbol{b}_2 + \nabla^2\boldsymbol{b}_4. \tag{7.40}$$

For $n \geq 4$, we then have

$$\varepsilon^n: \sum_{j+k=n}(\lambda_j\boldsymbol{b}_k) = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_{n-1}) + i\boldsymbol{K}\times(\boldsymbol{u}\times\boldsymbol{b}_{n-3}) - K^2\boldsymbol{b}_{n-4} + 2\mathrm{i}(\boldsymbol{K}\cdot\nabla)\boldsymbol{b}_{n-2} + \nabla^2\boldsymbol{b}_n. \tag{7.41}$$

To reduce the problem we use averaging, in a very similar sense to Moffatt (1978), though we define our method of averaging with respect to harmonics of $\boldsymbol{b}$ and $\boldsymbol{u}$. Because we are working in spectral space, all magnetic fields and flows are defined

as a sum of their respective harmonics $\hat{\boldsymbol{b}}_{\boldsymbol{k}}$ and $\hat{\boldsymbol{u}}_{\boldsymbol{k}}$ for $\boldsymbol{k} = (l, m, n)$. The use of caret ($\hat{}$) notation is to avoid confusion between fields defined in spectral space and real space; this convention is adopted in this section only. These fields are defined modulo 2π, so we define our averaging over this 2π period, thus have that

$$\langle \boldsymbol{b}_i \rangle = (\hat{\boldsymbol{b}}_i)_{0,0,0}, \tag{7.42}$$

as it is the case, for the $\boldsymbol{k}^{\text{th}}$ harmonic, that

$$\langle (\hat{\boldsymbol{b}}_i)_{\boldsymbol{k}} \rangle = \boldsymbol{0}, \quad \forall \boldsymbol{k} \neq \boldsymbol{0}. \tag{7.43}$$

Since our velocity field is purely fluctuating at its periodic scale, and contains no mean field, we know that

$$\langle \boldsymbol{u} \rangle = \hat{\boldsymbol{u}}_{0,0,0} = \boldsymbol{0}. \tag{7.44}$$

It also needs to be noted that terms involving a derivative of $\boldsymbol{b}_i$ will have zero mean. We can show that under averaging, the three terms in (7.41) involving derivatives will always have a zero mean. The Laplacian gives

$$\nabla^2 \boldsymbol{b}_i = \sum_{\boldsymbol{k}} k^2 \left(\hat{\boldsymbol{b}}_i \right)_{\boldsymbol{k}} \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}, \tag{7.45}$$

which, after averaging gives us

$$\langle \nabla^2 \boldsymbol{b}_i \rangle = 0 (\hat{\boldsymbol{b}}_i)_{0,0,0} = \boldsymbol{0}. \tag{7.46}$$

Similarly, the term $(\boldsymbol{K} \cdot \nabla)\boldsymbol{b}_i$ can be shown to have a zero average for all fields $\boldsymbol{b}_i$. We know that the derivatives ∂_x, ∂_y and ∂_z when acting on $\boldsymbol{u}$ or $\boldsymbol{b}$ produce coefficients l, m and n respectively, for the harmonics. The term can therefore be written as a sum of its harmonics by

$$(\boldsymbol{K} \cdot \nabla)\boldsymbol{b}_i = \sum_{\boldsymbol{k}} (\boldsymbol{K} \cdot \boldsymbol{k}) \, (\hat{\boldsymbol{b}}_i)_{\boldsymbol{k}} \, \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}, \tag{7.47}$$

and so averaging gives

$$\langle(\boldsymbol{K}\cdot\nabla)\boldsymbol{b}_i\rangle = (\boldsymbol{K}\cdot\boldsymbol{0})(\hat{\boldsymbol{b}}_i)_{0,0,0} = \boldsymbol{0}. \tag{7.48}$$

The same process applies to the term $\nabla\times(\boldsymbol{u}\times\boldsymbol{b}_i)$ and is written as a sum of its harmonics by

$$\nabla\times(\boldsymbol{u}\times\boldsymbol{b}_i) = \sum_{\boldsymbol{k}} \boldsymbol{k}\times\left[\hat{\boldsymbol{u}}_{\boldsymbol{k}}\times\left(\hat{\boldsymbol{b}}_i\right)_{\boldsymbol{k}}\right] \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}, \tag{7.49}$$

which again averages to $\boldsymbol{0}$, due to the factor of $\boldsymbol{k}$ operating on each harmonic.

Now the results from (7.46)–(7.49) can be used when averaging the equations for each power of ε. Taking the average of (7.36), we see two possibilities. The first is that $\lambda_0 = 0$ and the second is that $\langle\boldsymbol{b}_0\rangle = \boldsymbol{0}$. Of course, there is no necessity for $\langle\boldsymbol{b}_0\rangle$ to be zero so it must be that $\lambda_0 = 0$. Using this result, we conclude that

$$\nabla^2\boldsymbol{b}_0 = \boldsymbol{0}, \tag{7.50}$$

which implies that $\langle\boldsymbol{b}_0\rangle = \boldsymbol{b}_0$; $\boldsymbol{b}_0$ is purely a mean field (contains no fluctuating part), which is taken to be non-zero.

In (7.37), we have terms involving $\boldsymbol{b}_0$ and $\boldsymbol{b}_1$. Averaging gives us

$$\lambda_1\langle\boldsymbol{b}_0\rangle + \lambda_0\langle\boldsymbol{b}_1\rangle = \langle\nabla\times(\boldsymbol{u}\times\boldsymbol{b}_0)\rangle + \langle\nabla^2\boldsymbol{b}_1\rangle = \boldsymbol{0}, \tag{7.51}$$

and since $\lambda_0 = 0$ and $\boldsymbol{b}_0$ does not have a zero average (as it is constant), then we conclude that $\lambda_1 = 0$ also. Through this we see that

$$-\nabla^2\boldsymbol{b}_1 = \nabla\times(\boldsymbol{u}\times\boldsymbol{b}_0), \tag{7.52}$$

and when we consider that $\boldsymbol{b}_0$ is constant and $\boldsymbol{u}$ is non-divergent, we can simplify to

$$-\nabla^2\boldsymbol{b}_1 = (\boldsymbol{b}_0\cdot\nabla)\boldsymbol{u}, \tag{7.53}$$

which we can then resolve, for each harmonic, as

$$(\boldsymbol{b}_1)_{\boldsymbol{k}} = \frac{(\boldsymbol{b}_0 \cdot \mathrm{i}\boldsymbol{k})\hat{\boldsymbol{u}}_{\boldsymbol{k}}}{k^2}, \tag{7.54}$$

so we find that $\boldsymbol{b}_1$ is simply a function of $\boldsymbol{b}_0$ (constant) and $\boldsymbol{u}$ (a purely fluctuating, zero-mean flow), therefore $\boldsymbol{b}_1$ must be purely fluctuating and contains no mean field, that is $\langle \boldsymbol{b}_1 \rangle = \boldsymbol{0}$.

In (7.38), all terms of the right hand side will have zero average. Knowing that $\lambda_0 = \lambda_1 = 0$, we deduce that $\lambda_2 = 0$. Similarly, in (7.39), all terms on the right-hand side average to zero, including $\mathrm{i}\boldsymbol{K} \times (\boldsymbol{u} \times \boldsymbol{b}_0)$ as a consequence of (7.44) and $\boldsymbol{b}_0$ being constant. This gives that $\lambda_3 = 0$ also.

It is in (7.40) that the first non-zero terms emerge after averaging and we are left with

$$\lambda_4 \langle \boldsymbol{b}_0 \rangle = \mathrm{i}\boldsymbol{K} \times \langle \boldsymbol{u} \times \boldsymbol{b}_1 \rangle - K^2 \langle \boldsymbol{b}_0 \rangle. \tag{7.55}$$

The term $(\boldsymbol{u} \times \boldsymbol{b}_1)$ is not zero when averaged, due to taking the product of $\boldsymbol{u}$ and a purely fluctuating field $\boldsymbol{b}_1$ which produces a fluctuating field with a non-zero mean field component; we resolve the term $\langle \boldsymbol{u} \times \boldsymbol{b}_1 \rangle$ using the result from (7.54). It follows that

$$\boldsymbol{u} \times \boldsymbol{b}_1 = \sum_{\boldsymbol{k},\boldsymbol{k}'} \frac{\hat{\boldsymbol{u}}_{\boldsymbol{k}'} \times (\boldsymbol{b}_0 \cdot \mathrm{i}\boldsymbol{k})\hat{\boldsymbol{u}}_k}{k^2} \, \mathrm{e}^{\mathrm{i}(\boldsymbol{k}+\boldsymbol{k}')\cdot\boldsymbol{x}}, \tag{7.56}$$

with $\boldsymbol{k}' = (l', m', n')$ identifying wavenumbers distinct from $\boldsymbol{k} = (l, m, n)$. When averaging, the only non-zero harmonics are those of wavenumber $\boldsymbol{k}' + \boldsymbol{k} = \boldsymbol{0}$ and so $\boldsymbol{k}' = -\boldsymbol{k}$. Taking the average of (7.56), we therefore obtain

$$\langle \boldsymbol{u} \times \boldsymbol{b}_1 \rangle = \sum_{\boldsymbol{k}} \frac{\hat{\boldsymbol{u}}^*_{\boldsymbol{k}} \times \hat{\boldsymbol{u}}_{\boldsymbol{k}}}{k^2} (\boldsymbol{b}_0 \cdot i\boldsymbol{k}), \tag{7.57}$$

which uses the fact that $\boldsymbol{u}$ is real and so $\hat{\boldsymbol{u}}^*_{\boldsymbol{k}} = \hat{\boldsymbol{u}}_{-\boldsymbol{k}}$.

We note that (7.57) is identical to the e.m.f. of Moffatt (1978) for a flow with no

time-dependence. Therefore

$$\langle \boldsymbol{u} \times \boldsymbol{b}_1 \rangle_i = \alpha_{ij} \boldsymbol{b}_{0j}, \tag{7.58}$$

and for a rotationally invariant flow, such as the 1:1:1 ABC flow, α_{ij} is isotropic and takes the form

$$\alpha_{ij} = \alpha \delta_{ij}, \tag{7.59}$$

where α is a scalar constant. Taking the definitions of (7.58) and (7.59), we return to the equation for λ_4 (7.55) and rewrite it as

$$\lambda_4 \langle \boldsymbol{b}_0 \rangle = \mathrm{i}\boldsymbol{K} \times \alpha \boldsymbol{b}_0 - K^2 \langle \boldsymbol{b}_0 \rangle. \tag{7.60}$$

Recalling the original large-scale field definition (7.28), we write (7.60) in dimensional units with

$$\varepsilon^{-3} \lambda \langle \boldsymbol{b}_0 \rangle = \mathrm{i}\varepsilon^{-2} \boldsymbol{K} \times \alpha \boldsymbol{b}_0 - \varepsilon^{-4} K^2 \langle \boldsymbol{b}_0 \rangle, \tag{7.61}$$

and remembering that we set $R_m = \varepsilon$, we finally obtain

$$\lambda \boldsymbol{b}_0 = \mathrm{i} R_m \boldsymbol{K} \times \alpha \boldsymbol{b}_0 - R_m^{-1} K^2 \boldsymbol{b}_0. \tag{7.62}$$

For an individual Bloch wave, say $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, we have $\boldsymbol{b}_0 = (0, 1, \pm\mathrm{i})$ without loss of generality, with (7.62) becoming

$$\lambda = \pm \alpha R_m \mathcal{L} - R_m^{-1} \mathcal{L}^2, \tag{7.63}$$

as $\mathrm{i}(\mathcal{L}, 0, 0) \times \boldsymbol{b}_0 = \pm \mathcal{L} \boldsymbol{b}_0$ for the helical waves $\boldsymbol{b}_0 = (0, 1, \pm\mathrm{i})$. Of course, for the 1:1:1 ABC flow we can analytically resolve α from (7.57), with each wave of $\boldsymbol{u}$ contributing an element of α_{ij}. We find that $\alpha = -1$, thus

$$\lambda = \mp R_m \mathcal{L} - R_m^{-1} \mathcal{L}^2. \tag{7.64}$$

For $R_m \ll 1$, and under the assumption that the contribution from fluctuating

(non-large-scale) field is negligible, we expect to see growth rates with the profile given above. In theory, this holds only for $R_m \ll 1$ and in section 7.7, we aim to test whether the growth rate of the dominant eigenmode is described accurately by (7.64), and for what range of R_m this holds. From the form of (7.64), we expect to see only a narrow window of $\mathcal{L}$ in which the fields grow, with this window becoming smaller for $R_m \to 0$. For very small R_m, only the very largest-scale fields (i.e. those for $\mathcal{L} \to 0$) will see growth.

7.4 A comparison with Galanti *et al.* (1992)

In the introduction, an alternative definition of the ABC flows, originally given by Galanti *et al.* (1992), allows for the periodicity of the flow to be modified. The definition involves introducing a term $k_0 \in \mathbb{N}$ that we use to rescale x, y and z, and is displayed as

$$\boldsymbol{u} = (C \sin k_0 z + B \cos k_0 y, A \sin k_0 x + C \cos k_0 z, B \sin k_0 y + A \cos k_0 x), \quad (7.65)$$

which is modified to fit with our own definition of the ABC flow, in that A, B and C are associated (respectively) with x, y and z. For $k_0 = 1$, we have the original 2π-periodic ABC flows. For $k_0 \geq 2$ the flow's period is reduced to $2\pi/k_0$. In solving the induction equation with the flow (7.65), the parameter k_0 effectively controls the separation between the scale of the flow and the largest scale of the magnetic field, which is simply the scale of the domain $[0, 2\pi]^3$. This method of rescaling the flow with respect to the magnetic field is conceptually simpler than the fluctuating and mean components of Moffatt (1978), and differs in that there is no true mean field, rather it is periodic on a much greater scale than the flow (for large k_0). To investigate the 1:1:1 ABC flow, Galanti *et al.* set $A = B = C = k_0$ to give

$$\boldsymbol{u} = k_0(\sin k_0 z + \cos k_0 y, \sin k_0 x + \cos k_0 z, \sin k_0 y + \cos k_0 x), \quad (7.66)$$

so that R_m remains relative to the scale of the flow. This approach allows the authors to find the most effective scale separation for dynamo growth by investigating a range of scales k_0 over a range of R_m. However, there are two distinct disadvantages of the definition (7.65). Firstly, resolving the magnetic field in the periodic domain $[0, 2\pi]^3$ means that there are effectively k_0^3 copies of the flow (7.66), with each copy needing to be well-resolved. Thus the resolution required scales with k_0; if the period of the flow is halved (by doubling k_0), the number of grid points required to properly resolve the flow increases two-fold in each direction, or eight-fold overall, making the problem massively computationally demanding for high k_0. Secondly, k_0 must be a positive integer, which severely restricts the magnetic field scales that can be studied, as the magnetic field's frequency can only be $1/2, 1/3, 1/4, \ldots$ of the frequency of the flow. From the results of Galanti *et al.* (1992), it appears that the growth rate peaks in the range $2 < k_0 < 3$ at $R_m = 12$ and 20, near to (but not necessarily at) $k_0 = 2$.

The disadvantages of the definition (7.65) are overcome in our method of rescaling the magnetic field through the introduction of Bloch waves. It is advantageous to relate these two very different methods so that we have a basis for comparing results. The first step is to take the induction equation (1.13), with the original ABC flow (1.18), and rescale the variables to make it analogous to the problem that Galanti *et al.* are solving. We do this by rescaling $\boldsymbol{x}$ by setting $\boldsymbol{x} = k_0\boldsymbol{x}'$, which is the obvious choice given the form of the flow in (1.25). We also rescale the differential operator ∇, so that the derivatives are with respect to the rescaled spatial variable $\boldsymbol{x}'$. This produces $\nabla = (1/k_0)\nabla'$ as

$$\frac{\partial}{\partial x} = \frac{\partial}{\partial x'}\frac{\partial x'}{\partial x} = \frac{\partial}{\partial x'}\left(\frac{1}{k_0}\right). \tag{7.67}$$

Other variables also need to be rescaled so we can identify arbitrary scalings, for which we make appropriate choices once we have more information. We let

$$t = a_1 t' \quad R_m = a_2 R_m' \quad \boldsymbol{u} = a_3 \boldsymbol{u}', \tag{7.68}$$

though we need not rescale $\boldsymbol{b}$ as this occurs linearly in all terms of the induction equation. Expressing the induction equation in terms of transformed variables, we have

$$\frac{1}{a_1}\partial_{t'}\boldsymbol{b}' = \frac{a_3}{k_0}\nabla' \times (\boldsymbol{u}' \times \boldsymbol{b}') + \frac{1}{a_2 R_m'}\frac{1}{k_0^2}\nabla'^2\boldsymbol{b}'. \tag{7.69}$$

To make appropriate choices for a_1, a_2 and a_3, we balance the k_0 terms in (7.69) to obtain

$$\frac{a_3}{k_0} = \frac{1}{a_2 k_0^2}, \quad \frac{1}{a_1} = \frac{a_3}{k_0}. \tag{7.70}$$

We have three scaling parameters and two equations, so there is an element of arbitrariness in the choices made. Ideally, for the most useful comparison, we want $R_m' = R_m$ so we set $a_2 = 1$. This solidifies the other parameters giving $a_3 = k_0^{-1}$ and $a_1 = k_0^2$. We can immediately understand the choice of Galanti *et al.* to set $A = B = C = k_0$, as this incorporates the rescaling into the amplitude of the flow itself. It is also apparent that the timescale on which their magnetic fields evolve is much smaller, with $t' = k_0^{-2}t$. To compare the eigenvalues obtained from the original eigenproblem to those of Galanti *et al.*, we require them to be on the same timescale and since we know how the relative temporal variables scale, we know how to scale the eigenvalues. If we call the eigenvalue obtained by Galanti *et al.* λ', then we have that

$$\lambda = k_0^2\lambda', \tag{7.71}$$

with λ being the eigenvalue from the original (unscaled) eigenproblem. Galanti *et al.* have rescaled all of their growth rates $\gamma = \mathrm{Re}(\lambda)$ by k_0^2 for ease of comparison with standard ABC dynamo results.

The next step is to relate rescaling of the magnetic field by Bloch waves to rescaling the flow through multiplying the wavenumbers by a factor of k_0. One way to achieve this is to impose Bloch waves on the magnetic field of Galanti *et al.* and ask what scale is needed for the large scale field to counteract the rescaling of the flow, through making an appropriate choice for $\boldsymbol{K}$. In the modified coordinates $\boldsymbol{x}'$, a

large-scale magnetic field would be expressed as

$$\boldsymbol{B}' = \mathrm{e}^{\mathrm{i}k_0 \boldsymbol{K}\cdot\boldsymbol{x}'}\boldsymbol{b}, \tag{7.72}$$

and so to restore the field to its 2π-periodic state, we can simply let $\boldsymbol{K} = (n_1/k_0, n_2/k_0, n_3/k_0)$, with $n_i \in \mathbb{N}$. This is not particularly surprising and we can give an example to clarify: if we let $k_0 = 3$, this means that we would be solving for a magnetic field that is periodic on 3 times the scale (in each spatial direction) as that of the flow. In this case, we could set $\boldsymbol{K} = (1/3, 1/3, 1/3)$ which in our problem corresponds to a magnetic field that is 6π-periodic. Since our flow remains 2π-periodic, our magnetic field is therefore on a scale 3 times that of the flow.

Before attempting a direct comparison, however, we need to be aware of the numerical differences of the two methods. Figure 7.1 attempts to describe this, though some explanation is also required. In the method of Galanti *et al.*, which we will

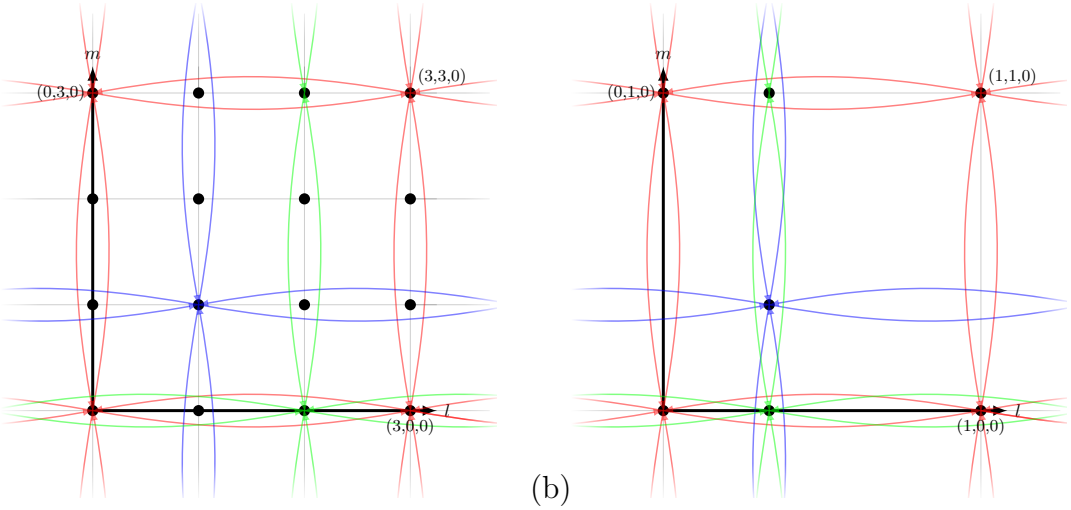

(a) (b)

Figure 7.1: Fourier modes $\boldsymbol{b}_{l,m,0}$ showing 'interactions' in (a) three of 27 independent eigenmodes of GPS, with $k_0 = 3$ and (b) three Bloch waves with $\boldsymbol{K} = \boldsymbol{0}, (1/3, 0, 0)$ and $(1/3, 1/3, 0)$. Note the different mode indices in each scheme.

abbreviate GPS, the Fourier modes of the magnetic field are no longer time stepped using their immediate neighbours but instead modes that are of distance k_0 in each index. In (4.9)–(4.11), the indices $l \pm 1$, $m \pm 1$ and $n \pm 1$ are replaced by $l \pm k_0$, $m \pm k_0$ and $n \pm k_0$ respectively. In figure 7.1(a), for example, $k_0 = 3$ and so mode

$\boldsymbol{b}_{0,0,0}$ depends on $\boldsymbol{b}_{3,0,0}, \boldsymbol{b}_{-3,0,0}, \boldsymbol{b}_{0,3,0}$, etc. This does not mean, however, that only modes of the form $\boldsymbol{b}_{nk_0,0,0}$ and permutations contain values and the rest are zero. Similarly, we have modes $\boldsymbol{b}_{2,0,0}$ communicating with $\boldsymbol{b}_{5,0,0}, \boldsymbol{b}_{-1,0,0}, \boldsymbol{b}_{2,3,0}$, etc. These two sets of modes are completely numerically independent and so will never interact. In the figure, three of these sets of modes are identified but for $k_0 = 3$, there are in fact 27 independent eigenmodes (without involving symmetries), though it is only possible to see up to nine in the figure's cross-section. A given simulation will resolve all 27 of these eigenmodes simultaneously, though the growth rate of only the most dominant can be found.

Using Bloch waves, each magnetic field mode still depends only on its immediate neighbours. Three eigenmodes are shown in figure 7.1(b): $\boldsymbol{K} = \boldsymbol{0}$ in red, $\boldsymbol{K} = (1/3, 0, 0)$ in green and $\boldsymbol{K} = (1/3, 1/3, 0)$ in blue. Like GPS, the eigenmodes for different Bloch waves are independent from one another. Unlike GPS, however, these eigenmodes are resolved in separate simulations by varying $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$, and as the periodicity of the magnetic field is controlled by the Bloch wave, the eigenvalue obtained is that of the particular Bloch wave imposed; no other periodicities are present. We recall that the space of $\boldsymbol{K}$ that requires investigating has been reduced and so we need not simulate for $\mathcal{L}, \mathcal{M}$ or $\mathcal{N}$ equal to 2/3, as each is symmetric about 0.5. In addition $\boldsymbol{K} = (0, 1/3, 0)$ is not shown in the figure, as it is possible to rotate the Bloch wave to $\boldsymbol{K} = (1/3, 0, 0)$, which is displayed. One further eigenmode is present, $\boldsymbol{K} = (1/3, 1/3, 1/3)$, though cannot be displayed in the cross-section. This leaves three simulations, one for each Bloch wave, to produce equivalent results to GPS, each with the same computational demand as the non-large-scale problem. In contrast, time stepping in GPS requires carrying out k_0^3 times as many computations as the $k_0 = 1$ problem. It should be noted that many of the independent eigenmodes resolved by GPS will evolve in an identical fashion, equivalent to the way that Bloch waves $\boldsymbol{K} = (1/3, 0, 0), (0, 1/3, 0)$ and $(0, 0, 1/3)$ are identical. It is not possible to avoid simulating these duplicate eigenmodes through GPS.

The eigenvalue obtained via GPS will therefore be the eigenvalue with the largest

real part of all the individual Bloch wave simulations, as the magnetic energy of GPS will grow as fast as the largest eigenvalue permits. This highlights another advantage of the Bloch wave approach, in that we are able to isolate the optimal (fastest growing) magnetic field scaling relative to the flow, without fear of other periodicities 'polluting' the problem.

For $k_0 > 3$, the comparison between growth rates is a little more complicated. In (7.72), we can set $\boldsymbol{K} = (n_1/k_0, n_2/k_0, n_3/k_0)$, with $n_i \in \mathbb{N}$ and for $k_0 = 4$, this allows for $n_i = 0, 1$ or 2 without overstepping the restrictions placed on $\boldsymbol{K}$. Therefore to achieve the same results as GPS, we would need to impose nine different Bloch waves: $\boldsymbol{K} = (1/4, 0, 0)$, $(1/4, 1/4, 0)$, $(1/4, 1/4, 1/4)$, $(1/2, 0, 0)$, $(1/2, 1/4, 0)$, $(1/2, 1/4, 1/4)$ $(1/2, 1/2, 0)$, $(1/2, 1/2, 1/4)$ and $(1/2, 1/2, 1/2)$. This still represents a huge saving, owing to us carrying out nine simulations at the same cost of the original problem, as opposed to one simulation at 64 times the cost of the original, as with GPS.

Overall, this means that we can only use the growth rates of GPS as an upper bound for the Bloch wave growth rates. Formally, the (Bloch) eigenvalue is a function of R_m and $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$, so $\lambda = \lambda(\mathcal{L}, \mathcal{M}, \mathcal{N}, R_m)$, whereas for GPS λ' is a function of R_m and k_0. Comparing growth rates only, we have

$$k_0^2 \, \gamma_{\max}(R_m) = \gamma'(R_m, k_0), \tag{7.73}$$

with

$$\gamma_{\max} = \max_{n_1, n_2, n_3} \{\gamma \left(R_m, n_1/k_0, n_2/k_0, n_3/k_0\right) \,|\, n_j \in \mathbb{N}, 0 \leq n_3 \leq n_2 \leq n_1 \leq k_0/2\}, \tag{7.74}$$

and with $\gamma = \mathrm{Re}(\lambda)$ and $\gamma' = \mathrm{Re}(\lambda')$.

7.5 Numerical methods

We are solving (7.5) as an eigenvalue problem, with eigenvalues $\lambda = \lambda(R_m, \mathcal{L}, \mathcal{M}, \mathcal{N})$. This section covers the main numerical considerations, with the time stepping scheme discussed for general $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$.

A major advantage of imposing large-scale fields, and the most important feature for this investigation, is that the same numerical methods can be employed as with the non-scale separated case with only minor modifications. This avoids lengthy rewriting of codes and has the added benefit of not having to re-explore parameter values for eigenvalue convergence (i.e. minimum resolution, maximum time step size); see table 4.1. We use the same three-step Adams-Bashforth numerical scheme, though owing to the addition of Bloch waves, there are some minor differences in the individual terms.

It should be recalled from section 7.1, that we are resolving $\boldsymbol{b}$, which is 2π-periodic, as opposed to $\boldsymbol{B}$ itself, which would require a domain that changes size, depending on the values of $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$. The Fourier mode indices are unchanged, as clearly we cannot have non-integer indices, but a general mode $\boldsymbol{b}_{l,m,n}$ now corresponds to wavenumbers $(l + \mathcal{L})$, $(m + \mathcal{M})$ and $(n + \mathcal{N})$.

The effects of the imposed Bloch wave are still seen however, as the ∇ operator acting on the field, pulls out an additional factor of $i\boldsymbol{K}$, as seen in (7.5). From the non-large-scale magnetic problem, anywhere that a factor of $\boldsymbol{k}$ appears will be replaced with $\boldsymbol{K} + \boldsymbol{k}$; all coefficients of l, m and n become $l + \mathcal{L}$, $m + \mathcal{M}$ and $n + \mathcal{N}$ respectively. For completeness, the full numerical scheme is given. The three-step Adam's Bashforth scheme is displayed as

$$\boldsymbol{b}^{j+1}_{l,m,n} = E_{l,m,n}\left[\boldsymbol{b}^{j}_{l,m,n} + \tfrac{1}{12}\Delta t\left(23\boldsymbol{\Lambda}^{j}_{l,m,n} - 16E_{l,m,n}\boldsymbol{\Lambda}^{j-1}_{l,m,n} + 5E^2_{l,m,n}\boldsymbol{\Lambda}^{j-2}_{l,m,n}\right)\right], \quad (7.75)$$

with $\boldsymbol{b}^{j}_{l,m,n} = (X^{j}_{l,m,n}, Y^{j}_{l,m,n}, Z^{j}_{l,m,n})$ and $E_{l,m,n}$ modified by the Bloch waves so that

$$E_{l,m,n} = \exp(-\eta\left[(l + \mathcal{L})^2 + (m + \mathcal{M})^2 + (n + \mathcal{N})^2\right]\Delta t). \quad (7.76)$$

The term $\mathbf{\Lambda} = (\Lambda_x, \Lambda_y, \Lambda_z)$ differs from the original dynamo problem also, with Λ_x written as

$$\begin{aligned} 2\left(\Lambda_x\right)^j_{l,m,n} = &\mathrm{i}B(m+\mathcal{M})\left(Y^j_{l,m-1,n} + Y^j_{l,m+1,n}\right) - \mathrm{i}C(m+\mathcal{M})\left(X^j_{l,m,n-1} + X^j_{l,m,n+1}\right) \\ &+ C(m+\mathcal{M})\left(Y^j_{l,m,n-1} - Y^j_{l,m,n+1}\right) - A(m+\mathcal{M})\left(X^j_{l-1,m,n} - X^j_{l+1,m,n}\right) \\ &+ \mathrm{i}B(n+\mathcal{N})\left(Z^j_{l,m-1,n} + Z^j_{l,m+1,n}\right) - \mathrm{i}A(n+\mathcal{N})\left(X^j_{l-1,m,n} + X^j_{l+1,m,n}\right) \\ &+ C(n+\mathcal{N})\left(Z^j_{l,m,n-1} - Z^j_{l,m,n+1}\right) - B(n+\mathcal{N})\left(X^j_{l,m-1,n} - X^j_{l,m+1,n}\right). \end{aligned} \tag{7.77}$$

Similarly the expression expression for Λ_y is given as

$$\begin{aligned} 2\left(\Lambda_y\right)^j_{l,m,n} = &\mathrm{i}C(n+\mathcal{N})\left(Z^j_{l,m,n-1} + Z^j_{l,m,n+1}\right) - \mathrm{i}A(n+\mathcal{N})\left(Y^j_{l-1,m,n} + Y^j_{l+1,m,n}\right) \\ &+ A(n+\mathcal{N})\left(Z^j_{l-1,m,n} - Z^j_{l+1,m,n}\right) - B(n+\mathcal{N})\left(Y^j_{l,m-1,n} - Y^j_{l,m+1,n}\right) \\ &+ \mathrm{i}C(l+\mathcal{L})\left(X^j_{l,m,n-1} + X^j_{l,m,n+1}\right) - \mathrm{i}B(l+\mathcal{L})\left(Y^j_{l,m-1,n} + Y^j_{l,m+1,n}\right) \\ &+ A(l+\mathcal{L})\left(X^j_{l-1,m,n} - X^j_{l+1,m,n}\right) - C(l+\mathcal{L})\left(Y^j_{l,m,n-1} - Y^j_{l,m,n+1}\right), \end{aligned} \tag{7.78}$$

and Λ_z is given by

$$\begin{aligned} 2\left(\Lambda_z\right)^j_{l,m,n} = &\mathrm{i}A(l+\mathcal{L})\left(X^j_{l-1,m,n} + X^j_{l+1,m,n}\right) - \mathrm{i}B(l+\mathcal{L})\left(Z^j_{l,m-1,n} + Z^j_{l,m+1,n}\right) \\ &+ B(l+\mathcal{L})\left(X^j_{l,m-1,n} - X^j_{l,m+1,n}\right) - C(l+\mathcal{L})\left(Z^j_{l,m,n-1} - Z^j_{l,m,n+1}\right) \\ &+ \mathrm{i}A(m+\mathcal{M})\left(Y^j_{l-1,m,n} + Y^j_{l+1,m,n}\right) - \mathrm{i}C(m+\mathcal{M})\left(Z^j_{l,m,n-1} + Z^j_{l,m,n+1}\right) \\ &+ B(m+\mathcal{M})\left(Y^j_{l,m-1,n} - Y^j_{l,m+1,n}\right) - A(m+\mathcal{M})\left(Z^j_{l-1,m,n} - Z^j_{l+1,m,n}\right). \end{aligned} \tag{7.79}$$

For the 1:1:1 ABC flow, we simply set $A = B = C = 1$.

With regards to the actual computation, we are required to define a vector for each parameter $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$. For instance for $\mathcal{L}$, the array "kvecx", which has indices $-N$ to N, stores the value of its index with the value of $\mathcal{L}$ added on to each. This provides an array which has entries $-N + \mathcal{L}$ at index $-N$ to $N + \mathcal{L}$ at index N. Similarly, arrays "kvecy" and "kvecz", corresponding to Bloch waves in the y and z

directions respectively, are defined with $\mathcal{M}$ and $\mathcal{N}$. Then taking the original (full) code, each reference to l, m and n that is not used in the context of an array index, can be replaced by a reference to index l of "kvecx", m of "kvecy" and n of "kvecz". On setting $\mathcal{L} = \mathcal{M} = \mathcal{N} = 0$, the values of the entries in each of these vectors will be equal to the index of those entries, thus returning the large-scale problem to that of the original. Additionally, we no longer enforce that $\boldsymbol{b}_{0,0,0} = \mathbf{0}$ (assuming $\boldsymbol{K} \neq \mathbf{0}$), as this mode now represents the lowest frequency Bloch wave and not a mean field. The large-scale magnetic field problem is, otherwise, numerically identical to the original problem without the ABC flow symmetries or Hermitian symmetry. Note that as in the original dynamo problem, we have still not set the values of A, B or C, so this numerical scheme is capable of simulating all ABC flows.

The biggest alterations are seen in the routines used to convert the field to real space. The fast Fourier transform routine now converts the set of $\boldsymbol{b}_{l,m,n}$ to complex $\boldsymbol{b}$ as opposed to a Hermitian symmetric complex set of $\boldsymbol{b}_{l,m,n}$ to real $\boldsymbol{b}$ in the original problem. To calculate $\boldsymbol{B}$ we use (7.1) and add on its complex conjugate:

$$\boldsymbol{B} = \boldsymbol{b}\,\mathrm{e}^{\mathrm{i}\boldsymbol{K}\cdot\boldsymbol{x}} + \boldsymbol{b}^*\,\mathrm{e}^{-\mathrm{i}\boldsymbol{K}\cdot\boldsymbol{x}}. \tag{7.80}$$

For ease, we split the complex field $\boldsymbol{b}$ into its real and imaginary parts such that $\boldsymbol{b} = \boldsymbol{b}_\mathrm{r} + \mathrm{i}\boldsymbol{b}_\mathrm{i}$. The real large-scale magnetic field is then given as

$$\boldsymbol{B} = 2\cos\left(\boldsymbol{K}\cdot\boldsymbol{x}\right)\boldsymbol{b}_\mathrm{r} - 2\sin\left(\boldsymbol{K}\cdot\boldsymbol{x}\right)\boldsymbol{b}_\mathrm{i}. \tag{7.81}$$

The field can be reconstructed for any range of $\boldsymbol{x}$ required. To visualise a field for one period of the Bloch wave, the field should be reconstructed for $\boldsymbol{x} = (0,0,0)$ to $\boldsymbol{x} = (2\pi/\mathcal{L}, 2\pi/\mathcal{M}, 2\pi/\mathcal{N})$. For producing large-scale fields in the domain of the original problem, we simply reconstruct the field for $\boldsymbol{x} = (0,0,0)$ to $(2\pi, 2\pi, 2\pi)$. For a one-dimensional Bloch wave, e.g. $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, (7.81) becomes

$$\boldsymbol{B} = 2\cos\left(\mathcal{L}x\right)\boldsymbol{b}_\mathrm{r} - 2\sin\left(\mathcal{L}x\right)\boldsymbol{b}_\mathrm{i}. \tag{7.82}$$

7.6 Symmetries and representations for large-scale fields

So far, we have discussed the use of Bloch waves for general $\boldsymbol{K}$. On the discussion of the presence of symmetries in a large-scale field, however, we need to consider different ranges of $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$, as the level of symmetry in the system is controlled by the values of these three parameters. Unlike the case when varying A, B and C, in which for $A \neq B \neq C$ some symmetries still remain, varying $\mathcal{L}$, $\mathcal{M}$ and $\mathcal{N}$ is much more destructive with regards to the level of symmetry.

Whilst we develop theory for the use of symmetries in the Bloch-wave problem, we must make it clear that these concepts are not used to obtain the results presented in this chapter; for completeness, full (non-symmetric) codes are used.

7.6.1 Loss of Hermitian symmetry

From section 7.1, we understand that only a limited parameter space needs to be explored due to the arbitrariness of the field's orientation. We can then identify three general situations in which we need to consider the symmetries, these being $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, 0)$ and $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$. What is important when using the Bloch wave formulation for large-scale fields, is that the magnetic fields are no longer Hermitian symmetric, except for certain special cases: if $\mathcal{L}, \mathcal{M}$ or $\mathcal{N}$ are either 0 or 1/2, then Hermitian symmetry is still present, as the Fourier modes of $\boldsymbol{b}$ and $\boldsymbol{b}^*$ coincide. To briefly explain why this is the case, we give an example, with $\boldsymbol{K} = (1/2, 0, 0)$. We recall that for $\boldsymbol{K} = \boldsymbol{0}$, we have the original problem, which is Hermitian symmetric, with $\boldsymbol{b}_{l,m,n} = \boldsymbol{b}^*_{-l,-m,-n}$. With $\boldsymbol{K} = (1/2, 0, 0)$, we have

$$\boldsymbol{B} = \mathrm{e}^{\mathrm{i}x/2}\,\boldsymbol{b} + \mathrm{e}^{-\mathrm{i}x/2}\,\boldsymbol{b}^*, \tag{7.83}$$

which, in its Fourier decomposition, gives

$$\boldsymbol{B} = \sum_{l,m,n} \boldsymbol{b}_{l,m,n}\, \mathrm{e}^{\mathrm{i}x/2}\, \mathrm{e}^{\mathrm{i}lx+\mathrm{i}my+\mathrm{i}nz} + \sum_{l,m,n} \boldsymbol{b}^*_{-l,-m,-n}\, \mathrm{e}^{-\mathrm{i}x/2}\, \mathrm{e}^{\mathrm{i}lx+\mathrm{i}my+\mathrm{i}nz}. \tag{7.84}$$

By changing the indices in the first sum from l to $l-1$, we have

$$\boldsymbol{B} = \sum_{l,m,n} \boldsymbol{b}_{l-1,m,n}\, \mathrm{e}^{-\mathrm{i}x/2}\, \mathrm{e}^{\mathrm{i}lx+\mathrm{i}my+\mathrm{i}nz} + \sum_{l,m,n} \boldsymbol{b}^*_{-l,-m,-n}\, \mathrm{e}^{-\mathrm{i}x/2}\, \mathrm{e}^{\mathrm{i}lx+\mathrm{i}my+\mathrm{i}nz}, \tag{7.85}$$

and so $\boldsymbol{b}$ and $\boldsymbol{b}^*$ both correspond to the same set of modes, giving

$$\boldsymbol{b}_{l-1,m,n} = \boldsymbol{b}^*_{-l,-m,-n}. \tag{7.86}$$

Similarly for $\boldsymbol{K} = (1/2, 1/2, 0)$, we have $\boldsymbol{b}_{l-1,m-1,n} = \boldsymbol{b}^*_{-l,-m,-n}$ and $\boldsymbol{K} = (1/2, 1/2, 1/2)$ gives us $\boldsymbol{b}_{l-1,m-1,n-1} = \boldsymbol{b}^*_{-l,-m,-n}$. For general $\boldsymbol{K}$, however, there is no equivalent solution, and as such $\boldsymbol{b}$ and $\boldsymbol{b}^*$ correspond to two non-interacting sets of modes. As our simulations involve varying $\boldsymbol{K}$, or more specifically $\mathcal{L}$, we do not take advantage of Hermitian symmetry and merely note that it is present for these cases.

7.6.2 Symmetries of the 1:1:1 ABC flow in large-scale fields

For $\mathcal{L} \neq 0$, it is the case that a transformation that maps x to $-x$ will not map a field back to itself and so will not be a valid symmetry of this system. We can see this by taking such a transformation (e.g. a^2) and applying it to the large-scale field

$$\mathsf{a}^2\boldsymbol{B} = \mathrm{e}^{\mathrm{i}\pi/2(-2m+2n)} \sum_{l,m,n} \boldsymbol{J}_{\mathsf{a}^2} \boldsymbol{b}_{-l,-m,n} \mathrm{e}^{\mathrm{i}(l-\mathcal{L})x+\mathrm{i}my+\mathrm{i}nz}, \tag{7.87}$$

which maps magnetic field modes with Bloch wave $\mathrm{e}^{\mathrm{i}\mathcal{L}x}$ to those of another eigenmode with Bloch wave $\mathrm{e}^{-\mathrm{i}\mathcal{L}x}$.

Any transformation that reverses the sign of x, therefore, will map the field to another field with a Bloch wave with wavenumber $-\mathcal{L}$ and not back to itself. In fact, the same will occur for any transformation that maps any variable, x, y or z to

any other negated variable $-x$, $-y$ or $-z$. Further to this, changing the direction of the large-scale field is not permitted, unless the field itself is invariant under the transformations. An example of this would be if $\mathcal{L} \neq 0$ but $\mathcal{M} = 0$, then any transformation that maps x to $\pm y$ or y to $\pm x$ would not be a valid symmetry of the system.

First, we will consider the case $\boldsymbol{K} = (\mathcal{L}, 0, 0)$. In such a field, we explore the list of symmetries (2.50) to identify fields that map x to x but are free to map y to $\pm z$ and z to $\pm y$, as there is no Bloch wave defined in these directions. Four transformations match these conditions. These are

$$\mathsf{i}(\boldsymbol{x}) = (x, y, z), \tag{7.88}$$

$$\mathsf{b}(\boldsymbol{x}) = (x - \pi/2, \pi/2 - z, y + \pi/2)\,, \tag{7.89}$$

$$\mathsf{b}^2(\boldsymbol{x}) = (x + \pi, -y, \pi - z), \tag{7.90}$$

$$\mathsf{b}^3(\boldsymbol{x}) = (x + \pi/2, z - \pi/2, \pi/2 - y)\,, \tag{7.91}$$

and so the large-scale field with $\boldsymbol{K} = (\mathcal{L}, 0, 0)$ has four-fold helical symmetry.

Considering a more general Bloch wave of the form $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, 0)$ then we are looking to identify transformations that map x to x, y to y and z to $\pm z$. Unsurprisingly, only one transformation fits these criteria, this being the identity, $\mathsf{i}(\boldsymbol{x})$. The only other non-identity transformation that fits the description would be a reflection, of which we know there are none in the group G. A subset of these waves, however, does contain some symmetry of the original 1:1:1 ABC flow. The case $\boldsymbol{K} = (\mathcal{L}, \mathcal{L}, 0)$ is invariant under transformations of the form $\mathsf{g} : (x, y, z) \mapsto (y, x, \pm z)$. Again, $(x, y, z) \mapsto (y, x, z)$ is a reflection in the plane $x = y$ and so is not an element of G, however the transformation k is valid. So for $\boldsymbol{K} = (\mathcal{L}, \mathcal{L}, 0)$, we have

$$\mathsf{i}(\boldsymbol{x}) = (x, y, z), \tag{7.92}$$

$$\mathsf{k}(\boldsymbol{x}) = (y + \pi/2, x - \pi/2, \pi/2 - z), \tag{7.93}$$

and thus two-fold helical (rotation composed with translation) symmetry.

For the very general case $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$, it is clear that no transformation, except the identity $\mathsf{i}(\boldsymbol{x})$, will map the field to itself. If we consider subsets of these Bloch waves, as with the previous case, we can identify valid transformations. Considering $\boldsymbol{K} = (\mathcal{L}, \mathcal{L}, \mathcal{N})$, the only possible transformations that map this type of field to itself are again the identity or a reflection through $x = y$, which is not valid. This case, therefore, has no valid symmetries of the 1:1:1 ABC flow. In the very specific case $\boldsymbol{K} = (\mathcal{L}, \mathcal{L}, \mathcal{L})$, we regain some symmetry due to the invariance of $\boldsymbol{K}$ under permutation of its elements. It is immediately clear that d and d^2 meet the desired permuting requirement, so for this case we have

$$\mathsf{i}(\boldsymbol{x}) = (x, y, z), \tag{7.94}$$

$$\mathsf{d}(\boldsymbol{x}) = (z, x, y), \tag{7.95}$$

$$\mathsf{d}^2(\boldsymbol{x}) = (y, z, x), \tag{7.96}$$

which represents three-fold rotation about the axis $x = y = z$.

7.6.3 Symmetries and representations for $\boldsymbol{K} = (\mathcal{L}, 0, 0)$

So far in this chapter, we have developed the problem for general $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$, despite our eventual aim of focusing on the case $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, and so now we include specifics for this particular case. As shown earlier in this section, the case $\boldsymbol{K} = (\mathcal{L}, 0, 0)$ contains four symmetries of the original 1:1:1 ABC flow, (7.88) to (7.91). In tandem with the label applied to the group of 24 symmetries of the 1:1:1 ABC flow, we have labelled this group of transformations G_{LSF}. This group is isomorphic to the cyclic group $\mathbb{Z}_4 = \{\mathsf{i}, \mathsf{g}, \mathsf{g}^2, \mathsf{g}^3\}$, which is abelian and so each element forms its own conjugacy class. The conjugacy classes of G_{LSF} are, therefore, given by

$$E = \{\mathsf{i}\}, \quad C_4 = \{\mathsf{b}\}, \quad C_4^2 = \{\mathsf{b}^2\}, \quad C_4^3 = \{\mathsf{b}^3\}. \tag{7.97}$$

From Chapter 3, we already understand that the number of symmetry classes is equal to the number of conjugacy classes. Therefore, it can be deduced that there

are four symmetry classes, which we will label I-IV. The character table of any group is invariant under a change of basis and so the character table for the well-known group $\mathbb{Z}_4$ will be identical to the character table for G_{LSF} (see Table 7.1). We know that the fundamental domain for G_{LSF} will be one quarter of the full

Table 7.1: Character table for both $\mathbb{Z}_4$ and G_{LSF}.

	E	C_4	C_4^2	C_4^3
Elements	1	1	1	1
I	1	1	1	1
II	1	-1	1	-1
III	1	i	-1	$-$i
IV	1	$-$i	-1	i

domain $\mathcal{F}_N$, as there are four transformations. We have symmetries that preserve x but map y to $-y$ or $\pm z$ and z to $-z$ or $\pm y$. Therefore in our frequency space, our domain will require the full range of l but only half the range each of m and n. We can therefore define the fundamental domain as

$$\mathcal{G}_{\text{LSF}} = \{(l, m, n) \,:\, |l| \leq N,\, 0 \leq m \leq N,\, 0 \leq n \leq N\}. \tag{7.98}$$

Unlike with the original dynamo problem, the fundamental domain cannot be reduced further through the application of Hermitian symmetry as it no longer holds for general $\mathcal{L}$. The fundamental domain is visualised in figure 7.2.

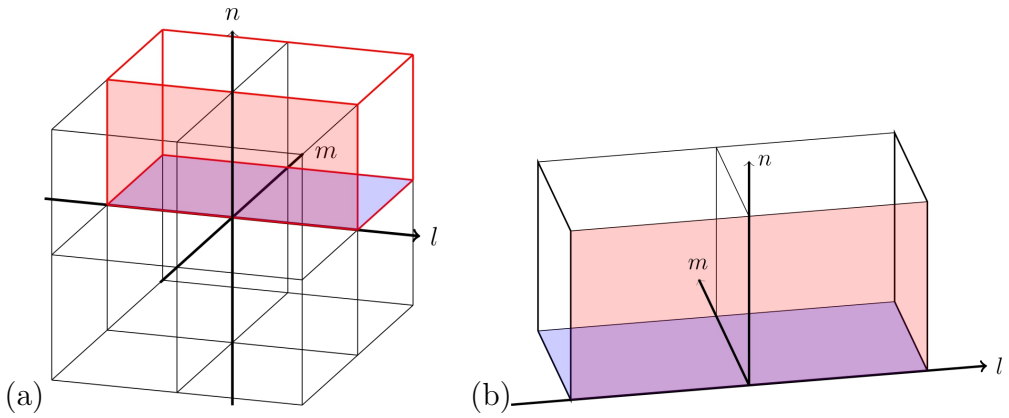

Figure 7.2: Depictions of fundamental domain $\mathcal{G}_{\text{LSF}}$ (a) within $\mathcal{F}_N$ and (b) isolated from $\mathcal{F}_N$.

The representations, M^α, can be read straight from the character table, as for one-

dimensional symmetry classes the irreducible representations and their characters are identical. The following is a list of representations for G_{LSF}:

$$M^{\mathrm{I}}(\mathsf{i}) = M^{\mathrm{I}}(\mathsf{b}) = M^{\mathrm{I}}(\mathsf{b}^2) = M^{\mathrm{I}}(\mathsf{b}^3) = 1, \tag{7.99}$$

$$M^{\mathrm{II}}(\mathsf{i}) = M^{\mathrm{II}}(\mathsf{b}^2) = 1, \quad M^{\mathrm{II}}(\mathsf{b}) = M^{\mathrm{II}}(\mathsf{b}^3) = -1, \tag{7.100}$$

$$M^{\mathrm{III}}(\mathsf{i}) = 1, \quad M^{\mathrm{III}}(\mathsf{b}) = \mathsf{i}, \quad M^{\mathrm{III}}(\mathsf{b}^2) = -1, \quad M^{\mathrm{III}}(\mathsf{b}^3) = -\mathsf{i}, \tag{7.101}$$

$$M^{\mathrm{IV}}(\mathsf{i}) = 1, \quad M^{\mathrm{III}}(\mathsf{b}) = -\mathsf{i}, \quad M^{\mathrm{IV}}(\mathsf{b}^2) = -1, \quad M^{\mathrm{IV}}(\mathsf{b}^3) = \mathsf{i}. \tag{7.102}$$

As all irreducible representations are one-dimensional, each symmetry class will consists of just one magnetic field. Due to the simplicity of the form of the representations, we can construct the magnetic fields explicitly as

$$\boldsymbol{B}^{\mathrm{I}} = \boldsymbol{B} + \mathsf{b}\boldsymbol{B} + \mathsf{b}^2\boldsymbol{B} + \mathsf{b}^3\boldsymbol{B}, \tag{7.103}$$

$$\boldsymbol{B}^{\mathrm{II}} = \boldsymbol{B} - \mathsf{b}\boldsymbol{B} + \mathsf{b}^2\boldsymbol{B} - \mathsf{b}^3\boldsymbol{B}, \tag{7.104}$$

$$\boldsymbol{B}^{\mathrm{III}} = \boldsymbol{B} + \mathsf{ib}\boldsymbol{B} - \mathsf{b}^2\boldsymbol{B} - \mathsf{ib}^3\boldsymbol{B}, \tag{7.105}$$

$$\boldsymbol{B}^{\mathrm{IV}} = \boldsymbol{B} - \mathsf{ib}\boldsymbol{B} - \mathsf{b}^2\boldsymbol{B} + \mathsf{ib}^3\boldsymbol{B}. \tag{7.106}$$

7.6.4 Numerical methods for $\boldsymbol{K} = (\mathcal{L}, 0, 0)$

The numerical considerations for the symmetrically reduced case $\boldsymbol{K} = (\mathcal{L}, 0, 0)$ can now be discussed. We have identified four (identity and three helical) transformations which are based around mapping the variable x to itself, with a translation that becomes a phase-shift in Fourier space. The transformations b, b^2 and b^3 serve to rotate the magnetic field through $\pi/2$ around the x-axis. In the frequency domain, this serves to rotate the modes through $\pi/2$ in the l-axis, giving the fundamental domain defined in (7.98) as 1/4 of the full set of harmonics; see figure 7.2. For modes to be time stepped in this fundamental domain, every mode requires the value of its immediate neighbours.

There are two 'faces' of modes that have neighbours lying outside the fundamental domain. As with the original problem, we are required to store an extra layer of

modes. Through the symmetries, their values can be calculated using only modes within the fundamental domain. The first face consists of modes of the form $\boldsymbol{b}_{l,m,0}$, which requires the layer of modes of the form $\boldsymbol{b}_{l,m,-1}$ to be filled. The second face consists of modes of the form $\boldsymbol{b}_{l,0,n}$ and relies on a layer of modes of the form $\boldsymbol{b}_{l,-1,n}$. Figure 7.3 shows visual representations of a cross-section of $\mathcal{G}_{\mathrm{LSF}}$. In Figure 7.3(a), the shaded area represents modes that are not time stepped (index N), whereas the yellow mode represents $\boldsymbol{b}_{l,0,0}$, which is mapped to itself under each of the symmetries and so may be restricted to certain values. Figure 7.3(b) indicates the location of the additional mode-layer outside of the fundamental domain.

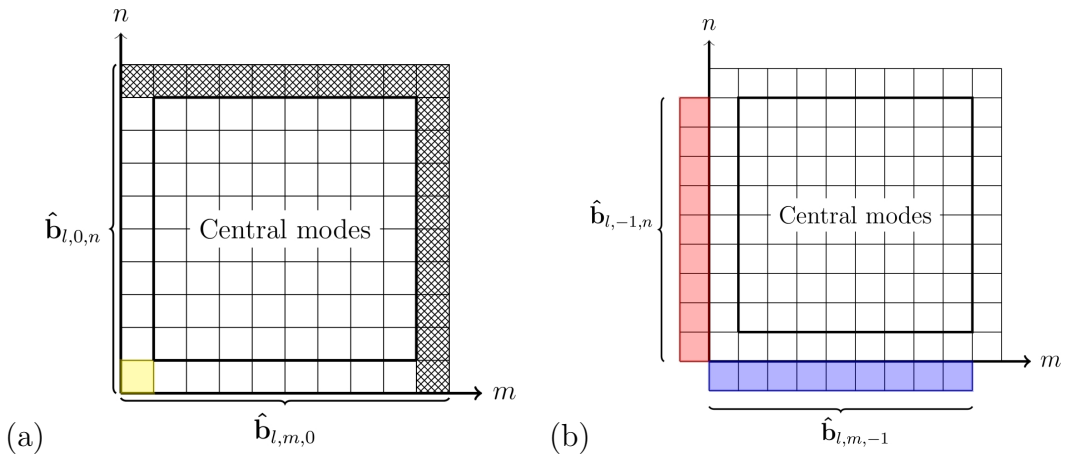

Figure 7.3: Cross section of $\mathcal{G}_{\mathrm{LSF}}$ for unspecified l, identifying (a) modes that are not time stepped (hashed) and (b) extent of extra mode layers.

Transformations b^3 and b are used to calculate modes $\boldsymbol{b}_{l,m,-1}$ and $\boldsymbol{b}_{l,-1,n}$ respectively. These relations are given below.

$$\begin{aligned} \boldsymbol{b}_{l,-1,n} &= M^{\alpha}(\mathsf{b})\, \mathrm{e}^{\mathrm{i}\pi/2(l+\mathcal{L}+1-n)}\, \boldsymbol{J}_{\mathsf{b}}\, \boldsymbol{b}_{l,n,1} \\ &= \mathrm{e}^{\mathrm{i}\mathcal{L}\pi/2}\, M^{\alpha}(\mathsf{b})\, \mathrm{e}^{\mathrm{i}\pi/2(l+1-n)}\, \boldsymbol{J}_{\mathsf{b}}\, \boldsymbol{b}_{l,n,1}, \end{aligned} \tag{7.107}$$

$$\begin{aligned} \boldsymbol{b}_{l,m,-1} &= M^{\alpha}(\mathsf{b}^3)\, \mathrm{e}^{\mathrm{i}\pi/2(-l-\mathcal{L}+m+1)} \boldsymbol{J}_{\mathsf{b}^3}\, \boldsymbol{b}_{l,1,m} \\ &= \mathrm{e}^{-\mathrm{i}\mathcal{L}\pi/2}\, M^{\alpha}(\mathsf{b}^3)\, \mathrm{e}^{\mathrm{i}\pi/2(-l+m+1)} \boldsymbol{J}_{\mathsf{b}^3}\, \boldsymbol{b}_{l,1,m}. \end{aligned} \tag{7.108}$$

In comparison, the transformations behave in the same way as with the original problem, except for an extra phase-shift with respect to $\mathcal{L}$ and, numerically, the

mode copying is carried out in an identical fashion. As $\mathcal{L}$ is fixed for each simulation, the phase-shift terms $e^{i\mathcal{L}\pi/2}$ and $e^{-i\mathcal{L}\pi/2}$ should be calculated and stored at runtime for later reference. In practical terms, one-dimensional representations mean that when simulating these fields, we would only need to store one field at any one time. If we were to resolve all four symmetry classes, we would be computationally solving four subproblems, each at 1/4 of the cost of the full problem. This means that there is no net reduction in computing time when solving all four subproblems, though we would be revealing far more about how effective different types of structure are at generating large-scale magnetic fields.

7.7 Results

This section communicates the findings of the numerous simulations for $\boldsymbol{K} = (\mathcal{L}, 0, 0)$. We carry out simulations for $\mathcal{L} = 0$ to $\mathcal{L} = 0.5$ in steps of 0.01, totalling 51 simulations per R_m investigated. The eigenvalues for every $\mathcal{L}$ at each R_m are determined, though we only supply information on the growth rates. We note that the eigenvalues for $\mathcal{L} = 0$ are equal (to within 3 s.f.) to those in the original problem, as expected. Simulations are carried out for R_m up to only 10^3, owing to the number of individual simulations per R_m studied. Section 7.7.1 summarises the optimal growth rates and their corresponding values of $\mathcal{L}$. The following section, section 7.7.2, briefly considers the resulting field structures and supplies some visualisations. A brief comparison with results extracted from Galanti *et al.* (1992) is made in 7.7.3 and a discussion on the applicability of the α-effect is given in section 7.7.4.

As already stated, these results are obtained using full numerical codes, with no assumption of symmetry.

7.7.1 Growth rates and optimal scales

In the original magnetic problem, the growth rate is found from the linear trend in $\ln(E_{\mathrm{M}})$, which is calculated from all modes in the fundamental domain. This remains unchanged for the large-scale field, with the exception that all modes within $\mathcal{F}_N$ are now used. For oscillating fields, minimum to minimum linear trends are calculated. Once growth rates for all values of $\mathcal{L}$ are obtained for each R_m, we fit a cubic spline so that a better estimate of the optimal growth rate (and $\mathcal{L}$) can be made. Assuming that there are no sudden changes in the growth rate for very small (i.e. $\sim O(10^{-3})$) changes in $\mathcal{L}$, then our cubic spline provides a reasonably accurate representation of the continuous profile of growth rate for changing $\mathcal{L}$.

Figure 7.4 displays the maximum growth rate obtained for each R_m. Where possible, a second set of results is displayed which identifies the growth rate at a distinct secondary peak. Alongside these growth rates, those of the non-large-scale problem (i.e. $\mathcal{L} = 0$) are given, for a direct comparison. Counter-intuitively, we see that the greatest growth rates occur at low R_m, particularly when the non-large-scale field is not a dynamo. The growth rate peaks at $R_m = 20$ with $\gamma = 0.144$ and $\mathcal{L} = 0.36$. From here, the growth rate decreases and at $R_m = 30$, a second (distinct) peak is seen in the profile of $\gamma = \gamma(\mathcal{L})$. The optimal growth rate steadily decreases, whilst the growth rate of the secondary peak slowly increases as R_m is increased. The two peaks of optimal growth converge at $R_m \approx 260$ after which the peak growth rate increases, with this emerging branch appearing to be a continuation of the path of the secondary peak. The growth rates reach a local maximum at $R_m = 400$, with $\gamma = 0.109$, and with a steady decrease to $\gamma = 0.106$ at $R_m = 10^3$.

Figure 7.5 displays the optimal value of $\mathcal{L}$ that the growth rates, seen in the previous figure, correspond to. Again we highlight the growth rates of the second peak and see that after the convergence of these two peaks, the branch of optimal $\mathcal{L}$ appears to be a continuation of the secondary peak's optimal $\mathcal{L}$. The optimal $\mathcal{L}$ increases from approximately 0.31 and peaks with $\mathcal{L} = 0.444$ (or $\mathcal{L} = 0.44$ without interpolation) at $R_m = 150$. Interestingly, the secondary peak of solutions reaches a

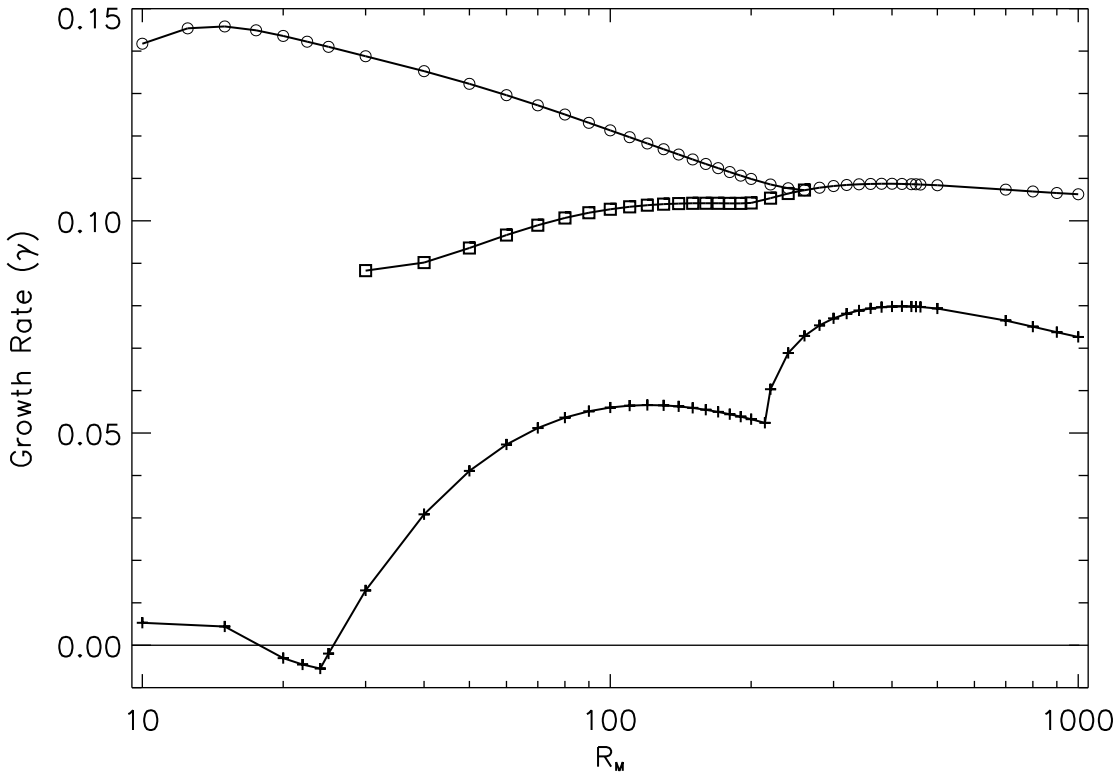

Figure 7.4: Maximum growth rates obtained for large-scale fields for R_m up to 10^3. Circles represent peak growth rates, squares identify a secondary growth rate peak (when present) and the growth rate for non-large scale fields (i.e. $\mathcal{L} = 0$) is shown with plus symbols for comparison.

minimum $\mathcal{L}$ of 0.0731 at the same R_m. As the Bloch waves that are amplified most effectively by the flow become shorter in period for increasing R_m, the secondary peak sees its most effective Bloch waves lengthening in their periods. At $R_m \approx 200$, a change occurs in that both branches tend towards a value of $\mathcal{L} \approx 0.2$ over a short interval of R_m. By $R_m = 280$, at which point the two peaks have converged, the scale is $\mathcal{L} = 0.203$ and as R_m increases, the optimal $\mathcal{L}$ fluctuates in the interval $\mathcal{L} \in [0.20, 0.21]$. This may possibly be an artefact of the simulations, owing to the discretisation of $\mathcal{L}$, though cubic-spline interpolation was implemented to avoid such problems. It is also possible that the fluctuations are real. For $R_m > 500$, $\mathcal{L}$ appears to be asymptotically approaching 0.2. For large R_m, therefore, the most effective magnetic field for dynamo action has a period five times that of the flow.

To explain the primary and secondary peaks, we give the results for twelve values

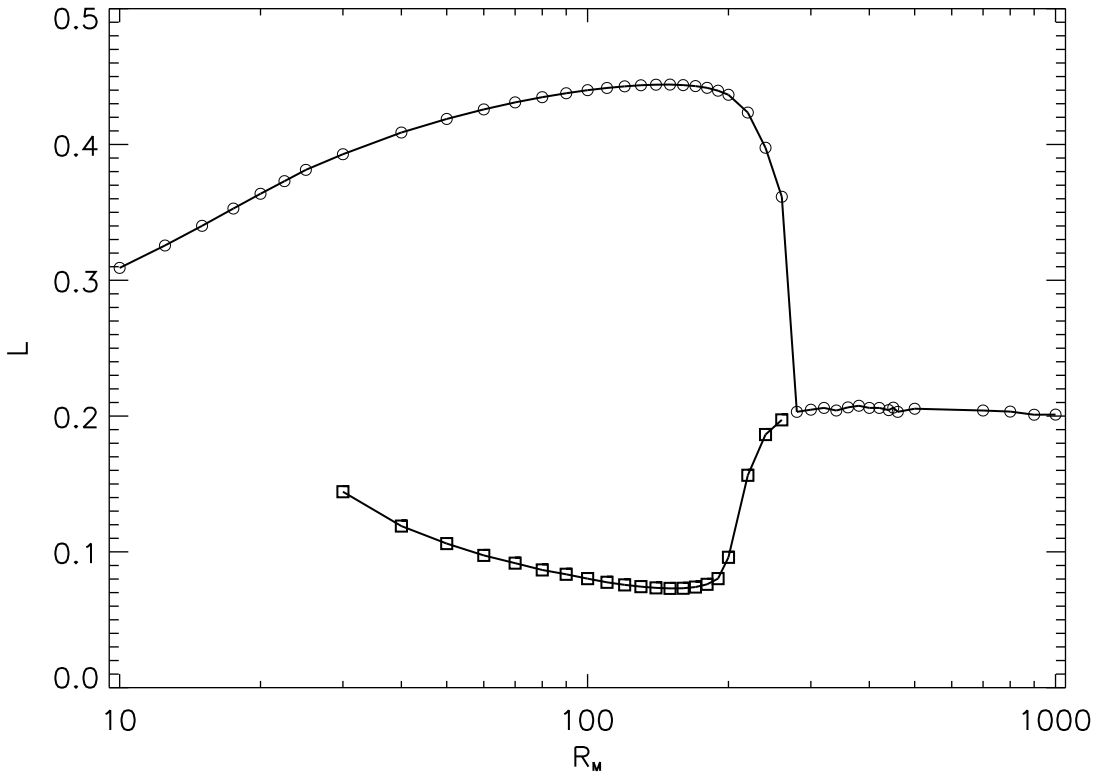

Figure 7.5: Value of $\mathcal{L}$ with largest growth rate for R_m up to 10^3. Circles represent optimal $\mathcal{L}$, with squares identifying $\mathcal{L}$ for secondary peak (when present).

of R_m, split into four plots. The distinct peaks identify that more than one scale exists for which the Bloch wave is effectively amplified. For $30 \leq R_m \leq 130$, it is not necessarily clear whether the two peaks belong to the same branch, due to the appearance of other dominant eigenmodes. For $R_m > 130$, we are able to see that the fields at the growth rate peaks do indeed belong to the same branch. Interestingly, we note that there are no large jumps in the growth rate for small changes in $\mathcal{L}$; the growth rate changes gradually as we vary $\mathcal{L}$. The growth rate profiles change significantly for increasing R_m, however, and a single growth rate peak for low R_m becomes two distinct peaks at intermediate R_m, with growth rate kinks seen suggesting contributions from multiple eigenmodes. Further increasing R_m sees the kink disappearing and the resulting smooth eigenfunction being convex (a line segment joining any two points will always lie above the curve) for a limited range of $\mathcal{L}$. These functions gradually flatten as R_m increases until at $R_m = 260$, the two peaks are roughly level and for a broad range of $\mathcal{L}$, the Bloch waves are

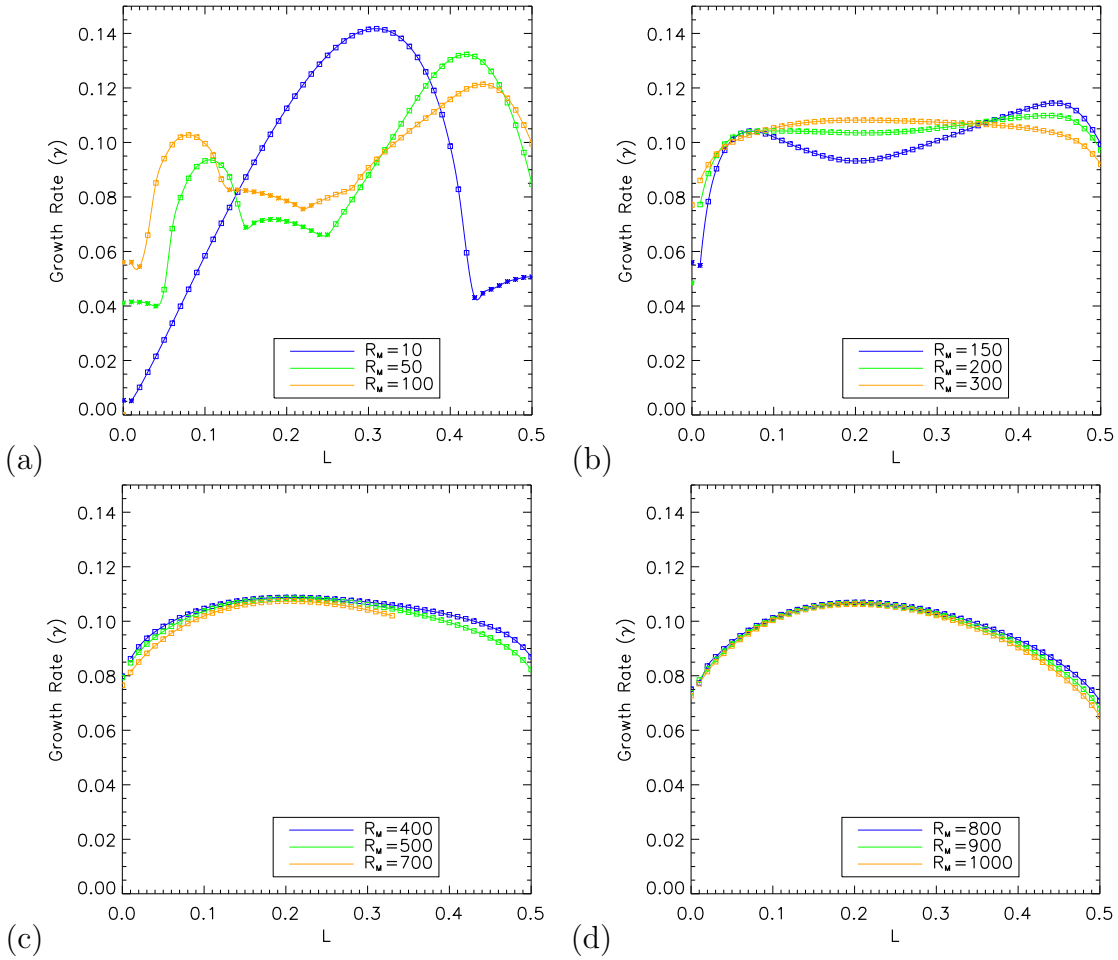

Figure 7.6: Growth rates for varying $\mathcal{L}$ with cubic spline overlaid at (a) $R_m = 10, 50, 100$, (b) $R_m = 150, 200, 300$, (c) $R_m = 400, 500, 700$, and (d) $R_m = 800, 900$ and 1000. Squares (asterisks) highlight steady (oscillatory) modes.

amplified with roughly equal effectiveness. For $R_m \geq 300$, we see a purely concave function and a single growth rate peak. For $R_m > 300$, both the peak growth rates and the profiles over $\mathcal{L}$ do not change significantly, in that γ is very similar for all $\mathcal{L}$ at these values of R_m; see figure 7.6(d). This seems to suggest that the dominant eigenmode is close to converging to some limit, which is consistent with the optimal $\mathcal{L}$ (and its corresponding growth rate) approaching an asymptotic limit.

In figure 7.6, we also see growth rates of oscillatory fields, identified by asterisks, as opposed to the squares indicating steady growth. These appear as transient branches between the peaks, at intermediate R_m, and for $\mathcal{L} \to 0.5$ for low R_m. Most importantly, the fields at the optimal $\mathcal{L}$ are always steady, along with the

majority of the resulting fields as $\mathcal{L}$ is varied. We clearly see that for $\mathcal{L} = 0$ (i.e. the non-large-scale field), the field is the slowest growing of all possible values of $\mathcal{L}$ explored.

Finally we compare the growth rates for fixed values of $\mathcal{L}$ over the range of R_m investigated. These are displayed in figure 7.7. In the first figure, we show the

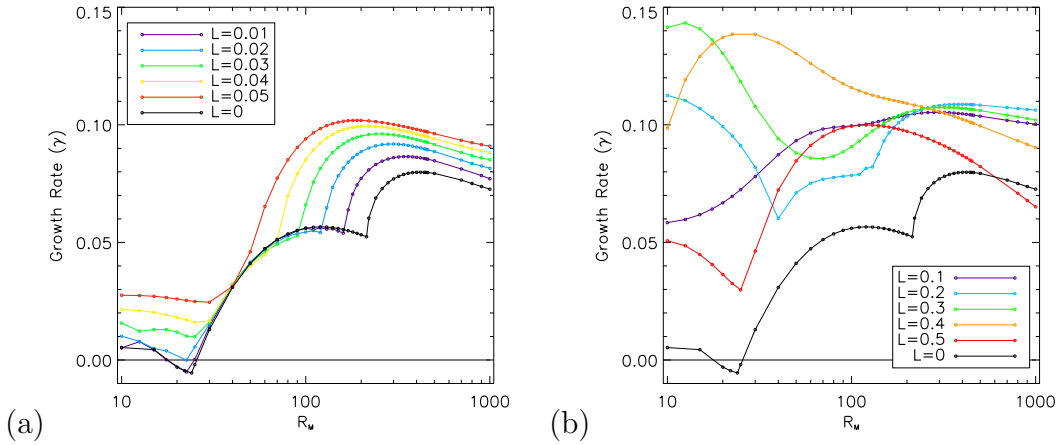

Figure 7.7: Growth rate of eigenmodes for (a) $\mathcal{L}$ from 0.01 to 0.05 and (b) $\mathcal{L}$ from 0.1 to 0.5, compared against growth rate at $\mathcal{L} = 0$.

growth rates for small $\mathcal{L}$, with results for the non-large-scale field also provided for comparison. In terms of periodicity, these Bloch waves approach a mean field-like situation as $\mathcal{L} \to 0$. Unsurprisingly, we see that as $\mathcal{L} \to 0$, the behaviour of the eigenmodes approaches that of the original problem. For these small $\mathcal{L}$, the same growth rate kinks are seen as in the original problem, though as $\mathcal{L}$ increases, these kinks become less obvious. This does suggest that for at least $\mathcal{L} \leq 0.05$, the eigenmodes and thus fields found in these cases are equivalent to those found in the original problem, with similar eigenvalue coalescences and mode crossings taking place. It is possible that the fields for low and high R_m resemble those of classes II and V, respectively, of the original problem; this will be dealt with in the next section when visualising the fields. For more moderate $\mathcal{L}$, we see very different growth rate profiles with, for example, $\mathcal{L} = 0.3$ and $\mathcal{L} = 0.4$ showing no sign of internal mode changes over the whole range of R_m. For large R_m, all eigenmodes seem to tend to an asymptotic decrease, similar to that seen after the local maximum at $R_m \approx 420$ in the original problem. Since there is a clear

relationship between the growth rate at $\mathcal{L} = 0$ and for $\mathcal{L} \neq 0$ (particularly at small $\mathcal{L}$), it is likely that the increase seen in the non-large-scale field for $R_m > 2000$ will also be mirrored in the large-scale fields.

7.7.2 Field structures

In this section, we visualise the resulting magnetic fields in the familiar cross-sections and three-dimensional field strength plots. These visualisations differ in that we either reconstruct the field for the entire length of the Bloch wave, which is not practical for small $\mathcal{L}$, or visualise in the original domain $x, y, z \in [0, 2\pi]$, so that only part of the field is seen. Most visualisations will be of the latter type but for demonstration purposes, we include an example of the former. We refer to the individual segments in which a single copy of the flow exists, as *cells*; the first cell is the original domain.

We start by visualising (in figure 7.8) the fastest growing field at $R_m = 1000$, which is that of $\mathcal{L} = 0.2$. To show the full periodic field, we have plotted x up to $(2\pi/\mathcal{L}) = 10\pi$ (5 cells); the plots are periodic in the sense that the field at one end matches up with the field at the other end. Although the field structures

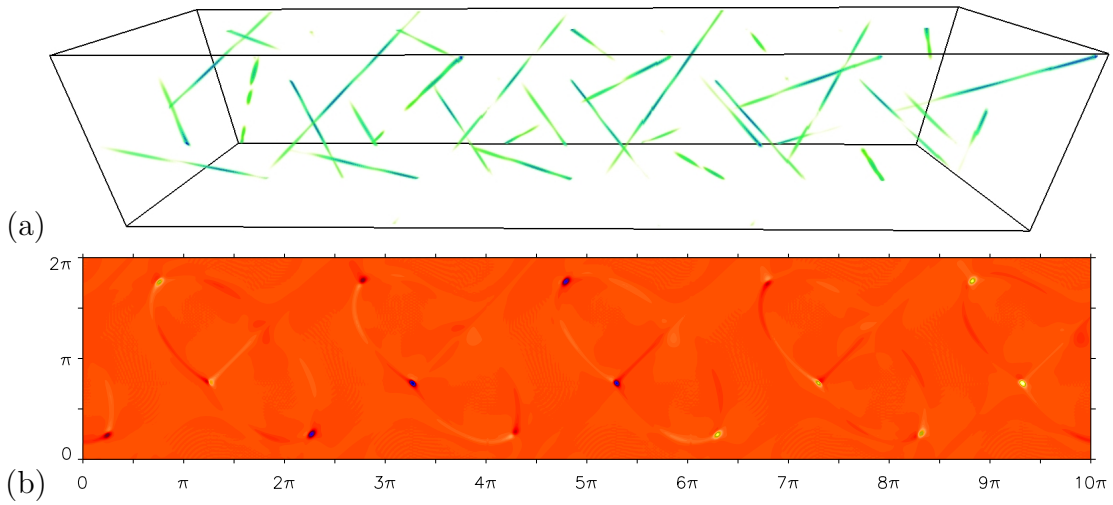

Figure 7.8: Field visualisations of (a) $|\boldsymbol{B}|$ and (b) B_z at $z = \pi/4$ for a full Bloch wave at $R_m = 10^3$ with $\mathcal{L} = 0.2$.

are very fine at this R_m, it is clear that this field closely resembles the equivalent

non-large-scale field, with the same cigar-like field concentrations. The cigars lie on the separatrices of the flow, as with the structures of class II and V of the original problem and seem to lack the three-fold rotation symmetry seen in class II, suggesting fields in class V are the closest analogue. Unlike the field in class V, however, there are no secondary cigars. Instead, the cigars form 'chains' in which the cigars are arranged consecutively along the separatrix, but with a gap at the β-type stagnation point as it disperses the field here. From further field visualisations (not shown here), it is found that neighbouring cigars have opposing polarity. We see that the arrangement is the same in each cell, but the field strength and polarity differ. For instance, in the first cell the cigars are weak but are strong in successive cells. This is not unexpected, however, as in reconstructing the large-scale field, we have $\boldsymbol{b}_{\mathrm{r}}$ and $\boldsymbol{b}_{\mathrm{i}}$ defined only for $x \in [0, 2\pi]$ and must construct $\boldsymbol{B}$ using a Bloch wave; we are seeing the effects of the periodic functions $\cos(Lx)$ and $\sin(Lx)$ in (7.82). We note that the field strength $|\boldsymbol{B}|$ does not change considerably from cell to cell, despite the changing polarity of the field.

Figure 7.9 shows the field structure of the eigenmode with the highest growth rate over the whole range of R_m and $\mathcal{L}$ studied. This occurs in the low-R_m regime, at $R_m \approx 15$ and with Bloch wave frequency $\mathcal{L} = 0.34$, which is just under three times the length of the original domain. The 3-dimensional plots are now displaying B_z, with positive values in green and negative in purple. Though the structures are very

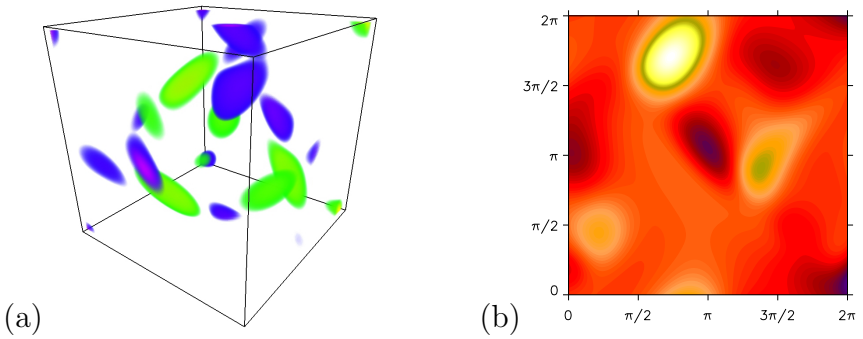

Figure 7.9: Visualisations in first cell of (a) field component B_z (green is -ve, purple is +ve) and (b) field component B_z at $z = \pi/4$, for $R_m = 15$ and optimal $\mathcal{L} = 0.34$.

diffuse, we see flux tubes on the separatrices. The strong areas of field are no longer centred at the α-type stagnation points, but instead at the separatrix midpoints.

The individual flux tubes that form either side of the α-type points are of opposite polarity. In the cross-section, we see these flux tubes do extend to the stagnation point but are very weak at this location. For comparison, we also visualise the field at $R_m = 300$, for $\mathcal{L} = 0.2$ (optimal scale) in figure 7.10. The structure appears to have a very similar configuration as in the previous figure, though less diffuse in this high-R_m regime. Interestingly, the cross-section of b_z clearly resembles that of

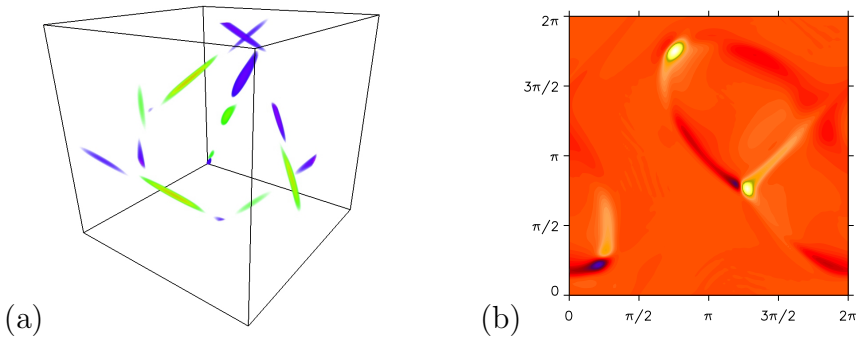

Figure 7.10: Visualisations in first cell of (a) magnetic field component B_z and (b) field component B_z at $z = \pi/4$ for $R_m = 300$ and optimal $\mathcal{L} = 0.2$.

class V in the original problem, with two cigars of opposing polarity at the α-type point. Instead of forming side-by-side (as in class V) at the stagnation point, these cigars remain centred on the separatrix midpoints, though it is clear that one of these cigars has a higher field strength than the other.

We briefly study the structures for small $\mathcal{L}$ in comparison to those of the original problem. Chosen are $R_m = 15$ and 100, as these lie before and after growth rate kinks, respectively. Figure 7.11 shows fields at $\mathcal{L} = 0$ (a,d), $\mathcal{L} = 0.01$ (b,e) and $\mathcal{L} = 0.05$ (c,f) at $R_m = 15$. At $\mathcal{L} = 0$, we clearly see the familiar cigar structures of class II, lying on the separatrices, though in the cross-section we see more complicated yet weaker field structure surrounding the cigar. At $\mathcal{L} = 0.01$, the field has not changed significantly, and remains in roughly cigar-like flux tubes. Strong field is also arranged around the cigar in the cross-section, though we note that at $R_m = 15$ the fields are oscillatory, making it more difficult to compare them directly. Further to this, for $\mathcal{L} = 0.01$, there are 100 cells (only the first is shown here) and from how the field is reconstructed, it is likely that one cell (other than the first) will more closely represent the field seen at $\mathcal{L} = 0$. Regardless, the field

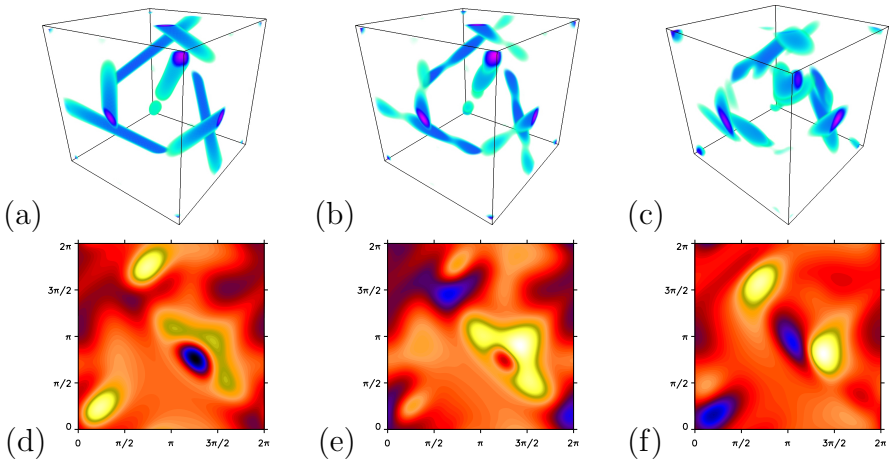

Figure 7.11: Visualisations of (a,b,c) magnetic field strength $|\boldsymbol{B}|$ and (d,e,f) field component B_z at $z = \pi/4$ for $R_m = 15$, with (a,d) $\mathcal{L} = 0$, (b,e) $\mathcal{L} = 0.01$ and (c,f) $\mathcal{L} = 0.05$.

structure has changed from $\mathcal{L} = 0$ to $\mathcal{L} = 0.01$, though clearly not significantly. By $\mathcal{L} = 0.05$, however, the field appears to be drastically different with twisted flux tubes of opposing polarity centred at the α-type stagnation point. In fields not visualised here, we also see these twisted flux tubes at $\mathcal{L} = 0.1$ but by $\mathcal{L} = 0.2$, the structures merge and appear more cigar-like. In terms of field configuration, no major changes occur until $\mathcal{L} > 0.4$, for which the areas of strong field are more diffuse and far less organised.

For the comparison after the mode change, we compare structures at $R_m = 100$. The familiar double-cigar configuration of class V is visible in figure 7.12(a,d). Again the structure does not change significantly with $\mathcal{L} = 0.01$, though the two cigars are not clearly distinguished in the figure, but by $\mathcal{L} = 0.05$, there are at least two twisted flux tubes formed on the separatrices. For larger $\mathcal{L}$, the structures are fairly homogeneous, with one cigar whose strength varies along its length.

7.7.3 Comparison with GPS

By varying only $\mathcal{L}$, we are testing a subset of the possible Bloch wave configurations and so we cannot be sure that these field scalings are the most effective. One way to check this is to compare our results directly with the limited results of Galanti *et*

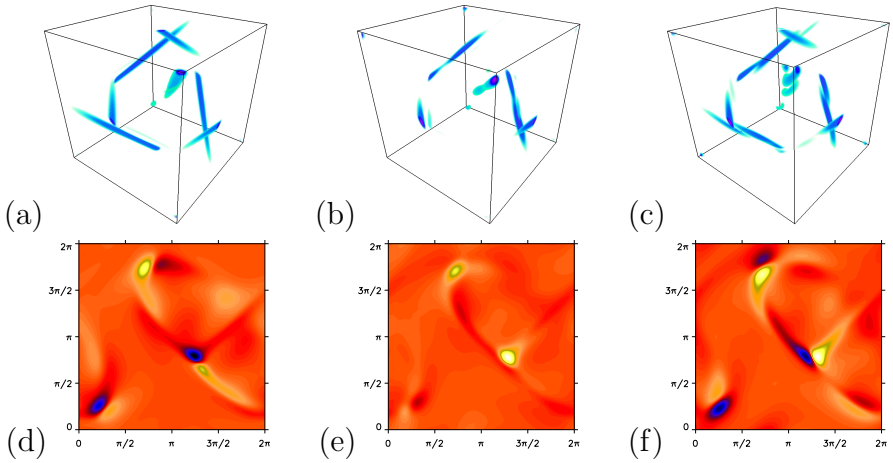

Figure 7.12: Visualisations of (a,b,c) magnetic field strength $\boldsymbol{B}$ and (d,e,f) field component B_z at $z = \pi/4$ for $R_m = 100$, with (a,d) $\mathcal{L} = 0$, (b,e) $\mathcal{L} = 0.01$ and (c,f) $\mathcal{L} = 0.05$.

al. (1992). This also allows us to assess whether our results are reliable by verifying that they do not exceed the growth rates of GPS. Figure 7.13 compares the GPS growth rate with $k_0 = 2$ and the equivalent Bloch wave in one direction ($\mathcal{L} = 0.5$). The GPS result is rescaled to create a fair comparison and yet it is clear that their growth rate far exceeds that of our large-scale field. Thus for $R_m \leq 40$, the Bloch

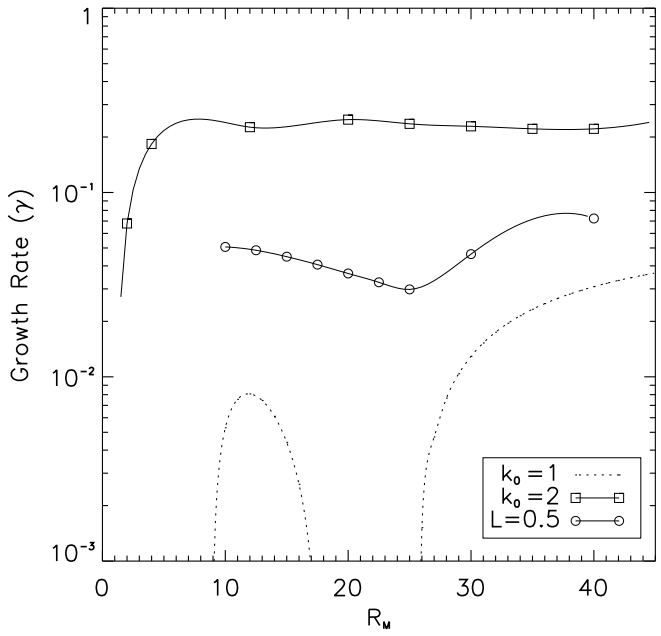

Figure 7.13: Growth rates as a function of R_m with GPS results (squares) and equivalent Bloch wave results (circles). For comparison, non-large scale growth rate is shown.

wave with $\boldsymbol{K} = (1/2, 0, 0)$ is far from the most effective scale for field amplification. Out of all three Bloch waves equivalent to those that GPS resolves at $k_0 = 2$, we speculate that the wave with wavenumber $\boldsymbol{K} = (1/2, 1/2, 1/2)$ results in the fastest dynamo growth.

As another point of comparison, we study the growth rates for fixed R_m and varying k_0, given by Galanti *et al.* (1992). For a fair comparison, we calculate $\gamma_{\max}$ as described by (7.73), with the exception that $\mathcal{M} = \mathcal{N} = 0$, as we only vary $\mathcal{L}$. For instance, to draw comparison with $k_0 = 7$ of GPS, we find the growth rates of $\mathcal{L}$ at $1/7, 2/7$ or $3/7$, rounded to the nearest 0.01 (this being the step in $\mathcal{L}$), giving $\mathcal{L} = 0.14, 0.29$ and 0.43 respectively; we then use the largest growth rate of all equivalent $\mathcal{L}$. (This most closely represents the method of GPS when varying k_0.) Figure 7.14 displays growth rate versus k_0 at $R_m = 12$ and 20 for GPS, with the results for the closest equivalent one-dimensional Bloch wave also given. The optimal Bloch wave scale and growth rate is also indicated by an asterisk, for contrast. At both $R_m = 12$ and 20, it is again clear that our large-scale fields are

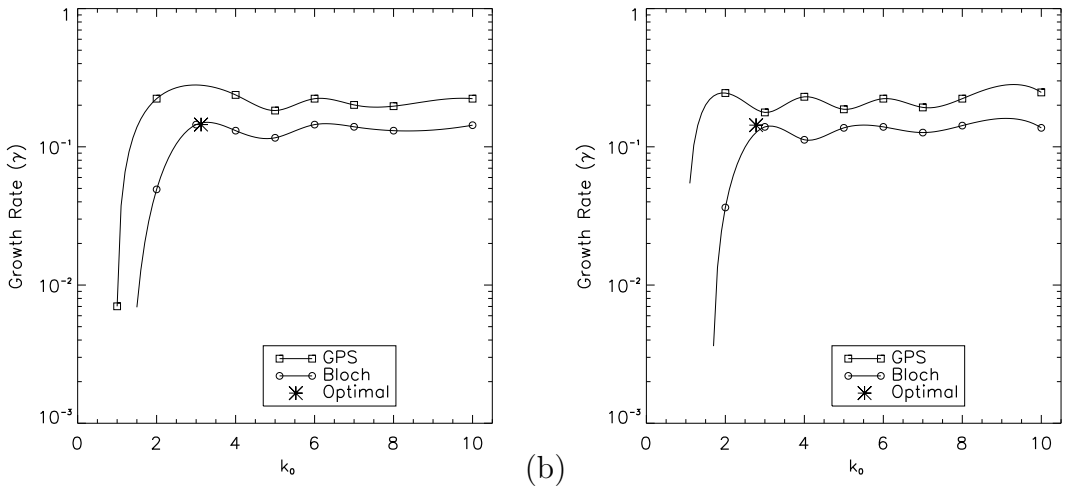

Figure 7.14: Growth rate versus k_0 for GPS and equivalent Bloch waves at (a) $R_m = 12$ and (b) $R_m = 20$. The optimal Bloch wave scale and growth rate is also shown.

less effective dynamos than those of GPS; this is to be expected. We see the same fluctuations in γ as k_0 is increased and it is now much more obvious why this is the case. Taking the Bloch wave results at $R_m = 20$, the optimal Bloch wavenumber is $\mathcal{L} = 0.36$. In the comparison, we use Bloch waves with $\mathcal{L} = (n/k_0) < 0.5$ for

$n \in \mathbb{N}$ and so as k_0 increases, there are more possible choices for the equivalent $\mathcal{L}$, but this does not guarantee that the closest choice to the optimal $\mathcal{L}$ at a given k_0 is closer than for the optimal choice at $k_0 - 1$. An example would be $k_0 = 6$, where the equivalent optimal Bloch wave has $\mathcal{L} = 2/6 = 1/3$ and at $k_0 = 7$, where the equivalent wave has $\mathcal{L} = 2/7$, which is further from $\mathcal{L} = 0.36$ than at $k_0 = 6$ and so has a lower growth rate. This causes the growth rate to fluctuate as the optimal scales oscillate either side of the 'true' optimum as it is approached. Clearly, on increasing k_0 and so allowing more possible choices for $\mathcal{L} = n/k_0$, it is possible to reduce the distance between our optimal choice and the true optimum. This also occurs for GPS, except that the values of all three of $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$ in the equivalent Bloch waves oscillate around the true optimum as k_0 increases.

7.7.4 Bloch waves and the α-effect

Finally, we address the extent to which the α-effect is seen in the problem of large-scale fields in the ABC flow. Section 7.3 culminated in (7.64), which describes how the analytically determined α would manifest itself in our results with $\boldsymbol{K} = (\mathcal{L}, 0, 0)$. Though we expect two curves, one for each helical wave $\boldsymbol{b}_0 = (0, 1, \mathrm{i})$ and $\boldsymbol{b}_0 = (0, 1, -\mathrm{i})$, only one of these grows whilst the other decays. With $\alpha = -1$ for the 1:1:1 ABC flow, we expect $\boldsymbol{b}_0 = (0, 1, -\mathrm{i})$ to be the growing field. For this wave, (7.64) becomes

$$\lambda = R_m \mathcal{L} - R_m^{-1} \mathcal{L}^2. \tag{7.109}$$

In figure 7.15, we present results for $R_m = 0.05, 0.1, 0.25, 0.5, 1.0$ and 5.0, in an effort to provide a sufficient range of R_m for comparison. At $R_m = 0.05$ and 0.1, we have near-perfect agreement, with all growth rates obtained numerically lying on the (dashed) parabolic curve given by λ in (7.109). By $R_m = 0.25$ and 0.5, the two curves are clearly separate with growth rates of our large-scale fields lower than predicted by (7.109). For $R_m \geq 1$, approximating growth rates with (7.109) clearly does not work; we expect this to be the case as this description should only hold for $R_m = \varepsilon \ll 1$. Despite the lack of agreement between the curves at $R_m = 1.0$, it

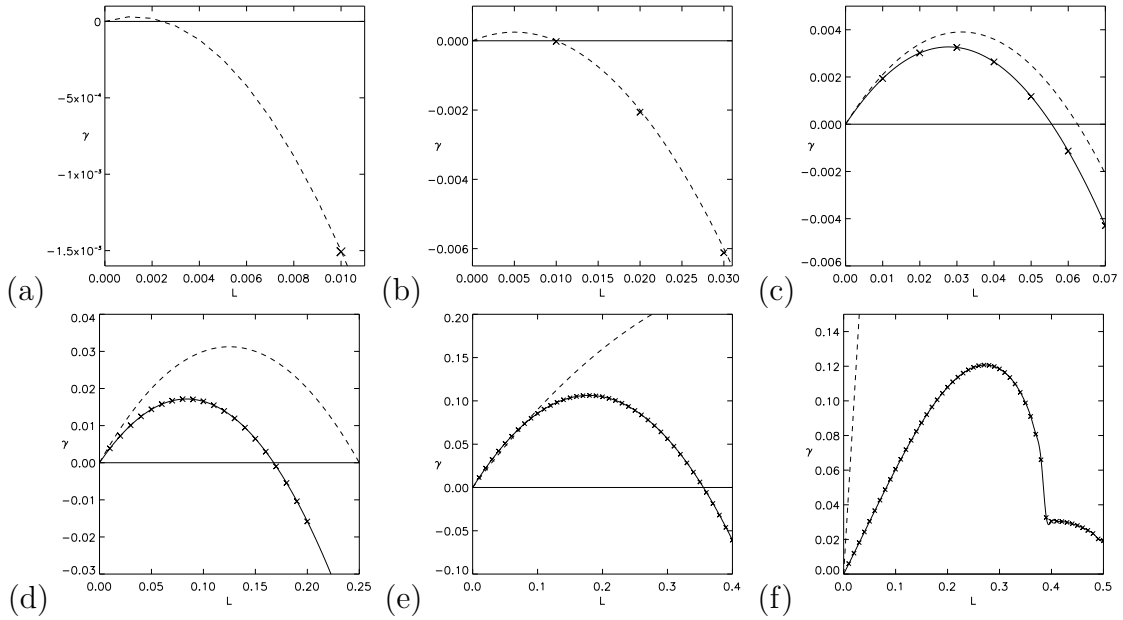

Figure 7.15: Growth rate versus $\mathcal{L}$ (crosses and solid line) in comparison to λ of (7.109) (dashed line) for (a) $R_m = 0.05$, (b) $R_m = 0.1$, (c) $R_m = 0.25$, (d) $R_m = 0.5$, (e) $R_m = 1.0$ and (f) $R_m = 5.0$.

is interesting to note that the gradients of both are very close for $\mathcal{L} \to 0$.

In summary, the numerical results are in good agreement with the expected contribution from the α-effect. As R_m increases, the description given by (7.109) becomes increasingly inaccurate, as expected.

7.8 Summary

In this chapter, we construct the problem of large-scale fields in the 1:1:1 ABC flow through the use of Bloch waves, as an alternative to the definition (7.66) of Galanti *et al.* (1992). We build the problem for general $\boldsymbol{K}$, though we eventually only simulate for $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, as exploring all possible $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$ would be computationally costly, even after considering the restrictions on these wavenumbers. Despite not numerically imposing the symmetries, we spend some time detailing the symmetries and representations of large-scale field, particularly for the case $\boldsymbol{K} = (\mathcal{L}, 0, 0)$, with the relevant numerical methods also given. Overall, the Bloch wave formulation is more fruitful as it isolates the most effective eigenmode through

imposing a fixed field scale, as opposed to GPS which numerically resolves multiple modes but cannot specify which is most effective. Moreover, we are able to rescale the magnetic field without increasing the computational cost from that of the original problem.

Numerically, we solve the large-scale problem for R_m up to 10^3, with $\mathcal{L}$ from 0 to 0.5 in steps of 0.01. We see the largest growth rate of $\gamma = 0.144$ with $\mathcal{L} = 0.36$ at $R_m = 20$. On increasing R_m, γ steadily decreases while the corresponding optimal $\mathcal{L}$ increases. For low R_m, there is a single clear peak near to $\mathcal{L} = 0.3$, though by $R_m = 30$, a second peak emerges, suggesting that there are multiple effective scales for the magnetic field. On increasing R_m, the periods of the optimal field become shorter (the opposite is seen for the secondary scale peak) and between $R_m = 200$ and 260, the optimal scale $\mathcal{L}$ changes dramatically. In later analysis we see that this is due to the eigenmode branch for the higher $\mathcal{L}$ being overtaken in quick succession by the lower $\mathcal{L}$ branches. For $R_m > 260$, the optimal scale remains at $\mathcal{L} \approx 0.20$ and again, there is a single peak growth rate. The growth rate then decreases steadily up to $R_m = 10^3$.

The structures of these large-scale fields are not significantly different to those of the original problem, with cigars forming along the separatrices. However, we now see chains of cigars of alternating polarity along the separatrices, though these generally become more twisted and irregular for the larger $\mathcal{L}$. Although the fields cannot be classified as belonging to the symmetry classes of the original problem, as these no longer exist in the same sense, we note that the field structures generally resemble those of classes II and V for small $\mathcal{L}$, owing to the similarity of eigenmodes to the original problem.

We draw comparison with both the methods and the results of GPS. We find that our growth rates, for fields rescaled in only one direction, are significantly lower than those of GPS, who allow fields to be rescaled in all three spatial directions, suggesting that the optimum scale exists for some $\boldsymbol{K} = (\mathcal{L}, \mathcal{M}, \mathcal{N})$ with at least one of $\mathcal{M}$ and $\mathcal{N}$ non-zero. In view of the fact that GPS resolve many eigenmodes in a given simulation, we can expect our equivalent Bloch wave formulation to produce

at most the same growth rate and no higher; this provides limited validation for the growth rates we obtained.

A considerable portion of this chapter concerns determining the contribution of the α-effect for Bloch wave formulation. For the 1:1:1 ABC flow, we identify $\alpha = -1$ and for a single-direction Bloch wave, we establish an equation for the expected growth rate profile. We find good agreement between growth rates obtained numerically and the expected growth rates for $R_m \leq 0.1$. For increasing R_m, there is an increasingly large disagreement between the two, in line with the predicted growth rates being valid under our assumption that $R_m \ll 1$.

Chapter 8

Discussion

Why do you go away? So that you can come back. So that you can see the place you came from with new eyes and extra colours. And the people there see you differently, too. Coming back to where you started is not the same as never leaving.

Terry Pratchett

In this final chapter, we summarise the main findings and conclusions of the multiple investigations detailed in this thesis. We relate our achievements back to our initial goals and objectives, given in chapter 1, addressing whether or not they have been met. The major obstacles faced in solving these problems are also discussed, with a final brief discussion on the possible future directions of these investigations.

8.1 Results and conclusions

In this thesis, the classic problem of kinematic dynamo action in the 1:1:1 ABC flow is revisited, with the additional approach of imposing the flow's symmetries on

the magnetic field. We have been motivated by the exploratory study of Galloway and Frisch (1986), who discover two different types of structure and two distinct windows of dynamo action, and from subsequent investigations. The methods used in this thesis are inspired by the work of Arnold (1984), in which the symmetries of the 1:1:1 ABC flow are identified and utilised in simple analytical approximations of the evolution of a magnetic field driven by the flow. As the use of symmetries and irreducible representations is generally unfamiliar in the context of Fluid Dynamics and MHD, a good portion of this thesis concerns developing the analytical framework and the numerical methods to numerically solve the kinematic induction equation for each of the five independent symmetry classes.

In terms of our initial goals, we have successfully achieved almost all that we set out to do. For the 1:1:1 ABC dynamo, we identify the two distinct windows of dynamo action as classes II and V, showing that the steady field of class V dominates up to $R_m = 10^4$. Interestingly, however, we find that all symmetry classes are capable of dynamo action, given a sufficiently large R_m. The most effective fields (those of V, II and IV) are concentrated into cigar-like flux tubes; we identify these as fast dynamo candidates. The less effective fields (those of classes I and III) form in the principal vortices; these we identify as slow dynamo candidates. We infer from the growth rates obtained for $R_m \leq 10^4$ that these values of R_m are not sufficiently large to establish these fields as truly fast or slow. To address the problem in the high-R_m regime, parallelisation was achieved with a limited degree of optimisation. The codes are found to scale well with the number of nodes used, provided the resolution is sufficiently large. Most importantly, we have successfully demonstrated the methods for decomposing the linear 1:1:1 ABC dynamo, with the methodology in place to deal with other symmetric linear problems.

In addition to time stepping, we present an alternative approach to solving the linear dynamo problem. The Implicitly Restarted Arnoldi method (IRAM) is used to provide independent confirmation of the time stepping results; we aimed to reproduce as many results as possible, as it was unclear whether this approach would be successful. We succeed in reproducing the dominant time stepping eigenvalue

for $R_m \leq 100$ in all symmetry classes, and even explore subdominant modes, due to the Arnoldi method's capacity to locate multiple eigenvalues. We have not attained eigenvalues for $R_m > O(10^2)$, owing to the increasing difficulty of convergence as R_m is increased, though we demonstrate one of the strengths of this approach, in that multiple eigenmodes can be tracked past mode crossings within a symmetry class. We are even able to follow the eigenmode pairs beyond eigenvalue coalescences, as is demonstrated at $R_m \approx 190$ in class II. Additionally, we extend the use of the Arnoldi method to constructing the magnetic fields of the eigenvalues obtained, though we find that the resultant fields (estimated from the Ritz vectors) only loosely resemble those of the time stepping. When comparing the two approaches, the Arnoldi method is quicker and more accurate for $R_m = O(10)$ but quickly becomes computationally demanding and less predictable as R_m increases, so that time stepping remains the best approach available for general exploration of eigenvalues and magnetic field structures.

In the analogous hydrodynamic stability problem, we take the same approach of dividing the full problem into five symmetry classes. We classify the eigenmode branch of Galloway and Frisch (1987) as belonging to symmetry class V, though we are only able to follow this branch to $R_e \approx 160$, after which a new mode emerges (belonging to class I). We calculate eigenvalues for all classes up to $R_e = 10^3$ and see at least one mode change for every class, with all growth rates remarkably similar at $R_e = 10^3$. On comparing the structures of the instabilities, we find the cigar-like structures of the magnetic problem, along with weaker but still significant instabilities in the principal vortices. It is clear that although the symmetries determine the structures of the respective instabilities, the rate of growth of an instability is not necessarily dependent on the symmetries present, unlike in the magnetic problem for which the magnetic fields of the individual symmetry classes can grow at very different rates. We expect the same pattern of mode changes and growth rate jumps seen for $R_e \leq 10^3$ will continue for $R_e > 10^3$.

The final objective of this thesis is to investigate the problem of magnetic fields that have a larger period than the 1:1:1 ABC flow; so-called large-scale fields.

We draw comparison with Galanti *et al.* (1992), whose methods we improve upon through the use of Bloch waves. Like Galanti *et al.*, we observe that magnetic fields of longer period than the $[0, 2\pi]^3$ domain are amplified more effectively, with massive increases in dynamo growth rate. The Bloch wave formulation allows us to accurately locate the most effective magnetic field scale, which we attain for $R_m \leq 10^3$. However, owing to the large parameter-space, we restrict our attention to rescaling the field in only the x-direction. The optimal scale changes dramatically for $R_m < 260$ but for $R_m \geq 260$, we ascertain that a magnetic field with 5 times the period of the flow (in the x-direction) makes the most effective dynamo. We also establish whether our growth rates resemble those predicted by the α-effect and derive an equation that describes the α-effect's contribution in terms of R_m and Bloch wavenumber $\mathcal{L}$. We see excellent agreement between expected and numerically obtained growth rates for R_m up to 0.1, good agreement for $R_m = 0.25$ and increasingly poor agreement as R_m increases further. This is entirely expected, as the derivation of the α-effect's contribution makes the assumption that $R_m \ll 1$. Though we develop the numerical methods for inclusion of symmetries of such large-scale fields, we choose to solve the full (non-reduced) problem in the interest of obtaining more general results. On comparing growth rates with those of Galanti *et al.* (1992), we observe that even our fastest growing fields are sub-optimal, indicating that for at least low-R_m regimes, the optimal field involves rescaling in at least two spatial directions, if not all three.

8.2 Obstacles and limitations

On investigating the various problems, we have identified problematic cases specific to the methods we adopted. In the time stepping results, we are unable to automatically calculate the eigenvalues because of the unpredictability of ranges of R_m in which the numerous difficult cases arise. Such cases include approaching (leaving) eigenvalue coalescences (decoalescences) where the oscillatory fields can have extraordinarily long periods, making their eigenvalues difficult to resolve

properly. Additionally, mode crossings can be problematic when the growth rates of two modes are very close. With the crossing of two oscillating modes, such as those in the high-R_m regime for classes II and IV, it takes an impractical length of time for the field to settle to one mode or another, and we often see overlapping of two frequencies in the oscillations. We see this in class III for $R_m > 1000$, with very complicated patterns of oscillations that settle over increasingly substantial timescales; we were thus unable to resolve class III in the high-R_m regime. In the crossing of an oscillatory and steady mode, we have a similar problem, in that oscillations may take so long to be damped that even over long timescales, it appears that the dominant mode is the oscillatory one; we see this at the crossing in class II at $R_m = 190$. Compounding such cases is the existence of spurious eigenmodes, found only in the results of the symmetry-reduced codes and appearing when mode values of heavily decaying eigenmodes approach machine precision. The spurious modes are always steady and have a decay rate that is much lower than the eigenmodes they supersede. On testing whether the spurious fields adhere to the symmetry requirements of the relevant class (as it appears for at least classes II and III), we find that symmetry-breaking has occurred. Furthermore, the Arnoldi method is not able to locate eigenvalues that correspond to these spurious eigenmodes, supporting the idea that they are erroneous. In an effort to combat the appearance of such eigenmodes, we rescale the field before it becomes sufficiently small and use the restrictions calculated for the Arnoldi method to impose the specific symmetries on the Fourier modes. However, we find that the appearance of the spurious eigenmode is merely delayed as opposed to avoided entirely. We speculate that rounding error builds up in such a way as to destroy the symmetries of modes within the fundamental domain, with this being avoided in the Arnoldi method by the mode restrictions imposed. This requires further investigation.

The relative success of the Arnoldi method is marred by the difficulty in ascertaining stable parameters n_{bv} and maximum iterations required. Because of the complicated underlying dynamics of the numerical problem, small changes in R_m can lead to large jumps in n_{bv}, making it very difficult to estimate the minimal

number of basis vectors required for a given simulation. Moreover, if n_{bv} is set too low, the routine may still locate eigenvalues but not those with the largest real part (or magnitude). The NAG user's manual states that 300 iterations is sufficient for most problems but we find that for $R_m \geq 200$, approximately 2500 to 3000 iterations are required to calculate 10 eigenvalues, though again there is no clear relationship between R_m and minimum iterations. These uncertainties seem to suggest that the Arnoldi method is not well suited for resolving the ABC dynamo problem, allowing us to investigate only a very limited range of R_m.

The major limiting factor of these investigations is computing power. We have access to a computer cluster of approximately thirty networked machines, which are used for running multiple serial codes, and to the University of Exeter supercomputer 'Zen' for parallel simulations. This allows for a substantial quantity of serial simulations, such as those for the large-scale field problem where 51 simulations are run for every R_m studied. For the main dynamo problem, however, the larger codes are parallelised as the power-law relationship between computing time and R_m means that desktop computing power is no longer sufficient. Even so, access to Zen is restricted on account of the large number of users and very limited resources. We do not manage to achieve the same level of R_m as Bouya and Dormy (2013), though they have access to a far larger supercomputer. We also are not able to invest much time in optimising the parallel codes, as parallelising the original dynamo problem is a secondary objective.

8.3 Future work

First of all, the successful demonstration of deconstructing the linear 1:1:1 ABC dynamo in its constituent symmetry classes, opens up the opportunity of using these methods for other symmetric flows. Immediate examples are the Kolmogorov flow (Galloway and Proctor, 1992) and the G.O. Roberts flow (Roberts, 1972), though these methods are applicable to any flow with discrete symmetries.

Secondly, there are several interesting questions raised, in the original dynamo

problem, that would benefit from further analysis. One is the behaviour of the symmetry classes at $R_m > 10^4$ and whether it is possible that another mode crossing occurs between classes II and V; there are some indications of competing modes at $R_m = O(10^5)$ (E. Dormy, private communication) that can be better determined (with greater computational resources) by the symmetry-reduced codes. In addition, field structures in chaotic regions are a potential wealth of information into how the Lagrangian chaos amplifies the fields.

Thirdly, the weaknesses of the Arnoldi solver could potentially be addressed. With more time and resources, a more detailed investigation into the convergence criteria of the Arnoldi method could be undertaken and a wider range of R_m investigated. This method could also be extended to the fluid stability and large-scale field problems to provide an alternative method of verifying the results of these problems. In the wider scope, the Arnoldi method could be adopted for completely different hydrodynamic or magnetohydrodynamic problems, provided that they are linear and have distinct eigenvalues.

Finally, we have only explored a subset of the possible large-scale fields using the Bloch wave formulation. Given more time and computing power, it would be useful to investigate the effect of two- and three-dimensional Bloch waves in order to ascertain the most effective dynamo scales. We envisage this as an optimisation problem at each R_m, given that the growth rate for varying $\mathcal{L}, \mathcal{M}$ and $\mathcal{N}$ is generally a smooth function in the absence of mode changes. More generally, Bloch waves could be adopted in any fluid or magnetic problem where one field is defined on scale that varies relative to that of another field.

Appendix A

Irreducible representations for symmetry classes III–V

The irreducible representations that we choose to use for symmetry classes IV and V are given below, with the upper sign for α as IV and the lower sign for α as V.

$$\boldsymbol{M}^{\alpha}(\mathsf{a}) = \pm\begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{a}^2) = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{a}^3) = \pm\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{a}) = \pm\begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{a}^2) = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{a}^3) = \pm\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{b}) = \pm\begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{b}^2) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{b}^3) = \pm\begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{c}) = \pm\begin{pmatrix} 0 & 0 & -1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{c}^2) = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{c}^3) = \pm\begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ -1 & 0 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{d}) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\alpha}(\mathsf{d}^2) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\alpha}(\mathsf{e}) = \begin{pmatrix} 0 & 0 & -1 \\ -1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{e}^2) = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & 1 \\ -1 & 0 & 0 \end{pmatrix}, \qquad \boldsymbol{M}^{\alpha}(\mathsf{f}) = \begin{pmatrix} 0 & 0 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{f}^2) = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ 1 & 0 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{g}) = \begin{pmatrix} 0 & 0 & -1 \\ 1 & 0 & 0 \\ 0 & -1 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{g}^2) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & -1 \\ -1 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{h}) = \pm\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & -1 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{j}) = \pm\begin{pmatrix} 0 & 0 & -1 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{k}) = \pm\begin{pmatrix} 0 & -1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{l}) = \pm\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix},$$

$$\boldsymbol{M}^{\alpha}(\mathsf{m}) = \pm\begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{n}) = \pm\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \boldsymbol{M}^{\alpha}(\mathsf{i}) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The irreducible representations that we use for symmetry class III are given explicitly by

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{i}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{a}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{b}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{c}^2) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{a}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{a}^3) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{k}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{n}) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{b}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{b}^3) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{h}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{l}) = \frac{1}{2}\begin{pmatrix} 1 & -\sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{c}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{c}^3) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{j}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{m}) = \frac{1}{2}\begin{pmatrix} 1 & \sqrt{3} \\ \sqrt{3} & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{d}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{e}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{f}) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{g}) = \frac{1}{2}\begin{pmatrix} -1 & \sqrt{3} \\ -\sqrt{3} & -1 \end{pmatrix},$$

$$\boldsymbol{M}^{\mathrm{III}}(\mathsf{d}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{e}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{f}^2) = \boldsymbol{M}^{\mathrm{III}}(\mathsf{g}^2) = \frac{1}{2}\begin{pmatrix} -1 & -\sqrt{3} \\ \sqrt{3} & -1 \end{pmatrix}.$$

Appendix B

Degrees of freedom for symmetry classes I–V

We work with the modes in $\mathcal{G}_N$ given in (4.31) and here we let M be the number of degrees of freedom for a given value of N. In other words M is the minimum number of independent complex quantities needed to specify any field $\boldsymbol{b}$ in $\mathcal{G}_N$ for symmetry classes I and II, and to specify the fields $\boldsymbol{b}_j$ for classes III–V. In each case we include the constraint that $\nabla \cdot \boldsymbol{b} = 0$, but we do not impose Hermitian symmetry.

We consider a slice of constant l depicted in figure 4.2. We first note that in view of the discussion of section 4.2.2, the left-hand column $\boldsymbol{b}_{l,0,n}$, $0 < n < l$ and top row $\boldsymbol{b}_{l,m,l}$, $0 < m < l$ can be obtained from $\boldsymbol{b}_{l,n,0}$ and $\boldsymbol{b}_{l,l,m}$ respectively, using (4.19) and (4.20). The remaining modes are treated as follows for $\boldsymbol{b} \equiv \boldsymbol{b}_1$ in classes I and II, and the $\boldsymbol{b}_j$ for III–V:

- Modes $(\boldsymbol{b}_j)_{l,m,n}$, $(\boldsymbol{b}_j)_{l,m,0}$ and $(\boldsymbol{b}_j)_{l,l,n}$ with $0 < m, n < l$ are unconstrained and each corresponds to 2 degrees of freedom using (4.60).

- The modes $(\boldsymbol{b}_j)_{l,l,l}$ map to themselves under d and this constrains $\boldsymbol{b}_{l,l,l} = \boldsymbol{0}$

for classes I and II. For III, these modes are constrained to the form

$$(\boldsymbol{b}_1)_{l,l,l} =(2\alpha, -\alpha + \sqrt{3}\beta, -\alpha + \sqrt{3}\beta), \tag{B.1}$$

$$(\boldsymbol{b}_2)_{l,l,l} =(2\beta, -\sqrt{3}\alpha - \beta, \sqrt{3}\alpha - \beta). \tag{B.2}$$

For IV and V the fields must take the form

$$(\boldsymbol{b}_1)_{l,l,l} = (\alpha, -\alpha-\beta, \beta), \quad (\boldsymbol{b}_2)_{l,l,l} = (\beta, \alpha, -\alpha-\beta), \quad (\boldsymbol{b}_3)_{l,l,l} = (-\alpha-\beta, \beta, \alpha). \tag{B.3}$$

- The modes $(\boldsymbol{b}_j)_{l,l,0}$ map to themselves under k and this constrains $\boldsymbol{b}_{l,l,0} = \boldsymbol{0}$ for I and

$$\boldsymbol{b}_{l,l,0} = (\alpha, -\alpha, \beta) \tag{B.4}$$

for II. For class III, however, the two fields take the form

$$(\boldsymbol{b}_1)_{l,l,0} = (\alpha, -\alpha, \beta), \qquad (\boldsymbol{b}_2)_{l,l,0} = \boldsymbol{0} \tag{B.5}$$

For IV the fields must take the form

$$(\boldsymbol{b}_1)_{l,l,0} = (\boldsymbol{b}_2)_{l,l,0} = (\alpha, -\alpha, \beta), \qquad (\boldsymbol{b}_3)_{l,l,0} = \boldsymbol{0}. \tag{B.6}$$

and for V

$$(\boldsymbol{b}_1)_{l,l,0} = (\alpha, -\alpha, \gamma), \qquad (\boldsymbol{b}_2)_{l,l,0} = (-\alpha, \alpha, -\gamma), \qquad (\boldsymbol{b}_3)_{l,l,0} = (\beta, -\beta, \delta). \tag{B.7}$$

- The modes $(\boldsymbol{b}_j)_{l,0,0}$ map to themselves under b. For classes I (upper sign) and II (lower sign) this constrains $\boldsymbol{b}_{l,0,0} = \boldsymbol{0}$ for $l = 2k$ even, and

$$\boldsymbol{b}_{l,0,0} = (0, \alpha, \mp(-1)^k \mathrm{i}\alpha), \tag{B.8}$$

for $l = 2k-1$ odd. For class III, the fields are of the form

$$(\boldsymbol{b}_1)_{l,0,0} = \mathbf{0}, \qquad (\boldsymbol{b}_2)_{l,0,0} = \mathbf{0} \tag{B.9}$$

for $l = 2k$ even and

$$(\boldsymbol{b}_1)_{l,0,0} = (0, 2\alpha, (-1)^k \mathrm{i}(\sqrt{3}\beta - \alpha)), \qquad (\boldsymbol{b}_2)_{l,0,0} = (0, 2\beta, (-1)^k \mathrm{i}(\sqrt{3}\alpha + \beta)) \tag{B.10}$$

for $l = 2k-1$ odd. For IV (upper sign) and V (lower sign) the fields must take the form

$$(\boldsymbol{b}_1)_{l,0,0} = \mathbf{0}, \qquad (\boldsymbol{b}_2)_{l,0,0} = (0, \alpha, \pm(-1)^k \beta), \qquad (\boldsymbol{b}_3)_{l,0,0} = (0, \beta, \mp(-1)^k \alpha) \tag{B.11}$$

for $l = 2k$ even, and

$$(\boldsymbol{b}_1)_{l,0,0} = (0, \alpha, \pm(-1)^k \mathrm{i}\alpha), \qquad (\boldsymbol{b}_2)_{l,0,0} = \mathbf{0}, \qquad (\boldsymbol{b}_3)_{l,0,0} = \mathbf{0} \tag{B.12}$$

for $l = 2k-1$ odd.

Counting of modes gives the the total number of degrees of freedom as

$$M = \tfrac{1}{6}(4N^3 - 6N^2 - 7N + 12), \tag{B.13}$$

$$M = \tfrac{1}{6}(4N^3 - 6N^2 + 5N), \tag{B.14}$$

$$M = \tfrac{1}{3}(4N^3 - 6N^2 + 5N), \tag{B.15}$$

$$M = \tfrac{1}{2}(4N^3 - 6N^2 + N), \tag{B.16}$$

$$M = \tfrac{1}{2}(4N^3 - 6N^2 + 5N - 4), \tag{B.17}$$

for I, II, III, IV and V respectively for N even, while for N odd,

$$M = \tfrac{1}{6}(4N^3 - 6N^2 - 7N + 9), \tag{B.18}$$

$$M = \tfrac{1}{6}(4N^3 - 6N^2 + 5N - 3), \tag{B.19}$$

$$M = \tfrac{1}{3}(4N^3 - 6N^2 + 5N - 3), \tag{B.20}$$

$$M = \tfrac{1}{2}(4N^3 - 6N^2 + N + 1), \tag{B.21}$$

$$M = \tfrac{1}{2}(4N^3 - 6N^2 + 5N - 3). \tag{B.22}$$

Bibliography

Alexakis, A. 2011 Searching for the fastest dynamo: laminar ABC flows. *Phys. Rev. E* **84**, 026321.

Alfvén, H. 1950 Discussion of the origin of the terrestrial and solar magnetic fields. *Tellus* **2**, 74–82.

Archontis, V. and Dorch, B. 1999 Numerical simulations of dynamos associated with ABC flows *Workshop on Stellar Dynamos, ASP Conference Series* (ed. Nunez, M. and Ferriz-Mas, A.), vol. 178, pp. 1–11.

Archontis, V., Dorch, S.B.F and Nordlund, A.A. 2003 Numerical simulations of kinematic dynamo action. *Astron. & Astrophys.* **397**, 393–399.

Arnold, V.I. 1965 Sur la topologie des écoulements stationnaires des fluides parfaits. *C.R. Acad. Sci. Paris* **261**, 17–20.

Arnold, V.I. 1984 On the evolution of a magnetic feld under the action of transport and diffusion. In *Some questions in modern analysis* (ed. V.M. Tikhomirov), *Mekh.-Mat. Fak. Moskov. Gos. Univ.*, pp. 8–21. Translated in *Fourteen papers translated from the Russian*, *Am. Math. Soc. Transl.*, series 2, vol. 137, pp. 119–129 (1987).

Arnold, V.I. and Korkina, E.I. 1983 The growth of a magnetic field in a three-dimensional steady incompressible flow. *Vest. Mosk. Un. Ta. Ser. 1, Matehm. Mekh.*, no. 3, 43–46.

Arnold, V.I., Zeldovich, Y.B., Ruzmaikin, A.A. and Sokolov, D.D. 1981 A magnetic

field in a stationary flow with stretching in Riemannian space. *Sov. Phys. JETP* **54**, 1083–1086.

Beltrami, E.1889 *Opera matematiche*, vol. 4, p. 304.

Bouya, I. and Dormy, E. 2013 Revisiting the ABC flow dynamo. *Phys. Fluids* **25**, 037103

Childress, S. 1970 New solutions of the kinematic dynamo problem. *J. Math. Phys.* **11**, 3063–3076.

Childress, S. and Gilbert, A.D. 1995 *Stretch, twist, fold: the fast dynamo.* Springer-Verlag

Clyne, J., Mininni, P., Norton, A., and Rast, M. 2007 Interactive desktop analysis of high resolution simulations: application to turbulent plume dynamics and current sheet formation. *New J. Phys.* **9**, 301.

Cowling T.G. 1933 The Magnetic Field of Sunspots. *Mon. Not. R. Astron. Soc.* **94**, 39–48.

Dombre, T., Frisch, U., Greene, J.M., Hénon, M., Mehr, A. and Soward, A.M. 1986 Chaotic streamlines in the ABC flows. *J. Fluid Mech.* **167**, 353–391.

Dorch, S. B. F. 2000 On the Structure of the Magnetic Field in a Kinematic ABC Flow Dynamo *Physica Scripta* **61**, 717–722.

Finn, J.M. and Ott, E. 1988 Chaotic flows and magnetic dynamos *Phys. Rev. Lett.* **60**, 760–763.

Finn, J.M. and Ott, E. 1988 Chaotic flows and fast magnetic dynamos *Phys. Fluids* **31**, 2992–3011.

Gailitis, A., Lielausis, O., Platacis, E., Gerbeth, G. and Stefani, F. 2003 The Riga Dynamo Experiment. *Surveys in Geophysics* **24**, 247–267.

Galanti, B., Pouquet, A. and Sulem, P.-L. 1992 Linear and nonlinear dynamos associated with ABC flows. *Geophys. Astrophys. Fluid Dynam.* **66**, 183–208.

Galanti, B., Pouquet, A. and Sulem, P.-L. 1993 Influence of the Period of an ABC Flow on its Dynamo Action. In: *Proctor, M.R.E., Matthews, P.C., Rucklidge, A.M. (Eds.), Solar and Planetary Dynamos*. Cambridge University Press

Galloway, D. and Frisch, U. 1986 Dynamo action in a family of flows with chaotic streamlines. *Geophys. Astrophys. Fluid Dynam.* **36**, 53–83.

Galloway, D.J. and Frisch, U. 1987 A note on the stability of a family of space-periodic Beltrami flows. *J. Fluid Mech.* **180**, 557–564.

Galloway, D.J. and O'Brian, N.R. 1994 Numerical calculations of dynamos for ABC and related flows In: *Solar and planetary dynamos* (ed. Proctor, M.R.E., Matthews, P.C. and Rucklidge, A.M.), Cambridge University Press

Galloway, D.J. and Proctor, M.R.E. 1992 Numerical calculations of fast dynamos in smooth velocity fields with realistic diffusion. *Nature* **356**, 691–693.

Gilbert, A.D. 1988 Fast dynamo action in the Ponomarenko dynamo. *Geophys. Astrophys. Fluid Dynam.* **44**, 214–258.

Gilbert, A.D. 1992 Magnetic field evolution in steady chaotic flows. *Phil. Trans. R. Soc. Lond. A* **339**, 627–656.

Gilbert, A. D. 2003 Dynamo Theory *Handbook of Mathematical Fluid Dynamics* (ed. Friedlander, S. and Serre, D.), **2**, 355–441. North-Holland

Gilbert, A.D., Ponty, Y. and Zheligovsky, V. 2011 Dissipative structures in a non-linear dynamo. *Geophys. Astrophys. Fluid Dyn.* **105**, 629–653.

Glatzmaier, G.A. and Roberts, P.H. 1995 A three-dimensional convective dynamo solution with rotating and finitely conducting inner core and mantle. *Phys. Earth Planet. Inter.* **91**, 63–75.

Golub, G.H. and Van Loan, C.F. 1996 *Matrix Computations*. Johns Hopkins University Press

Hamermesh, M. 1962 *Group theory and its application to physical problems.* Addison-Wesley

Hartmann, J. 1937 Theory of the laminar flow of an electrically conductive liquid in a homogeneous magnetic field. *K. Dan. Vidensk. Selsk. Mat. Fys. Medd.* **15**, 1–28.

Hénon, M. 1966 Sur la topologie des lignes de courant dans un cas particulier. *C.R. Acad. Sci. Paris A* **262**, 312–314.

Johnson, D. L. 2001 *Symmetries.* Springer-Verlag

Larmor, J. 1919 How could a rotating body such as the Sun become a magnet? *Reports of the British Association* **87** 159–160.

Lau, Y.-T. and Finn, J.M. 1993 Fast dynamos with finite resistivity in steady flows with stagnation points. *Phys. Fluids B* **5**, 365–375.

Lehoucq, R.B., Sorenson, D.C. and Yang, C. 1998 *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods.* SIAM, Philadelphia

Matthews, P.C. 1999 Dynamo action in simple convective flows. *Proc. R. Soc. Lond. A* **455**, 1829–1840.

Metropolis, N., Rota, G.-C. and Stein, J.A. 1991 Theory of symmetry classes. *Proc. Natl. Acad. Sci. USA* **88**, 8415–8419.

Moffatt, H.K. 1978 *Magnetic field generation in electrically conducting fluids.* Cambridge University Press.

Moffatt, H.K. 1985 Magnetostatic equilibria and analogous Euler flows of arbitrarily complex topology. Part 1, Fundamentals. *J. Fluid Mech.* **159**, 359–378.

Moffatt, H.K. 1986 Magnetostatic equilibria and analogous Euler flows of arbitrarily complex topology. Part 2, Stability considerations. *J. Fluid Mech.* **166**, 359–378.

Moffatt, H.K. and Proctor, M.R.E. 1985 Topological constraints associated with fast dynamo action. *J. Fluid Mech.* **154**, 493–507.

Monchaux, R., Berhanu, M., Aumaitre, S., Chiffaudel, A., Daviaud, F., Dubrulle, B., Ravelet, F., Fauve, S., Mordant, N., Petrelis, F., Bourgoin, M., Odier, P., Pinton, J.-F., Plihon, N. and Volk, R. 2009 The von Kármán Sodium experiment: Turbulent dynamical dynamos. *Phys. Fluids* **21**, 035108.

Newton, I., Cohen, I.B. and Whitman, A.M. 1999 *The Principia: Mathematical Principles of Natural Philosophy*. University of California Press.

Nussbaum, A. 1971 *Applied Group Theory for Chemists, Physicists and Engineers*. Prentice-Hall.

Oxford English Dictionary (2013), *symmetry, n.* [Online], http://www.oed.com/view/Entry/196242?redirectedFrom=symmetry (as of July 19, 2013, 18:30 GMT)

Parker, E.N. 1955 Hydromagnetic dynamo models, *Astrophys. J.* **122**, 293–314.

Podvigina, O. and Pouquet, A. 1994 On the non-linear stability of the 1:1:1 ABC flow. *Physica D* **75**, 471–508.

Podvigina, O. 1999 Spatially-periodic steady solutions to the three-dimensional Navier–Stokes equation with the ABC force. *Physica D.* **128** 250–272.

Ponomarenko, Y.B. 1973 On the theory of hydromagnetic dynamos. *Zh. Prikl. Mekh. & Tekh. Fiz.* **6**, 47–51.

Roberts, G.O. 1970 Spatially periodic dynamos. *Phil. Trans. R. Soc. Lond.* A **266**, 535–558.

Roberts, G.O. 1972 Dynamo Action of Fluid Motions with Two-Dimensional Periodicity. *Phil. Trans. R. Soc. Lond.* A **271**, 411–454.

Ruzmaikin, A.A., Sokoloff, D.D. and Shukurov, A.M. 1988 A hydromagnetic screw dynamo. *J. Fluid Mech.* **197**, 39–56.

Saad, Y. 2011 *Numerical Methods for Large Eigenvalue Problems*. SIAM, Philadelphia.

Soward, A.M. 1987 Fast dynamo action in a steady flow. *J. Fluid Mech.* **180**, 267–295.

Steenbeck M., Krause F. and Rädler, K.-H. 1966 A calculation of the mean electromotive force in an electrically conducting fluid in a turbulent motion, under the influence of Coriolis forces. *Z. Naturforsch* **21**a, 369–376.

Takezawa, S., Tamama, H., Sugawawa, K., Sakai, H., Matsuyama, C., Morita, H., Suzuki, H. and Ueyama, Y. 1995 Operation of the Thruster for Superconducting Electromagnetohydrodynamic Propulsion Ship "YAMATO 1". *Bulletin of the M.E.S.J* **23**, 46

Trefethen, L.N. and Bau, D. 1997 *Numerical linear algebra*. SIAM, Philadelphia.

Vainshtein, S.I. and Zeldovich, Y.B. 1972 Origin of magnetic fields in astrophysics (turbulent "dynamo" mechanisms). *Sov. Phys. Usp.* **15**, 159–172.

Zeldovich, Y.B. 1957 The magnetic field in the two-dimensional motion of a conducting turbulent fluid. *Sov. Phys. JETP* **4**, 460–462.