DOCTOR OF PHILOSOPHY

Braiding of Coronal Loops and Consequences for Coronal Heating

Ritchie, Miriam

Award date:
2018

DOCTOR OF PHILOSOPHY

Braiding of Coronal Loops and Consequences for Coronal Heating

Ritchie, Miriam

Award date:
2018

Awarding institution:
University of Dundee

Braiding of Coronal Loops and Consequences for Coronal Heating

By

Miriam Louisa Ritchie

A Thesis submitted for the degree of Doctor of Philosophy

Division of Mathematics

University of Dundee

Dundee

May 2018

Contents

List of figures

List of tables

Acknowledgements

Thank you to everyone at Dundee! I would like to thank both supervisors I worked with over the course of this project, Dr Antonia Wilmot-Smith and Professor Gunnar Hornig, for all of their help and support. Thank you also to my second supervisor, Dr David Pontin. To Dr Alex Russell, the helpful discussions and supportive chats are so appreciated, thank you. To Dr Simon Candelaresi, thank you also for your patience with many emails and questions and practical help with magneto! Thank you to Dr Nick Dawes - you were introduced to me as 'the most important person in the department' and I definitely relied on your expertise more than few times!

Thank you to Dr Herbert Fruchtl at St Andrews for continued access to the cluster, help when it was not working, and all of his work maintaining it for as long as possible! For helpful advice on topological entropy calculations, thank you to Dr Jean-Luc Thiffeault. For advice on Lare3d, thank you to Dr Tony Arber.

There are not enough thank yous I can make to my endlessly supportive and patient family! I would like to dedicate this to all of you, because it would not have happened without you. You are a constant source of strength and I am eternally grateful to you all. Mum, Dad and Hannah, you have all been there, on the phone, text, video call or in person, every minute of every day as I have gathered both myself and this thesis together over the last few years, and of course long before, through the happiest times

and the hardest. Duncan, even when we have been apart you have been there with a joke to keep me going! For so much fun, love, laughter and support, thank you. To my inspirational and loving grandparents, Gran, Papa, Wee Gran and Granda, knowing you have never stopped rooting for me has helped keep me going when it has been toughest. Gran and Papa, I hope you know. Thank you, all.

I have wanted to do this PhD in Solar Physics since I discovered the theory of Solar MHD looking at university prospectuses when I was in school. I will always remember the excitement of discovering what looked like beautiful light sculptures of glowing loops adorning our still mysterious star, and how lucky I am to have been able to do this.

Declaration

I declare that the following thesis is my own composition and that it has not been submitted before in application for a higher degree.

Miriam Louisa Ritchie

Certification

This is to certify that Miriam Louisa Ritchie has complied with all the requirements for the submission of this Doctor of Philosophy thesis to the University of Dundee.

Prof. Gunnar Hornig

Abstract

This thesis is inspired by one of the greatest mysteries surrounding our Sun: the coronal heating problem. Scientists have worked for decades to explain the extraordinary temperature increase which occurs as we move up through the solar atmosphere into the corona. Theories abound, including that of heating via reconnection, following the formation of current layers, which occur as the coronal field is braided by photospheric footpoint motions. In this work we designed different driving functions to act as braiding motions on a numerical photosphere, and sought to establish how each one affected overlying coronal fields. We used the concepts of topological entropy and helicity to categorise the drivers and applied them to a uniform field representing a simple coronal loop. We found that both complexity of driving and ability to inject helicity are key to the type of evolution which takes place, with direct consequences due to differences in these properties. The findings from the uniform field experiments also lend support to the idea that frequent nanoflares triggered along a loop could provide significant heating to the plasma. We then considered a field also involving parasitic polarities, creating a magnetic carpet, and attempted to find whether the same driving motions induced similar behaviour in a different environment. This research led us to challenge how we approach numerical work, and how the practical choices made in codes can drastically affect the success of a simulation. We found that one must consider carefully when choosing a numerical scheme to use, and be prepared potentially to change either elements of the project or the scheme itself.

Copyright

Chapter 1

(1) Figure 1.1 courtesy of NASA/SDO and the AIA, EVE, and HMI science teams, `hmi.stanford.edu`.

(2) Figure 1.2 courtesy Stanford-Lockheed Institute for Space Research Palo Alto, and NASA Goddard Space Flight Centre, Greenbelt, MD, `umbra.nascom.nasa.gov`.

(3) Figure 1.3 reprinted from:

E. R. Priest, D. I. Pontin, 'Three Dimensional Null Point Reconnection Regimes', Physics of Plasmas, December 2009, Volume 16, Issue 12, Page 12201, DOI 10.1063/1.3257901,

with the permission of AIP Publishing, license number 4315040928232.

(4) Figure 1.4 reprinted from:

E.R. Priest, D.W. Longcope, J. Heyvaerts, 'Coronal Heating at Separators and Separatrices', The Astrophysical Journal, May 2005, Volume 624, Issue 2, Pages 1057 – 1071, DOI 10.1086/429312.

©AAS. Reproduced with permission.

(5) Figure 1.5 reprinted from:

E. R. Priest, J. F. Heyvaerts, A. M. Title, 'A Flux-tube Tectonics Model for Solar Coronal Heating Driven by the Magnetic Carpet', The Astrophysical Journal, September 2002, Volume 576, Issue 1, Pages 533 – 551, DOI 10.1086/341539.

(6) Figure 1.8 courtesy of NASA/SDO and the AIA, EVE, and HMI science teams, `sdo.gsfc.nasa.gov`.

Chapter 2

(7) Figures 2.1 and 2.2 reprinted from:

J-L. Thiffeault, 'Braids of Entangled Particle Trajectories', Chaos, March 2010, Volume 20, Issue 1, DOI 10.1063/1.3262494,

with the permission of AIP Publishing, license number 4315050124268.

Chapter 3

(8) Figure 3.3 reprinted from:

Journal of Computational Physics, Volume 171, Issue 1, T. D. Arber, A. W. Longbottom, C. L. Gerrard, A. M. Milne, 'A Staggered Grid Lagrangian-Eulerian Remap Code for 3D MHD Simulations', Pages 151 – 181, July 2001, DOI 10.1006/*jcph*.2001.6780,

with permission from Elsevier, license number 4314881218611.

(9) Figures 3.5 and 3.6 from:

S. Candelaresi, D. I. Pontin, G. Hornig, 'Effects of Fieldline Topology on Energy Propagation in the Corona', The Astrophysical Journal, December 2016, Volume 832, Issue

2, DOI 10.3847/0004 – 637*X*/832/2/150.

Chapter 1

Introduction

1.1 Thesis Overview

This thesis is concerned with the decades old mystery of how the Sun's corona is heated to temperatures so much higher than the photosphere below. We look in particular at photospheric footpoint motions and how they can distort the overlying atmospheric fields of the Sun due to the typically low plasma beta nature of solar atmospheric plasma. We investigate the idea that this braiding of fields can lead to a build up of current sheets in the typically close to ideal conditions, which could be dissipated through magnetic reconnection, and thus release heat into chromospheric and coronal plasma. We consider the development of braiding and reconnection theories since their inception with Parker in 1972 [68] and construct driving functions which will simulate idealised photospheric motions to induce braiding of various magnetic field configurations. These drivers will themselves have different characteristics which we quantify in various ways. The evolution of test environments constructed with these features is modelled using the Lagrangian-Eulerian Remap code, Lare3d [2]. Ultimately we

are interested in the different types of photospheric motions' abilities to induce current sheet formation and significant reconnection and energy release. We examine the extent to which our different cases could realistically contribute to the heating of the solar corona.

In this introduction, we begin by describing the different regions of the Sun, in particular its atmosphere, and the model of magnetohydrodynamics we will use to understand it. We will briefly discuss the network of theories being developed to explain coronal heating and some of the literature providing evidence for and against. The theories of magnetic reconnection and helicity are discussed. Observations of braiding are also considered.

1.1.1 The Sun

We will first give a general overview of the structure and pertinent (to our work) processes at work on the Sun, to provide context for our area of interest. An excellent starting point for general information is [83], but we will only touch on a few properties here.

Internal Structure and Surface

Through the development of helioseismology, we can describe the internal regions of the Sun. Helioseismology uses observed vibrations of the solar surface to determine how waves are traveling through the interior, and therefore what the interior structure must be. This is in analogy to the study of seismology on Earth by using vibrations caused by earthquakes to learn about the Earth's interior. For example, patches of the photosphere are observed to be oscillating vertically over periods averaging 5 minutes

(known therefore as 5 minute oscillations), due to trapped acoustic waves, produced by convective motions. The waves must travel through and be reflected between different depths of plasma with different properties, such as density and temperature, to explain the vibrations observed. Therefore, studying the nature of the vibrations provides information on the internal structure, and there are many processes producing millions of vibrations. For a full review of the development and details of the theory see [51]. The Sun is revealed to be a highly complex and dynamic structure - in each region there are different dominant forces, energy transport mechanisms, temperatures and densities.

The core of our Sun is a nuclear fusion reactor with a radius of about 0.25 solar radii. Temperatures reach 15 MK with densities of 1.6×10^5 kg m^{-3}. Neutrinos released as a result of the fusion process pass through the rest of the solar interior and escape easily, however the gamma rays produced in the core interact significantly with the material in the radiative zone. This is a slightly cooler region with a depth of around 0.45 solar radii where energy transport is by radiative diffusion. This region is still very dense - photons can take millions of years to travel through it due to repeated absorption and re-emission. Their energy is decreased each time, thereby reducing the wavelength even to within the visible light spectrum, to be observed by us at the surface.

The transport mechanism in the final interior layer is by convection. This convection zone has a depth of 0.3 solar radii, with a temperature at its base of about 5×10^5 K. The density is estimated to be around 8×10^{-5} kg m^{-3}, a large decrease from that at the core.

Convective instability is at work in this region, contributing to the dynamism of the solar surface. A quantity of plasma at high temperatures at the base of the convective zone becomes buoyant and rises up to the surface, cooling as it does so. At the surface

the cooler plasma spreads out, and begins to fall back down into the interior to be re-heated again. This is what forms convection cells. The tops of these cells are referred to as granules, separated by intergranular lanes where cooled plasma is descending. These granules vary in diameter from less than 1 Mm to a few megametres, and in fact supergranules can measure tens of megametres. This continuous evolution contributes to the tangling of the overlying field.

The photosphere is the visible surface of the Sun, a highly active and variable layer. It has a depth of just a few hundred kilometres and temperatures varying from about 4000K to 6000K. This has been shown to be a highly dynamic region, by observations which have improved in resolution with ever improving telescopes. Ropes of strong magnetic field breaking through the surface result in the lower-temperature sunspots. Papers such as [17] discuss how observations from TRACE show that sunspots rotate, twisting loops associated with them. We have also known since, for example, the work of Leighton [53], that the granules observed on the surface involve substantial plasma flows over the photosphere. This flow and the breaking through of magnetic field creates a complex network on the photosphere called the magnetic carpet. Even in quiet Sun regions, where more obvious structures such as large coronal loops, prominences and sunspots are absent, smaller structures are created by the constant emergence, cancellation, coalescence and fragmentation of magnetic flux. Small loops and arcades and vertical field cover the entire solar surface, anchored in the surface, and all of this activity has been proposed as a source for coronal heating, in work using observations such as [26]. In addition to rotating sunspots, there is also observational evidence of other vortical plasma motions on the photosphere which could contribute to the dynamic field behaviour. For example, the authors in [11] discuss specific vortical motions which they refer to as whirlpools, where plasma is advected down under the surface by the rotational flow. Due to the small scale nature of these motions (less than a megametre in diameter), they observed the behaviour indirectly thanks to the

influence on nearby bright points, using data from the Swedish Solar Telescope. The authors in [5] also use bright points to track the action of vortical flows on the solar surface, which would distort the overlying fieldlines. This all leads us to consider placing a vortical driver function at our numerical photosphere, in order to further investigate the potential effect of such activity on the twisting on fieldlines and reconnection opportunities. We will elaborate on this in Chapter 3, Section 3.1.1.

Chromosphere, Transition Region and Corona

As we leave the surface of the Sun, we move through a layer a few thousand kilometres thick called the chromosphere. Temperatures in the lower chromosphere are of the order of 10^3 K, but as we move up this rises to 10^4 K. Next we arrive in a very thin layer (around 100 km) called the transition region, with a sharp temperature gradient, rising to hundreds of thousands of kelvin. Finally we reach the solar corona, a layer which reaches far out into the solar system. It is incredibly diffuse, however temperatures can reach millions of kelvin. This was first observed by [41] and [31] who observed from a solar spectrum that highly ionised iron was present in the corona, something that can only occur with extremely high temperatures. Naturally this was an unexpected discovery - how can temperatures rise so suddenly and steeply as we move away from the primary heat source? Extensive research has been conducted ever since to investigate numerous possible mechanisms which could be responsible, and this is the umbrella of investigations under which this thesis sits.

Magnetic Field, Atmospheric Structures and Dynamic Processes

The creation of the Sun's magnetic field is thought to be due to a solar dynamo. The

flowing, ionised material within the Sun generates electric currents, which in turn generates a magnetic field. It appears to follow specific, ever repeating cycles. Dynamo theory tries to explain how the many different types of flows and rotation in and of the Sun generate and wind up the magnetic field, giving rise to these cycles. Reading on the solar dynamo cycles can be found in [20] but we shall not discuss it further, other than the following brief description of its role in sunspot creation, since the plasma flows around sunspots also contribute to tangling of the overlying field.

Sunspots are regions of strong magnetic field (a few thousand Gauss) where the convective region field has been twisted and distorted until a loop breaks through the surface. These areas consist of a dark central region called the umbra (with a temperature of about 4000K) surrounded by a slightly hotter and brighter ring called the penumbra, and can measure tens of megametres across. The cool temperature of a sunspot (relative to the surrounding regions) is due to the magnetic field strength there. The interaction of magnetic fields with convective processes, known as magnetoconvection, is complex and more detail can be found in [88] and references therein. Briefly, the strong magnetic fields can inhibit the convection process which transports heat to the photosphere from the convection zone. If hotter plasma from the base of the convection zone is no longer able to rise to the surface and spread out horizontally as in other regions, then the sunspot is not continuously heated to such temperatures and is therefore cooler. These sunspots typically occur in pairs of positive and negative magnetic field, (field emerging into or exiting from above the surface, respectively), with one polarity 'leading' in the Northern Hemisphere and the other in the Southern Hemisphere. The number of sunspots follows a continuous cycle - over 11 years the number of sunspots will go from a maximum down to a minimum and back up to a maximum again. They also tend to appear closer and closer to the equator as the cycle progresses. As field is twisted and folded within the Sun, more loops break through the surface, resulting in more sunspots. After 11 years it seems that the global field reaches some critical

point, resets, and the process begins again. Furthermore, every 11 years the magnetic North and South poles reverse, leading to a switch in which sunspot polarity is leading in each hemisphere. Therefore the complete sunspot cycle occurs over the course of 22 years.

In addition to sunspots we observe features such as coronal loops, spicules, filaments, prominences, coronal holes and the solar wind. In this thesis we are predominantly concerned with coronal loops. These are ropes of hot, glowing plasma flux tubes which follow the trajectories of field lines (since particles moving in an electromagnetic field are under the Lorentz force, which causes them to spiral around and along fieldlines), exiting and re-entering the photosphere at footpoints. Lengths vary from the very small up to hundreds of megametres. Temperatures can reach millions of kelvin. They have lifetimes ranging from a few hours to a day or two, moving through different stages labelled the rise, main and decay stages by the authors of [56]. They are long-lived in comparison to typical coronal cooling times so are typically hotter than their surroundings. Therefore many theories for coronal heating consider the heating mechanisms of loops to be the subject of much interest. At this point direct observations of the fine detail of loops and evidence of braiding are limited - see Subsection 1.2.3 for more on observations. However, evidence that loops consist of many individual strands of plasma too fine to be resolved currently is given in many papers, such as [97]. In the absence of any external driving from footpoint motions, a coronal loop would relax to an equilibirum. Many coronal loops sit at heights above the solar surface where the plasma beta is small, (at the surface plasma pressure dominates, and also at a certain height above it as magnetic pressure begins to fall off faster than plasma density, but many coronal features lie in between theses depths in a region where magnetic pressure dominates), so this equilibrium is said to be force-free; this means that magnetic forces dominate over gas pressure and gravity and the forces of magnetic tension and magnetic pressure are balanced. Also, since in general magnetic diffusion in the corona

is negligible (we elaborate on this in Section 1.1.6.1) plasma elements are confined to fieldlines, so there are no plasma flows between the strands comprising the loop even though the strands themselves may be being twisted and tangled together by the photospheric flows in which they are anchored. In this thesis we will be interested in situations where the equilibrium of loops is compromised by photospheric motions with varying characteristics and exchange of plasma between strands becomes possible via reconnection. Further general reading on loops can be found in [94].

Under certain circumstances which we will detail in Section 1.1.6.1, the topology of coronal loop elements can change, releasing free magnetic energy, into a different configuration of lower energy. This is known as the process of magnetic reconnection. In some cases these restructuring events are so large and release such large amounts amounts of energy that we observe a large solar flare, which can trigger a coronal mass ejection. However flares have also been noted, directly when large enough and indirectly otherwise, on a range of scales and are predicted to occur on currently unobservable scales too.

The authors in [42] present a review of flare research and we refer now to some of the discussions there and in references therein. Large flares, releasing energies up to 10^{26} Joules/ 10^{33} ergs, are seen to occur infrequently, mostly during solar maximum. Smaller flares called microflares release around 10^{20} J/10^{27} ergs. Smaller still are so called nanoflares, with energies of around 10^{18} J/10^{24} ergs. The existence of these types of flares were theoretically proposed by Parker in [71] who proposed that heating may also be provided by numerous smaller events. He suggested that the hot corona we see in X-ray observations is maintained by many of these lower energy, localised heating events, taking place at the sites of small current sheets constantly being formed by the continuous tangling of fieldlines by photospheric motions. For example, [110] looks at changes in X-ray emission as evidence of this and [102] considers rapid

changes in 'moss' regions to be a sign of nanoflares. Further, there is the suggestion of picoflares - even smaller events releasing around 10^{14}J/10^{21} ergs, which [89] claims to have observed in radio data.

Research considering the role of events of different energies has been carried out by many authors, such as [77], following on from proposals by [47]. The idea is to assume that the distribution of events follows a curve given by a power law for the heating power:

$$\int_{E_{min}}^{E_{max}} fE\mathrm{d}E = \int_{E_{min}}^{E_{max}} f_0 E^{-\alpha} E\mathrm{d}E$$

where E_{min} is the energy of the smallest event, E_{max} the largest event and f the frequency of events where f_0 and α are constants. Calculating the integral gives:

$$Power = \frac{f_0}{2-\alpha}[E_{max}^{2-\alpha} - E_{min}^{2-\alpha}]$$

Assuming that the largest event energy is much bigger than the smallest event energy, then either;

$$\begin{aligned} \alpha < 2, \text{ so that } Power &\approx \frac{f_0}{2-\alpha} E_{max}^{2-\alpha} - \text{ so large events dominate the heating, or;} \\ \alpha > 2, \text{ so that } Power &\approx \frac{f_0}{\alpha-2} E_{min}^{2-\alpha} - \text{ so small events dominate the heating.} \end{aligned}$$

So the value of α, giving the slope of the power law, is crucial in determining whether more power is supplied to the corona through many smaller events, like nanoflares, or less frequent but higher energy events, like micro and occasional large flares. Many

authors have found this sort of power law relation, indicating smaller events happening with more frequency, but calculated different values for the power law index, so disagreeing on the roles of different flare categories (see references in [42], composites of estimates by several authors in [3]). There is also the issue of the smallest events being beyond our limit of resolution. We can only assume the power law extends in the same way below the limit of our perception. The conclusion of [75] is that whether we look at the full solar disk or zoom in on smaller regions, the flux in that area follows the same power law with frequency. Whichever scale you look on, high flux becomes less frequent at the same rate, (so magnetic feature complexity must driven by the same mechanism - either a dynamo acting the same way on all scales or surface motions dominating the structure of magnetic features). It follows that flares also occur according to this scale invariant trend. Whatever the size of the feature, smaller events are happening most frequently.

1.1.2 Magnetic Carpet

Research and observations have shown that the solar atmosphere has an incredibly complex structure, including the lower atmosphere, near the photosphere. Even in Quiet Sun regions, there is a highly complicated network of magnetic features from the surface up through the chromosphere - the magnetic carpet. The constant flows of plasma up from the convection zone and across the surface of the photosphere, threaded by the ever changing magnetic field breaking through from the solar interior, leads to a highly dynamic surface which continuously distorts the overlying structures. With the SOHO launch in 1995, scientists gained new insight into the magnetic carpet, realising how complex it is, and also how quickly emerging flux is distributed and recycled through the solar corona. The authors in [98] calculated that photospheric flux is replaced around every 40 hours. Taking flux emergence and cancellation and

reconnection into account, the authors of [22] also calculate that it takes only 1.4 hours for the entire coronal field to be reconnected. Clearly, incorporating a magnetic carpet field with its associated topological features should provide an abundance of heating opportunities, and applying different drivers may result in interesting current sheet formation.

The work of [26] is among many others which investigate how the magnetic carpet could be a prime candidate for coronal heating mechanisms, therefore modelling of the magnetic carpet is a large area of ongoing research.

Magnetograms of the solar surface show that the whole photosphere is covered in clusters of mixed polarity magnetic flux concentrations. They map out the strength and direction of the line of sight magnetic field. Black and white represent regions of strong magnetic field concentrations, where the line of sight field is directed either away from or towards us, respectively. The most obvious are features such as sunspots, but smaller mixed polarity clusters are present all the way down to observable scales and probably beyond. Concentrations of magnetic flux are constantly being transformed through four main processes: emergence, fragmentation, coalescence and cancellation. Ephemeral regions are clusters of bipolar pairs of flux, thought to originate from motions below the surface, which trigger their emergence. These ephemeral regions are thought to consist 90% of quiet Sun flux (see [85] and references within). Smaller clusters of magnetic flux are also formed from the remains of other magnetic structures, such as a decaying active region, forming so called network concentrations. Even smaller clusters are referred to as intra-network concentrations. All of these concentrations are subjected to horizontal plasma flows, sweeping them towards granular and super granular boundaries, where material is subjected to downflows back under the surface. As they are advected, these clusters can fragment into smaller groups, and coalesce with other groups,. Further, a bipolar pair can cancel with another bipolar

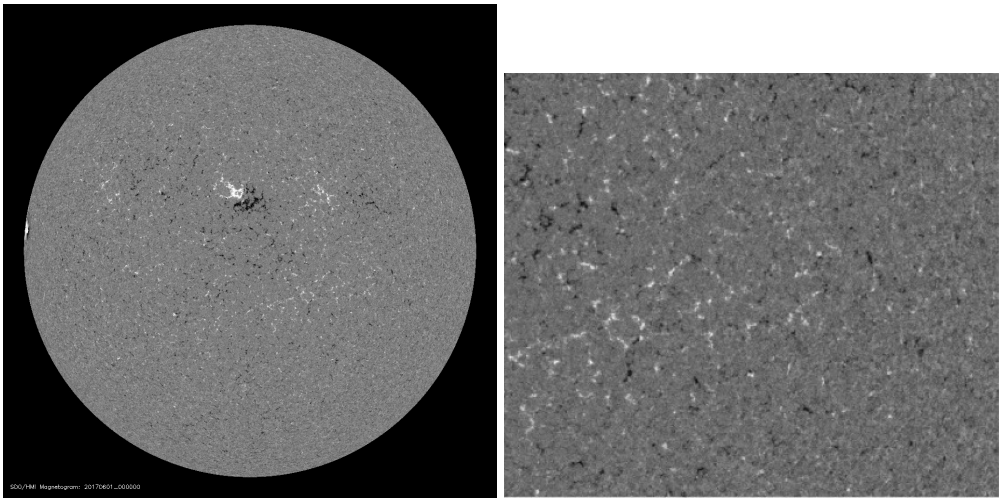

Figure 1.1: HMI magnetogram. Courtesy of NASA/SDO and AIA, HMI, and EVE science teams. This image is from SDO's HMI equipment, and is available from the HMI website, `hmi.stanford.edu`. It illustrates the varying magnetic field across the solar surface: in black areas the line of sight field is directed away from over viewpoint; in white areas it is pointing towards us.

pair. Figure 1.1 is a magnetogram from SDO's HMI apparatus and shows the variable polarity and strength of the magnetic flux across the surface. While we can see an active region near the centre of the disk, most of the rest of the disk is quiet, but speckled all over by field concentrations of varying scales. Zooming in conveys how variable the flux is, down to the smallest scales. We can in fact just make out some edges of granulation cells, outlined by flux concentrations which have been swept and gathered there by the plasma flows. Further reading can be found in [73] and references therein.

An illustration of the variable surface flux and how the atmospheric topology depends on it can be seen in Figure 1.2. The green base represents the magnetic carpet, the black and white pockets representing flux receding or emerging respectively. Fieldlines have been plotted using other collected data which trace some of the loops anchored in the carpet. These loops reach a range of heights, some remaining very low down in the atmosphere.

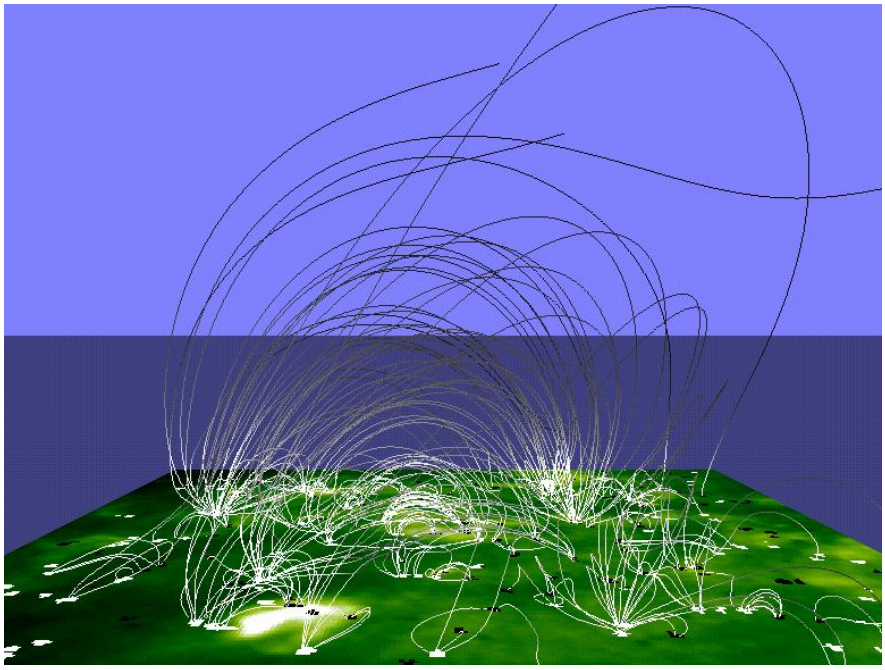

Figure 1.2: Figure from magnetic carpet simulations and observations: this is an image created by tracing fieldlines above the solar photosphere. The black and white patches are the regions of strong positive and negative line of sight field, and we can see how the overlying field is anchored at these concentrations. This image was created using data from SOHO and can be found at `umbra.nascom.nasa.gov`. Credit: Stanford-Lockheed Institute for Space Research Palo Alto, and NASA Goddard Space Flight Centre, Greenbelt, MD.

1.1.3 Magnetic Skeleton

The 'magnetic skeleton' consists of those fieldlines which when mapped, illustrate the topological features which are dictated by the characteristics of the magnetic carpet at that point in time. The structure of this skeleton and its connection to photospheric flux patterns are described by the authors of [85] as coronal tectonics. Everything is anchored in the photosphere and stems from the flux concentrations we see in the magnetograms. The largest loops can be anchored via multiple flux sources at each

footpoint, and the smallest loops via single sources. Even the individual strands consisting coronal loops can have their footpoints rooted in different sources. We are interested in how loops could be affected by the motions taking place on the surface and so the nature of field topology due to the magnetic carpet is of great interest. We proceed now to detail the fundamental elements which create this complex network.

Magnetic skeleton features can include various combinations and configurations of null points, separatrix surfaces and separator fieldlines. A great deal of research has shown that these features can be reconnection sites. Just two examples are the investigations of [76], who found reconnection occurring along separators, while [81] considers reconnection in a structure consisting a null sitting inside a separatrix dome (a concept we will consider ourselves in Chapter 3, Section 3.3.2.)

Null points are points where the magnetic field is zero, and the field around a null point has specific properties. The spine fieldline is an individual fieldline leading away or towards the null, while the surface created by the set of fieldlines running towards or away from the null together is called the fan. The exact geometry of the field around the null can vary - for example Figure 1.3 (taken from [87]) shows the fan and spine for a so called proper and improper radial null. Null points can provide reconnection opportunities - [87] propose that under stress from rotational motion, currents can build along the spin and/or fan, leading to reconnection and heating. Furthermore, [81] and the references therein discuss how reconnection at nulls can contribute significantly to dynamic events such as solar flares and CMEs, solar jets and particle acceleration.

Separatrix surfaces are the boundaries between topologically distinct regions. An example of a separatrix is the fan plan of a null point - it divides regions of differing magnetic connectivity. Separator fieldlines are those which lie along the intersection of separatrix surfaces, connecting null points. Both separatrices and separator curves

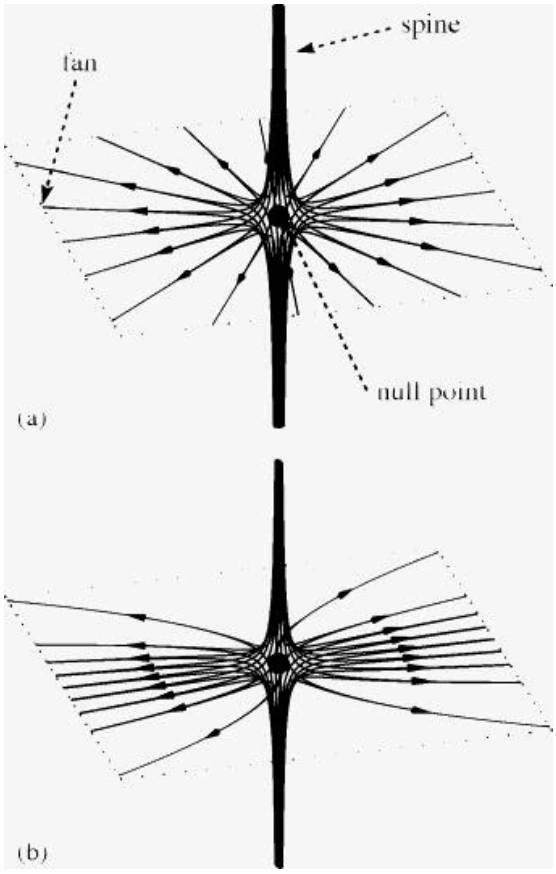

Figure 1.3: The fundamental elements of a magnetic nulls, showing the spine and fan. This is taken from [87]

are thought to be abundant in the corona, in the small scale, low altitude field and higher up in large scale structures. A diagram of topologies involving separatrices and separators is shown in Figure 1.4, taken from [86]. The first image in Figure 1.4 shows the field above the surface, constituting two dome shaped separatrix surfaces. Their intersection, the separator, is shown in bold. It connects two null points. The stars highlight points of opposite flux sitting on the photosphere in the different regions. The second imagines straightening a loop out so we have photosphere at both ends. This shows a scenario where an initial field has been stressed and magnetic energy has built up, such that the separatrices are planar surfaces and the separator takes the form of a 2D sheet. Figure 1.5 from [85] shows a further example of separatrices, this time

lying between nested coronal loops of different scales. Separatrices can also form inside loops between individual strands. The work of [85] concludes that current sheets form along separatrices and separators under the influence of photospheric motions and reconnection can occur. They also found that the nature of the motions dictated the position of the current sheets formed and their dimensions and strength.

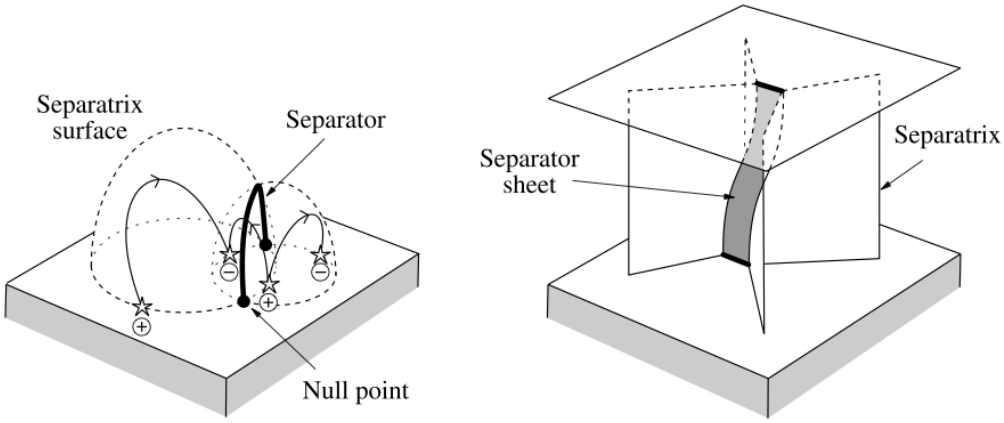

Figure 1.4: Diagram from [86], showing how the magnetic skeleton creates distinct regions.

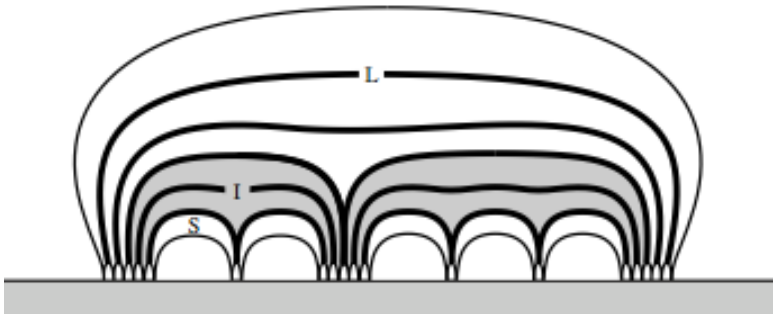

Figure 1.5: Again taken from [85], this diagram further illustrates how many separatice surfaces and separators can form many distinct regions in the field, with different scales, creating the highly variable magnetic carpet.

1.1.4 Magnetohydrodynamics

Magnetohydrodynamics, or MHD as we will generally refer to it from here on in, is a model describing a magnetised plasma. It is valid assuming typical plasma velocities are much less than the speed of light; we do not consider relativistic effects in the model. The MHD equations are constructed from Maxwell's equations, the fluid equations and Ohm's Law. Under the MHD approximation the displacement current term of Ampére's Law:

$$\nabla \times \mathbf{B} = \mu \mathbf{j} + \underbrace{\frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t}}_{\text{displacement current}},$$

can be neglected. With this approximation we can manipulate the aforementioned sets of equations to obtain the governing equations of magnetohydrodynamics:

$$\nabla \cdot \mathbf{B} = 0\text{: solenoidal constraint} \tag{1.1}$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0\text{: conservation of mass} \tag{1.2}$$

$$\rho\left(\frac{D\mathbf{v}}{Dt}\right) = -\nabla p + \mathbf{j} \times \mathbf{B} + \rho \mathbf{g} + \rho \nu \left(\nabla^2 \mathbf{v} + \frac{1}{3}\nabla(\nabla \cdot \mathbf{v})\right): \tag{1.3}$$

$$\text{equation of motion} \tag{1.4}$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) + \eta \nabla^2 \mathbf{B}\text{: induction equation} \tag{1.5}$$

$$\frac{\rho^\gamma}{\gamma - 1}\left\{\frac{D}{Dt}\left(\frac{p}{\rho^\gamma}\right)\right\} = -L\text{: energy equation} \tag{1.6}$$

$$\mathbf{j} = \frac{\nabla \times \mathbf{B}}{\mu}\text{: Ampére's Law} \tag{1.7}$$

$$p = \frac{1}{\tilde{\mu}}\rho R T\text{: equation of state- ideal gas law} \tag{1.8}$$

where **B** is the magnetic field, **v** is plasma velocity, **j** current density and **g** gravitational acceleration. Scalars p and ρ are plasma pressure and density. The variable η is the magnetic diffusivity $1/\sigma\mu_0$, where σ is the the conductivity. The constants are:

ν : coefficient of kinematic viscosity,

μ_0 : permeability of free space,

γ : ratio of specific heats, for us $= \frac{5}{3}$,

μ : magnetic permeability,

$\tilde{\mu}$: mean atomic weight,

R $= kN_A$: universal gas constant (Boltzmann's constant k, Avogadro's number N_A).

The variable L represents various processes, some of which the user may not wish to consider in a numerical simulation. For example, in many cases an adiabatic (no net energy loss or gain) environment is considered so this is taken to be 0. We expand upon it in the next section.

1.1.5 Coronal Heating Problem

Affecting the plasma energies are losses from optically thin radiation, energy transport via conduction, and contributions from Ohmic and viscous heating. Expanding L then we have:

$$-L = \underbrace{\nabla\cdot(\kappa\nabla T)}_{\text{isotropic conduction}} - \underbrace{\rho^2 Q(T)}_{\text{optically thin radiative loss}} + \underbrace{\frac{j^2}{\sigma}}_{\text{ohmic heating}} + \underbrace{H_\nu}_{\text{viscous heating}} \tag{1.9}$$

To explain coronal heating we must understand the processes which could result in ohmic and viscous heating, since they must balance radiative losses to result in the temperatures observed.

Paper [117] estimated that losses in the quiet chromosphere total around 4×10^6 erg cm^{-2}s^{-1}, 2×10^7 erg cm^{-2}s^{-1} in active chromosphere. In the corona, these losses are 3×10^5 erg cm^{-2}s^{-1} for quiet Sun and 1×10^7 erg cm^{-2}s^{-1} for active regions, so a heating mechanism must provide energy in excess of these requirements. Using observations, [116] estimated that the total energy input to the open magnetic field corona from non-radiative sources is $5 \pm 1 \times 10^5$ ergs cm^{-2}s^{-1}, on a par with the slightly earlier estimates.

1.1.5.1 Importance of Magnetic Field to Heating

In their paper [106], Vaiana, Krieger and Timothy analysed observations from newly developed satellites measuring X-ray emission and started to uncover the variation throughout the corona. Areas such as active regions, coronal loops and coronal holes and their relation to activity on the photosphere were discovered. It was first recognised that changes in the photospheric magnetic field had huge effect on the structure of coronal features. This paper helped open up a whole new frontier of research concerning the role of the magnetic field in the solar atmosphere and potential heating mechanisms. Paper [74] explains how the energy required to heat the solar atmosphere is sourced from the convection zone, through flux emergence and horizontal movement of footpoints. Theories for how this energy is then transferred to and dissipated in the chromosphere and corona generally fall under one of two categories - Alternating Current, (AC), or Direct Current, (DC), depending on the current driven by the mechanism under consideration.

AC versus DC and Proposed Mechanisms

Consider photospheric motions which cause perturbations in the field that propagate up through the atmosphere at the Alfvén speed $v_A = |B|/\sqrt{\mu_0 \rho}$ in that region. We class AC heating as that induced by motions which vary over timescales short compared to the Alfvén loop travel time. This is the time taken for a wave generated at a loop footpoint to travel the length of the loop. So a relatively fast photospheric motion would send up another wave before the previous wave has reached the other end of the loop, so the loop has not had time to adjust to an equilibrium within an Alfvén travel time. Alternating currents are flowing through the loop over one Alfvén time. A coronal field subject to AC motions is said to evolve dynamically. Typically AC theories involve wave heating.

Waves can be triggered by many different mechanisms - fast photospheric motions, reconnection, flares and CMEs. Research has explored the potential for waves to transport energy up through the solar atmosphere and indeed since observations have improved, more and more evidence of wave activity has been gathered. However only Alfvén waves are able to easily reach the corona - acoustic and magnetoacoustic waves may be easily damped or reflected before they reach the upper atmosphere. Consider that quiet Sun coronal fields may be dominated by vertical field. Alfvén waves propagate along fieldlines up through the atmosphere, while fast and slow waves may travel at an angle, such that two ends of a wavefront lie at different heights. As we move up through the chromosphere, the plasma density drops such that the Alfvén speed increases with height, jumping in particular across the transition region. As the magnetoacoustic wave travels, the part of the wavefront at the highest height is moving faster than the lower part, and the wave starts to turn. This effect is only intensified as the wave travels higher: the Alfvén speed continues to increase and the wavefront

tilts further until it has been reflected entirely and travels back down toads the surface, never reaching the corona in order to have the opportunity to release energy there. Phase mixing, resonant absorption and viscous effects are examples of mechanisms which could damp Alfvén waves and release their energy throughout the atmosphere. See the review of [74] and references therein for more on wave heating.

On the other hand, if currents are being generated by motions on longer timescales, then the field can react directly within an Alfvén time and we have a quasi-static evolution. The field evolves via a sequence of equilibria. This is referred to as Direct Current heating. In this case stresses may build up in the overlying field in the form of current sheets, which may lead to reconnection.

Current sheets can appear on many scales due to braiding, and emergence and cancellation of flux as a result of plasma rising through the photosphere. The details of reconnection will be discussed in Section 1.1.6.1, however, a general feature of magnetic reconnection is plasma heating. This is due to the release of free magnetic energy, which can go into generating electric fields, which in turn accelerates particles and generates thermal energy by Ohmic heating. In addition some of that free magnetic energy is also first converted to kinetic energy, which can result in viscous heating via shocks, and trigger further reconnection events, in an avalanche effect. For more information on the development of heating theories again see [74] and the references within.

Research suggests that different mechanisms falling under the AC and DC headings may be more or less influential in different atmospheric regions and also in the quiet Sun versus active regions. The authors in [27] proposed that Alfvén waves travelling

through the solar atmosphere could be largely responsible for heating the chromosphere and [59] went further to conclude that they could also provide the necessary heating to quiet Sun corona. In addition to Alfvén waves propagating up from the photosphere, [111] found that they could be generated higher up in the atmosphere too. The authors found that the vortical motions induced on the photosphere from convection are transferred to the overlying layers of plasma, creating 'tornadoes'. Such volumes of swirling plasma generate more waves higher up which can dissipate more energy. Paper [61] concludes after analytically considering a coronal loop model that DC heating is more important in the active region corona. Paper [84] supports the idea of active region coronal heating due mostly to turbulence induced by loop footpoint motions and subsequent reconnection.

Taking all of this into account it seems clear that the question of what heats the corona probably has many answers. The mechanisms investigated can be seen as coupled and not just acting in an isolated way. There are also factors such as the drop off in density from chromosphere to corona and the launching of the solar wind which any potential full model will also have to contend with. Therefore determining what is most important, if any one factor can be deemed so, is a formidable task.

1.1.6 Magnetic Reconnection and Helicity

1.1.6.1 Magnetic Reconnection

Magnetic reconnection is fundamental to our work. It is the process by which the connectivity of fieldlines changes and releases free magnetic energy. This leads to Ohmic dissipation, release of kinetic energy and acceleration of particles which can launch

flares. We consider a brief history of the theory, and how it can arise.

Theories of reconnection have been developed over decades. They began with papers such as [38] and [46] who considered the possibility that currents flowing at X-type magnetic null points could trigger flares and heating. Paper [24] added to the idea, concluding that thin current sheets could induce heating. The work of Dungey (see [28] among many others) was fundamental to the development of the theory - he was first to describe the breaking apart of fieldlines. Other landmarks include [67], [101] and [78] by Parker, Sweet and Petschek. In fact it was Parker that finally used the term reconnection. The Sweet-Parker model was the first MHD reconnection model. It considered reconnection in 2D due to current sheets at null points, which could trigger solar flares. However it was shown that in this initial set-up, the rate of reconnection (the rate at which flux is reconnected in time and space) was not fast enough to launch solar flares. Petschek's model on the other hand described a situation that does give rise to sufficiently fast reconnection. These were all still 2D models however, and the properties of reconnection were shown to be quite different in 3D. In 2D, reconnection only occurs at null points and the rate of reconnection depends only on the local electric field. In 3D, reconnection does not require nulls and the reconnection rate is obtained through integration of the electric field parallel to the fieldlines involved. Papers by Schindler and Hesse ([96], [44]) present a formalised model of general magnetic reconnection in 3D and can be consulted for more details.

To understand how reconnection may occur, consider the evolution in time of a magnetic field as described by the induction equation:

$$\frac{\partial \mathbf{B}}{\partial t} = \underbrace{\nabla \times (\mathbf{v} \times \mathbf{B})}_{\text{advection}} + \underbrace{\eta \nabla^2 \mathbf{B}}_{\text{diffusion}} \tag{1.10}$$

The advection term describes the movement of the field by plasma motion and the diffusion term relates to the smoothing out of any variation in the field. In general on the Sun the diffusion term is small. We can see this by letting B, l, t and v represent a typical field strength, length scale, timescale and velocity:

$$\frac{B}{t} = \frac{vB}{l} + \eta\frac{B}{l^2} \tag{1.11}$$

Let us look at the ratio of advection to diffusion terms, which is known as the magnetic Reynolds number, R_m:

$$\begin{aligned} R_m &= \frac{vB}{l} / \eta\frac{B}{l^2} \\ R_m &= \frac{lv}{\eta} \end{aligned} \tag{1.12}$$

Since length scales and velocity scales on the Sun are typically large, and $\eta = \frac{1}{\sigma\mu}$ is typically small, this means R_m is typically much larger than 1. In other words, advection has much more of an influence on the evolution of the field than diffusion (while quantities remain sufficiently smooth). On the other hand, in some special cases R_m can be smaller than one, in which case diffusion is important. Let us look at the limiting cases.

Advection - the Frozen-In Flux Theorem

Consider the ideal case when diffusion is negligible so that:

$$\frac{\partial \mathbf{B}}{\partial t} \approx \nabla \times (\mathbf{v} \times \mathbf{B})$$

In this case, fieldlines are 'dragged around' with the flow. This is known as frozen in flux - i.e. the flux through any open surface S, bounded by C, does not change. The change in time of magnetic flux is given by:

$$\frac{dF}{dt} = \frac{d}{dt}\int_s \mathbf{B}\cdot d\mathbf{S}$$

This flux could change in time due to a change in **B** or from movement of the boundary C of surface S . An element $\delta\mathbf{l}$ of C moving at **v** in an element of time δt covers an area of $\delta t\mathbf{v}\times\delta\mathbf{l}$. The flux passing through this area is

$$\mathbf{B}\cdot(\delta t\mathbf{v}\times\delta\mathbf{l}) = (-\delta t\mathbf{v}\times\mathbf{B}).\delta\mathbf{l}$$

by vector identity. Over the whole boundary then:

$$\begin{aligned}\frac{dF}{dt} &= \int_S \frac{\partial\mathbf{B}}{\partial t}\cdot d\mathbf{S} - \oint_C(\mathbf{v}\times\mathbf{B}\cdot d\mathbf{l}) \\ &= \int_S(\frac{\partial\mathbf{B}}{\partial t} - \nabla\times(\mathbf{v}\times\mathbf{B}))\cdot d\mathbf{S} \\ &= 0,\end{aligned} \tag{1.13}$$

using Stoke's Theorem and then the induction equation in the ideal limit we assumed at the start. In other words, magnetic field lines are frozen into the plasma and will not change their connectivity - the geometry of the field is altered by plasma motion but the topology is unaltered.

Since we are generally close to ideal on the Sun, we see plasma, following fieldlines, twisted and contorted in structures such as loops. Plasma elements remain largely with

the fieldline they start on. As fieldlines are tangled due to the action of photospheric motions on footpoints, their corresponding threads of plasma also become so. In the limit of vanishing resistivity, stresses introduced are able to build up. In paper [68] Parker poses that with sufficient tangling, these stresses can approach singular current sheets. We will elaborate on this later. In reality of course, we do not have zero resistivity and diffusion may eventually play a role.

Diffusion Limit

If we were to reach a situation where $R_m << 1$, we would have:

$$\frac{\partial \mathbf{B}}{\partial t} \approx \eta \nabla^2 \mathbf{B}$$

which we can see has the form of a diffusion equation. Here any spikes in any quantities would be smoothed out. By the same argument used above we can illustrate that here the flux through an area would be free to change.

$$\begin{aligned} \frac{dF}{dt} &= \int_S \frac{\partial \mathbf{B}}{\partial t} \cdot d\mathbf{S} - \oint_C (\mathbf{v} \times \mathbf{B} \cdot d\mathbf{l}) \\ &= \int_S (\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{v} \times \mathbf{B})) \cdot d\mathbf{S} \\ &= \int_S (\eta \nabla^2 \mathbf{B}) \cdot d\mathbf{S} \\ &\neq 0. \end{aligned} \tag{1.14}$$

So a particular plasma element would no longer be confined to its original fieldline and energy can be dissipated.

The solar corona is mostly highly electrically conducting. Diffusivity is always low.

For diffusion to become non-negligible then, R_m must become smaller. This can happen, for example, in current sheets - regions of high localised current where the typical length scale over which the field varies is very small. Then;

$$\begin{aligned} \nabla^2 \mathbf{B} &= \nabla(\nabla \cdot \mathbf{B}) - \nabla \times (\nabla \times \mathbf{B}) \\ &= -\nabla \times \mathbf{j}. \end{aligned}$$

So a region of high current within a region of relatively lower current would result in large gradients across that domain. Ideal conditions break down in that volume (and only that volume, globally the field may still be close to ideal) and fieldlines can reconnect in a lower energy state there. Reconnection has consequences beyond this restructuring of the field. It also converts some of the released magnetic energy into thermal energy through Ohmic heating, and some of it into kinetic energy via acceleration of plasma away from the site. This movement of plasma can generate waves which in turn can generate heating in other areas beyond the reconnection site.

1.1.6.2 Magnetic Helicity

The concept of magnetic helicity can be used to give a sense of the large scale topology of a field. We will use it throughout and so here give a brief summary of its properties and role in our work. Say the magnetic field **B** can be expressed as;

$$\mathbf{B} = \nabla \times \mathbf{A}$$

where $\mathbf{A}$ is a vector potential. Where $\mathbf{B}\cdot\mathbf{n}=0$ on the boundary of the volume, we say the field is magnetically closed. $\mathbf{A}$ is non-unique - to see this consider the gauge transformation

$$\mathbf{B}=\nabla\times(\mathbf{A}+\nabla\Phi)$$

where Φ is a scalar potential in the class of gauge functions whereby the curl of $\mathbf{A}+\nabla\Phi$ still equals the magnetic field. Then the helicity is given by:

$$\begin{aligned} H=\int_V(\mathbf{A}+\nabla\Phi)\cdot\mathbf{B}\,\mathrm{d}V &= \int_V\mathbf{A}\cdot\mathbf{B}\,\mathrm{d}V+\int_V\nabla\Phi\cdot\mathbf{B}\,\mathrm{d}V \\ &= \int_V\mathbf{A}\cdot\mathbf{B}\,\mathrm{d}V+\int_{\partial V}\Phi\mathbf{B}\cdot\mathbf{n}\,\mathrm{d}\,\partial V \\ \mathbf{B}\cdot\mathbf{n}|_{\partial V}=0 \;\Rightarrow\; H&=\int_V\mathbf{A}\cdot\mathbf{B}\,\mathrm{d}V, \end{aligned}$$

regardless of the choice of gauge function. Helicity is gauge invariant in this situation. This expression was developed using Gauss's equation for the linkage between a pair of closed fieldlines, then summing over all fieldline pairs in a volume - more details on the development of the concept can be found in [7]. High helicity must mean there is a high degree of twist and linkage in the field. However, we will be considering a straightened loop, where the footpoints do not return to their starting point. The field is not closed and $\mathbf{B}\cdot\mathbf{n}=B_n\neq 0$ as a general rule. Hence, for us helicity is no longer gauge invariant, and the physical interpretation is less clear. We turn to the concept of relative helicity to address this.

Relative Helicity

One way to adapt the idea of helicity for open fields was developed by the authors in [8]. They proposed, when dealing with open fields, the concept of relative helicity:

the helicity with respect to some reference field also defined within the volume V. This reference field, $\mathbf{B}_R$, must satisfy $\mathbf{B}_R \cdot \mathbf{n} = \mathbf{B} \cdot \mathbf{n}$ on the boundary. Both the original and reference field can be 'closed' by expanding the volume V to a larger volume, $V \cup \tilde{V}$. The field $\tilde{\mathbf{B}}$ is defined in $\tilde{V}$ and closes $\mathbf{B}$ and $\mathbf{B}_R$ within the volume $V \cup \tilde{V}$. So, within $\tilde{V}$, $\mathbf{B} = \mathbf{B}_R = \tilde{\mathbf{B}}$. Therefore the relative helicity $H_V(\mathbf{B}|\mathbf{B}_R)$ becomes

$$\begin{aligned} H_V(\mathbf{B}|\mathbf{B}_R) &= H_{V\cup\tilde{V}}(\mathbf{B},\tilde{\mathbf{B}}) - H_{V\cup\tilde{V}}(\mathbf{B}_R,\tilde{\mathbf{B}}) \\ &= \int_V \mathbf{A}\cdot\mathbf{B}dV + \int_{\tilde{V}} \tilde{\mathbf{A}}\cdot\tilde{\mathbf{B}}d\tilde{V} \\ &- \int_V \mathbf{A}_R\cdot\mathbf{B}_R dV - \int_{\tilde{V}} \tilde{\mathbf{A}}_R\cdot\tilde{\mathbf{B}}_R d\tilde{V} \\ &= \int_V \mathbf{A}\cdot\mathbf{B} - \mathbf{A}_R\cdot\mathbf{B}_R dV \\ &+ \int_{\tilde{V}} \tilde{\mathbf{A}}\cdot\tilde{\mathbf{B}} - \tilde{\mathbf{A}_R}\cdot\tilde{\mathbf{B}_R} d\tilde{V} \end{aligned} \tag{1.15}$$

Now, on the surface S between the two volumes, $\mathbf{B}\cdot\mathbf{n} = \mathbf{B}_R\cdot\mathbf{n} = \tilde{\mathbf{B}}\cdot\mathbf{n}$, therefore $\nabla\times\mathbf{A}\cdot\mathbf{n} = \nabla\times\mathbf{A}_R\cdot\mathbf{n}$ and so $\mathbf{A}-\mathbf{A}_R = \nabla\phi$, where ϕ is some scalar. Also recall that in $\tilde{V}$, $\mathbf{B} = \mathbf{B}_R = \tilde{\mathbf{B}}$. Then the final integral in 1.15 can be written:

$$\begin{aligned} \int_{\tilde{V}} (\tilde{\mathbf{A}} - \tilde{\mathbf{A}}_R)\cdot\tilde{\mathbf{B}}d\tilde{V} &= \int_{\tilde{V}} \nabla\phi\cdot\tilde{\mathbf{B}}d\tilde{V} \\ &= \int_S \phi\mathbf{B}\cdot\mathbf{n}dS \end{aligned}$$

Finally, using Stoke's Theorem and vector identities it is simple to show that :

$$\int_S \phi\mathbf{B}\cdot\mathbf{n}dS = \int_V (\mathbf{A}_R\cdot\mathbf{B} - \mathbf{A}\cdot\mathbf{B}_R)dV, \tag{1.16}$$

and so the relative helicity of $\mathbf{B}$ in V is:

$$H_V(\mathbf{B}|\mathbf{B}_R) = \int_V (\mathbf{A}+\mathbf{A}_R)\cdot(\mathbf{B}-\mathbf{B}_R)dV, \tag{1.17}$$

which is independent of the choice of extension field $\tilde{\mathbf{B}}$ and gauge invariant.

We now have established a valid way to discuss helicity even for an open field structure. In addition, in principle it can be possible to choose $\mathbf{A}_R$ such that the relative helicity reduces to $H(\mathbf{B})$. Expanding 1.17 gives;

$$\begin{aligned} H_V(\mathbf{B}|\mathbf{B}_R) &= \int \mathbf{A}\cdot\mathbf{B}\mathrm{d}V - \int \mathbf{A}\cdot\mathbf{B}_R\mathrm{d}V + \int \mathbf{A}_R\cdot\mathbf{B}\mathrm{d}V - \int \mathbf{A}_R\cdot\mathbf{B}_R\mathrm{d}V \\ &= H(\mathbf{B}) - H(\mathbf{B}_R) + \int (\mathbf{A}_R\cdot\mathbf{B} - \mathbf{A}\cdot\mathbf{B}_R)\mathrm{d}V \\ &= H(\mathbf{B}) - H(\mathbf{B}_R) + \int (\mathbf{A}\times\mathbf{A}_R)\cdot\mathbf{n}\mathrm{d}S \end{aligned}$$

If $\mathbf{A}_R$ is chosen such that $\mathbf{A}\times\mathbf{n} = \mathbf{A}_R\times\mathbf{n}$ on the surface S and $\mathbf{A}_R\cdot\mathbf{B}_R = 0$, then the reference field does not contribute to the relative helicity. In our work we will not have analytic expressions for $\mathbf{B}$ and $\mathbf{A}$ at $t > 0$, so we cannot specify exactly what this reference field would be, although as we will see in Chapter 3, Section 3.1.2 we will be able to make an approximation to the helicity. The crux of the matter is however, that we can still confidently discuss a meaningful helicity for our open field, without worrying about what the reference field would be.

Role of Boundary Motions in Helicity

It is our aim to characterise the types of systems created by different boundary driving,

and so here we show how helicity is a quantity which will vary due to the type of motions applied. Taking the dot product of vector potential $\mathbf{A}$ with Faraday's Law and manipulating with vector identities yields:

$$\begin{aligned}
\frac{\partial}{\partial t}(\mathbf{A}\cdot\mathbf{B}) - \mathbf{B}\cdot\frac{\partial \mathbf{A}}{\partial t} &= \nabla\cdot(\mathbf{A}\times\mathbf{E}) - (\nabla\times\mathbf{A})\cdot\mathbf{E} \\
\frac{\partial}{\partial t}(\mathbf{A}\cdot\mathbf{B}) + \mathbf{B}\cdot(\mathbf{E}+\nabla\psi) &= \nabla\cdot(\mathbf{A}\times\mathbf{E}) - \mathbf{E}\cdot\mathbf{B} \\
\frac{\partial}{\partial t}\underbrace{(\mathbf{A}\cdot\mathbf{B})}_{\text{helicity density}} + \nabla\cdot\underbrace{(\psi\mathbf{B}+\mathbf{E}\times\mathbf{A})}_{\text{helicity current}} &= \underbrace{-2\mathbf{E}\cdot\mathbf{B}}_{\text{helicity source}},
\end{aligned}$$

a balance equation for the evolution of the helicity density. Integrating over the whole volume and applying Gauss' divergence theorem:

$$\int_V \frac{\partial}{\partial t}(\mathbf{A}\cdot\mathbf{B})\,\mathrm{d}V + \int_S (\psi\mathbf{B}+\mathbf{E}\times\mathbf{A})\cdot\mathbf{n}\,\mathrm{d}S = -\int_V 2\mathbf{E}\cdot\mathbf{B}\,\mathrm{d}V, \tag{1.18}$$

we see that the evolution of the total helicity depends on the flux through the boundaries of the volume V, and the source. We can take this further if we re-write the electric field using Ohm's Law. Then:

$$\begin{aligned}
-\int_V 2\mathbf{E}\cdot\mathbf{B}\,\mathrm{d}V &= -\int_V 2(\eta\mathbf{j}-\mathbf{v}\times\mathbf{B})\cdot\mathbf{B}\,\mathrm{d}V \\
&= -\int_V 2\eta\mathbf{j}\cdot\mathbf{B}\,\mathrm{d}V \\
\int_S (\psi\mathbf{B}+\mathbf{E}\times\mathbf{A})\cdot\mathbf{n}\,\mathrm{d}S &= \int_S (\psi\mathbf{B}+\eta\mathbf{j}\times\mathbf{A}+(\mathbf{A}\cdot\mathbf{B})\mathbf{v}-(\mathbf{A}\cdot\mathbf{v})\mathbf{B})\cdot\mathbf{n}\,\mathrm{d}S \\
&= \int_S ((\psi-(\mathbf{A}\cdot\mathbf{v}))\mathbf{B}+\eta\mathbf{j}\times\mathbf{A}+(\mathbf{A}\cdot\mathbf{B})\mathbf{v})\cdot\mathbf{n}\,\mathrm{d}S
\end{aligned}$$

by vector identity, where $\eta = \frac{1}{\sigma}$ is the resistivity. Resistivity is small, so η terms are less significant, and for sufficiently small η would be negligible, so neglect these terms for now. We are left with the flux through the boundaries, therefore clearly the velocity on the boundaries is key to the helicity balance equation. If we set the x and y boundaries to be periodic (as indeed we will in later simulations), then flux crossing these boundaries cancels out. Turning now to the upper boundary, where we plan to have $\mathbf{v} = \mathbf{0}$ and where the normal into the domain is $\mathbf{n} = -\mathbf{e}_z$, so the flux is:

$$\int_S (\psi \mathbf{B} + \mathbf{E} \times \mathbf{A}) \cdot \mathbf{n} \, \mathrm{d}S = -\int_S \psi B_z \, \mathrm{d}S$$

Now on the lower boundary, $v_z = 0$ and $\mathbf{n} = \mathbf{e}_z$, the flux is:

$$\int_S (\psi \mathbf{B} + \mathbf{E} \times \mathbf{A}) \cdot \mathbf{n} \, \mathrm{d}S = \int_S (\psi - \mathbf{A} \cdot \mathbf{v}) B_z \, \mathrm{d}S$$

Furthermore, we are free to choose a ψ which is zero on boundaries so the flux through the top boundary disappears and the flux through the bottom boundary is then given by one term only:

$$-\int_S (\mathbf{A} \cdot \mathbf{v}) B_z \, \mathrm{d}S = -\int_S (A_x v_x + A_y v_y) B_z \, \mathrm{d}S \tag{1.19}$$

Interestingly then, in this case it would seem that the helicity in the domain depends only on the velocity and normal component of magnetic field on the lower z boundary (the tangential components of $\mathbf{A}$ correspond to B_z also) and not the tangential components of $\mathbf{B}$. B_x and B_y will depend on the evolution in the domain thus far. However, if $B_z = 1$ at $t = 0$ and the boundary driver is incompressible, then it will remain close

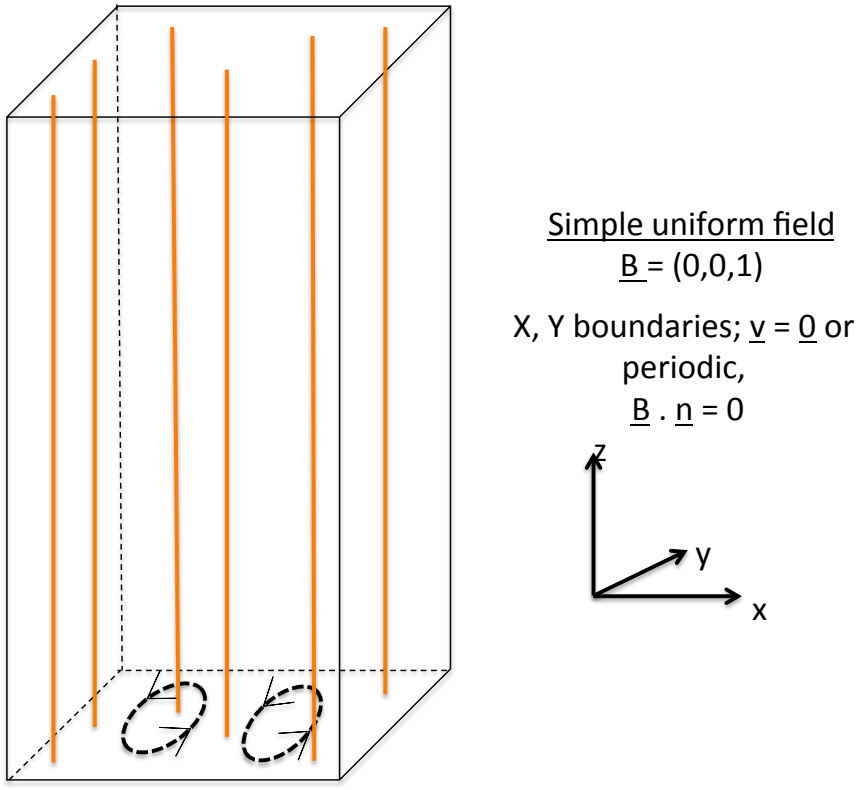

Figure 1.6: Cartoon of a uniform field under the influence of rotating boundary motions on $z = 0$.

to 1 for $t > 0$ regardless of the evolution inside the domain (in fact for ideal MHD it would be unchanged for all time). Compare this to the Poynting flux for example: Poynting flux across the boundary depends on both the normal and tangential components of **B**, and so the history of the system does impact the flux. In the specific set of circumstances discussed however, only the sign of the driving affects whether helicity is added or removed. Clearly then the specifics of our driver will have the potential to change the characteristics of the field greatly, as these vortical motions will be placed on the lower z boundary as pictured in Figure 1.6. In fact when we have specified the driver functions in Chapter 3 we will be able to calculate the helicity contributed by each one.

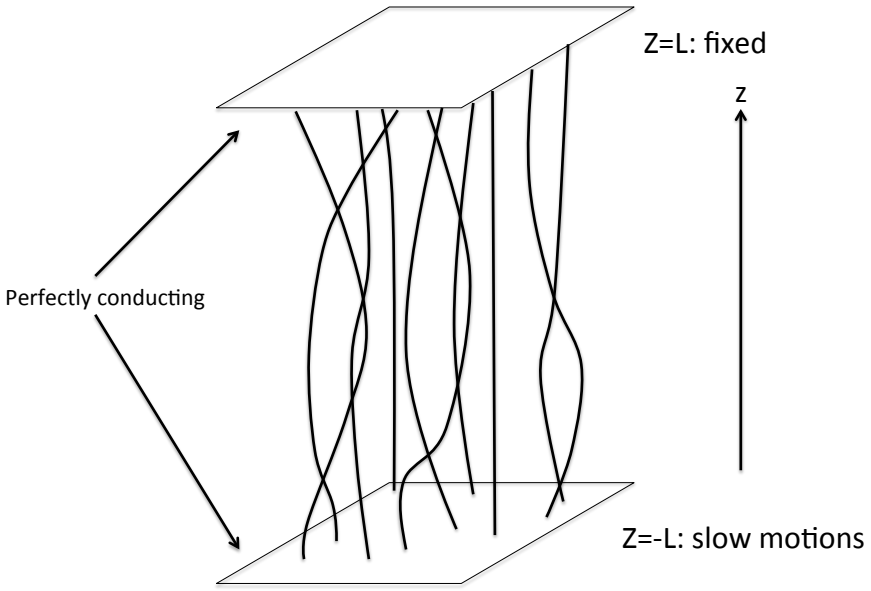

Figure 1.7: Parker set-up of fieldlines connecting to plates representing the photosphere. The sketch illustrates a few flux tubes and how they can be braided around each other.

1.2 Braiding

We now move on to focus on the development of braiding theory and topological dissipation.

1.2.1 Parker and the Braiding Theory

The landmark paper by Parker in 1972 (see [68]) was first to pose the idea of braiding and topological dissipation as a mechanism for coronal heating.

The scenario we now describe is illustrated in Figure 1.7. Parker considers an ideal, uniform field anchored between two perfectly conducting parallel plates, representing the photosphere, as if a loop has been stretched out. Now imagine arbitrary slow motions on one of the plates, moving the footpoints of that surface. At the other plate no

motions are applied so fieldlines remain anchored in their initial positions here. The movement of footpoints at one end will braid the field to some extent and some stress will build in the structure in the form of regions of higher current. Since the motions are slow, the field will try to respond to this by relaxing ideally to a new, smooth, force-free equilibrium with the same original topology. If the motions are applied again, the field responds accordingly again. The coronal field evolves by continuously responding to the photospheric motions with a sequence of force-free equilibria; this is referred to as quasi-static evolution. However, Parker posits that after arbitrary footpoint motions, in most cases the resulting field will be unable to relax ideally to a smooth equilibirum, because the space of continuous solutions to such activity is very limited. Parker proposes that there exists a smooth equilibirum only when the fieldline twist introduced by footpoint motions is very simple and uniform along the field - and braiding does not meet these criteria. Parker's magnetostatic theroem, developed throughout his work and described in [72], says that in general only equilibria containing tangential discontinuities in the form of current sheets exist. In an environment where $\eta = 0$, localised increases in current will not be smoothed out and tangential discontinuities will form - currents could become infinite within an infinitesimally thin sheet along the magnetic field. This could then trigger magnetic reconnection and release of magnetic energy. This is known as topological dissipation. This proposal has been a contentious one and is yet to be completely proved or disproved. In the following sections we will elaborate on some of the experiments which have attempted to do so. However note that in a system with finite resistivity such as ours, of course diffusion becomes non-negligible when currents layers are sufficiently thin. We may not require true discontinuities to form at all for the topology to be able to change through magnetic reconnection and release free magnetic energy.

Current Sheets

Current sheets are magnetic surfaces over which an electric current flows. Imagine

two magnetic flux tubes extending between two plates as before, where the tubes have been twisted around each other. The surface consisting the area where the flux tubes meet will have some current flowing along it. In an ideal case, where the degree of twist of the flux tubes reaches some critical point, (i.e the curl of the magnetic field is such that the lengthscale over which the field varies goes to zero), then we have a true discontinuity where an infinite current is flowing along the infinitesimally thin magnetic surface between the flux tubes. Physically, this is unrealistic, and however close to perfectly conducting the solar plasma often is, the current will not have the chance to become so concentrated before diffusion becomes non-negligible. Therefore in our work we will consider thin, but finite, current layers, which are not singular but can trigger reconnection regardless.

1.2.2 Numerical Braiding Experiments

Parker's idea triggered a new area of research which sought to understand the effect braiding as a result of photospheric motions can have on coronal structures, whether discontinuities can form, and what the consequences for heating could be. Parker himself published further papers (for example [71], [72]) on the theory. It remains a much debated topic. As we will discuss, observational evidence of braided structures is limited, but is improving and looks promising- see Section 1.2.3 for more on this. Furthermore, different authors considering different scenarios have found evidence both for and against the formation of current sheets as posed by Parker. Here we will review the key findings of some of the many numerical experiments conducted concerning braided fields. We will see that the type of photospheric motions can greatly affect the outcomes. There is also a broad area of research exploring this idea analytically: see for example [120], analytic parts of [25] (which find smooth equilibria in their cases) and [69]/ [70] (which find evidence for discontinuities).

The common starting point for most of these types of simulation is to consider a stretched out loop, so that both upper and lower z boundaries represent the photosphere. Simulations vary from using uniform initial fields to more complicated set-ups, setting different resistivities, and of course driving motions.

There have been few experiments thus far which fully concur with the hypothesis as first described. Of course, only in ideal conditions where resistivity is set to zero can true tangential discontinuities form, as otherwise diffusion will act and prevent currents building to the point of singularity. In [54], the authors took a uniform field under ideal conditions. Various shear motions were applied to both 'photospheric boundaries'. They found that as the strength of the shear was increased, the thinner the layers of current became. It was indeed possible to find smooth equilibria under a certain shear strength, but above a threshold the current appeared to become singular - at least at the grid resolution available at the time of the work. New simulations with a more efficient code, outlined in [18], have since concluded that in fact these apparent sheets were of finite thickness. The studies of [25] applied drivers to a uniform field and failed to find any evidence of current sheets in all but one experiment. The majority of runs settled to smooth equilibria free from discontinuities, and the one run that did form current sheets was the only one where the non-photospheric numerical boundaries were altered as a result of the driving.

The work of [107], [108] and [60] all took an ideal, uniform field and applied a sequence of motions to tangle the field. Velocity at footpoints was introduced and then the field was allowed to relax before the next application. The authors observed an increasingly filamentary nature to the field structure and current profiles, with narrow enough current layers to induce some dissipation due to resistivity had this not been

an ideal experiment. They refer to a cascade effect, where each successive iteration results in a decrease in current length scales at an exponential rate. It would therefore take infinite time to produce infinitesimal current sheets. They do not find evidence to support the spontaneous formation of true current sheets Parker proposed, even in this ideal environment. On the other hand though, if true current sheets are not necessary to trigger reconnection, we can consider an adjusted Parker's theory based on the formation of 'thin enough' current layers in a resistive case due to footpoint motions.

The authors in [55] ran simulations on an initially uniform, resistive field under the influence of continuous, slow photospheric motions and found that small regions of higher current were present in the statistically steady states reached. It was concluded in [49] that current sheets were 'inevitable' even in untwisted fields. Further supporting evidence was provided by [57], using continuous footpoint motions as part of a magnetostatic model and failing to find any smooth equilibria. One key paper which gives us a glimpse into the highly complex nature of field behaviour and the consequences footpoint motions can have is [35]. The authors again took a uniform field representing a straightened loop and subjected it to a variety of continuous shearing motions on both photospheric boundaries. The driving motions had random properties and the effects of driving speed, duration of driving, grid resolution and loop aspect ratio were shown to alter the resulting heating rate. Different resistivity values were also used. The authors found that even one driving cycle was sufficient to induce exponential growth in currents. Quantities reached statistically steady states after just a few periods and reconnection took place. Thin current layers were formed easily. No discontinuities in the true mathematical sense would have formed even with zero resistivity, but for all practical purposes we can count these runs as encouraging for the topological dissipation theory. In addition, the reconnection events set off waves of supersonic and super-Alfvénic speeds, which themselves gave rise to more current layers on even smaller scales. The current structure becomes filamentary and there is

plenty of opportunity for heating.

Other authors have modelled simplified versions of convection cells by applying rotational motions to uniform fields and found results offering further support to the possibility of current layer formation and subsequent heating. The authors of [43], [40], [64] and all see current layer formation under various different types of simulation, all applying continuous, rotational motions but under different MHD models, grid resolutions, aspect ratios and resistivities. Further, the work of [92], [93],[90],[91] all analyse the effect of footpoint motions on magnetic fields and find current layers leading to energy dissipation, instabilities and turbulence under reduced MHD. Reduced MHD is a simplified model which can be used for some simulations as it saves on computational expense, allowing large experiments to run faster. When studying a magnetic field which is said to be almost uniform and uni-directional, one can assume fairly homogeneous behaviour along the direction of the dominant field component and the equations of MHD simplify to a set of 6 with 6 unknowns. Often the field under consideration will be taken to be almost uniform in the z direction, such that it can be written;

$$\begin{aligned} \mathbf{B} &= \mathbf{B}_{\perp} + B_z \mathbf{e}_z \\ \text{where } \frac{\mathbf{B}_{\perp}}{B_z} &\sim \varepsilon << 1 \end{aligned}$$

Using this and assuming that:

- there is equipartition of energy in the x, y directions, so that perpendicular kinetic, internal and magnetic energy densities are approximately of the same size, and;

- variations are limited along the dominant direction, so $\partial/\partial z \sim \varepsilon$ and $v_z = 0$,

then each variable can be considered as some power of ε and only the most significant terms kept. The variables $\mathbf{B}_\perp$ and $\mathbf{v}_\perp$ can also be written in terms of scalar functions ψ and ϕ respectively. Putting all of this together leads to the reduced MHD equations:

$$\begin{aligned}
\frac{\partial \psi}{\partial t} + \mathbf{v}_\perp \cdot \nabla \psi &= \eta \nabla^2 \psi - B_z \frac{\partial \phi}{\partial z} \\
\rho_0 \left(\frac{\partial \omega_z}{\partial t} + \mathbf{v}_\perp \cdot \nabla \omega_z \right) &= \mathbf{B} \cdot \nabla j_z \\
\omega_z &= \nabla_\perp^2 \phi \\
\mu_0 j_z &= \nabla^2 \psi \\
\mathbf{v}_\perp &= \mathbf{e}_z \times \nabla \phi
\end{aligned}$$

The detailed derivation can be found in [9]. We outline the key points since RMHD is a common approach which has been used in research similar to ours, but we will see from our results later that it may not always be sufficient for accurately describing a system's evolution.

Braided Field Experiments

An extensive series of numerical experiments using initially braided fields have been carried out by the assortment of authors in [112], [113], [114], [82], [115]. In [112] the authors point out that many investigations have looked only at current layer thickness as the main indicator of how effective braiding could be at upsetting equilibrium and triggering reconnection. In fact as mentioned before, it is the integrated parallel electric field which determines the rate of reconnection in three dimensions. The authors consider an initially braided field with no net twist and allow it to relax ideally,

and find that the field reaches a force-free state with no small scale currents. However they also consider the integrated parallel electric current (as it is related clearly to the electric field) and find that the length scales over which this quantity varies actually decrease with increased braiding, and its maximum increases. Clearly equilibrium would eventually be lost as a field was continuously braided, without small scale currents necessarily occurring. This work was taken further in [114] and [82], using the final equilibrium state as their initial condition and evolving it under a resistive numerical simulation. The authors saw thinner and thinner current layers developing, which they refer to as a cascade to smaller scales. There was a global restructuring of the field during this phase, eventually bringing it to a smooth equilibrium with two oppositely twisted flux tubes spanning the length of the domain. The braided nature of the initial state appears to have triggered the evolution of small scale currents and subsequent field relaxation to a smooth state. In addition, the authors of [79] carried out ideal relaxation experiments on initial braided fields of varying complexities. These initial conditions were based on the same field as in [112], and the authors found finite current layers in the final states whose thickness depended on the complexity of the initial state. This is particularly interesting to us as it suggests that the degree of braiding can have an impact on heating opportunities.

Some of the latest work has focused on making the coronal loop model more realistic by including more solar effects, such as in [109] where the authors include stratification of the atmosphere. They find that taking variations in the lower atmosphere into account alters the way a coronal loop evolves, in that it changes from being via the DC heating an un-stratified model induces, to being via AC heating. Furthermore, researchers are looking now at more complex and realistic fields. Rather than just taking a coronal loop, later in this thesis we will attempt to apply driving motions to a structure with more similarities to the real magnetic carpet of the solar atmosphere. See also

[18], which applies photospheric motions under ideal conditions to more topologically complex initial fields. The various cases result in the formation of different features, from finite current layers to singular currents at null point locations. Clearly there remains great deal to consider.

Conclusion

The question of whether true tangential discontinuities can form in simple fields where $\eta = 0$ is somewhat academic - while not everyone is agreed on it, on the Sun we never have exactly zero resistivity anyway. Some dissipation will occur, particularly when currents build in narrow regions. A current layer need not be infinitesimally thin to trigger reconnection, and even without thin current layers the integrated parallel electric current can build. The take-away point is that finite but thin (i.e. small enough length-scales for resistivity to be important) current layers are generally found in numerical simulations where braiding is applied, with the potential to induce heating. We have also seen how the type of motion driving the braiding of the field can affect current build-up. We wish to take this subject further by quantifying types of motion in order to pinpoint the nature of heating arising.

1.2.3 Observations

Great strides have been made over the years with instruments and satellites designed to observe the Sun. Different instruments can look at different wavelengths and temperatures of structures from the photosphere up through the corona. Table 1.1 gives a brief overview of solar missions past and present which have contributed to our understanding. More details can be found at `https://www.nasa.gov/missions`.

Mission	Period	Instruments	Focus
Skylab	1973-74	8 different cameras, telescopes etc. in various wavelengths	First of its kind, initial look at solar surface/structures.
SOHO	1995 -	12 different instruments	Interior, heliosphere, solar wind
TRACE	1998-2010	Multi-wavelength telescope	Magnetic structures/ temperatures
Rhessi	2002 -	High energy spectroscopy equipment	Particle acceleration/ flares
STEREO	2006 -	SECCHI IMPACT PLASTIC SWAVES	CMEs/solar wind
Hinode	2006 -	SOT XRT EIS	Field structure
SDO	2010 -	AIA EVE HMI	Field generation/ energy storage, release
Hi-C	2012 & 2016 (separate short flights)	UV imager	High resolution coronal images

Table 1.1: Overview of solar observation missions.

With each mission we have been able to see finer and finer resolution in all wavelengths, although we are still far from indisputable observational evidence of braiding. The most exciting data for us comes from the Hi-C spacecraft. In just a five minute flight, this instrument resolved high temperature structures in the corona to a scale never before seen. Paper [21] details the specifics and shows the best evidence so far that coronal structures can be braided. In the images we see fine structure at lengthscales of around 0.2 arcseconds (150 km) and temperatures of around 1.5 million kelvin. Previously our best images in terms of spatial and temporal resolution came from the Solar Dynamics Observatory (SDO) instrument, the Atmospheric Imaging Assembly (AIA). This apparatus takes images with spatial resolution of 1 arcsecond in 10 different wavelengths around every 10 seconds. The resulting images are impressive in their own right and show huge tangled loops on a larger scale, as shown in Figure 1.8 (both images can be found at `https://sdo.gsfc.nasa.gov`).

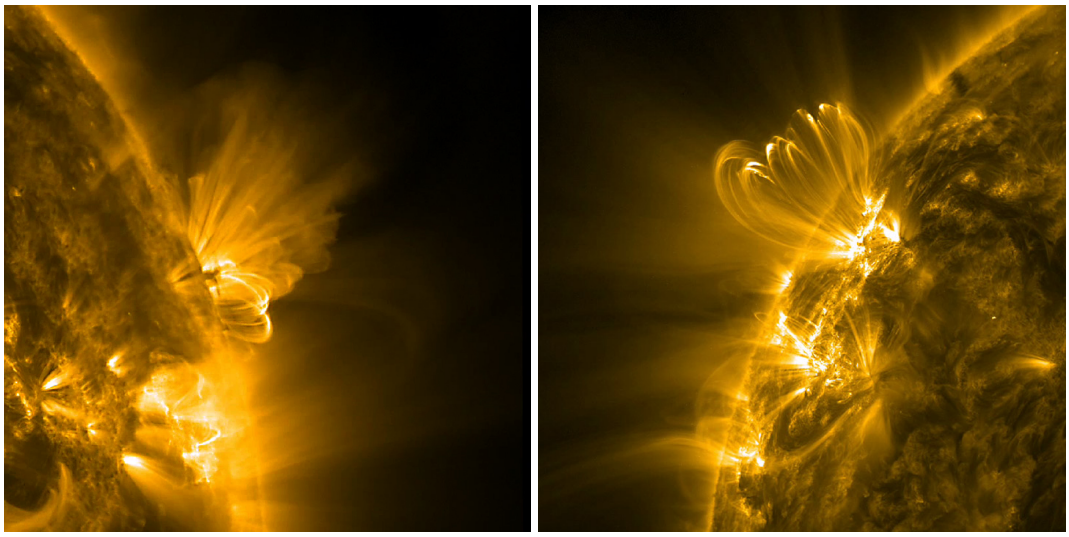

Figure 1.8: AIA active region loops. Courtesy of NASA/SDO and the AIA, EVE and HMI science teams.

The authors in [21] made comparisons of AIA and Hi-C images. The Hi-C images show groups of twisted, braided loops, reconnecting and releasing free magnetic energy, which cannot be seen in the lower resolution AIA images. However, comparing

AIA data at the same wavelengths shows even more. In one observation, a small flare and acceleration of plasma is seen in AIA shortly after Hi-C had stopped taking images, at the very location where strands of plasma had just been seen to interact in Hi-C. Furthermore, Hinode data allowed the authors to calculate that the maximum temperature of this loop hits 7 million kelvin. Now obviously this 5 minute run is not enough to definitively prove topological dissipation as a significant coronal heating mechanism and more research is still required, but it is very encouraging that our theories here may have huge potential.

Other authors have also examined the braided structures of Hi-C data, for example [105]. They found by looking at the data in conjunction with AIA and HMI data, that a small flare, or subflare, was triggered in the coronal braided structure by a separate reconnection event, involving loops underneath in the transition region. The authors found evidence that reconnection events sent a one small loop downwards and another larger one upwards, with this interaction with the overlying braided field triggering the flare event.

While this is all encouraging for the theory of braiding as a heating mechanism, one should also be cautious in interpreting such observations too. Intuitively we can understand that while two strands may appear to cross over in one image, they may in fact be quite separate as seen from a different angle. This is one of several conclusions drawn by the authors in [80], who ran relaxation simulations on braided field setups and compared their synthetic emission profiles to their magnetic field structure, to clarify what we can safely conclude from observations such as those studied in [21]. In addition to noting the potential for the viewing angle to be misleading, the research poses that the structure suggested by bright strands in observations does not necessarily match exactly with field structure, and that a bright crossing does not guarantee reconnection there. The findings do also state however that the lack of a braided appearance in observations does not mean that the underlying field is not highly tangled - a bright but

smooth patch in observed emission may still have a braided field topology.

Observations have also been used to determine how much Poynting flux (energy being directed up from the photosphere) is created by photospheric motions and whether it is sufficient, in theory, to heat the corona. For example, the ideal MHD treatment in [118] used data from Hinode to calculate upper and lower bounds for the Poynting flux by deriving photospheric motions from magnetograms. The authors were then also able to calculate an estimate of the real Poynting flux for one particular observation and found it fell between the calculated bounds. While this is only one experiment, it is encouraging to see observations supporting the idea that there is enough energy pumped into the solar atmosphere by photospheric motions to account for the temperatures we see, even if we cannot yet be sure of the dissipation mechanisms.

Chapter 2

Topological Entropy and Calculation Methods

2.1 Topological Entropy

2.1.1 Motivation

A key part of this work is to investigate how the complexity of photospheric motions affects the braiding of field lines in coronal loops, and subsequent heating of the plasma via magnetic reconnection. The quantity chosen to measure the complexity is the topological entropy. The concept of topological entropy was first developed by the authors in [1] in 1965 concerning continuous mappings, and has since been applied to situations such as quantifying the efficiency of mixing in [14] and later to flows where the stirring is created by point vortices in [13]. The mathematical definition of topological entropy makes it difficult to calculate directly. Instead, one can estimate it from the rate of stretching of a material line subject to the mapping. This concept has been

considered by many authors, including those in [63], and numerical algorithms developed, such as those described in [103]. The method considers a material line in a two-dimensional fluid flow, where particle motions on the surface mix the fluid, acting as stirrers. When referring to particles, this can mean physical particles in a fluid, fluid elements, or whatever is pertinent to the work. As the fluid is mixed, the original material line is stretched and folded repeatedly. The entropy is a measure of the maximum rate of stretching of the line, so the extent to which the fluid has been mixed, with the application of the mapping which describes the mixing motion. For certain maps, such as the 'horseshoe map' of chaos theory where a material is stretched then folded over and over, the stretching can be calculated analytically. The trajectories of particles can be plotted in a space-time diagram, with time increasing along the z-axis. This gives a 'spaghetti plot' where the trajectories are lifted up in time producing a braid. The greater the entanglement of the strands of the braid, the more thorough the mixing and the greater a material line on the surface is distorted. So the higher the entropy, the higher the complexity induced by the flow. This relationship between braids and topological entropy can be exploited to provide a way of calculating entropy, as described in [103] and here in Sections 2.1.6 and 2.1.7. We will expand on the mathematical theory behind topological entropy in Section 2.1.3.

2.1.2 Photospheric driver and entropy

In our work the mixing will be carried out by the driver consisting of vortical motions which will be described in more detail later. For the moment we will give a only brief explanation to see the context for entropy. We consider a coronal loop, represented by a uniform field, stretched out so that one set of footpoints is anchored in the lower z boundary and the other in the upper z boundary. We now apply our vortical motions

at the lower boundary while keeping the upper boundary fixed. The basic profile of the driver was chosen to produce braiding of field lines similar to that in the work of [112], [114] and [82]. Our vortices are a pair of 'blinking vortices', so one twists on, stirring the fluid, then switches off and the second takes over, stops and the whole process repeats for as many periods as one wishes. The aim is to alter properties of this driver and then measure the topological entropy given by each scenario. We can then determine if there is a correspondence between the entropy and the amount of heating taking place as field lines reconnect in attempt to restore an equilibrium to a field strained by braiding.

2.1.3 Growth Measurements

There are two main ways to measure the growth rate of a material line - the metric growth or the topological growth. We will define these as in [12]. The function ϕ_t is the fluid displacement due to the movement of stirrers, which in our case will be particle trajectories. Take a material line γ to be a smooth arc or a simple closed curve. If $l_0(\gamma)$ is the initial length of γ and $l_t(\phi_t \circ \gamma)$ is its length after a time t under the stirring, then the metric growth by time t is given by;

$$L_t^{met}(\gamma) = \frac{l_t(\phi_t \circ \gamma)}{l_0(\gamma)}. \tag{2.1}$$

The topological growth is more complicated. If an arc or simple closed curve is wrapped in some way around particle/rod stirrers, we can think of tightening it around the stirrers, without letting it pass through them, contracting the loop to its shortest possible length. To measure topological entropy a 2D flow will be considered over some area, so we can think of this flow being defined within boundaries. Now, two curves are homotopic if one can be continuously deformed into the other. We call an

arc essential if both of its endpoints are on a flow boundary, and essential arcs are homotopic if the homotopy between them keeps their end points on the same boundary. A simple closed curve is essential if there is no homotopy which deforms it to a single point or the boundary itself. If we defined the set of essential curves homotopic to γ as its homotopy class $[\gamma]$, then we take from this set the curve with the shortest length, i.e. the one wrapped most tightly around string particles, giving a topological length for γ of;

$$L^{top}(\gamma) = min(l(\sigma) : \sigma \in [\gamma]). \tag{2.2}$$

The topological growth rate can then be defined as;

$$L_t^{top}(\gamma) = \frac{L^{top}(\phi_t \circ \gamma)}{L^{top}(\gamma)} \tag{2.3}$$

Seeing as the topological growth rate uses the shortest possible length, it follows that

$$L_t^{met}(\gamma) \geq L_t^{top}(\gamma). \tag{2.4}$$

Now, if we consider all possible essential curves under the influence of a flow ϕ_t, the topological entropy is defined as the supremum of the metric growth rates of all of these possible curves in the limit of t approaching infinity. We note that the topological entropy uses the metric growth rate, not the topological growth rate. It is referred to as a topological quantity due to a variational principle link with the metric entropy: the topological entropy maximises the metric growth rate. Further details of these concepts are beyond the scope of this work, but can be found in [6]. The complex derivation of the formula for topological entropy is also unnecessary for our purposes, so we simply

present the formulae given in [12];

$$h_{top}(\phi_t) = \sup\left(\limsup_{t\to\infty} \frac{\log(L_t^{met}(\gamma))}{t} : \gamma \text{ is a smooth arc}\right) \tag{2.5}$$

The topological growth rate can however tell us about the nature of the topological entropy, a fact that is exploited by the Thurston-Nielsen Theorem.

Thurston-Nielsen Theorem

The Thurston-Nielsen theorem or trichotomy gives us a link between the topological growth rate and the topological entropy resulting from different fluid motion.

If for every essential curve there exist real constants $C_1, C_2 > 0$ and $\lambda > 1$ with $C_1\lambda^t \leq L_t^{top}(\gamma) \leq C_2\lambda^t$ then we can say the topological growth rate is exponential, and by Equation 2.3, that the metric growth rate is also exponential, and therefore:

$$\begin{aligned} h_{top}(\phi_t) &\geq \limsup_{t\to\infty} \frac{\log(C\lambda^t)}{t} \\ &= \limsup_{t\to\infty} \frac{\log(C) + t\log(\lambda)}{t} \\ &= \limsup_{t\to\infty} \frac{\log(C)}{t} + \log(\lambda) \\ h_{top}(\phi_t) &\geq \log(\lambda) > 0 \end{aligned}$$

This is known as the pseudoAnosov (or pA) case, where the mixing is highly complex. If however $L_t^{top} < Kt, K > 0$, we say the fluid motion is finite order, or fo, and mixing is not as complex. Here we cannot deduce if the metric growth rate is always linear or exponential, only that any essential curve γ can be deformed into a curve in the same homotopy class with linear topological growth. Also, although a pA motion means exponential growth and positive entropy, positive entropy does not necessarily mean

universal exponential growth and a pA motion. There could be exponential metric growth rate locally somewhere, resulting in a positive entropy, and still the underlying topological growth rate could be linear. If the topological entropy approaches zero in time then we can say the metric growth rate must be linear and therefore so is the topological growth rate and so we are finite order. If neither a finite order or pA regime hold, then the case is reducible, and the surface of motion must be able to be split into fo or pA subsurfaces.

2.1.4 Comparison to Lyapunov exponent

To further understand the concept of topological entropy we can compare it to the Lyapunov exponent, another quantity which indicates complexity and chaos in systems. This is a measure of how quickly trajectories, which are initially close together, move away from each other over successive iterations or time. Consider a discrete, one dimensional mapping whereby:

$$f(x_i) = x_{i+1},$$

so an orbit starting at x_1 is mapped to x_2, then x_3 and so on, such that the trajectory is given by the set of points $(x_1, x_2, x_3, ..., x_n)$. Letting $f'(x_i)$ represent the derivative at x_i, the Lyapunov exponent of the trajectory beginning at x_1 is given by:

$$\lambda(x_1) = \lim_{n\to\infty} \left[\frac{1}{n} (ln|f'(x_1)| + \cdots + ln|f'(x_n)|) \right]. \tag{2.6}$$

A positive value for this would indicate chaos, with the difference between successive points in the trajectory increasing at an exponential rate. One can study many system trajectories and if the largest Lyapunov exponent is positive then the system is chaotic.

The concept is generalised to analyse multidimensional and continuous mappings also. Further information can be found in [119] and references therein. As mentioned however, this is a quantity which considers trajectories starting close together and describes their rate of divergence, often over a small scale. We want to quantify the degree of entanglement of trajectories on a global scale, from particles covering a whole plane, which is what the topological entropy provides. The topological entropy estimation method presented in [103] needs only particle trajectories in order to supply us with a sense of the degree of mixing on the surface from any set of chosen starting points.

2.1.5 Braids of Trajectories and Braided Fieldlines

As mentioned before, a finite set of particle trajectories can be plotted in time to give a space-time diagram, which can also be described as a braid. Firstly, however, we should make a distinction between the braid describing the particle trajectories in space-time and the physical braid of field lines resulting from the action of particle motions which we have previously discussed.
Since the field is constructed of fieldlines anchored in the photosphere, each footpoint can be thought of as representing a seed particle. We are considering a coronal loop, where the fieldlines are of finite length and close at the surface. In this close to ideal environment, the overlying field will be imprinted in some way by photospheric motions. In this case the 2D footpoint motion mapping will translate to a 3D tangling of the associated fieldlines.

Assuming ideal conditions, the stirring motions on the photosphere would result in field lines being braided as per the braid of trajectories, giving the same topological picture to the space-time diagrams. In our simulations we will have resistivity therefore the picture will not be so similar - we expect reconnection to be taking place which obviously would lead to a different picture of crossings in the physical field line braid

than in the braid of trajectories.

Our approach is to apply vortices to carry out the mixing, in analogy to the stirrers in [13]. A question to consider is: where should we place our vortices? Intuitively we can see that there will be some configurations which induce either a maximum or minimum entanglement. Let us consider the limiting cases.

Initially consider placing a pair of vortices of opposing circulation on top of each other. The first vortex would twist strands and stretch a material line on the surface. The second vortex would then twist the other way, undoing the previous action: net mixing is zero and tangling is cancelled out. One vortex injects positive helicity and the other injects the same amount of negative helicity. If we start to move one vortex away, less mixing is undone and the field remains more braided: the driver has higher entropy. If the vortex centres are placed far enough apart, the action of each vortex may have very little effect on that of the other (each vortex will have an exponential profile so their influence falls away exponentially with radial distance); the field tangling is less complicated again; the driver has lower entropy. So there will be some point in between the extremes where we see a maximum tangling and highest topological entropy. If we think of the same argument for vortices twisting in the same sense, we can predict that we will have a similar optimum distance for highest entropy. In this case however, the action is more coherent - one vortex on top of the other does not undo, but continues to twist the field tighter, into a helical structure rather than one with strands crossing over each other. Instead, helicity into the system is always positive, always increasing, but we would expect generally lower entropies with fieldlines weaving in out less. The helix will become looser as the vortices are moved further apart, until once again we have vortical motions whose effects are not felt by the other. Varying the vortex placements in our entropy tests will give us an idea of how to achieve the greatest entanglement of trajectories. Once again we should keep in mind that this optimum

would be for an ideal environment where any crossing of trajectories is preserved and is therefore an estimate for our resistive framework.

2.1.6 Flow and Braid Entropy

So how can braids (now in the sense of entangled particle trajectories) be used in the determination of the topological entropy of a flow, which will subsequently be applied to a three dimensional magnetic field?

Mathematically, the topological entropy of a flow, h_{flow}, is calculated considering the growth rates of all possible material lines in the flow and taking the supremum. Clearly this is difficult in practice. Therefore we estimate the entropy of the flow using braids. We choose to follow certain particles and their trajectories become the strands of the braid, and calculating the topological entropy of the braid gives a lower bound estimate of that of the whole flow. Clearly then the more trajectories we consider, the closer we are to a continuous flow, and the closer to the true entropy value. If we follow n particles then;

$$h_{flow} \geq h_{braid}^{(n)} \tag{2.7}$$

As n increases the better our estimate to the entropy of the flow becomes (see [33]). There are various ways of calculating $h_{braid}^{(n)}$ but a simple method is first described in [62]. A numerical code for the entropy calculation developed from this idea is presented in [103]. This is the algorithm that we will implement. We first give an explanation of some theory behind the method before considering specifics of the code.

2.1.7 The method

The approach is to calculate the trajectories of particles seeded at random points and advected by a flow, giving our braid. We can split the braid into units, where each unit corresponds to one crossing between a pair of strands (strands being the two dimensional trajectories plotted in time). Each unit is called a 'generator', given by:

1: σ_i for a clockwise twist between strands i and $i+1$ (from the 2D viewpoint shown in diagrams this looks like strand i passing underneath strand $i+1$), or;
2: σ_i^{-1} for strands i and $i+1$ twisting around each other in an anti-clockwise sense (in this case the 2D visualisation shows strand i passing over strand $i+1$).

This is illustrated in Figure 2.1. The braid can be encoded by the set of generators which describe each unit of the braid, ordered in time.

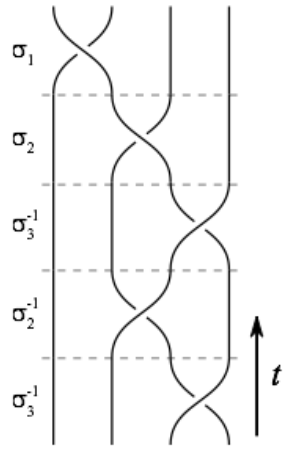

Figure 2.1: Generators σ representing crossings between strands. Taken from [103].

In practice, when considering a 2D flow, we project the trajectories on to the horizontal axis (or some other preferred projection line) and then calculate our string of generators according to how trajectories move on this line. If we have n particles, we will have

$n-1$ possible crossings, so $n-1$ generators. How do we now use these to calculate the effect on a material line?

We insert an arbitrary, non-self intersecting closed loop wrapped around the seed particles (again note the difference between referring to a loop in this context, and our discussions of coronal loops) whose braid of trajectories is now encoded by braid generators. We can then encode this loop by counting the crossings of the loop around the particles. Figure 2.2 is taken from [103] and illustrates the general set of 'crossing numbers' used to do this. The topology of a material loop being stretched in a flow is determined by some set of i particles moving according to the flow. The loop will be wrapped around and between these points in some pattern. The particles can be positioned along an axis, such that loop topology is unaffected, and numbered from left to right, from 1 to n. We then count crossings of the loop relative to the particles' positions. Breaking up the picture of the particles such as with the vertical lines in Figure 2.2, anytime the loop crosses a vertical line we count a loop crossing. There are two types of crossing:

1. Those labelled ν_i, $i = 1...n-1$, represent the number of times that the loop crosses the vertical line between particles i and $i+1$;

2. For particles $i = 2..n-1$, μ_{2i-3} is the number of crossings of the loop with the vertical line through that particle's position above that particle. The values of μ_{2i-2} are the number of crossings under particle i.

To further illustrate, see Figure 2.3, where we elaborate on the example given in [103]. A simple closed loop is shown in red, wrapped around 5 particles. In the top diagram we point out with arrows where crossings occur between particles, across the purple vertical lines. We can see that between particles 1 and 2, the red loop crosses the purple

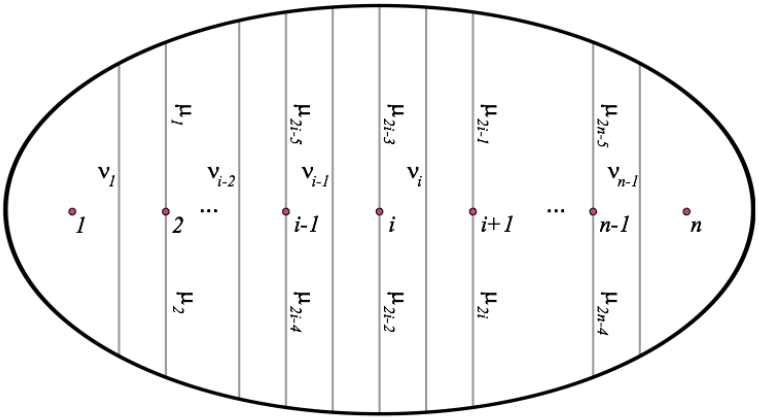

Figure 2.2: Crossing numbers. Figure taken from [103]. μ_i indicates a crossing of loop with the vertical line above or below a particle, while the $\nu_i s$ is a crossing between particles.

line twice. Therefore $\nu_1 = 2$. In the second plot we point out the crossings above and below particles, across the orange vertical lines. For example, above particle 2 the loop crosses the orange line twice, so $\mu_1 = 2$, but it does not cross the orange line below the particle, so $\mu_2 = 0$. The set of numbers $(\nu_1, \nu_2, \nu_3, \nu_4, \mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_6)$ then represent this loop.

Once we have obtained the crossing numbers we can use them to find the the minimal set describing the loop, known as the Dynnikov coordinates (see [30]), by the formulae:

$$a_i = \frac{1}{2}(\mu_{2i} - \mu_{2i-1}) \tag{2.8}$$

$$b_i = \frac{1}{2}(\nu_i - \nu_{i+1}) \tag{2.9}$$

$$\mathbf{u} = (a_1, ...a_{n-2}, b_1, ...b_{n-2}) \tag{2.10}$$

The vector $\mathbf{u}$ gives us the algebraic representation of a material line lying in a flow. We can use this to quantify how stretched it becomes in the flow by calculating the

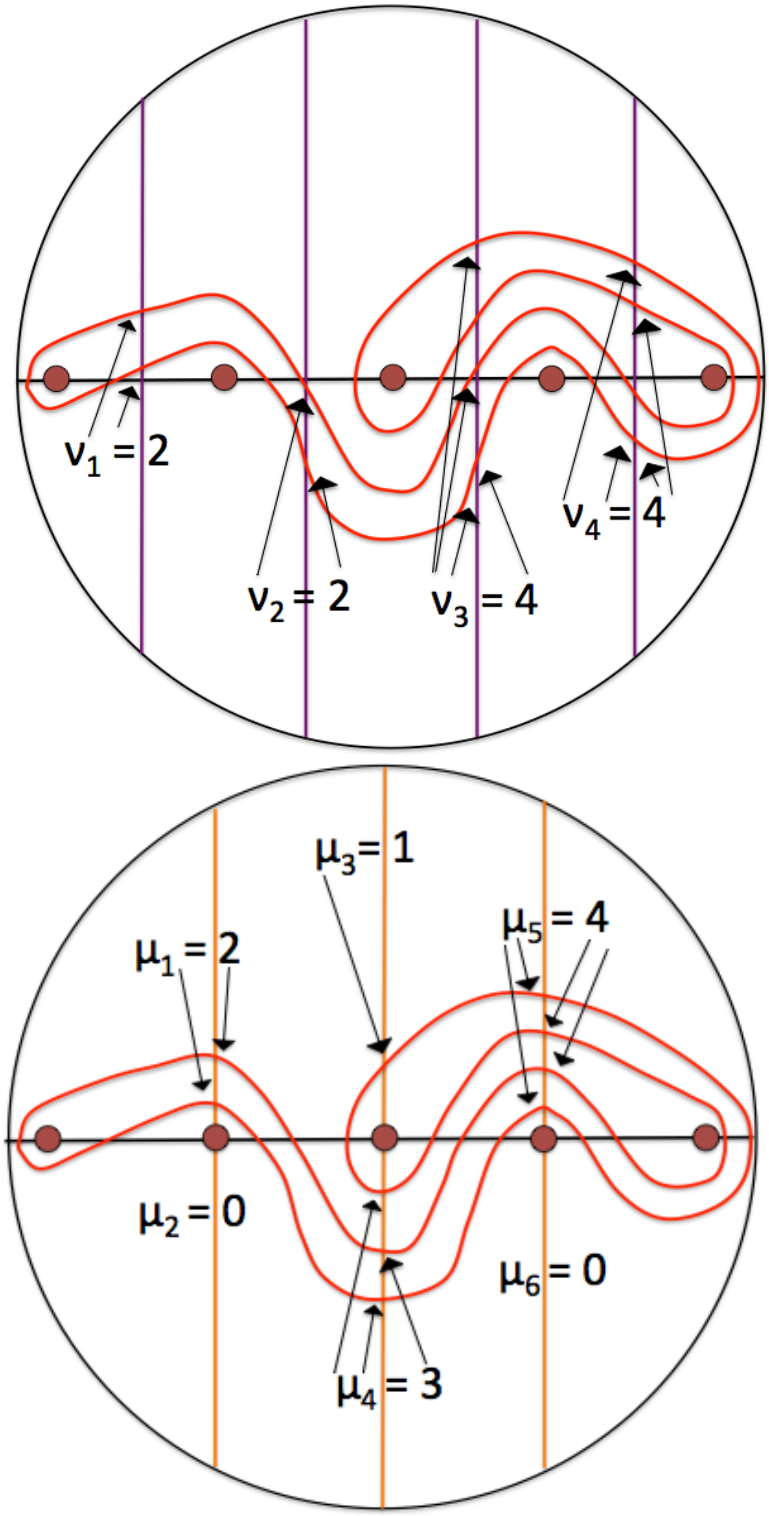

Figure 2.3: Illustrated crossings for a loop, with loop topology taken from the example in [103]. Arrows indicate the crossings that contribute to the value of each ν and μ.

number of intersections of the loop with the real axis (represented in Figure 2.3 by the horizontal line through the particles). The more crossings the more tangled the loop must be becoming. There is literature on the derivation of a formula for this (see [62]) but these specifics are not our main concern. The crux of the matter is that a formula for the minimum number of intersections of the loop around the particles can be obtained from **u**, as quoted in [103] is;

$$L(\mathbf{u}) = |a_1| + |a_{n-2}| + \sum_{i=1}^{n-3} |a_{i+1} - a_i| + \sum_{i=0}^{n-1} |b_i| \tag{2.11}$$

$$b_0 = -\max_{1\le i\le n-2} \left(|a_i| + b_i^+ + \sum_{j=1}^{i-1} b_j \right), \quad b_{n-1} = -b_0 - \sum_{j=1}^{n-2} b_i$$

$$b_i^+ = max((b_i, 0)$$

So now we have a way to encode the action of a flow by using particle trajectories to obtain our braid, a way to describe a material line lying in the flow, and a way to quantify how this loop is being stretched by the flow. In other words, we have all the ingredients we need to calculate the entropy of the braid. We can apply the action of the braid to an encoded material line, update the crossing numbers of the line (using the update rules presented in [103]), and then calculate the number of intersections now occurring. The rate of increase in intersections with many iterations gives us an estimate to the entropy.

2.2 MATLAB code and braidlab

We will briefly describe the numerical procedure used for estimating topological entropy, which uses MATLAB functions and the package braidlab. The braidlab software

and user guide (see [104]) is available through `www.math.wisc.edu/~jeanluc/software.php`.

We first generate a set of random seed particles in the domain of interest. Using the expression for the motion velocity (for us our blinking vortices) we can solve for each particle's position at regular time intervals up to a time t_{max}. With these trajectories, we use the braidlab function 'braid' to find the generators describing the braid created by them. These can then be passed to another braidlab function, 'entropy', which applies the braid to a random initial loop for a user-defined number of iterations, and counts the intersections of the loop with the real axis, using the techniques discussed previously. The estimate for the entropy is updated at each iteration: after an application of the braid, the estimate is set to the logarithm of the number of intersections minus the logarithm of the number of intersections after the previous application. Then the new estimate is compared to the previous estimate, and if the absolute value of the difference between them is less than some user-defined tolerance, the code stops and the latest estimate is taken to be the entropy. Once we reach a point where each iteration adds the same number of intersections (to within the requested tolerance) we can quantify how complex the flow is. The larger the number of intersections the flow is able to induce, the higher the degree of stretching of the loop, and the higher the entropy of the flow.

2.2.1 Issues to be Aware of

While the procedures described give good estimates of the topological entropy for many cases, it is possible to run into problems, such as those we found.

1) Initially, we calculated the braid created by particle trajectories over only one driving period. We found that our simplest, least tangling case was producing higher entropy values than what should have been the more complex cases, which intuitively can be determined to be incorrect. This was found to be due to the error introduced due to the fact that we do not have a true braid, in that the starting points do not match the finishing points. Running the driving for more periods and calculating a longer braid to apply repeatedly to the material line renders the error of closing the braid negligible, and we obtain an entropy per period. There is a balance to be found - how many periods gives a satisfactory accuracy without being too computationally expensive? This is something we will discuss further when we come to detail the calculations in Section 3.1.3 .

2) The above solution caused its own problems - running the driving for enough periods to obtain accurate estimates sometimes resulted in such long braids that we exceeded the computational limit. At this time the best option was to use the braidlab function 'complexity'. The entropy function applies the braid action to material line multiple times, iterating until the number of intersections with the real axis converges, which can lead to very large numbers. On the other hand, the complexity function only applies the braid once, thereby avoiding this problem.

2.2.2 Complexity

For the results used in the following chapter, we calculated the complexity of the driving action, where the complexity is as defined in [29]. In this description when we say complexity and entropy we are referring specifically to the complexity and entropy functions of the braidlab package. Take the braid generators representing the action of

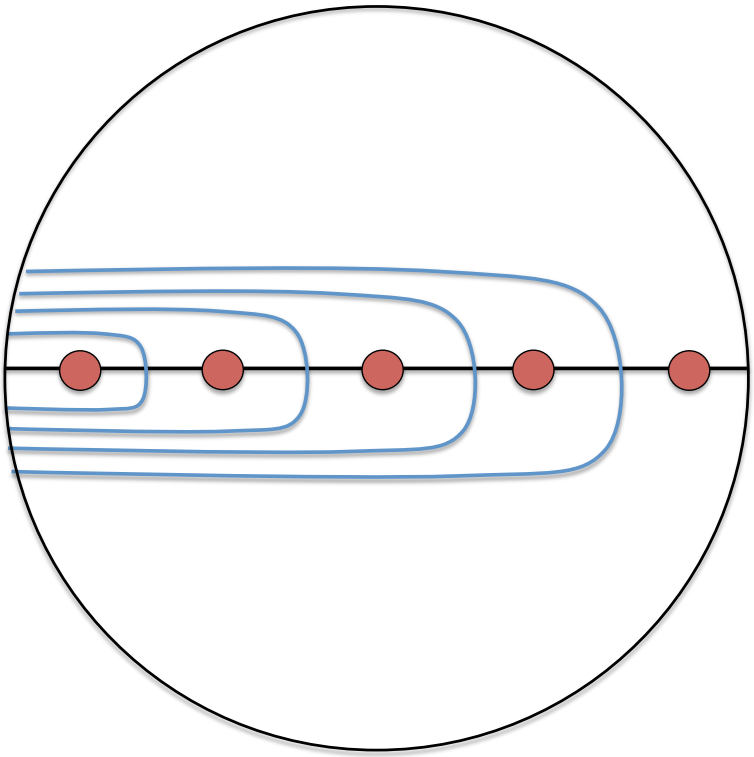

Figure 2.4: Curve diagram for the set E (blue lines) sectioning off the red seed particles.

the flow to be β. This braid will act on the set of curves E, illustrated for an example with 5 seed particles in Figure 2.4. For a system with n particles, E will have $n-1$ curves.

The five seed particles are the red circles along the horizontal axis. The set E consists of the blue lines and sections off each seed particle from the others, with each arc anchored on the boundary - these are essential arcs, rather than the random loop as in the entropy process. Each curve is referred to as being tight with respect to the real axis (the horizontal line through the seed particles), i.e intersects at right angles. Now, rather than applying braid generators repeatedly to one material line wrapped in some way around our initial seed particles as discussed for the entropy method, the complexity method applies β to E once, and counts the intersections of the transformed curves with the real axis. Since the number of intersections can increase exponentially even

with one application, the logarithm of this is taken. The logarithm of the number of intersections of E before braid action is deducted, leaving us with the intersections due specifically to the braiding. This is described mathematically in [29] by the formula:

$$c(\beta) = \log_2\|\beta \cdot E\| - \log_2\|E\| \tag{2.12}$$

where the norm represents the number of intersections. There is of course also the option of using entropy and applying just one iteration, to keep numbers down, but we would lose accuracy. We can instead apply the braid representing many periods of our driving flow to E and see how the number of intersections increases this way. Studied over many periods, both the complexity and entropy functions will describe the increase in the length of a loop in the long term, and become comparable, with the complexity function having the advantage in terms of efficiency. It is not pertinent to us to consider further details of the complexity concept - we simply need a way to quantify the relative degrees of braiding induced by different driving motions and we have one - but more information can be found in [29].

Illustration To further justify our use of the complexity, Table 2.1 shows the gradual convergence of the complexity and entropy to similar values for one of the braiding cases we will investigate in full in later sections. Running this particular driving profile on one set of seed points for many periods and examining the entropy and complexity after each period reveals that both eventually converge, and to values of the same order of magnitude. Note that the complexity estimate is larger particularly in the beginning; this is due to the fact that this function applies the motions to E, rather than to one loop only, so there are likely to be more intersections straight away. They will never give exactly the same values due to differences in the algorithms, but entropy calculations using braids are estimates anyway. Number of periods, number of seed particles,

Quantity	**After** 1 **period**	**After** 20 **periods**	**After** 80 **periods**
Complexity	5.598	0.576	0.321
Entropy	3.248	0.365	0.216

Table 2.1: Entropy versus complexity.

area driven, tolerance to convergence, number of points solved for along particle trajectories etc, can all affect the result. Furthermore, as long as we consistently use the complexity function for all drivers, then it is still valid to compare and contrast the numbers, even if those numbers do not match those which would be produced by the entropy function. The results will be consistent with each other. We will later detail the pragmatic parameter choices made in order to achieve reliable but practically attainable results. These numbers also further illustrate the earlier point of how applying a braid one period long does not provide an accurate result when the braid is not closed. It is only as we apply a braid of more and more periods that the quantities converge towards their true estimates.

Chapter 3

Photospheric Driver, Magnetic Environment and LARE3D

3.1 Photospheric Driver

3.1.1 Concept

The convective motions on and below the solar surface are understood to cause vortical plasma flows, as material which has emerged from the convection zone cools and falls back below the surface (see [65] for a full review of surface convection.) Plasma may swirl, like water down a drain, around downdrafts at intergranular lanes. Observations have also found evidence of this. Vortices on the scale of 5 Mm in diameter were observed and studied in [16] as granules were seen to swirl around each other for at least a duration of 1.5 hours. Vortical motions on a smaller scale have also been observed more as telescopes have become more and powerful. The authors of [11] discuss observing vortical motions on a smaller scale indirectly using their effects on magnetic bright points. By examining the trajectories of the bright points as they were dictated

by the 'whirlpools', the authors estimated a lower bound of 0.9×10^{-2} such vortices per square megametre, each one affecting the surface plasma over an area of around half a megametre squared, with lifetimes of around 5 minutes. The review [66] refers to a series of papers which highlight observations of the action of vortical flows, again flagged up by the motions of bright points which were influenced by the presence of the vortices, for example in [5], which considers motions lasting about 20 minutes. In fact, vortical flows have even been seen to occur spontaneously in the modelling of solar convection, such as in the experiments in [100]. Clearly then these types of motions are frequent and occur on a range of scales. This evidence, coupled with the intention to continue in the style of previous braided field simulations, led us to design the following driving functions.

The driver that will induce the braiding consists of 'blinking vortices', with incompressible velocity profiles. By 'blinking' we mean that each vortex spins alternately, each one blinking on and off: one vortex begins to spin, reaching a maximum speed, then slowing to a stop, at which point the second vortex follows this same pattern, after which one period of the driving is completed. This periodic action will be applied repeatedly throughout the simulations. On the Sun, photospheric motions are constantly distorting the atmospheric fields so we will also apply our mixing for the duration of our experiments.The vortices are positioned in the $x-y$ plane, which for us represents the photosphere. They were initially chosen to behave in this blinking fashion, with opposing directions of twist, in order to create a braid similar to those studied in [112]. In these papers the so-called pigtail braid was the initial condition for simulations. Here we wish to take further the ideas of this work, and create the braiding by boundary motions as part of the experiment to assess the impact of the braiding action. In an ideal case, imposing three full periods of this type of driving on a uniform $\mathbf{e}_z$ field would create a braid with six pockets of twist as in the E3 braid used in [112].

Note that in this work, since we are concerned with identifying the impact of particular types of driving motions, we are constructing a driver which can be adapted to create different quantifiable categories- not the most physically realistic driver based on observations. At this stage, we wish to be able to say if one category of motion seems to give rise to a certain type or level of heating. This is intended to be a somewhat idealised set-up. Nevertheless, we can point back to the observations of [11] where the authors concluded that vortices with clockwise and anti-clockwise twists occur with roughly the same frequency - this would correspond to our oppositely twisting case. Furthermore, the freedom to choose normalisation constants for simulation results (as we will elaborate on in Section 3.3) means that in theory, the vortices designed could represent flows of the sizes studied in the aforementioned observational papers.

Taking this idea further we also consider blinking vortices with equal circulation. Let us explore in detail the characteristics of systems under the influence of such types of motion. Recall the limiting cases we discussed in 2.1.5, i.e. the two extremes where we could have the vortices centred at the same point, or at an infinite distance from each other. Here we elaborate on these ideas.

Consider an ideal case with no diffusion. Intuitively we predicted that the case of oppositely twisting vortices will provide opportunity for more heterogeneous, complex mixing. With both vortices placed at the origin, the twisting of the second vortex will undo, in an ideal case, the action of the first vortex and any braiding will be undone. In addition, helicity injected by one vortex will be removed by the alternate twisting of the other. At an infinite distance from each other, the tangling induced by one vortex will be unaffected by the other, so the degree of complex tangling will be low. The overall helicity in the whole volume will be zero again since the total twist will cancel out. At some point between these two extremes we will see a maximum of complex

tangling, when vortices are at a distance apart where twisting of some fieldlines and counter twisting of others will induce maximum criss-crossing, creating a braid with the most intricate tangling. Still however, in an ideal scenario the total helicity in the domain will be zero. In a real life resistive case, diffusion and reconnection can occur, so a twist of one sense will not necessarily be followed by a twist of the opposite sense such that helicity is exactly cancelled out, but in general here we will have a very low helicity environment.

A further thought experiment can be considered for a scenario where both vortices twist with equal circulation. If both vortex centres are the same, the twisting of the second vortex will only add to the twist of the first. Each action will continue to twist the overlying field into a tighter spiral. Fieldlines will not criss-cross in different directions, but wind round together in a tight, coherent, helical structure. Therefore the tangling is less complex. Similarly for vortices placed at a distance of infinity - but this time we will see two helical structures formed in twice the time. In both cases, only twist in one sense will occur. In fact in a purely ideal scenario where all other boundaries are magnetically closed, the total helicity in the volume will increase indefinitely at the same rate for both cases. Even at points in between these two extremes we would still see the twist increase consistently. This scenario will always add helicity. The complexity of the braids created however may vary as we place the vortices at a distance in between the two extremes, and the tightness of the spiral will vary. Twisting will still be of the same sense, with fieldlines winding around in a consistent fashion, but not symmetrically around one axis, so there will still be some point of maximum complexity.

Summarising this then, we have formed Table 3.1. Taking the above arguments into account, we can say that the oppositely twisting case will have lower helicity and higher complexity than the equally twisting case.

Case	Circulation	Complexity	Helicity
1	Opposite	Complex	Low
2	Equal	Coherent	High

Table 3.1: Driver Properties. Circulation refers to the relative directions in which the blinking vortices will spin.

3.1.2 Driver Design

Now that we have explored the type of simulations we wish to create with a blinking vortex photospheric driver, we move on to the specifics of the driving profile.

We want to create a system whereby we can easily adjust the distance between vortex centres and the direction of twisting. In addition, we want to be able to control the rate of ramping up and down, the amplitude of driving and speed of driving. Let us choose now to centre one vortex, v_1, at $(0,0)$, and the second, v_2, at $(x_0,0)$. Also, set v_2 always to rotate clockwise, while v_1 will spin anticlockwise for Case 1 and clockwise for Case 2. Then the velocity components are given by:

$$v_x(x,y,t) = A_1(t)ky\exp\left(\frac{-x^2-y^2}{a}\right) + A_2(t)y\exp\left(\frac{-(x-x_0)^2-y^2}{a}\right),$$

$$v_y(x,y,t) = \underbrace{-A_1(t)kx\exp\left(\frac{-x^2-y^2}{a}\right)}_{\text{from } v_1} - \underbrace{A_2(t)(x-x_0)\exp\left(\frac{-(x-x_0)^2-y^2}{a}\right)}_{\text{from } v_2},$$

where A_1 is the time dependent factor for vortex 1, and A_2 the time dependence of vortex 2. We will choose $a = 2$, thereby making the circle of influence of the vortices comparable to the region of twist in [112]. The constant k will determine which way v_1 is spinning: $k = -1$ will set it anticlockwise for Case 1, while $k = 1$ will set it clockwise for Case 2. Now consider the time-dependent terms. Over the first period

$t = [0 : 2L]$, we wish to have only v_1 spin over the first half period L and only v_2 over the second half period L. We set:

$$A_1(t) = \begin{cases} a & 0 < t < \frac{\pi}{m} \\ 2b & \frac{\pi}{m} < t < \frac{n\pi}{m} \\ c & \frac{n\pi}{m} < t < L \\ 0 & L < t < 2L, \end{cases} \qquad A_2(t) = \begin{cases} 0 & 0 < t < L \\ a & L < t < L + \frac{\pi}{m} \\ 2b & L + \frac{\pi}{m} < t < L + \frac{n\pi}{m} \\ c & L + \frac{n\pi}{m} < t < 2L, \end{cases}$$

where:

$$\begin{aligned} 2b &: \text{amplitude of maximum driving} \\ a &= b\cos(mt - \pi) + b : \text{ramping up,} \\ c &= -b\cos(mt) + b : \text{ramping down,} \\ L &= (n+1)\pi/m : \text{half period.} \end{aligned}$$

The constants m and b affect the slope of the ramp up and down, n the duration of maximum driving and therefore the period, and additionally b also dictates the amplitude of the driving. Incorporating this all together then we have an expression for one period of driving.

The constants m, n and b were chosen using trial and error in order to meet loose criteria. We wish the maximum driving speed never to exceed $\sim 10\%$ of the initial Alfvén speed, maintaining relatively slow motions. The twists in the E3 braid of works such as [112] twist the field by a maximum of $\pm\pi$, and we are taking inspiration from this research. The driver does not twist by a uniform amount as we move radially away from vortex centres, so particles starting at different points will be mapped by different

angles. We chose constants such that the maximum twist induced by each vortex will be about $|\pi|$, also. This was judged simply by eye using plots of trajectories of seed particles - it is not exact. The maximum velocity occurs at radius of about 1 from the vortex centre and for a particle starting at $(1,0)$, the action of vortex 1 in Case 1 brings it to sit at about $(-1,0)$. In Case 2 a particle seeded at $(-1,0)$ will be moved to $(1,0)$ by the first vortex. Only particles starting on the circle of radius one centred at a vortex will be moved by $|\pi|$ by the action of that vortex. The resulting function has a total period of 24π time units. That is, the first vortex will ramp up in π time units, drive at its maximum for 10π, and ramp back down to a stop by 12π units, after which the second vortex does the same in the same time. The values giving these properties and those used in subsequent experiments are: $m = 1$, $n = 11$, $b = 0.075$.

Recall in Chapter 1 the discussions of relative helicity. We deduced that the helicity current can be calculated by:

$$-\int_S (\mathbf{A}\cdot\mathbf{v})B_z \mathrm{d}S. \tag{3.1}$$

We can calculate the helicity each driver induces using the presented formulae and setting $B_z = 1$. If we set:

$$A_x = -\frac{y}{2},\ A_y = \frac{x}{2},$$

then plugging these into Equation 3.1 and following through the indefinite integrals gives;

$$\text{Flux from Vortex 1} : -2\pi k A_1(t)$$

$$\text{Flux from Vortex 2} : -2\pi A_2(t).$$

Note that this does not depend on x_0, so we know the helicity wherever we choose to place the vortices. Over the course of one period then for the oppositely twisting vortices, $k = -1$, and:

$$\begin{aligned} \mathrm{H}_1 &= H_{v_1} + H_{v_2} && (3.2)\\ &= -\left(\int_0^{12\pi} +2\pi A_1(t)\mathrm{d}t - \int_{12\pi}^{24\pi} 2\pi A_2(t)\mathrm{d}t \right) && (3.3)\\ &= -32.57 + 32.57 = 0. && (3.4) \end{aligned}$$

For the equally twisting case, $k = 1$, and:

$$\begin{aligned} \mathrm{H}_2 &= H_{v_1} + Hv_2 && (3.5)\\ &= -\left(-\int_0^{12\pi} 2\pi A_1(t)\mathrm{d}t - \int_{12\pi}^{24\pi} 2\pi A_2(t)\mathrm{d}t \right) && (3.6)\\ &= 32.57 + 32.57 = 65.14. && (3.7) \end{aligned}$$

These are in the same non-dimensional units used in the simulations we will carry out later. We now know exactly the helicity each driver is able to inject over 1 period. All of the Case 1 drivers, regardless of x_0, will overall inject no net helicity. All Case 2 drivers, again regardless of x_0, can inject 65.14 units of helicity each period. The

last step is to extend this to an infinitely periodic function which can be applied at the lower z boundary for as long as we choose, by setting time t to be the time modulo $2L$ in the above expressions when programming this boundary condition in Lare3d, so this driving is repeated for the simulation duration.

3.1.3 Calculating Complexity

We return to the concept of topological entropy: our method of measuring the different drivers' abilities to create complex fields. It is pertinent to assess the impact of different types of motions on potential heating. Remember our work is concerned with the argument about whether reconnection of braided fields can contribute to a background level of heating in the solar atmosphere. Fields can be braided in different ways by different types of photospheric motion, so carrying out this analysis will allow us to describe how different categories of motion could lead to heating.

We have discussed how the complexity of the tangling induced by driving will depend on where the second vortex centre x_0 lies along the spectrum of $x = 0 \rightarrow \infty$. Therefore in our calculations we wish to examine each driving function for different values of x_0.

In this work we ran many tests to help decide the best course of action with regards to this and also due to the issues discussed in Section 2.2.1. Let us recap a little before we continue.

The code provides an estimate to the true entropy of a flow by estimating the entropy of a braid of trajectories, where the trajectories are the paths of seed particles in the flow across a two dimensional plane lifted up in time. Therefore, the more particles we follow, the closer our estimate to the true value. Further, the braid of trajectories from our flows are not true braids in the sense that they are not closed: at the end of a driving

period the particles have not returned to their starting point. Therefore there is an error which can lead to non-zero entropies where there should be zero entropy. We get around this by calculating the entropy of a braid constructed from more periods of the flow velocity to find an entropy per period, which reduces the error and gives a more accurate answer. We also talked previously about the fact that the entropy programme in the braidlab package we use iterates many applications of a braid, and in some cases repeatedly applying a long braid can lead to attainment of the computational limit and a failure of the code. We moved instead to using braidlab's complexity function, which also gives a sense of the degree of tangling induced with values comparable to the entropy function over many periods, without leading to overflow since it only applies braid action once.

With all of this in mind, we summarise the tests carried out to determine the parameters to use.

We constructed tables of data, whereby we calculated the complexity of a braid constructed by n particles over m periods. Time and computational power must be taken into account - we could run for more periods or for more particles but the increased accuracy may not be great enough to justify the extra calculations. Tests were carried out where both values ran from 10 upwards by 10, such that $m = n = 10$ gave the least accurate answer. From previous tests using the entropy function we judged that the value $x_0 = 2$ for oppositely twisting vortices would be around the point where we would have the most complex braids. Therefore, if we could carry out analysis for this point with our chosen m and n, we would hopefully not run into difficulties for other, simpler braids. For $x_0 = 2$, Case 1, then, we judged the payoff between accuracy and computational expense to be around $n = m = 110$.

Now, the seed particles are randomly generated by the code from an area which changes with x_0. We want to fairly capture braiding created by each set-up, and as x_0 increases,

the influence of the driving will be felt over a larger area. Vortex 1 will remain at $(0,0)$, and the influence of each vortex has tailed off beyond a radius from the centre of 2.5. Therefore the x co-ordinate of a seed particle will be $x \in [-2.5, x_0 + 2.5]$. In y we select random particles between $y \in [-2.5, 2.5]$.

Clearly, because we are dealing with randomly selected seed particles, we can run the same experiment twice and get a slightly different answer. To account for this, and since the accuracy of an experiment is always improved by repeating and averaging, for each x_0 we run the code for the same five sets of randomly selected particles and average the results over the five realisations. This also allows us to calculate a confidence interval to show the variation depending on the initial set of seed particles.

Once our seed particles are chosen, their trajectories under the influence of each type of driver described previously are calculated, and the complexity of the braid determined. Finally, we can present our results. Varying x_0 between 0 and 4 at intervals of 0.25, we end up with Figure 3.1.

The two main features of Figure 3.1 are:

1) In general, Case 1, with oppositely twisting vortices and zero net helicity injection, gives rise to larger complexity (with the exception of $x_0 < 0.5$, which we will discuss shortly). This is as we predicted - the crossing over of fieldlines induced by Case 1 leads to complex tangling, whereas Case 2 leads to more coherent twisting and therefore simpler braids.

2) For both cases, the entropy reaches a peak at some value $x_{0_p} > 0$ before declining towards zero for $x_0 > x_{0_p}$. Once again, this is as expected from thought experiments about the interaction between vortical motions at various distances.

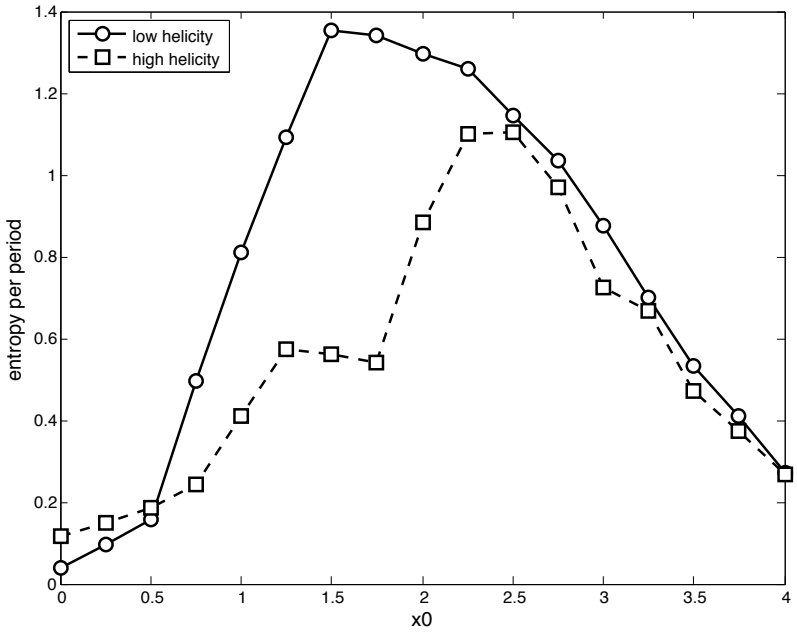

Figure 3.1: Entropy comparison for varying x_0 for the high and low helicity cases. Points are plotted at the mean value calculated for each run from multiple tests. The solid line corresponds to oppositely twisting case where helicity injection is zero, while the dashed line represents the equally twisting vortex case, which injects helicity consistently.

We address the behaviour for $x_0 = 0$. When both vortices lie on $x_0 = 0$, the entropy should mathematically be zero. In Case 1, a material line in flow will be stretched one way and then stretched by the same amount again in the opposite direction. Repeating this stretching does increase the length of a line, but the growth rate is linear. For Case 2, each twist stretches by the same amount in the same direction, so the growth rate again is linear. A linear growth rate puts both cases in the finite order category of the Thurston-Nielsen theorem discussed in Section 2.1.3, therefore the topological entropy is zero. It is only as the vortices are placed further apart and the distortion of a material line is less predictable and uniform that the opportunities arise for exponential growth rates.

Number of Periods	1	10	100	150	200
Case 1 **entropy**	0.412	0.198	0.044	0.029	0.023
Case 2 **entropy**	5.927	0.923	0.124	0.087	0.068

Table 3.2: Convergence of entropy values when $x_0 = 0$ towards zero when calculations are made over more periods.

Clearly our calculations have a positive value for both scenarios which should have zero entropy. However we can see that as we run the code for more periods, the value does converge towards zero, and the values we have plotted here are due to the choice of when to stop. We can check this by running calculations for these two scenarios for even longer, as in Table 3.2. Notice that for Case 2 the estimate starts off much larger and converges more slowly, explaining why we have a larger value for this coherent case plotted in 3.1. After the first period of driving, the oppositely twisting case has undone tangling - the net growth of a material line is small, and the error is mostly in closing the braid as discussed in section 2.2.1. In the equally twisting case, twisting is not undone by the second vortex, but added to, and the length of a material line increases, with the addition of the error in closing the braid. It is only with repeated application over many periods that it becomes clear that the growth rate of a material line in the flow is linear, for both cases, as shown by the estimates' direction of travel .

Taking all of this into account, we now select and categorise specific set-ups to run full Lare3d simulations on. Since the coherent twisting of the high helicity case leads to simpler braiding, we can refer to Case 1 as the complex, zero helicity case, and Case 2 as the coherent, zero helicity case. Using Figure 3.1 as a guide, we choose values of 0.5, 1 and 2 for x_0 and run a Case 1 and 2 simulation for each. In total we will have 6 runs to examine. This is all summarised in Table 3.3, where we quote the

Run	**Circulation**	**Helicity**	x_0	**Entropy** (95% **C.I)**	**Category**
1A	Opposite	0	0.5	0.161 ± 0.004	Complex
1B	Opposite	0	1	0.786 ± 0.017	Complex
1C	Opposite	0	2	1.321 ± 0.028	Complex
2A	Equal	65.14	0.5	0.265 ± 0.104	Coherent
2B	Equal	65.14	1	0.444 ± 0.046	Coherent
2C	Equal	65.14	2	0.888 ± 0.039	Coherent

Table 3.3: Entropy calculations with 95% confidence intervals for each driver and their characteristics.

entropy values as the mean result of the multiple runs for each scenario with errors corresponding to a 95% confidence interval.

Recall that this work aims to build on the findings of [112] which studied the 'E3' braid. Our choice of set-ups does in fact include one which is analogous to the topology of the E3 braid unit. For ease of explanation, consider tracing 3 fieldlines of E3 as Figure 2a of [112]. The braid is constructed of 3 units, each unit containing one positive twist and one negative twist. This could be described by the braid generators $\sigma_1\sigma_2^{-1}$. The three fieldines are rooted at $x = -2, 0, 2$. The regions of twist are centred at $x = -1$ or $x = 1$, i.e a distance of 2 units apart along the x-axis. The photospheric velocity which will be applied in Run 1C has oppositely twisting vortices, centred at $(0,0)$ and $(2,0)$, making this the analogous case. Here the braid of trajectories of seed particles placed initially at $x = -1, 1, 3$ (echoing the anchor positions of the E3 sample fieldlines) after 1 period of the flow can also be described as $\sigma_1\sigma_2^{-1}$. We plot these trajectories, with time along the y-axis, of these seed particles over the course of one driving period in Figure 3.2. The only difference between this braid and the first unit of the E3 fieldines in [112] (apart from a shift along the x axis) is that due to the order of twisting (vortex 1 followed by vortex 2 as opposed to vortex 2 followed by vortex 1), the order of interaction between fieldlines is altered.

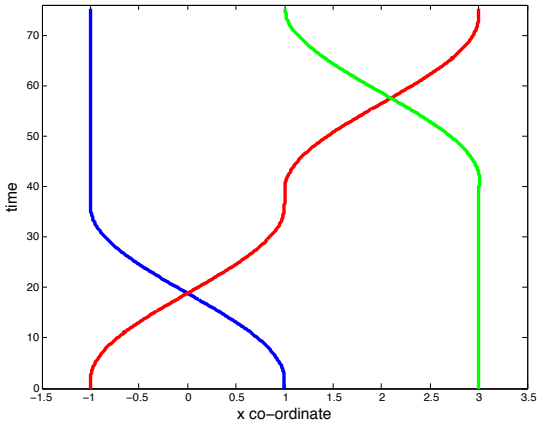

Figure 3.2: E3 braid unit illustration. Three illustrative fieldlines are plotted, showing the one anti-clockwise and one clockwise twist which make up each fundamental unit. E3 consists three of these units.

3.2 LARE3D

3.2.1 Code Features

This work uses the Lagrangian Remap code Lare3d. The code was developed by Arber, Longbottom, Gerrard and Milne, [2] and is available from `http://ccpforge.cse.rl.ac.uk/gf/`. Having previously presented the MHD equations in a comprehensive, non-normalised form in Chapter 1, Section 1.1.4, we now consider the Lare3d specific equations, including only the processes we are interested in for this work. The code uses normalised MHD equations where quantities can be specified in terms of a length-scale L_0, magnetic field strength B_0 and plasma density ρ_0. The normalised 3D resistive equations solved are;

$$\frac{D\rho}{Dt} = -\rho\nabla\cdot\mathbf{v}, \tag{3.8}$$
$$\frac{D\mathbf{v}}{Dt} = \frac{1}{\rho}(\nabla\times\mathbf{B})\times\mathbf{B} - \frac{1}{\rho}\nabla P, \tag{3.9}$$
$$\frac{D\mathbf{B}}{Dt} = (\mathbf{B}\cdot\nabla)\mathbf{v} - \mathbf{B}(\nabla\cdot\mathbf{v}) - \nabla\times(\eta\nabla\times\mathbf{B}), \tag{3.10}$$
$$\frac{D\varepsilon}{Dt} = -\frac{P}{\rho}\nabla\cdot\mathbf{v} + \frac{\eta}{\rho}j^2, \tag{3.11}$$
$$\mathbf{j} = \nabla\times\mathbf{B}, \tag{3.12}$$

where $\mathbf{B}$ is the magnetic field, $\mathbf{v}$ is the plasma velocity, $\mathbf{j}$ is the current density, ρ is the plasma density, and j^2 is the current density squared. The variable P is the pressure which takes the form $P = \rho\varepsilon(\gamma-1)$, where γ is the ratio of specific heats, and ε is the internal energy density. We will return to the value for η in the next section.

To deal with potential shocks or gradient build ups, Lare3d uses artificial, or shock viscosity. The formula for this is not included in the quoted equations since in a 3D, resistive scenario in particular it is far from trivial, and involves many details that we do not need to be concerned with for this work. The technique consists of adding a tensor shock viscosity term into the momentum equation and energy equation, which is only non-negligible where gradients become large enough. More details can be found in [2] and the references therein.

The code allows the user to choose to include effects such as conduction, radiation and gravity. At this point in our work we will neglect all of these: the coronal pressure scale height is much larger than typical length scales so gravitational effects are negligible; and we can assess the relative heating inputs of the different driving scenarios without involving radiation and conduction, which would be a non-trivial task.

It is possible to have uniform grid spacing, or to stretch the grid in any or all directions. In this work we keep the grid spacing uniform. The user can deploy pre-coded boundary conditions of periodic, open and line-tied boundaries, or alternatively write their own. Also user defined are the initial conditions, normalisation constants, resistivity, resolution and physical grid - we will elaborate on our choices for these in the following sections.

The code will output energy data along with the velocity and magnetic field components, density, pressure and the numerical grid values.

Like many numerical codes, Lare3d uses finite differencing to solve the MHD equations. An important feature of Lare3d to be aware of is the fact that variables are set up on a staggered Eulerian grid. This is most clearly illustrated in Figure 3.3, which is taken from [2]. The magnetic field components are specified in the centre of cell faces - B_x is defined at the centre of each cell face lying in the $y-z$ plane, B_y at the centre of cell faces along $x-z$ and B_z in the middle of cell faces in the $x-y$ plane. This allowed the code writers to use existing numerical methods for maintaining the initial value of $\nabla \cdot \mathbf{B}$. More straightforwardly, each of the velocity components are prescribed at cell vertices, while all other quantities are prescribed at cell centres; further details can be found in [2], but the reason for this is to avoid 'checkerboard' instabilities .

Lare3d follows a predictor-corrector Lagrangian procedure, where each timestep is split into two. Updated variables are then mapped back onto the original Eulerian grid at the end of the full time step. The advantage this method has over other MHD codes is that it does not use equations in their conservative form which can lead to inaccurate temperature calculations (further details on this can be found again in [2], but for now more explanation is not pertinent to our work). Lare3d can deal with shocks in ideal and resistive environments. The code is second order accurate in time and space.

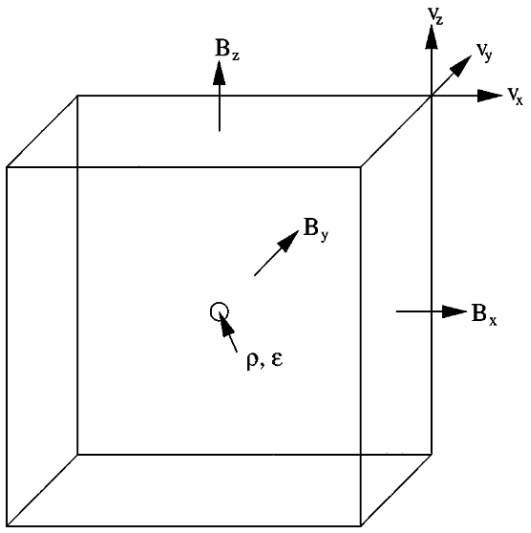

Figure 3.3: Illustration of the staggered grid used in Lare3d. This taken from [2]. For a three dimensional grid, all scalar variables are defined at the centre and all velocity components at the cell vertex. Each magnetic field component is defined in the middle of the cell face perpendicular to its direction.

3.3 Magnetic Environment

3.3.1 Simple Uniform Field

Our first set of experiments will each take a uniform field, $\mathbf{e}_z$, as the initial condition. We will apply our six drivers to the lower z boundary and investigate the outcomes of each after long duration driving. Figure 3.4 shows some fieldlines plotted over the domain of our initial condition.

We will keep the upper z boundary fixed, as in Parker's topological dissipation set-ups. The x and y boundaries will be periodic, so that any flux lost across one boundary will be replaced across the opposite boundary.

Unlike in entropy calculations where we varied the area of driving with x_0 to take in a proportionately fair amount of tangling, here we will keep the same physical domain

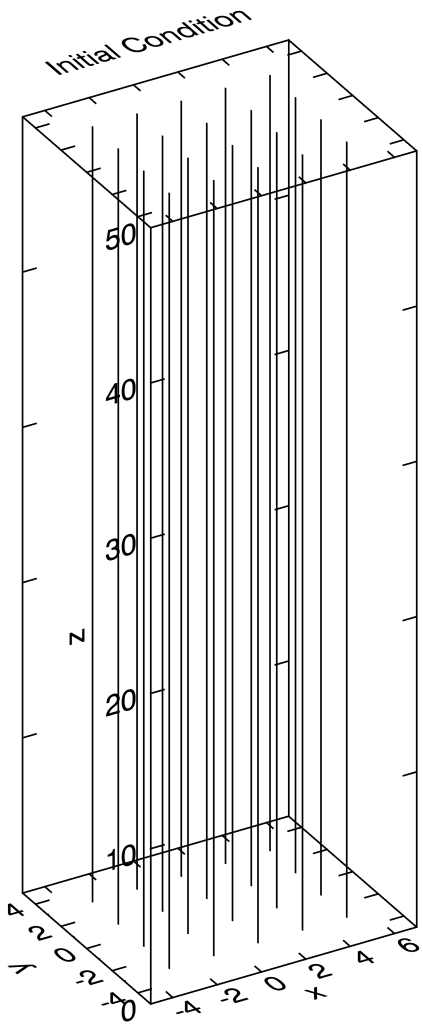

Figure 3.4: Some field lines traced showing the initial condition for this set of runs. We have a simple uniform field, $\mathbf{e}_z$.

for each run, but ensure the $x-y$ plane is large enough for all cases to include fieldines under significant influence of the driving. Therefore, $x \in [-5,7]$ and $y \in [-5,5]$. The z axis ranges from 0 to 50 - this is somewhat arbitrary and gives an aspect ratio smaller than is physically realistic, but computational expense and grid resolution must be taken into account.

The normalisation process of the MHD equations is such that all normalisation constants drop out and theequations are effectively unchanged. We can choose what to set normalisation constants B_0, L_0, and ρ_0 (and in turn all other normalisation constants which are defined in terms of these) to later, as it makes no difference to the numbers produced by the simulations and we may wish to interpret the simulation quantities for different solar scenarios.

The length of the domain in simulations units is 50. The simulation Alfvén speed (where hat variables are non-dimensional simulation variables) is as follows:

$$\begin{aligned} \hat{\mathbf{v}}_A = \frac{\mathbf{v}_A}{v_0} &= \frac{B_0\hat{\mathbf{B}}}{\sqrt{\mu_0\rho_0\hat{\rho}}} \Big/ \frac{B_0}{\sqrt{\mu_0\rho_0}} \\ &= \frac{\hat{\mathbf{B}}}{\sqrt{\hat{\rho}}}. \end{aligned}$$

So the time it takes for an Alfvén wave to travel from the lower z boundary to the top of the domain is:

$$t_A = 50 \Big/ \frac{\hat{\mathbf{B}}}{\sqrt{\hat{\rho}}}.$$

Initially, we have a uniform field $\mathbf{e}_z$, so setting $\rho = 1$ gives us an Alfvén speed of 1 and an Alfvén time of 50. Therefore our driving speed never exceeds 10% of the initial

Property	Value/Nature
x, y boundaries	periodic
Upper z boundary	line-tied, fixed
Lower z boundary	line-tied, driven
Initial **B**	$\mathbf{e_z}$
Initial **v**	**0**
Resolution	256^3
Domain	$[-5,7]\times[-5,5]\times[0,50]$
η	0.0005, uniform
Initial ρ,ε	1.0,0.01
viscosity	shock capturing form
Duration	45000 time units 596 driver periods

Table 3.4: Simulation details for the uniform field experiment.

Alfvén speed since we designed it to have a maximum speed of 0.1. We wish to run for many Alfvén times, and so set the simulations to run for 45000 time units, which is 900 initial Alfvén times and 596 full driver periods.

Now we come to the resistivity. Due to the normalisation process, the η specified in the code is not the resistivity $1/\sigma$, but the inverse Lundquist number. Therefore we do not actually set a value for resistivity. The Lundquist number $S = L_0 v_A/\eta$ (where here η is the magnetic diffusivity) is a particular case of the magnetic Reynolds number R_m, where **v** is the Alfén speed. This means that the larger we set $\eta = 1/S$, the more dissipation will take place. The experiments of [114] tested dissipation constants of $0.01, 0.005$ and 0.0002. We chose to take $\eta = 0.0005$ for our tests.

Lastly, the only viscosity we will include is shock viscosity, and the specific internal energy density is initialised to 0.01. Table 3.4 summarises and sets out other properties of this field.

3.3.2 Magnetic Carpet Field

In our second set of simulations, we will again apply our drivers on the photosphere, but we will use a magnetic carpet initial condition. We wish to include null points, and one particular simple null topology observed frequently on the Sun is referred to as a 'parasitic polarity'. This is the basic unit which we will include in our field. It has been investigated before, for example in [81], who considered the nature of reconnection at these points. In this configuration, the null point sits just underneath the top of a dome structure which is created by its fan. These separatrix fieldlines will form a cage around the null, anchored back down in the photosphere, so that the null is isolated in a region of one polarity, outside of which the local field has opposite polarity. The strength of the field enclosed in the fan is smaller than that of the external field. We will include several of these parasitic polarities, sitting in an otherwise uniform $\mathbf{e}_z$ field, and apply our drivers at the photosphere again.

The structure is designed by first creating a vector potential $\mathbf{A}$, to ensure that the resulting field satisfies $\nabla \cdot \mathbf{B} = 0$ exactly. This will consist of the sum of the expressions for each point source which will create a parasitic polarity region;

$$\mathbf{A} = \sum_{i=1}^{n} \varepsilon_i \left(\frac{(y-y_i)}{((x-x_i)^2+(y-y_i)^2+(z-z_i)^2)^{3/2}}, \frac{(x-x_i)}{((x-x_i)^2+(y-y_i)^2+(z-z_i)^2)^{3/2}}, 0 \right),$$

where n is the number of sources, (x_i, y_i, z_i) is the position of a point source and ε_i dictates the sign and strength. In our field we place three point sources beneath $z = 0$ in order to create three parasitic polarities above the surface. We choose the set of points $\mathbf{x}_1 = (-2, 0, -0.85)$, $\mathbf{x}_2 = (0, 0, -0.85)$ and $\mathbf{x}_3 = (2, 0, -0.85)$ and each ε equal to -2.0. Taking the curl of this gives us an expression for $\mathbf{B}$. Lastly, we add the

uniform background field $\mathbf{e}_z$ so that:

$$\mathbf{B} = \nabla \times \mathbf{A} + \mathbf{e}_z.$$

This particular field is also a potential one, so our initial condition should have no current. This basic unit is illustrated in Figures 3.5 and 3.6, both taken from [19]. Figure 3.5 shows the magnetic skeleton of the basic unit, consisting the three nulls and their separatrix domes and spines. The red spheres mark the positions of the nulls, while the green tubes depict the spine fieldlines and the blue mesh the fan separatrices. Figure 3.6 plots some fieldlines, with the nulls again denoted by red spheres. In the region under the nulls, some field lines descend back down under the photosphere. Outside of these positions fieldlines trace up the box and become part of the dominant $\mathbf{e}_z$ field.

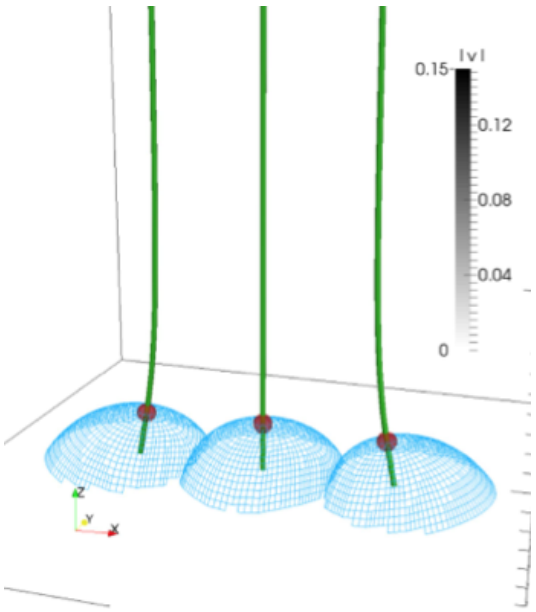

Figure 3.5: Magnetic carpet initial condition, taken from [19]. This figure illustrates the magnetic skeleton of the basic unit field. The nulls are the red spheres, enclosed in their blue separatrix cages, with the green spines reaching down below the photosphere and up into the uniform field region.

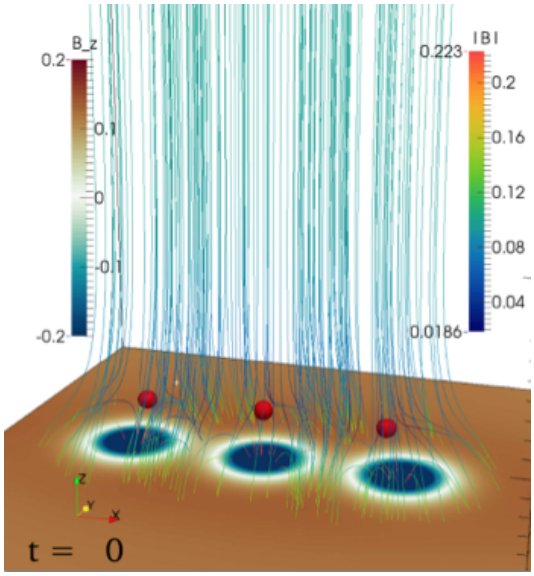

Figure 3.6: Some fieldlines of the magnetic carpet initial condition, taken again from [19]. We can see how some fieldlines curl under the separatrix domes, while others flow over the top of these fan surfaces. Null points are again denoted by the red spheres.

The x and y components of the magnetic field on the photosphere are plotted in the vector field in Figure 3.7. The arrow size is proportionate to field strength. On $z = 0$, B_x and B_y are small outwith the non-uniform region - the field is dominated by $\mathbf{e}_z$ away from the parasitic polarities. We can see that on the photosphere, underneath where the domes would sit, the horizontal magnetic field is oriented towards the centre of the circle which would be created by the separatrices' intersections with the photosphere. This is where we could see fieldlines receding under the surface in 3.6. In this region, vertical flux is negative.

This basic unit can easily be replicated over a larger area. Mirroring the point sources first in x, then in y, and then in x and y leads to copies of the configuration positioned around the original. In addition to the previous three point sources at (x_i, y_i, z_i) we position eight new groups of three surrounding them, at points;

$(x_i \pm 8, y_i, z_i),\ (x_i, y_i \pm 8, z_i)$

$(x_i + 8, y_i - 8, z_i),\ (x_i - 8, y_i + 8, z_i)$

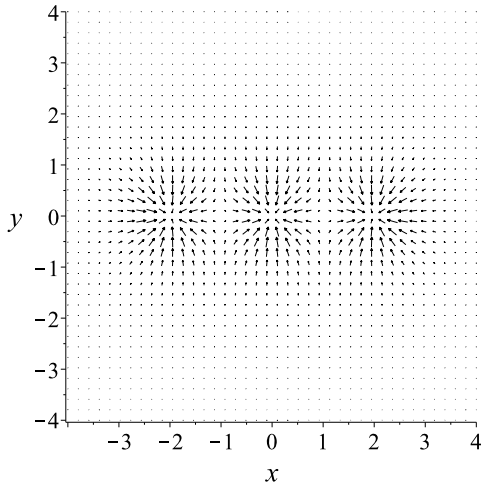

Figure 3.7: Vector representation of the horizontal magnetic field on the photosphere. This component is relatively weak beyond the region lying under the dome structures. It is concentrated inside the circles formed by the intersection of the fan meshes with $z = 0$.

$(x_i \pm 8, y_i \pm 8, z_i)$.

Adding these in gives us a field with the parallel component on the photosphere as shown in Figure 3.8.

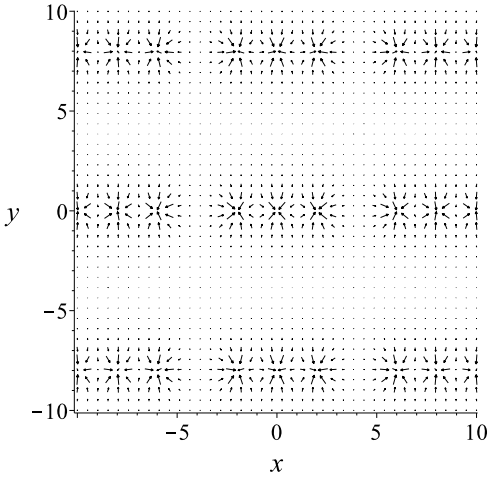

Figure 3.8: Mirrored horizontal vector field at $z = 0$. The original three structures centred at the origin are now surrounded by replicas.

Over the domain we used in previous simulations, we shall take in a fourth parasitic polarity along $y = 0$. This can be seen in Figure 3.9 where we plot the vector representations firstly of the horizontal field on $z = 0$, and secondly the vertical field at $y = 0$. The first is simply the same view as before but taking in only our domain of $[-5,7] \times [-5,5]$. In the second we can see the difference in topology underneath where the separatrix surfaces sit compared to above. Being a parasitic polarity setup, the field strength around the nulls is weaker than the rest of the field, to the point that plotting the vector field with proportionate arrow size produces a field where the detail around the null is not visible. Therefore, in this one plot the arrow size is fixed. Figure 3.10 shows the vertical magnetic field component on the lower z boundary over our domain. This is similar to a magnetogram - the black area shows negative vertical magnetic field, while the white area is positive. The negative patches sit under the separatrix domes, illustrating again the distinct regions present in the domain.

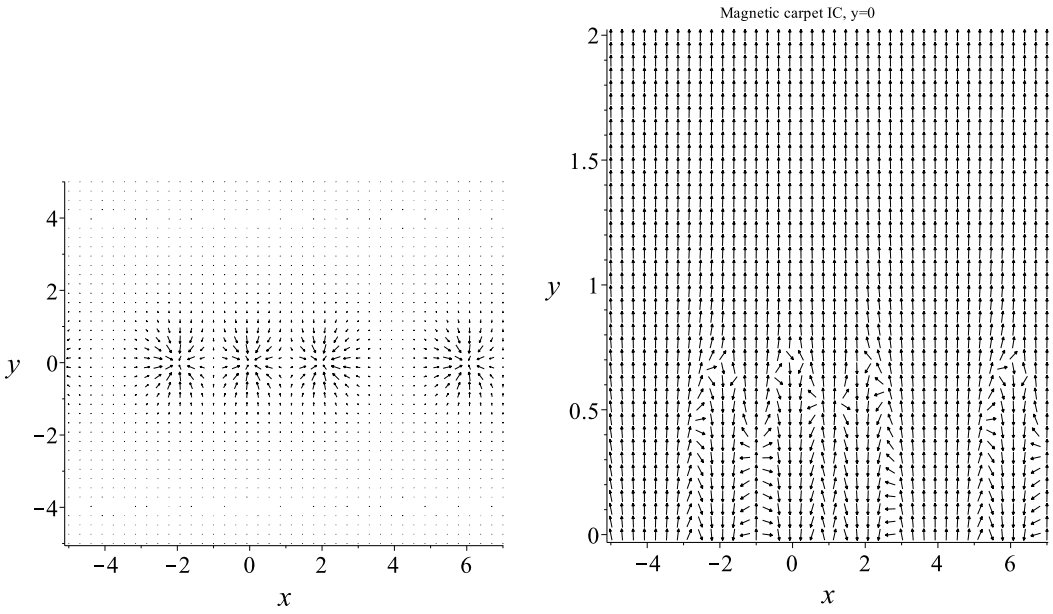

Figure 3.9: Vector plots for horizontal field on $z = 0$ and vertical field at $y = 0$. We observe the original three dome configurations and one extra from one of the copies in this size of domain.

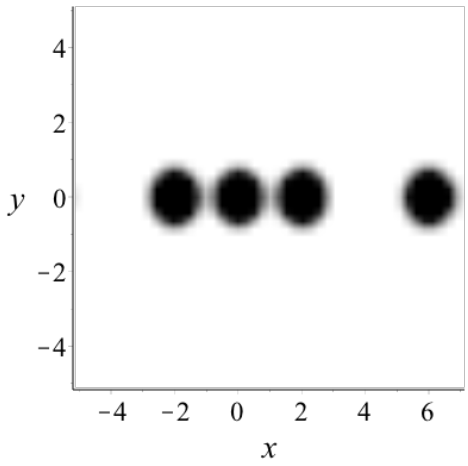

Figure 3.10: B_z on $z = 0$. The four black circles sit under the domes, where B_z is negative. The rest of this area is dominated by positive, uniform vertical field

Chapter 4

Uniform Field Simulations

4.1 Results

We now turn to the results of the numerical simulations set up with differing drivers acting on an initially uniform field. This section contains results also discussed in the paper [95].

4.1.1 Simple Uniform Field

4.1.1.1 Case 1

First we shall discuss the properties of the zero helicity, high complexity, oppositely spinning vortices runs 1A, 1B and 1C, whose entropies and properties we summarised in Chapter 3, Table 3.3.

Let us consider what we might expect to see and how to approach investigations. Run 1C had the highest entropy value out of the three runs, so we would expect this to

induce more complex tangling and provide more opportunities for current layers and reconnection. As the least complex, Run 1A will perhaps induce fewer opportunities for reconnection. We will look for signals of magnetic energy releases and how the kinetic energy and heating quantities are affected by this. Examining the structure of currents will provide an insight into reconnection events.

Comparison plots of quantities for the three runs 1A, 1B and 1C

Magnetic Energy

In Figure 4.1 we present the evolution of the magnetic energy over: a) the full simulation time and; b) over a shorter interval where the extent of the oscillatory behaviour is revealed. At the beginning of each simulation, the initial potential field has a total magnetic energy of 3000 simulation units. For each of these runs, we observe a broadly similar 'big picture': all three see an overall increase in magnetic energy from the initial potential state, until they settle into quasi-periodic states. This is particularly apparent for runs 1B and 1C ($x_0 = 1, 2$). After about $22,000$ time units (440 Alfvén times), the magnetic energy average over 100 driving periods remains roughly constant. For example for run 1B we split the last 300 full driving periods of the profile into 3, since we can see from the full plots about 3 repeated sections. The average magnetic energy values in these three sections are 3073.59, 3074.18, 3077.74 .

A common feature for all three profiles is their highly variable, 'bursty' character. Enlarging sections, as in Figure 4.1 (b), reveals the extent of the oscillatory behaviour. There appear to be many short-lived energy releases on a variety of scales. Zooming in on each of the runs at random time intervals showed that some of the smallest waves created by a build up and release of energy occur on a timescale similar to that of the driver; we will expand upon this shortly.

We also note the clear ordering of magnetic energy levels above potential with topological entropy: recall, $x_0 = 2$ gave the highest complexity, $x_0 = 0.5$ the lowest. The magnetic energy profiles match this ordering. More complex driving has given rise to higher magnetic energy, both early in the simulation and as the runs settle into their statistically steady states. The initial potential field has a magnetic energy of 3000 units. Taking the magnetic energy after $25,000$ time units (during the statistically steady state) we see that the average energy in excess of the potential level for each run, as presented in Table 4.1, increases with complexity. Clearly this illustrates that higher complexity photospheric motions have lead to higher injection of free magnetic energy. This is a new finding.

So far we have discussed quantities in terms of the non-dimensional units of the simulation, which allow us to compare and contrast simulation data directly. However it is obviously useful to think about what these numbers mean in real terms. Magnetic energy is given by $\int_V B^2/2\mu_0 dV$. At potential level, the magnetic energy in simulation units is 3000. If we take a length scale of $L_0 = 1$ Mm, then the domain in physical space becomes $[-5\text{ Mm}, 7\text{ Mm}] \times [-5\text{ Mm}, 5\text{ Mm}] \times [0\text{ Mm}, 50\text{ Mm}]$. Taking $B_0 = 5$ G and $\mu_0 = 4\pi \times 10^{-7}$ Hm^{-1}, we obtain a potential magnetic energy in physical units of 5.97×10^{20} J. Therefore, one simulation unit of magnetic energy is equal to 1.99×10^{17} J.

A typical drop in the magnetic energy profile of run 1C covers about 10 units: this means that the real energy release would be 1.99×10^{18} J. To put this is context, a nanoflare releases energy on the order of $10^{17} - 10^{19}$ J, so many of the events seen here could correspond to nano flares. The largest drop seen for run 1C in plot (b) of Figure 4.1 spans about 30 units - an event releasing 5.97×10^{18} J. So we see that the energy discharges are far from trivial in real terms. Even smaller releases could be interpreted

Averages of Quantities for Case 1

Run	x_0	**Excess**	**Viscous/Ohmic**	**Heating Rate**	**|j| Filling Factor**
1A	0.5	1.6 %	1.12	0.09	$0.026\% > 1$ $0\% > 4$
1B	1	2.3 %	0.66	0.12	$0.13\% > 1$ $0.00025\% > 4$
1C	2	2.9 %	0.46	0.18	$0.41\% > 1$ $0.006\% > 4$

Table 4.1: Case 1 average values: average magnetic energy in excess of potential after $t = 25000$; average value of viscous/Ohmic heating for full simulations; average values of heating rate (rate of change of Ohmic + viscous heating per timestep) for full simulations; average filling factor over thresholds of 1 and 4 for full simulations.

as the signatures of events on the scale of pico flares (with energies of less than 10^{17} J) - see some of the smaller drops in the zoomed in plot of Figure 4.1. We discussed the proposal that continuous occurrence of small scale events may be a contributing factor to the basal coronal heating rate in Chapter 1, so these regular small drops are very interesting. Say some of the smallest drops cover just one unit. This translates to between 1.99×10^{17} J or 1.99×10^{24} ergs. This gives a range in event sizes from 10^{24} to 10^{25} ergs. The authors in [4] estimated, based on observations, that the largest flares can release up to 10^{33} ergs, while research discussed in [42] calls events of energies at least 10^{27} ergs microflares and 10^{24} ergs nanoflares. Our range puts events for Case 1 runs in the nanoflare category, with some even smaller picoflares. This suggests that in a solar environment similar to ours, nanoflare heating could be significant.

Kinetic Energy

The kinetic energy displays particularly erratic behaviour, as seen in Figure 4.2. Again we have plotted the full simulation data for all three runs (top 3 plots) and then focused in on a limited interval (between 14800 and 15500 time units) of run 1B. As the magnetic energy, the kinetic energy seems to settle into a statistically steady state. Taking the last 300 full driving periods of run 1C and averaging as before gives values

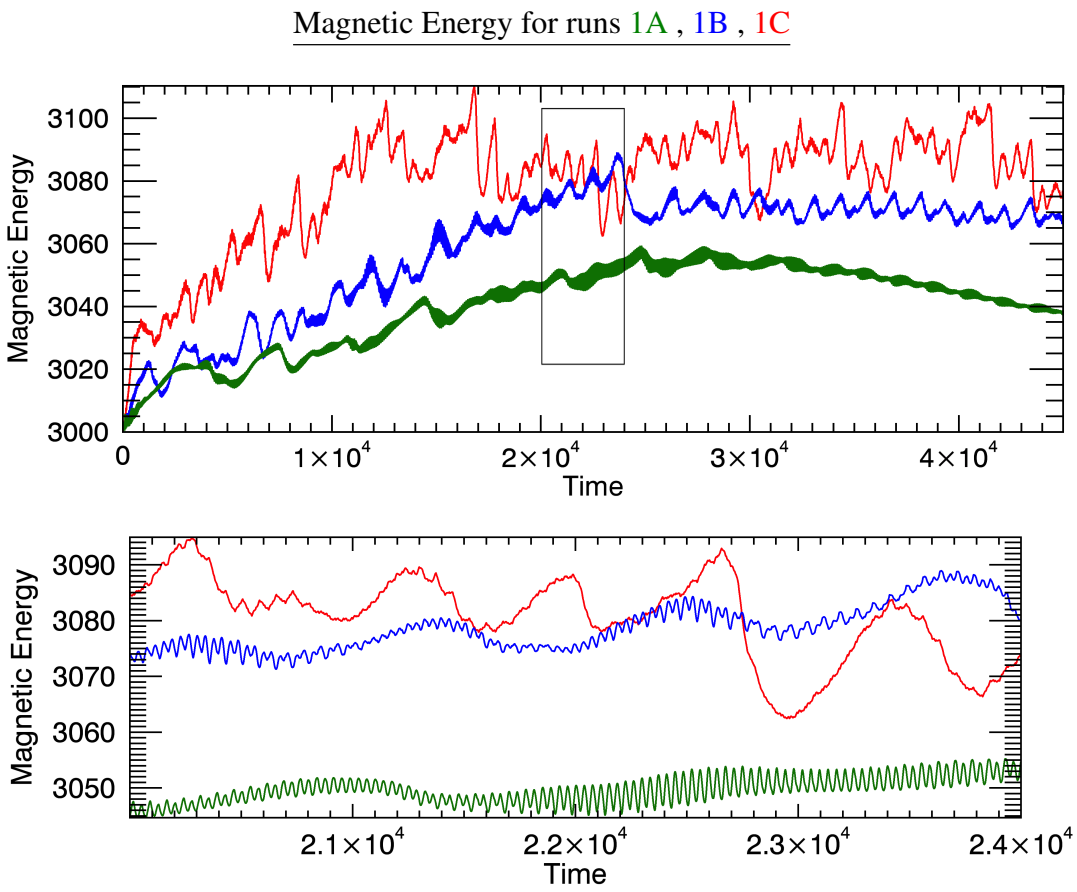

Figure 4.1: Magnetic energy evolution for runs 1A, ($x_0 = 0.5$), 1B, ($x_0 = 1$), and 1C, ($x_0 = 2$). The top plot, (a), shows the full simulation up to 45000 time units. The section within the black box is plotted in the bottom plot, (b). Here we have zoomed in on a section between 20000 and 24000 time units, part of the statistically steady state section, to show the variation in the sizes of oscillations in and between the runs. Plots are ordered with the entropy levels of the scenarios, with the largest magnetic energy input above potential due to the highest entropy run 1C and so on. Profiles exhibit bursty behaviour and runs 1B and 1C in particular appear to reach statistically steady states with easily recognisable patterns.

of 1.22, 1.07 and 0.90. Upon inspection of the enlarged plot in Figure 4.2, we see that the time between troughs at the level in the 1B case is fairly even. The half period of the driver, i.e. the time taken for one vortex to spin, is $12\pi \approx 38$. A rough estimation of the time between troughs in the plotted section is also around this value. This periodicity appears in the kinetic energy profile of this run into the statistically steady

state section (there are wavelengths over even shorter times earlier in the simulation). However, there is also clearly further periodicity on a larger time scale, as is visible in the last 15000 time units.

We can examine the periodicity beyond that resulting from the driver. Consider splitting the simulation into intervals, each the duration of one vortical twist of the driver. Starting from $t = 0$, the first interval or half period runs from $t = 0..12\pi$, the time for the first vortex to spin. The second half period runs from $t = 12\pi..24\pi$, the time taken for the second vortex to spin, and so on. In this way the half periods correspond to simulation time for the full simulation duration. We can average the kinetic energy during each half period, thereby smoothing those variations due just to the action of each twist. Plotting the average kinetic energy for each half period then, it is clearer to see those larger scale oscillations which have influence over many periods. We can see where plasma motion is perhaps a consequence of reconnection rather than driving. We illustrate this technique for run 1B in Figure 4.3.

We can do the same for the magnetic energy. Figure 4.5 compares the two averaged quantities during the steady state section while Figure 4.4 examines these variables over a peak at around 15000 time units (about 400 half periods). The larger scale variations which remain after averaging occur over many driving periods. We can also see by comparing the two figures that these variations become more uniform during the steady state section. This indicates that the system has settled into a steady pattern of energy build up and release: this could be promising as a potential mechanism for consistent background heating.

There is a clear difference to be seen in the magnitudes of magnetic and kinetic energy. An average of the two quantities for run 1B shows that the ratio of excess magnetic

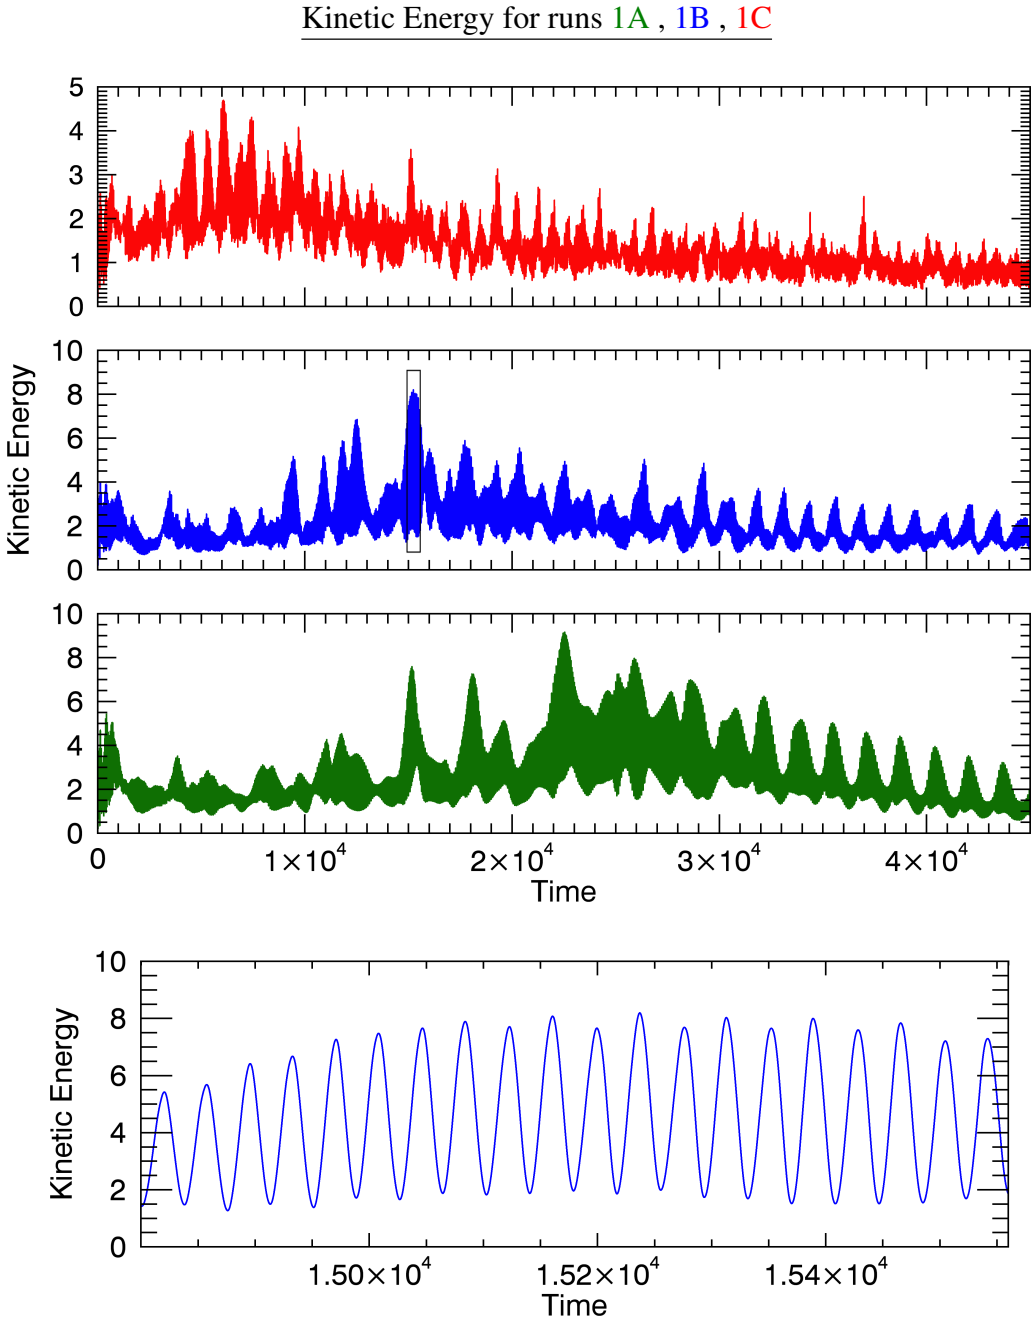

Figure 4.2: Kinetic energy evolution for runs 1A, 1B and 1C again with colours corresponding as before. The top three plots show the full simulations for all three runs, with a black box again around the section of the middle plot which we enlarge in the bottom plot. These profiles are so highly oscillatory that separate plots are required to glean information and still we must zoom in, as in the bottom plot, for a clearer picture.

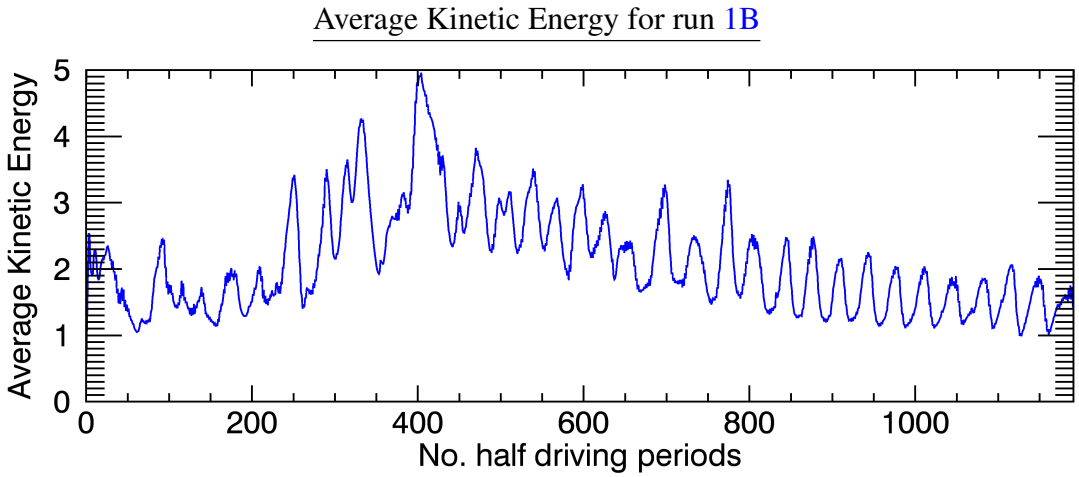

Figure 4.3: Kinetic energy is averaged in each half driving period in run 1B. This leaves more clearly visible the larger scales variations.

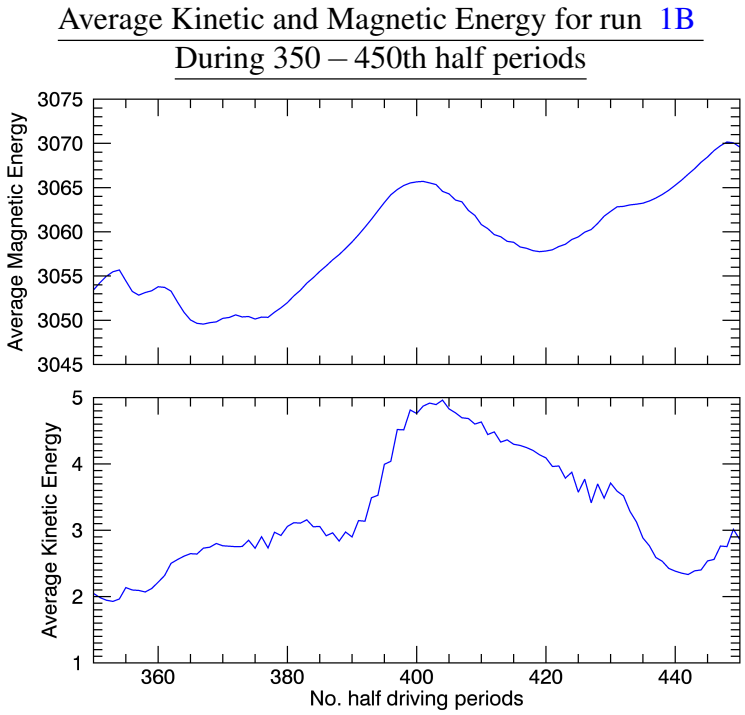

Figure 4.4: Comparison of averaged kinetic and magnetic energies in run 1B during an obvious peak in kinetic energy at around 15000 time units.

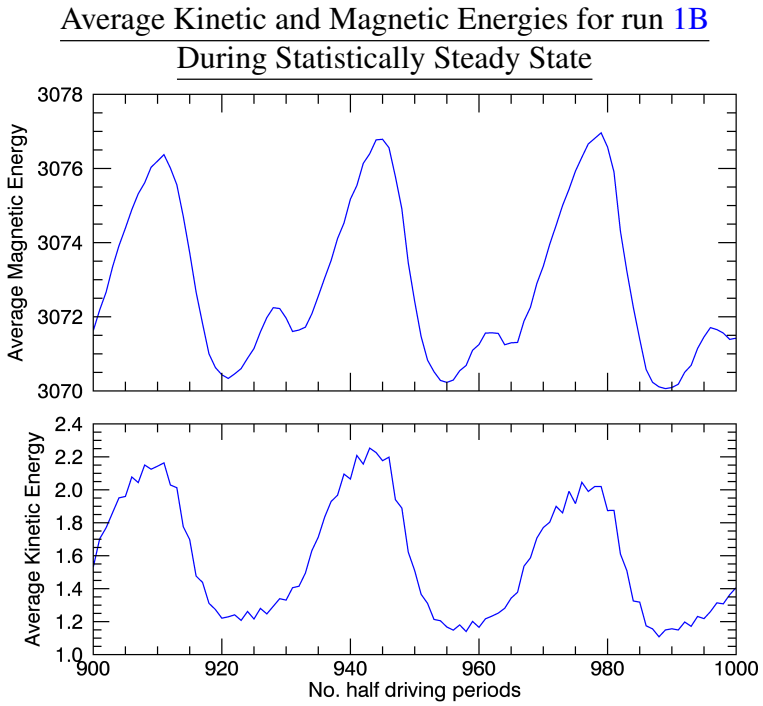

Figure 4.5: Comparison of averaged kinetic and magnetic energies in run 1B during the statistically steady state section.

energy to kinetic energy is 30 : 1. These results are in agreement with previous simulations of continuously driven systems, for example [92], [93]. The highest entropy run 1C, which had the highest energy injection, has the lowest kinetic energy. This run settles into oscillating around roughly half the value of the other two lower complexity runs: taking averages after 25000 gives values of 2.7, 1.7 and 1.0 for runs 1A, B and C respectively, so here there is inverse ordering with complexity.

As seen from Figure 4.1, run 1C had the highest levels and largest drops in magnetic energy. Looking ahead to the heating for this run, (Figure 4.8) it also has the the highest Ohmic heating. These two facts suggests that there is more reconnection occurring, which would trigger plasma flows and boost velocities and therefore kinetic energy, so from this alone we would expect the kinetic energy to be higher than the other two runs. However 1C also has the highest levels of viscous heating, as also seen in Figure 4.8. Viscous heating removes kinetic energy from a system, therefore this

could explain the lower levels for this run. More reconnection in turn may provide more opportunities for shocks, leading to viscous effects, and the removal of kinetic energy. Furthermore, from Figure 4.6, where we plot two snapshots of the magnitude of velocity along $y = 0$ for runs 1C and 1A, we see that run 1C appears to induce velocity fields which vary more over small scales. It may be that the more complex driving function, as well as being able to create smaller scales in the magnetic field and form more reconnection sites, can bring about smaller scales in the velocity field too - particularly in the region closest to the driving action. These areas are noticeably more variable, with strand-like features of varying strength close together. The same site in the other run's snapshots shows a much more homogenous velocity strength. Finer velocity fields can lead to more conversion of kinetic energy to heating by viscosity, and we have already identified this run as having the highest viscous heating.

Internal Energy

The internal energy evolution here is steadily increasing for all three runs, as seen in Figure 4.7. This is expected, due to the fact that we have no radiation, so there is no dissipation of internal energy. This consistent increase in internal energy will have an effect on the plasma beta. At each end of a loop the plasma beta is large, which is why we have line tying. As mentioned in Chapter 1, Section 1.1.1, the plasma beta is generally taken to be small for a certain height window above the photosphere. The unavoidable consequence of neglecting radiation gives us this increasing internal energy which means the plasma pressure will also increase, until some point where the plasma pressure dominates the magnetic pressure, and we no longer have a small plasma beta regime. Including radiation may be something to consider in different simulations and provide interesting alternative results. Again here there is an ordering as per the complexity. The highest complexity run, which had the highest magnetic energy and the largest scale energy releases, builds up the largest internal energy.

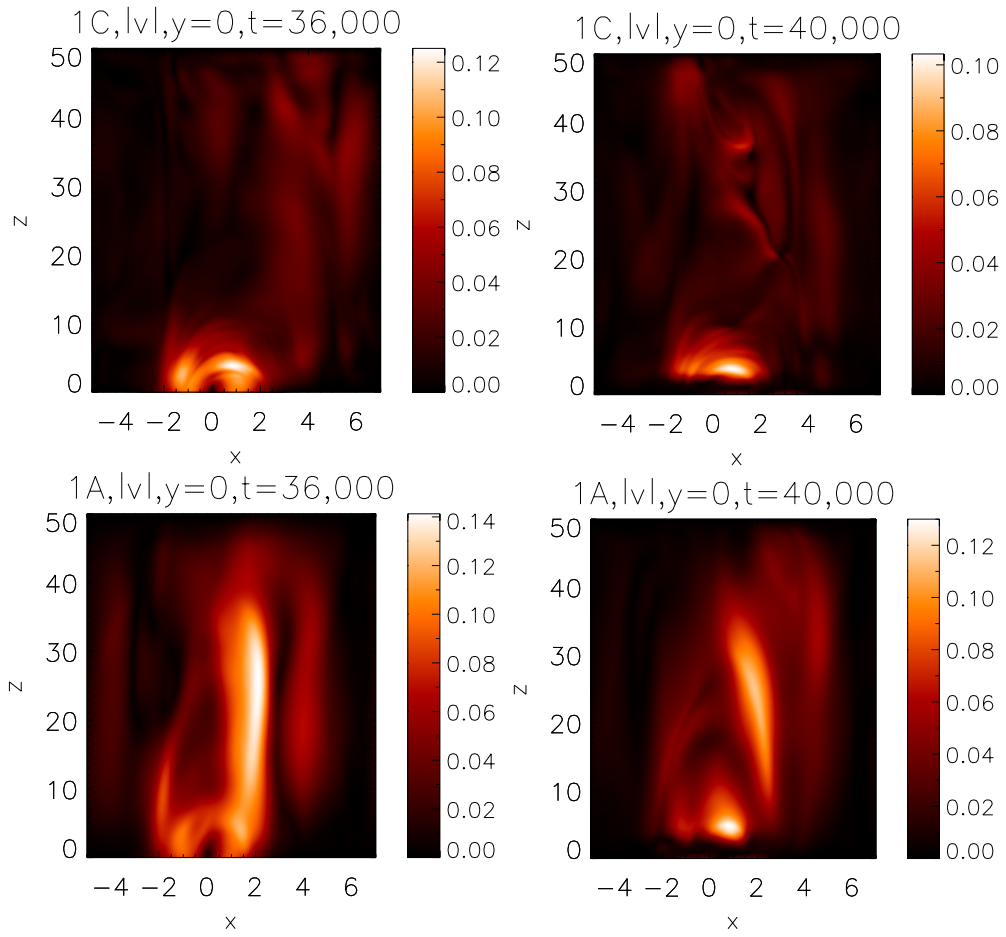

Figure 4.6: Magnitude of velocity along $y = 0$ at times $t = 36,000$ and $t = 40,000$ for runs 1C (top) and 1A. This quick check does appear to show slightly more variable velocity in the 1C run. We can see finer detail for 1C. Being the most complex driving motion, we have seen it tangle up the magnetic field and lead to smaller scales and more reconnection, and it also seems that this case is able to create variation over smaller scales in the velocity field too. This helps explain the lower kinetic energy in this experiment, as the smaller scales in velocity give rise to higher levels of viscous heating.

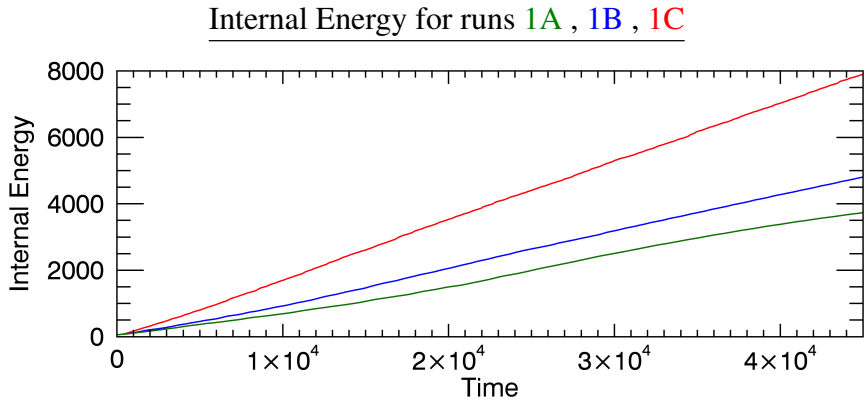

Figure 4.7: Internal energy evolution for runs 1A, 1B and 1C, with colours corresponding as before. The simulation under the influence of the highest entropy driver builds up the largest internal energy, by some way.

Heating

Figure 4.8 displays graphs of the heating in the system: Ohmic, viscous and total (ohmic+viscous). In these plots the heating at each time is cumulative. The heating at each stage is calculated by integrating over the domain and the timestep, and adding this to the previous heating value. Ratios of the average viscous to average Ohmic heating for each run are quoted in Table 4.1 and show that viscous heating dominates only in run 1A, and only by a small margin; the values are 2220 average viscous heating: 1980 average ohmic heating. It appears that the higher the entropy of the run, the more dominant the Ohmic heating over viscous heating. Remember that for us, the viscous heating is an artificial shock viscosity. It comes into play where quantities build to a certain level in a localised region. So while our measure of viscous heating is not the most physical concept, the presence of viscous heating implies occasions where quantities must have built sufficiently to trigger the mechanism in the first place. Higher viscous heating must mean that sufficiently large gradients must form more frequently in that simulation. The viscous heating in each run, while still ordered with complexity, varies less between each run than the Ohmic heating. Ohmic heating is triggered by the flow of electric currents. It is sufficiently large electric currents in

thin layers which can trigger reconnection. Therefore, the increasing value of Ohmic heating and its dominance over viscous heating as the domain is subject to increasingly complex tangling mechanisms suggests that highly braided fields are more prone to reconnection.

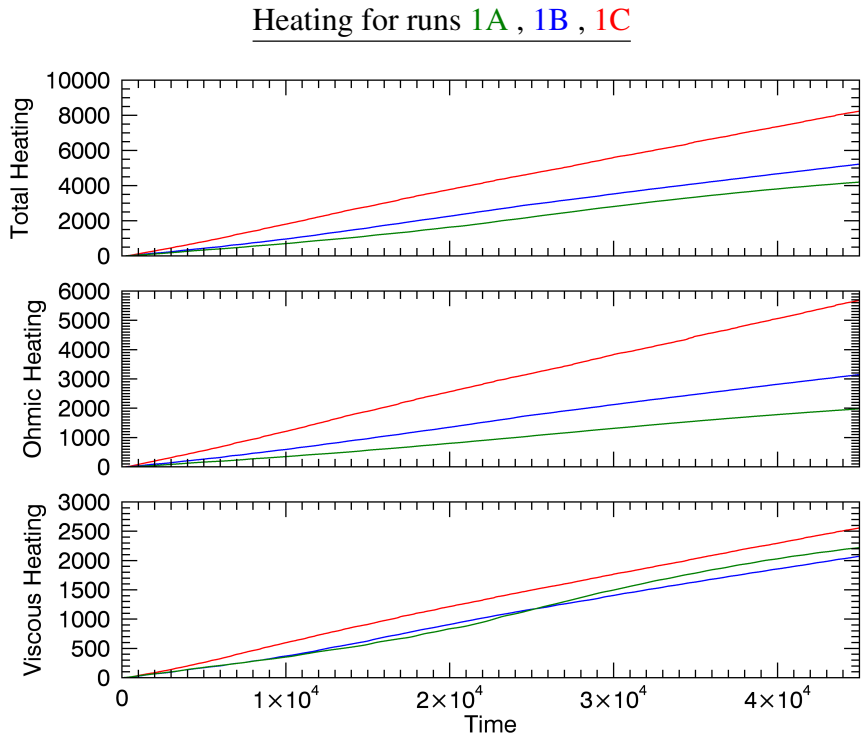

Figure 4.8: Cumulative heating evolution for runs 1A, 1B and 1C. Total heating is the sum of the viscous and Ohmic heating. Clearly the more complex braiding in Run 1C has produced higher levels of heating.

We can examine where the Ohmic heating is occurring by looking at the Ohmic 'heating density' at certain times. Unlike the cumulative values considered in Figure 4.8, here we take the quantity ηj^2 over the whole domain at one particular time to assess how heating is distributed at that time. We compare this to the current density and find that often the Ohmic heating takes place where currents have built up. Figure 4.9 shows the Ohmic heating (left) and the current density (right) for $t = 30,000$ at $y = 0$ for run 1A. The heating is greatest where the current density, and therefore magnetic field gradients, are largest. Figure 4.10 displays these quantities at $z = 0$, $z = 20$ and

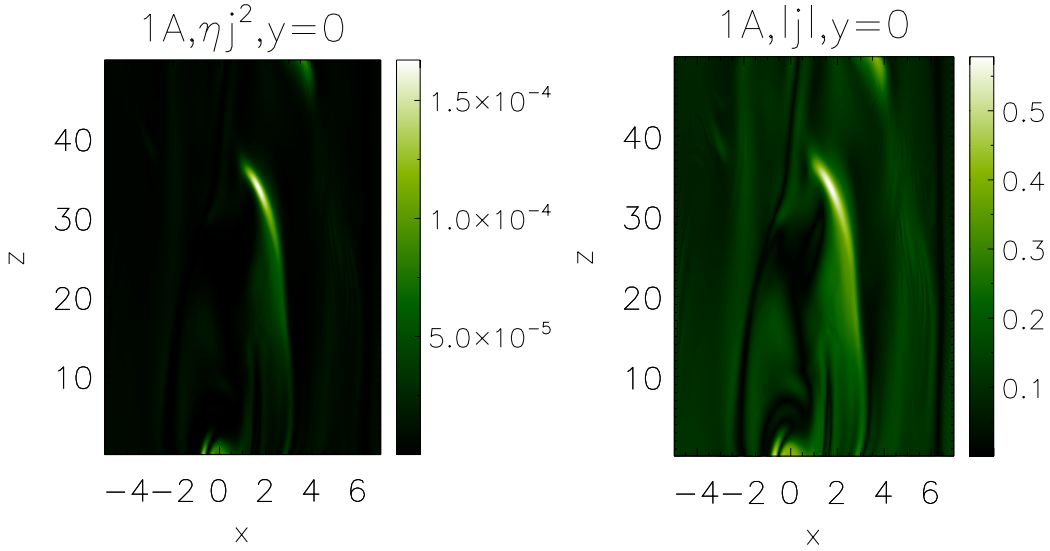

Figure 4.9: Contour slices at $y = 0$ for run 1A at $t = 30,000$. The left hand plot shows the Ohmic heating distribution while the right hand plot shows the current density.

$z = 40$ at the same time. Again the heating echoes the pattern of the current, concentrated in regions of largest current. For example at $z = 0$ we can see a thin curve of high current around $(0,0)$ and a corresponding curve in the heating plot.

Figures 4.11 and 4.12 display contours of Ohmic heating (left again) and current density (right again) for run 1B, at the same time $t = 30,000$ and at $y = 0$ and $z = 0,20,40$ respectively. There are clearly long, thin structures to be seen in the distribution of heating and again the pattern of heating corresponds to that of the current density. More diffuse regions of current translate into regions of lower heating.

Lastly, we do the same for run 1C. Figure 4.13 and Figure 4.14 show that at time $30,000$ this highest complexity run results in more highly localised Ohmic heating, in regions corresponding to thin current layers observed in the current density plots. The maximum values are also larger. In all three runs then, it appears that the build up of currents in long, thin layers triggers Ohmic heating in these layers, with the level of heating becoming more intense with increasing complexity.

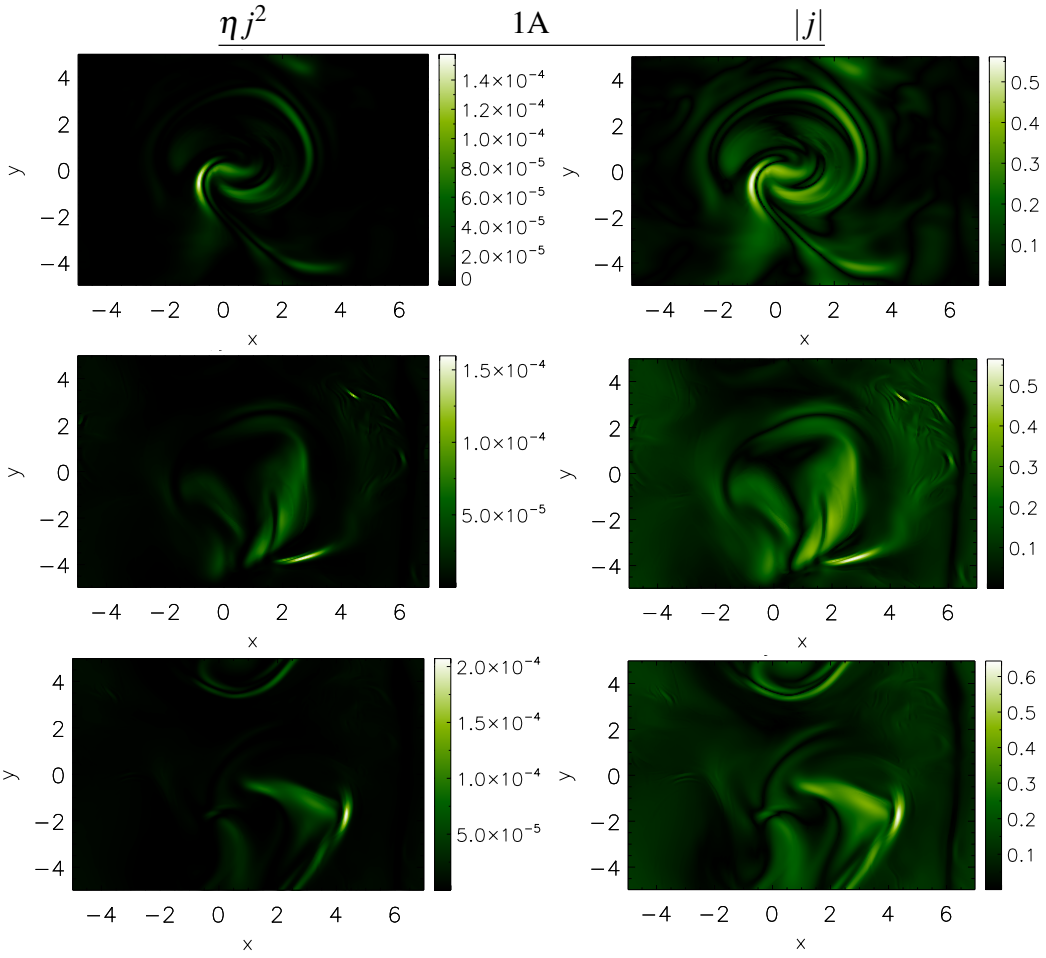

Figure 4.10: Slices at $z = 0, 20, 40$ for run 1A at $t = 30,000$. Left hand plots display the Ohmic heating at the three heights and the right hand plots show the current densities in those planes.

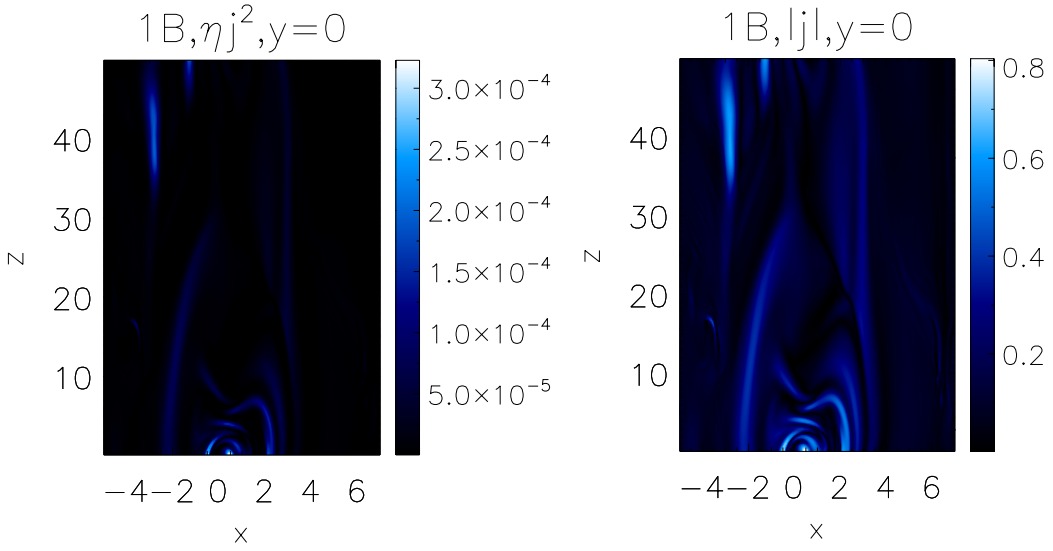

Figure 4.11: For run 1B at $t = 30,000$ we see contours of Ohmic heating (left) and current density (right) through the slice $y = 0$.

This assessment also holds when looking at another time, $t = 10,000$, in Figure 4.15. The top three plots show the distribution of Ohmic heating for runs 1A, 1B and 1C (from left to right) at $y = 0$, while the bottom three plots displays the current density over the same slices. Again we see the heating distribution echoing the current density, with higher complexity drivers inducing more current layers with greater heating levels.

Heating Rate

We can also quantify the heat being supplied to the system by calculating the heating rate; by this we mean the rate of change in total heat (Ohmic+viscous) in the system per timestep. (Note the difference between time step and time unit: there are 45000 full time units in the simulation, but the timestep *dt* in which the code updates all variables varies and multiple time steps can make up one time unit. We will be elaborating on timestep factors in the next chapter, but it is interesting to note that the highest energy run, 1C, took the most timesteps to reach $t = 45,000$. The number of timesteps also is

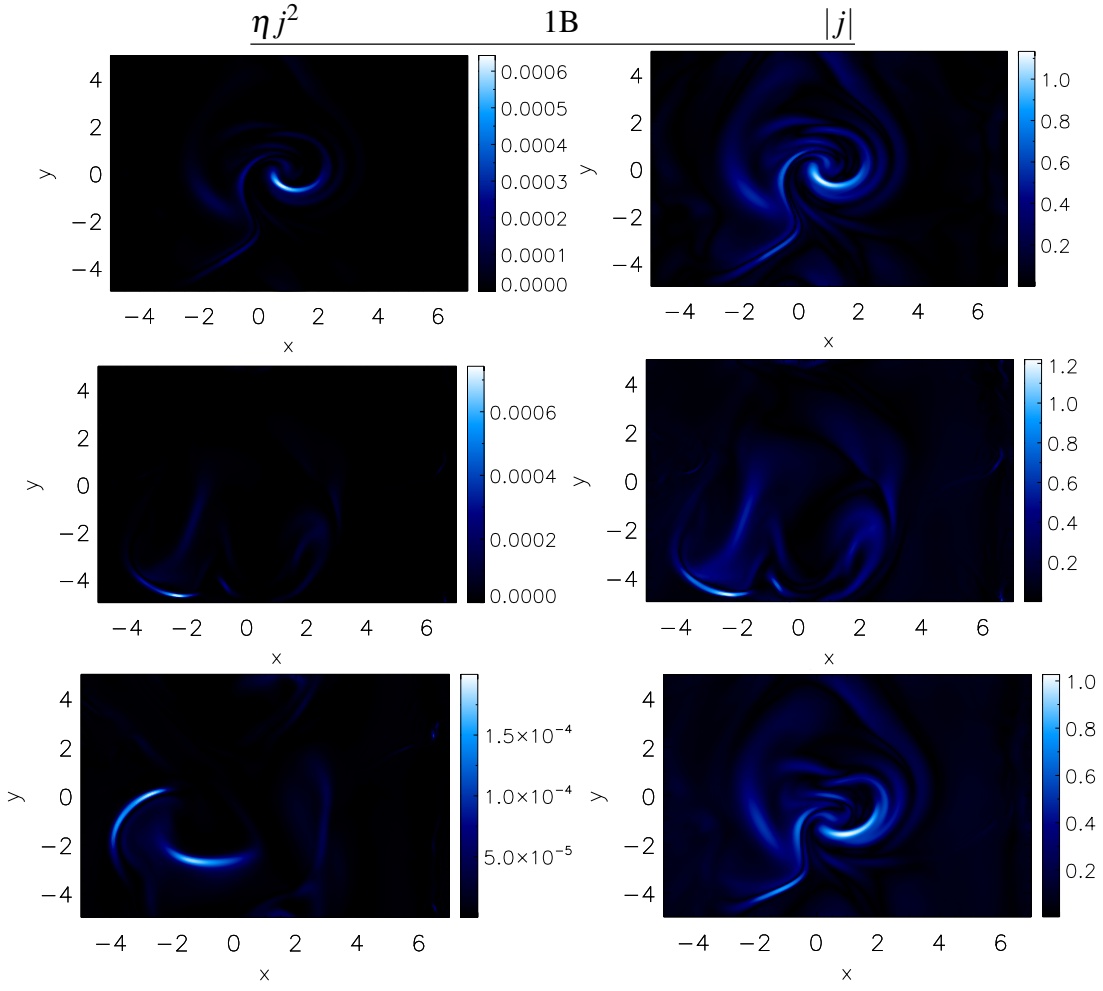

Figure 4.12: The left hand plots show the Ohmic heating at $z = 0, 20, 40$ for run 1B at $t = 30,000$. The right hand plots show the same slices through the current density.

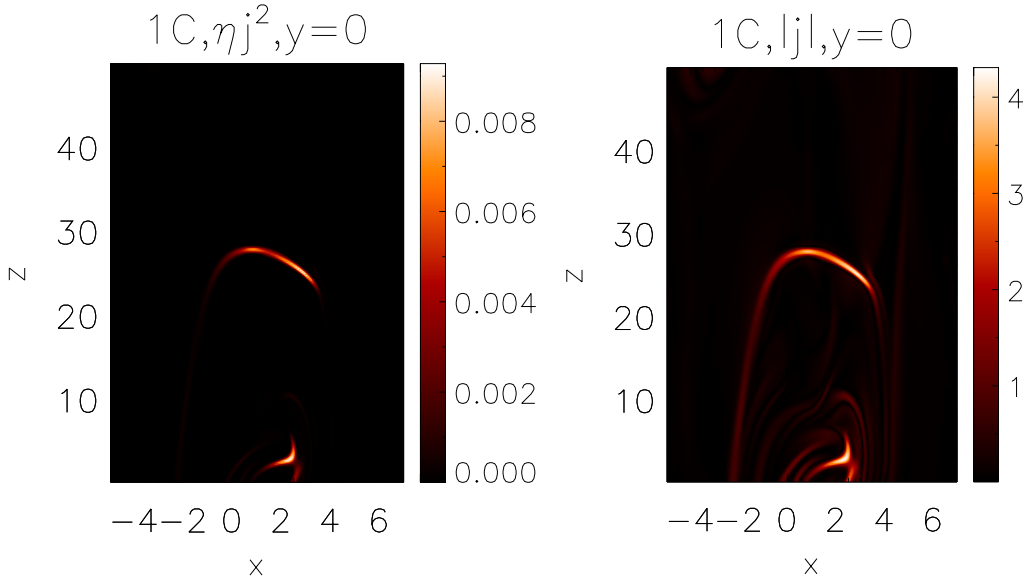

Figure 4.13: For run 1C at $t = 30,000$ the plot on the left displays the Ohmic heating and the plot on the right the current density, at $y = 0$.

an indicator of the most dynamic, variable systems. All energy diagnostics are given by the code on every timestep).

Figure 4.16 shows how the heating rate changes over the course of the simulation. While appearing highly variable, the magnitude of the variations in these profiles are extremely small relative to the values of total heating themselves. Table 4.1 shows the average values of the heating rates, and illustrates that this quantity is also ordered with topological entropy. The Case 1 setup appears to give rise to a nicely consistent pattern. It is satisfying that again we have the ordering of heating levels with ordering of complexity. This would support the idea that the degree of complexity of a photospheric motion could be key variable for heating.

Current Properties

Identification of finite current layers would be further support of reconnection. The knowledge we have gained from energy diagnostics leads us to suspect that most complex setup will provide more frequent and thinner current layers. Figure 4.17 shows

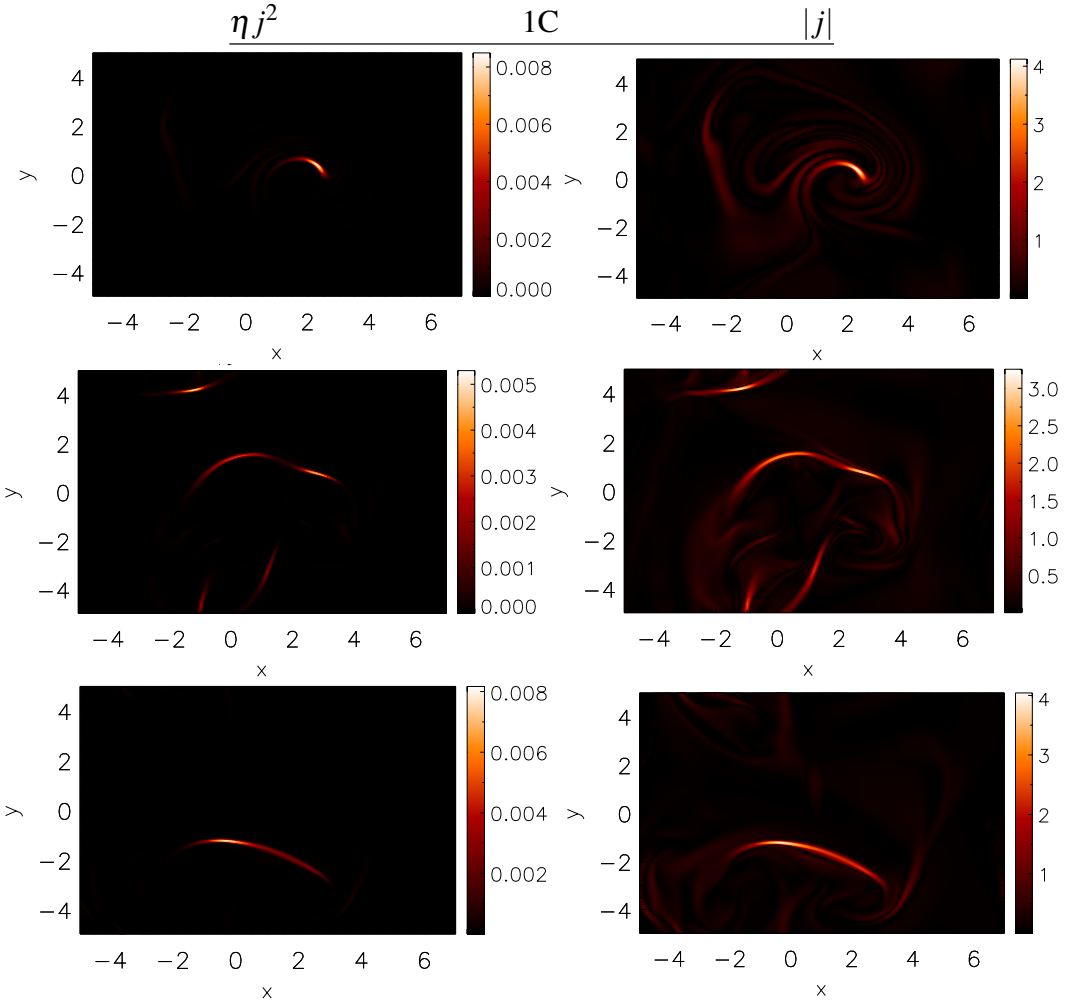

Figure 4.14: The Ohmic heating distribution (left) and the current density (right) for 1C, $t = 30,000$, at $z = 0, 20, 40$.

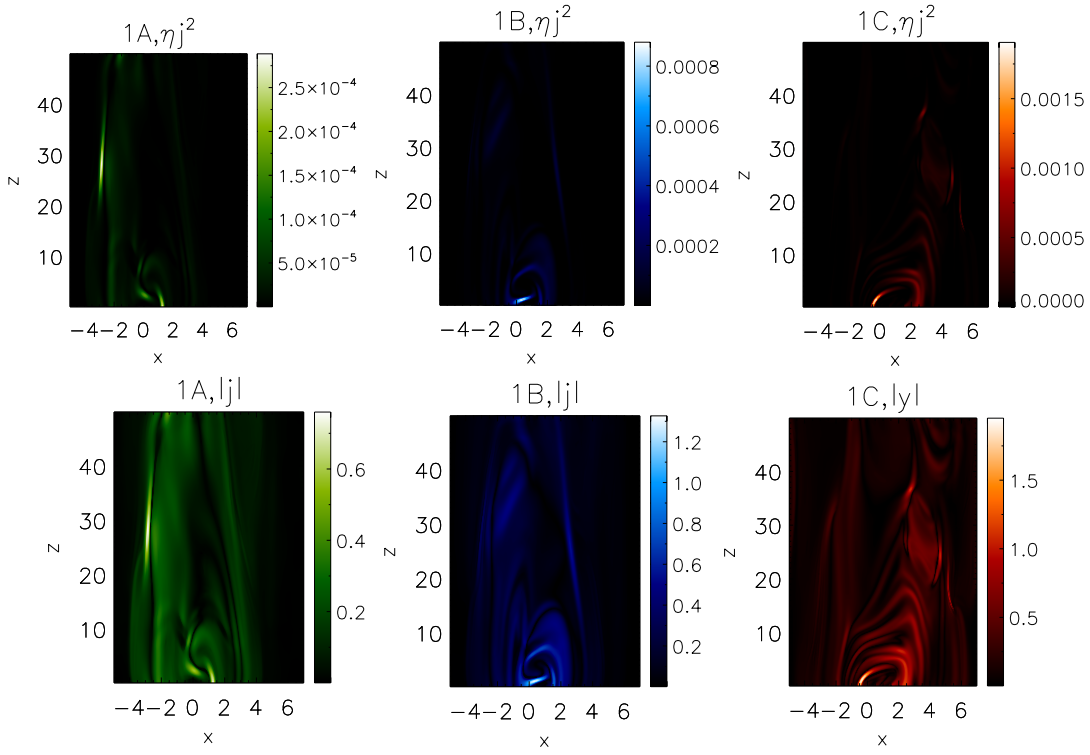

Figure 4.15: For time $t = 10,000$ we take another look at the Ohmic heating and the correspondence with the current structure. For $y = 0$, the top three plots show the heating distribution for runs 1A, 1B and 1C, while the bottom three plots show the current density.

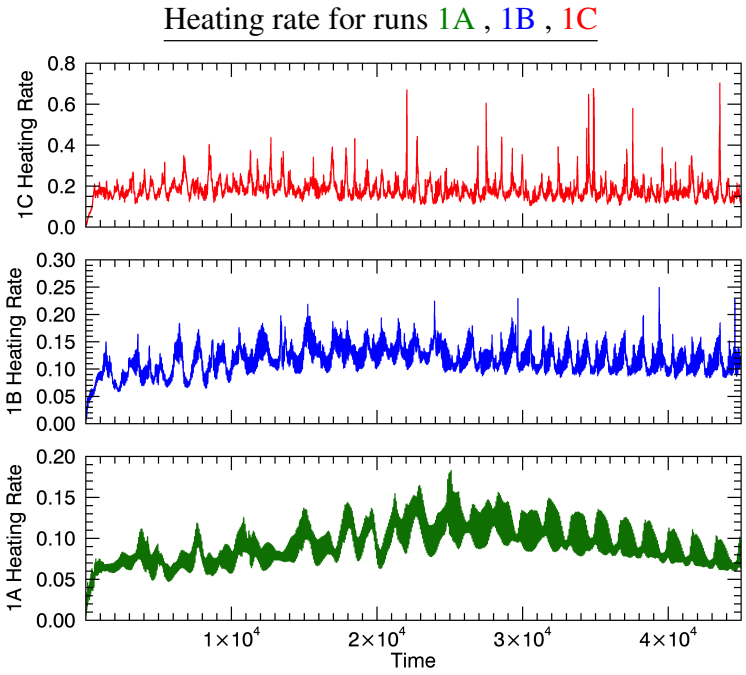

Figure 4.16: Heating rate for runs 1A, 1B and 1C. This is the difference in total heating per timestep.

the maximum current in the domain at points in time across the simulations. The simulation with the highest entropy driver generally provides the largest current densities. Further, we observe larger drops in the maximum current value for the highest entropy simulation than in the others; this suggests that a higher degree of braiding has resulted in larger reconnection events. Not only does this run have the largest values, but we see the most significant dissipation of those currents in the form of large decreases in the maximum current. Furthermore, Figures 4.18 and 4.19 illustrate the percentage of the domain filled with current densities larger than 1 and 4 units, respectively. We observe that run 1C also has larger currents filling more of the domain than the other two lower complexity simulations. Run 1A does not even have any instances of currents larger than 4. These analyses suggest that the higher complexity run, with its ability to induce more intricate braiding, leads to more regions of higher currents.

Current Structures

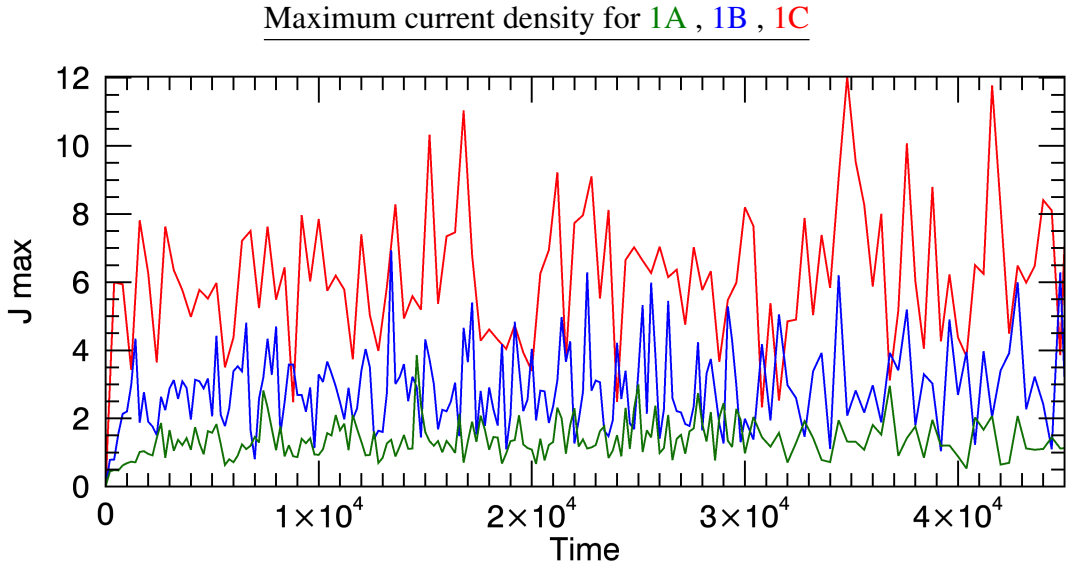

Figure 4.17: Comparison of the maximum current density j in the domains, for runs 1A,B and C. The highest complexity/largest magnetic energy run has the largest maximum currents.

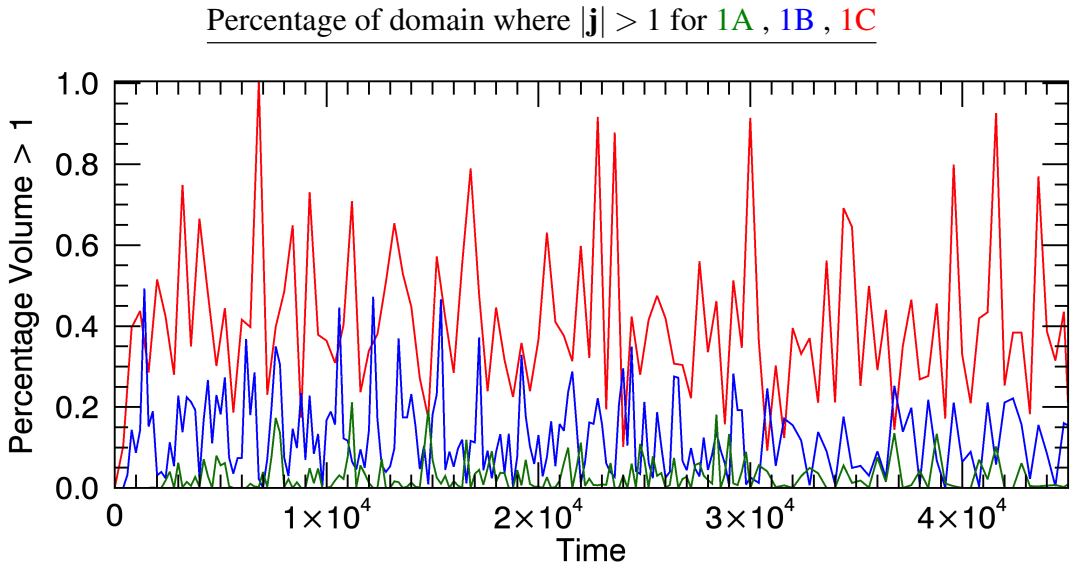

Figure 4.18: Here we see the percentage of the domain at points in time which is filled with current densities larger than 1 unit, again for runs 1A,B and C. As well as having the largest current maxima, run 1C also fairly consistently has the largest volume of currents above this level.

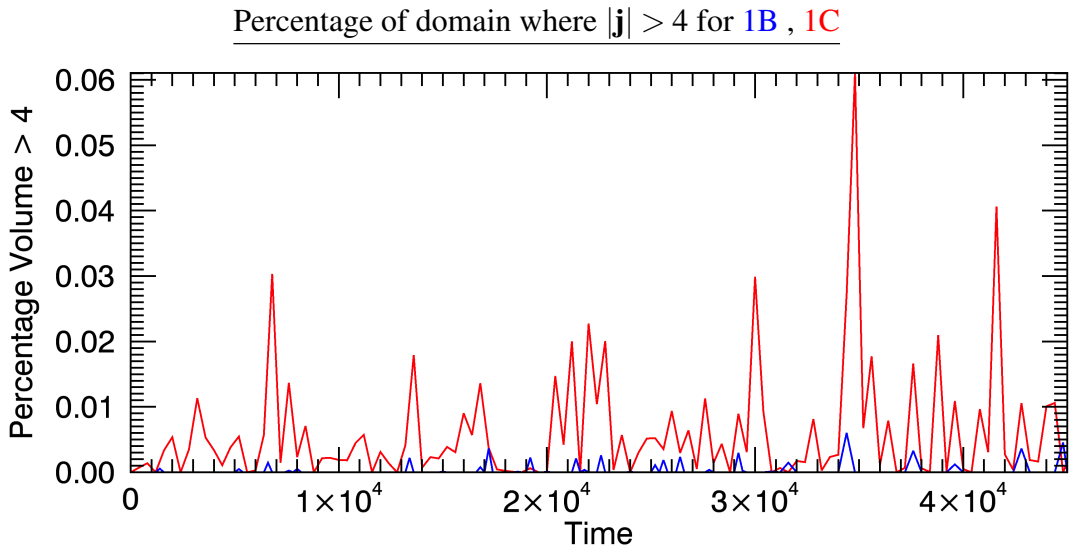

Figure 4.19: Only runs 1B and C have any parts of the domain with currents above 4 units.

We now turn to the structure of currents within the domain - do we see thin current layers forming? We will focus on run 1C since this looks to be the best candidate for reconnection. Figure 4.20 displays isosurface plots of current for run 1C. The plots show current structures at different levels, both before and after one of the largest visible magnetic energy release, between 22800 and 23200 time units. Over this interval we see a release of around 13% of free magnetic energy. The top two plots show the current density at a value of 2 units, the middle plots are at 3 units, and the last plot is at 6 units (the domain maximum current at the earlier time is 8.51). Before the drop, we see some thin current layers scattered through the domain at current density levels of 2, 3 and 6. Afterwards we see fewer layers, covering a smaller surface area, in the case of each threshold - in fact there are no current values of 6 at all at this later time. The volume filling factor has been reduced, indicating that larger currents have been dissipated.

It is also instructive to look at current structures at lower levels, where most of the

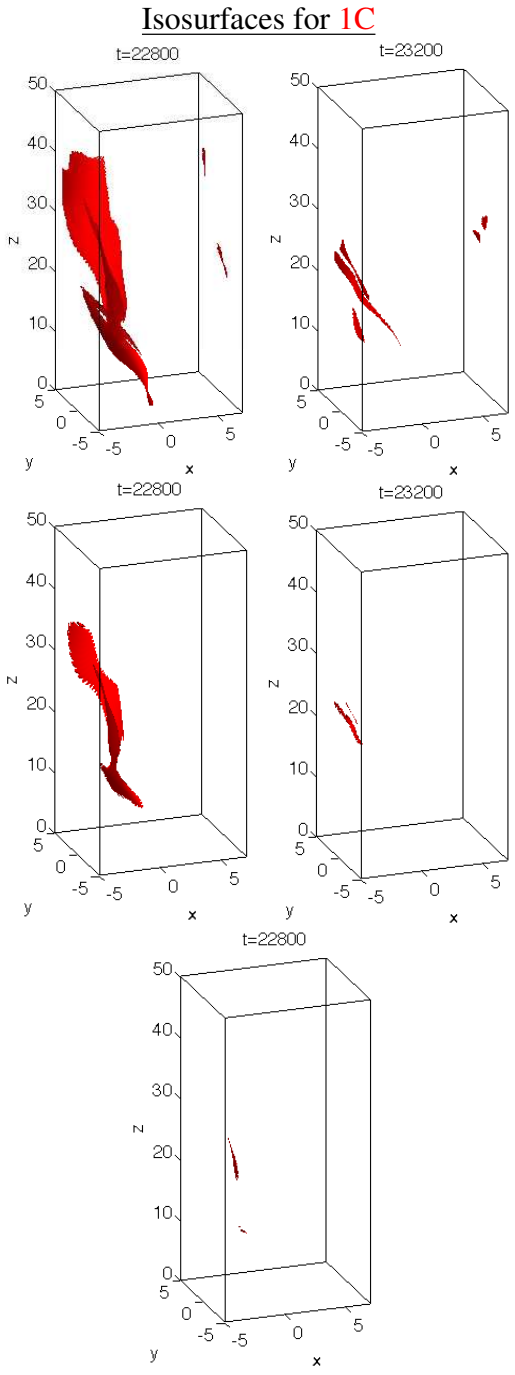

Figure 4.20: Isosurface plots for run 1C at times $t = 22,800$ and $t = 23,200$. The top two plots are at a $|\mathbf{j}|$value of 2 units, the middle two at 3 units, and the last at 6 units.

current activity lies. Figure 4.21 shows current density at 10% of the domain maxima for run 1B and 1C respectively at time $t = 35200$ (during the steady state section). The lower complexity and magnetic energy run, 1B, has a more volume filling current structure at this point. The second plot, for run 1C, the most complicated simulation, appears to be less volume filling at this level, but we have already seen how this run has a higher fraction of current density at higher levels. Both show thin but three dimensional layers, agreeing with the results presented in [36].

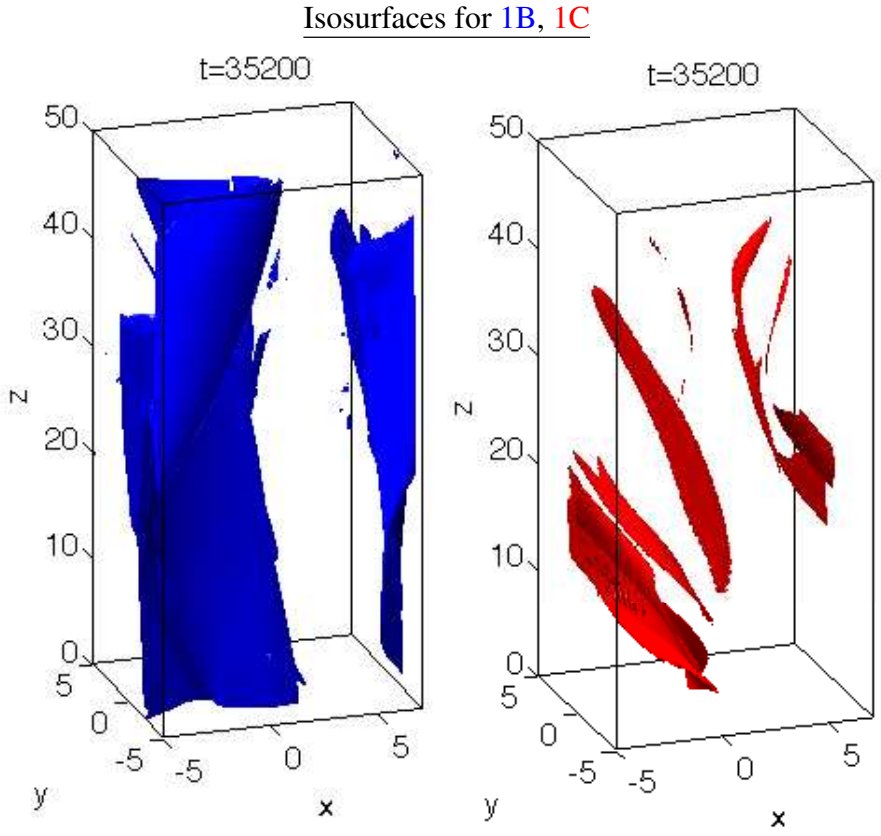

Figure 4.21: Isosurfaces for runs 1B and C at $t = 35200$ at 10% domain maximum.

We note also the difference to reduced MHD simulations. In RMHD, $|B_x|$ and $|B_y|$ are taken to be much less than $|B_z|$. Looking over snapshots from all three runs (ignoring $t = 0$ since we know the field at that time), the minima and maxima of the absolute value of each of the components cover the ranges:

$$
\begin{aligned}
1A: \quad & |B_x| : 2.0 \times 10^{-12} \rightarrow 0.17 \\
& |B_y| : 1.2 \times 10^{-11} \rightarrow 0.21 \\
& |B_z| : 0.82 \rightarrow 1.0 \\
1B: \quad & |B_x| : 4.2 \times 10^{-12} \rightarrow 0.17 \\
& |B_y| : 2.4 \times 10^{-11} \rightarrow 0.20 \\
& |B_z| : 0.77 \rightarrow 1.0 \\
1C: \quad & |B_x| : 1.7 \times 10^{-12} \rightarrow 0.28 \\
& |B_y| : 8.3 \times 10^{-11} \rightarrow 0.36 \\
& |B_z| : 0.63 \rightarrow 1.0
\end{aligned}
$$

These ranges show B_z to be larger, but we cannot say we are in a situation where $|B_z| >> |B_x|, |B_y|$. Other models using reduced MHD and similar field set-ups have seen long ribbons of current flowing from the bottom to the top of the domain. Seeing as many of our current isosurfaces, particularly for the most complex run, show more localised current layers, we can say it appears that a difference is made by using a fully 3D magnetohydrodynamic model.

The plots of Figure 4.22 show contours of current density at these same times, in the z mid plane (top) and y mid plane (bottom). The maximum current in these slices decreases across the drop, with the structure in the region above the driver in the $y = 0$ plots appearing to simplify slightly.

Fieldlines

Examining some sample fieldlines traced from seed points on $z = 0$ up through the domain further illustrates the degree of braiding. We provide a 'bird's eye' view, looking

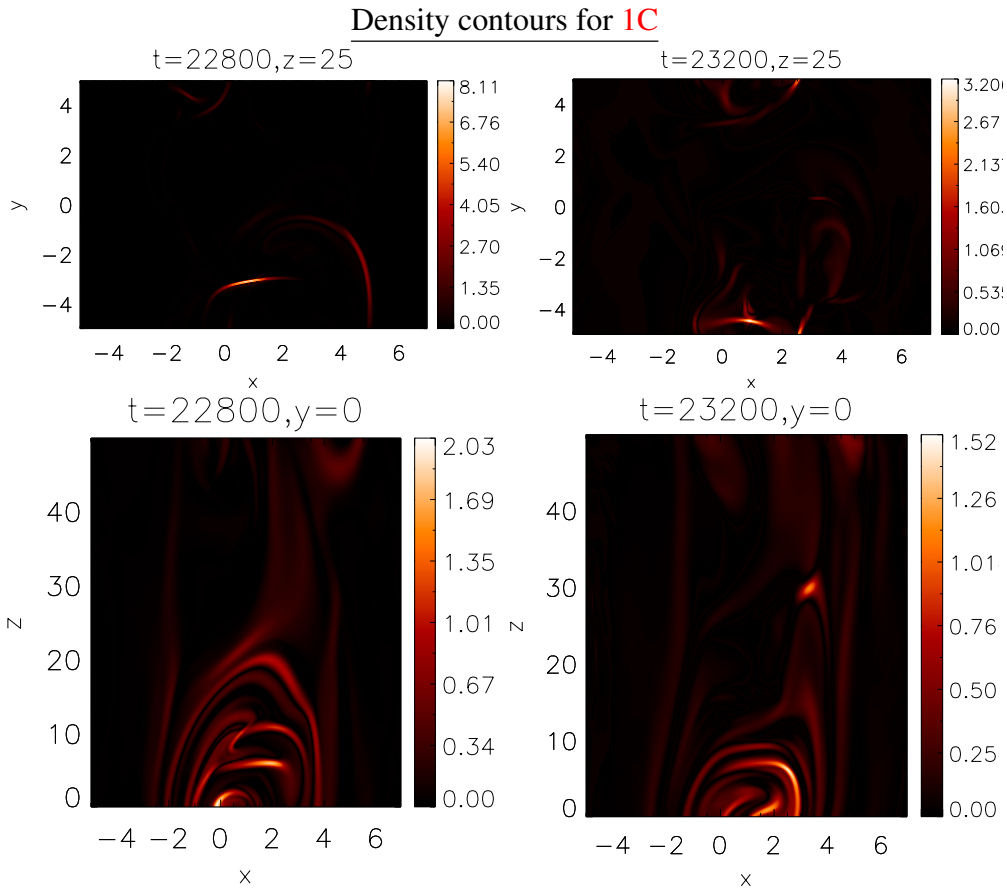

Figure 4.22: Contours for run 1C at the same times as the isosurface plots, $t = 22,800$ and $t = 23200$.

down from the top of the box. As we move across the base of the box from minimum to maximum x value, the colour of the lines darkens. This makes it clearer to see the braiding; the more lines of differing shades interacting, the more tangled the field. In addition, we have taken 2D slices at constant z and over-plotted the magnetic field vectors too fieldlines to show the structure throughout the domain.

Figure 4.23 displays these concepts at $t = 10,000$ for run 1A. In the top plot, we have traced fieldlines up the box. At the extremities of the x-axis, the fieldlines maintain similar x and y coordinates; in between they swirl around as they travel up the domain. We recognise the action of the vortices creating twist in the field, but it does not appear to be particularly tangled at this time. The lower plots show the structure in detail at $z = 20$ and $z = 40$. The field is dominated by two large, swirling structures, which create strong field where they interact.

Figure 4.24 shows the same information at the same time for Run 1B. Again we see large vortical structures. The fieldline plot shows a slightly more complicated structure: there is a little more interaction in the middle of the domain between fieldlines of varying shades. The slices illustrate two large vortical forms with opposing circulation, with the field strengthened where they interact.

Figure 4.25 corresponds to run 1C at $t = 10,000$. This is the most complex run of the three, with the highest levels of magnetic energy and heating. At this snapshot the fieldlines appear to be dominated by one large vortical formation and an area of sheared field; more variable than the previous runs.

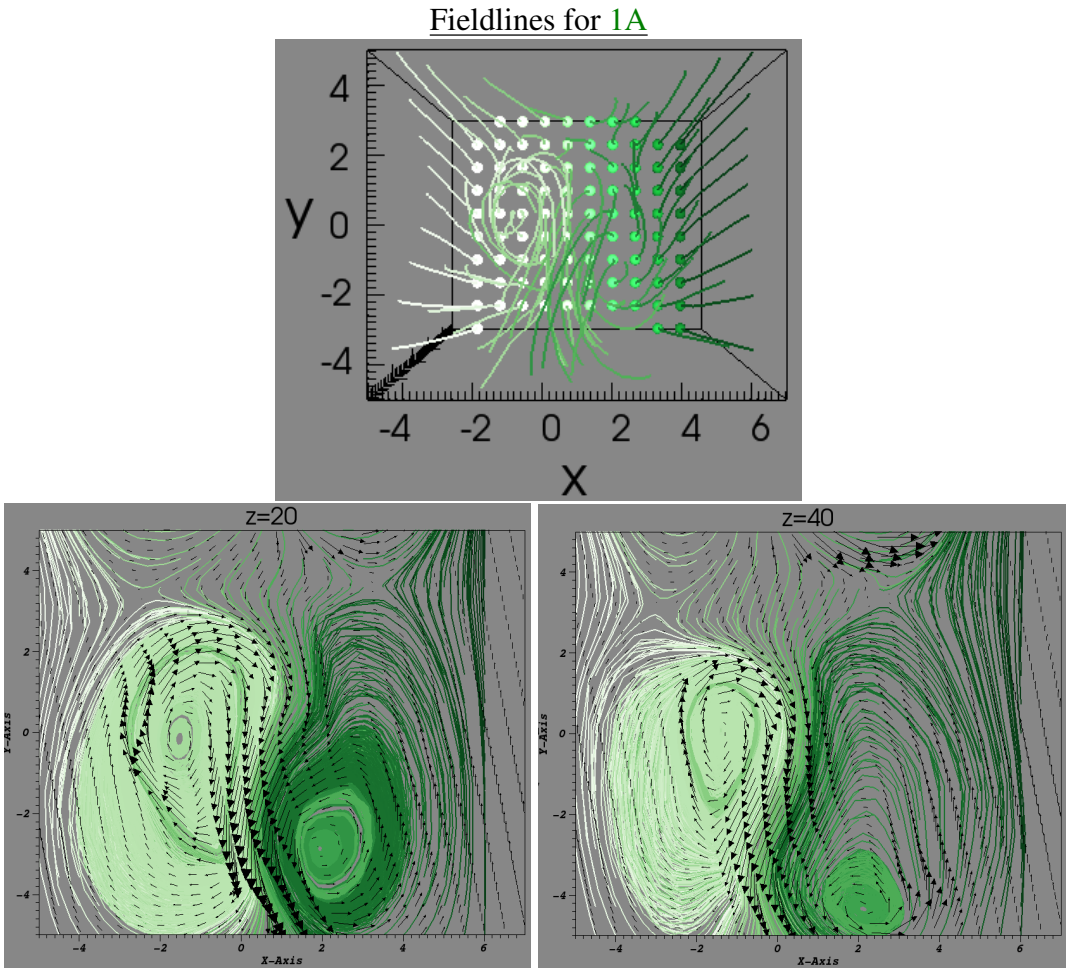

Figure 4.23: Field structures for run 1A at $t = 10,000$. The top plot shows a bird's eye view of the domains fieldlines trace up the box. The lower two plots show 2 slices through $z = 20, 40$, where the fieldlines and vector field are traced.

Fieldlines for 1B

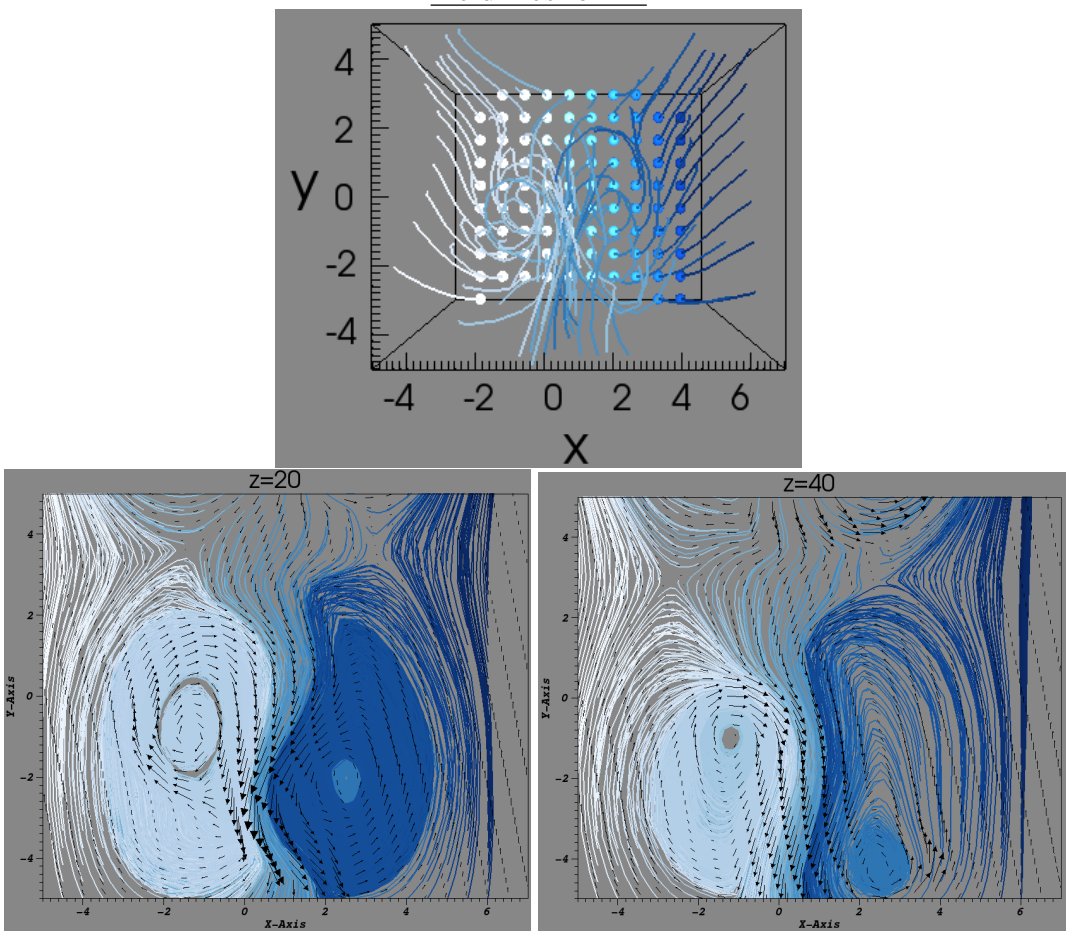

Figure 4.24: Fieldlines for run 1B at $t = 10,000$.

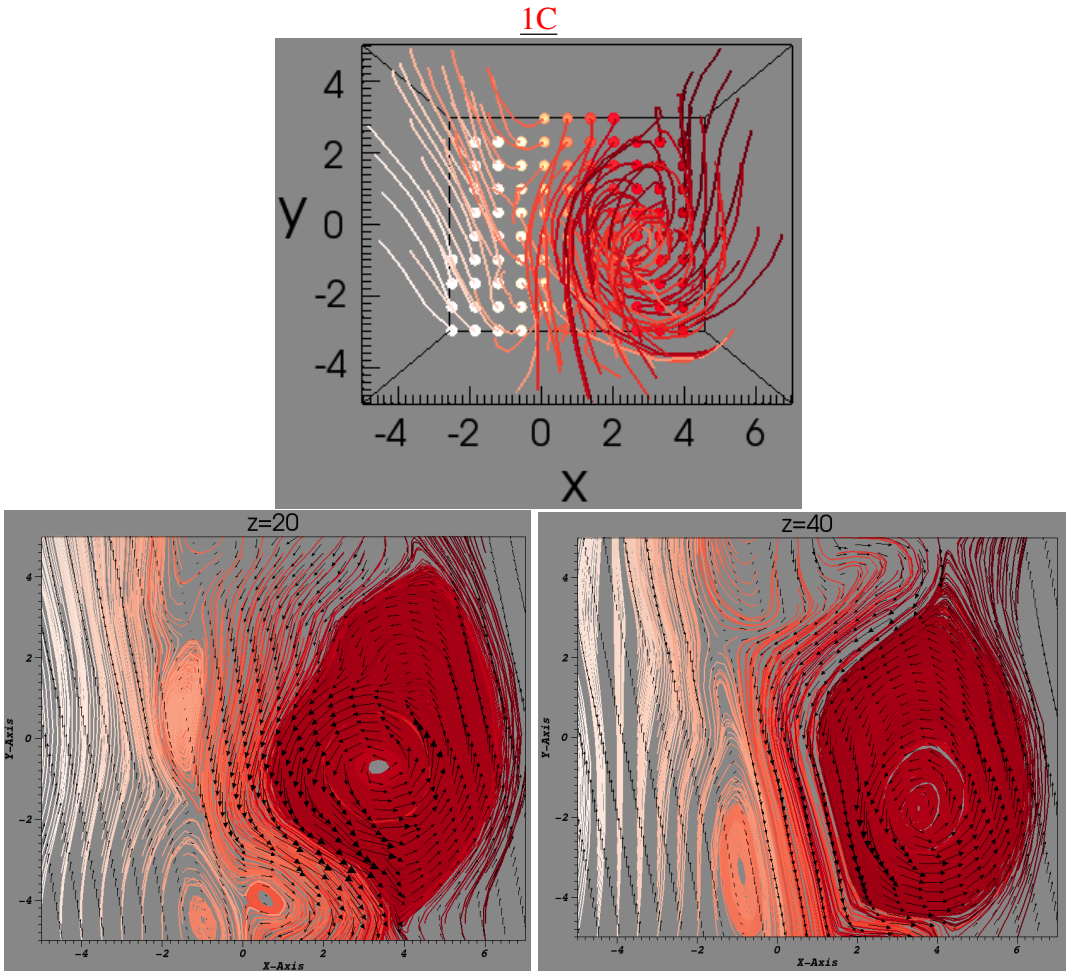

Figure 4.25: Fieldlines for run 1C at $t = 10,000$.

Figures 4.26 to 4.28 display the same plots for all three runs at time $30,000$. Again the most obvious features are the large swirling forms. Run 1C shows the most interaction between fieldlines originating in different sections of the numerical photosphere, and the vortical forms are more distorted. The most complex driver appears to create the most complexity in the magnetic field.

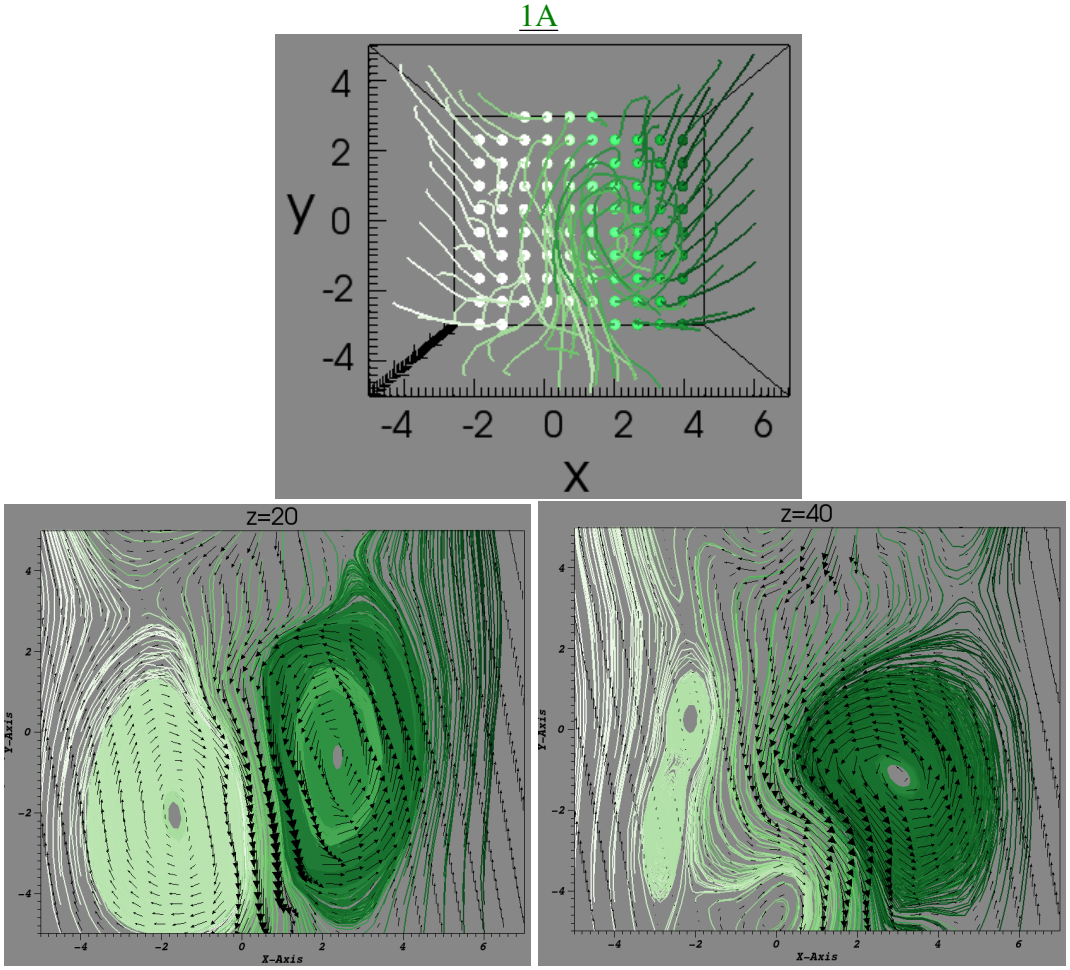

Figure 4.26: Fieldlines for run 1A at $t = 30,000$.

Fieldline plots can also illustrate the change in the system with the release of free magnetic energy. Figure 4.29 takes a snapshot at $t = 22,800$ during run 1C, during a drop in magnetic energy. The fieldlines are particularly messy, with lines of differing shades reaching across the domain. We see distorted swirling features, and also sheared field regions. By Figure 4.30 at time $t = 23,200$, the swirling is more ordered: fieldlines of different shades remain closer to the regions they were seeded in. The 2D streamlines also appear less variable,often either in a more spherical forms. The field seems to have been untangled somewhat, with the twisted regions more contained.

Overall, the highest complexity driver has consistently produced evidence for more

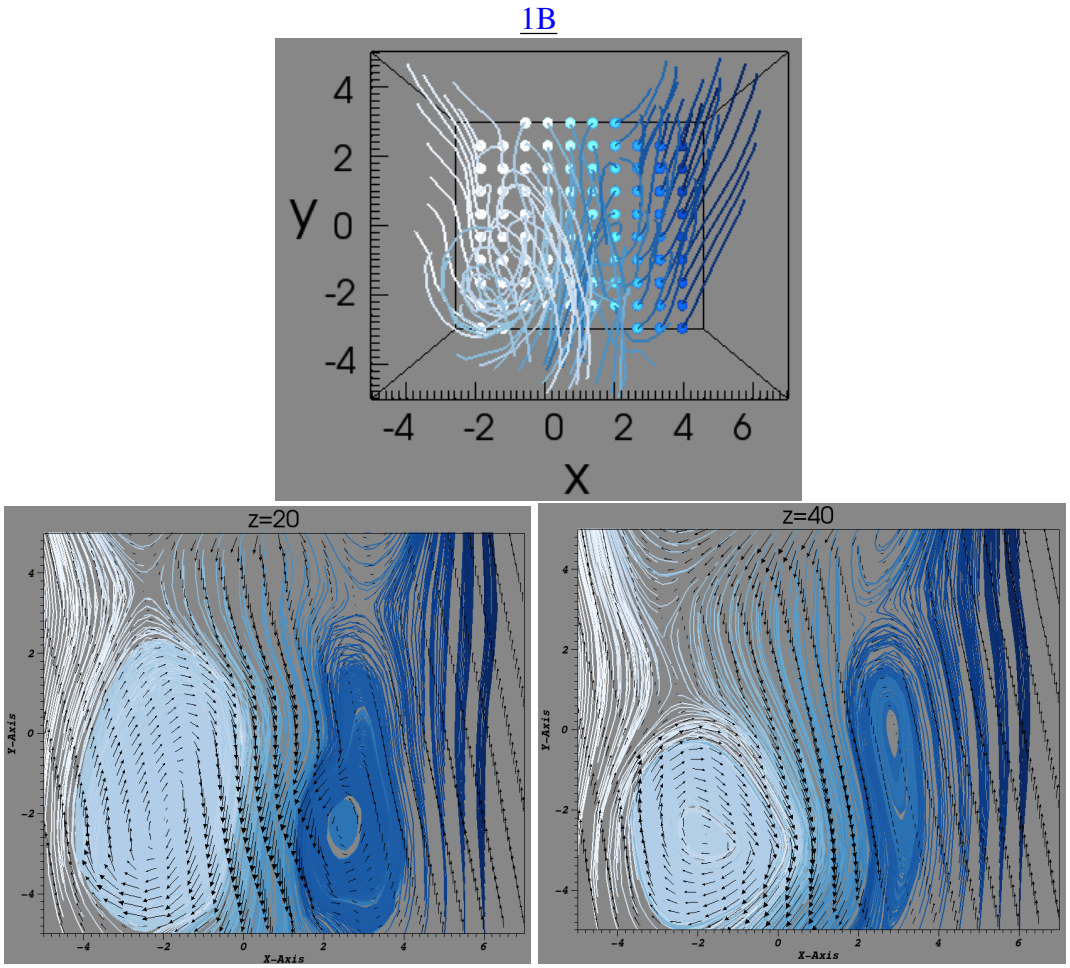

Figure 4.27: Fieldlines for run 1B at $t = 30,000$.

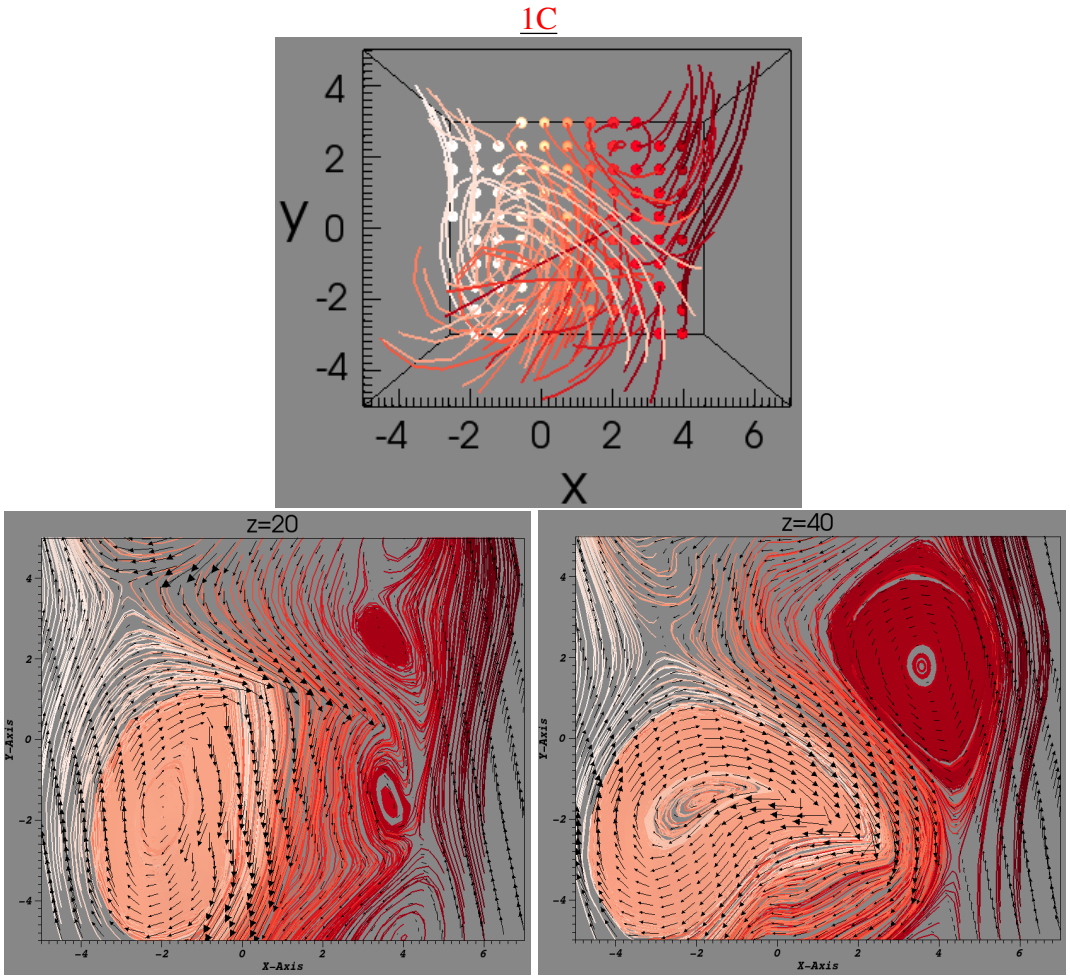

Figure 4.28: Fieldlines for run 1C at $t = 30,000$.

1C, $t = 22,800$

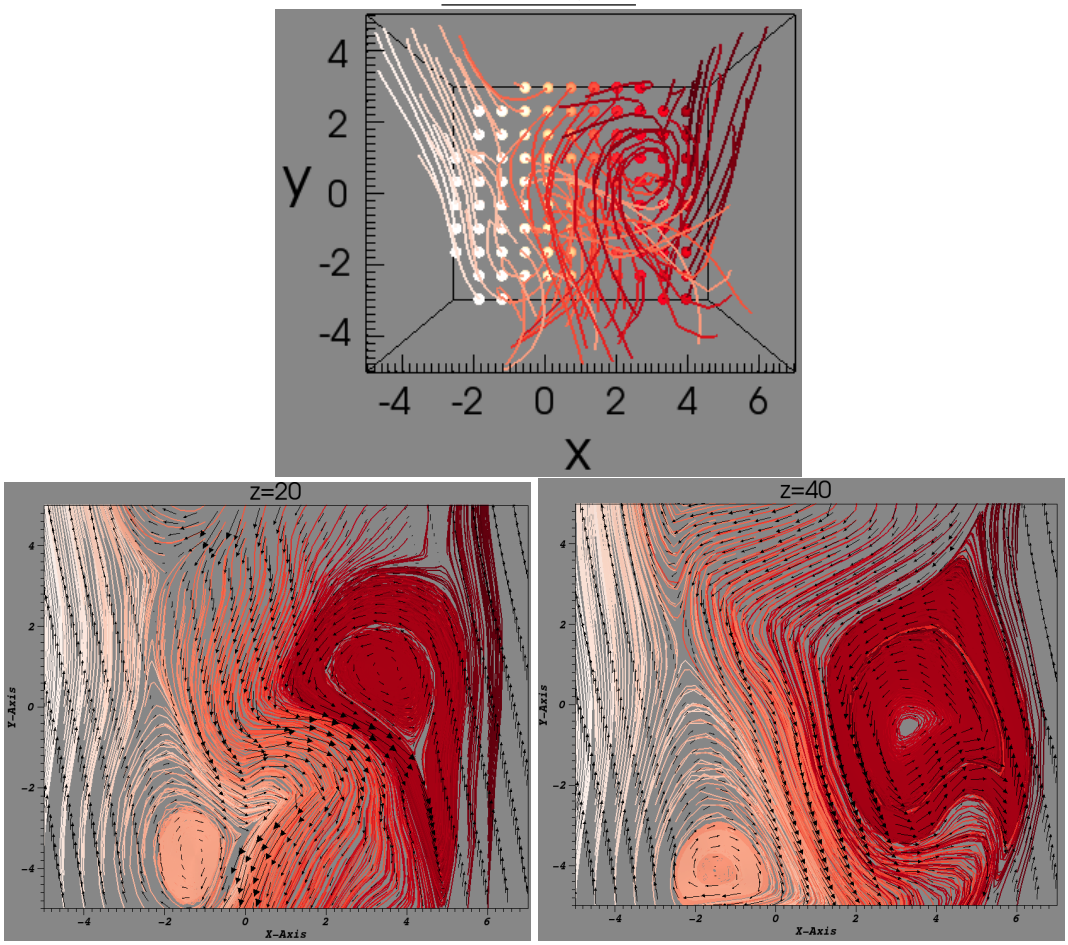

Figure 4.29: Fieldlines for run 1C during a drop .

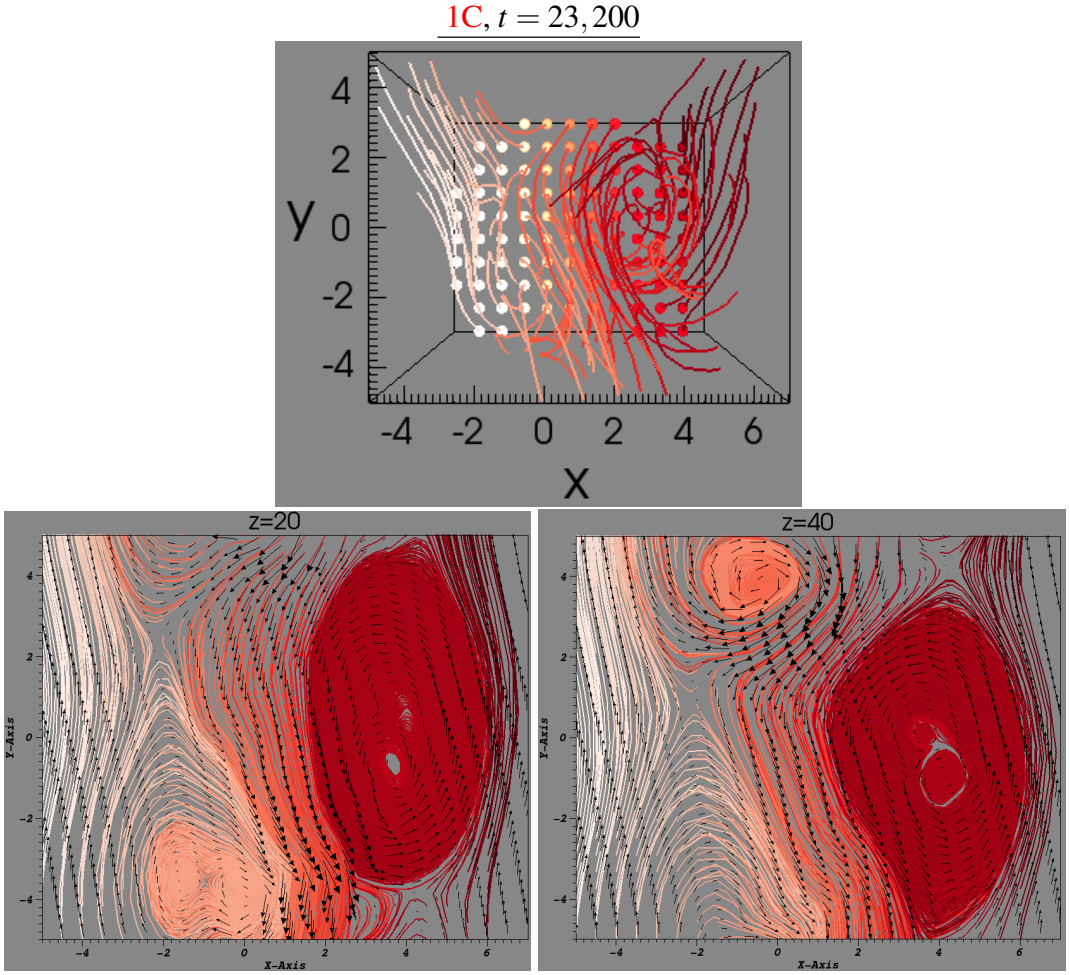

Figure 4.30: Fieldlines for run 1C after a drop.

reconnection opportunities and heating. In this net zero helicity driver case, the complexity dominates the behaviour, and the more complex the driver the more complex the magnetic and velocity fields, leading to higher levels of heating.

4.1.1.2 Case 2

We now turn to the high helicity, low complexity runs with driving vortices spinning with equal circulation. We may see an ordering with complexity as before, but cannot predict how the higher helicity in these systems will come into play and affect potential heating. Again we would expect to see the periodicity of the driver appearing in some quantities.

Comparison plots of quantities for the three runs 2A, B and C

Magnetic Energy

Figure 4.31 (a) illustrates quite a different picture of magnetic energy evolution to that seen previously. Immediately we see that we do not reach a statistically steady state for any of the simulations in this time frame, even though we are running for 596 full periods of driving. The profiles do again have an overall jagged appearance with many small oscillations, but the most intriguing features are drops, much larger than previously seen, over short time periods. At several points in the 2A and 2B runs, there appear to be release events much larger than anything we have observed before in this work. Plot (b) of Figure 4.31 displays an enlarged section of the run during which the largest energy release occurs. This event is part of the evolution of run 2A: the lowest complexity run. In this set of tests, it appears that the ordering of magnetic energy levels is inverted, with the most complex driving leading to lower magnetic energy and vice versa. Driver 2A would twist the field in to the tightest helical structure. The

coherence of the twisting seems to be more important than the complexity: the change in the driving action has made a significant difference. We suspect there is some instability being triggered after sufficient build up of energy. This has not been seen before. Previously, authors such as [90] and [91], when applying shearing motions and coherent twisting, saw tearing mode and kink instabilities take place. However, there the authors only saw a single occurrence, even though these simulations ran for some time; the kink instability case ran for on the order of 1000 Alfvén times and the tearing mode case for 600 Alfvén times. The main difference with our experiments is that the authors used reduced MHD. Again we see that fully 3D MHD may be necessary to capture system behaviour. We will analyse in detail the properties of one of these events and assess the nature of the instability in Section 4.1.2. For now note that here we appear to have the first demonstration of how significant instabilities can actually build up multiple times and lead to reconnection in this type of system.

The peak in magnetic energy for the highest energy run, 2A, 3473 units, is 15.8% in excess of potential. The result of this event is to return the system almost back to its potential level. Here the set-up has allowed a large amount of energy to build up in the field and the release event has successfully dissipated almost all of that energy in a short time span. The magnetic energy drops from 3473 units to 3029 units in a time interval of 909.6 units. This is a release of 93.9% of free magnetic energy over 18.2 Alfvén times, or 12.1 driving periods. From previous calculations we have one unit of magnetic energy representing 1.99×10^{17} J for a potential field strength $B_0 = 5$ G and a length scale of 1 Mm. Therefore in real terms, this release would correspond to a flare of 8.84×10^{19} J, or 8.84×10^{26} ergs - a large nanoflare, not far from microflare category. In addition, the excitement of the large events still should not overshadow the more numerous, smaller events. There are regular discharges of magnetic energy, on a scale comparable to the previous, zero helicity/higher complexity runs. We may

have similar small scale activity, potentially sufficient to maintain a background level of heating, supplemented by larger flare events. The smaller drops here look around the same size as some of the larger drops in 1C, suggesting that these coherent cases induce larger events overall.

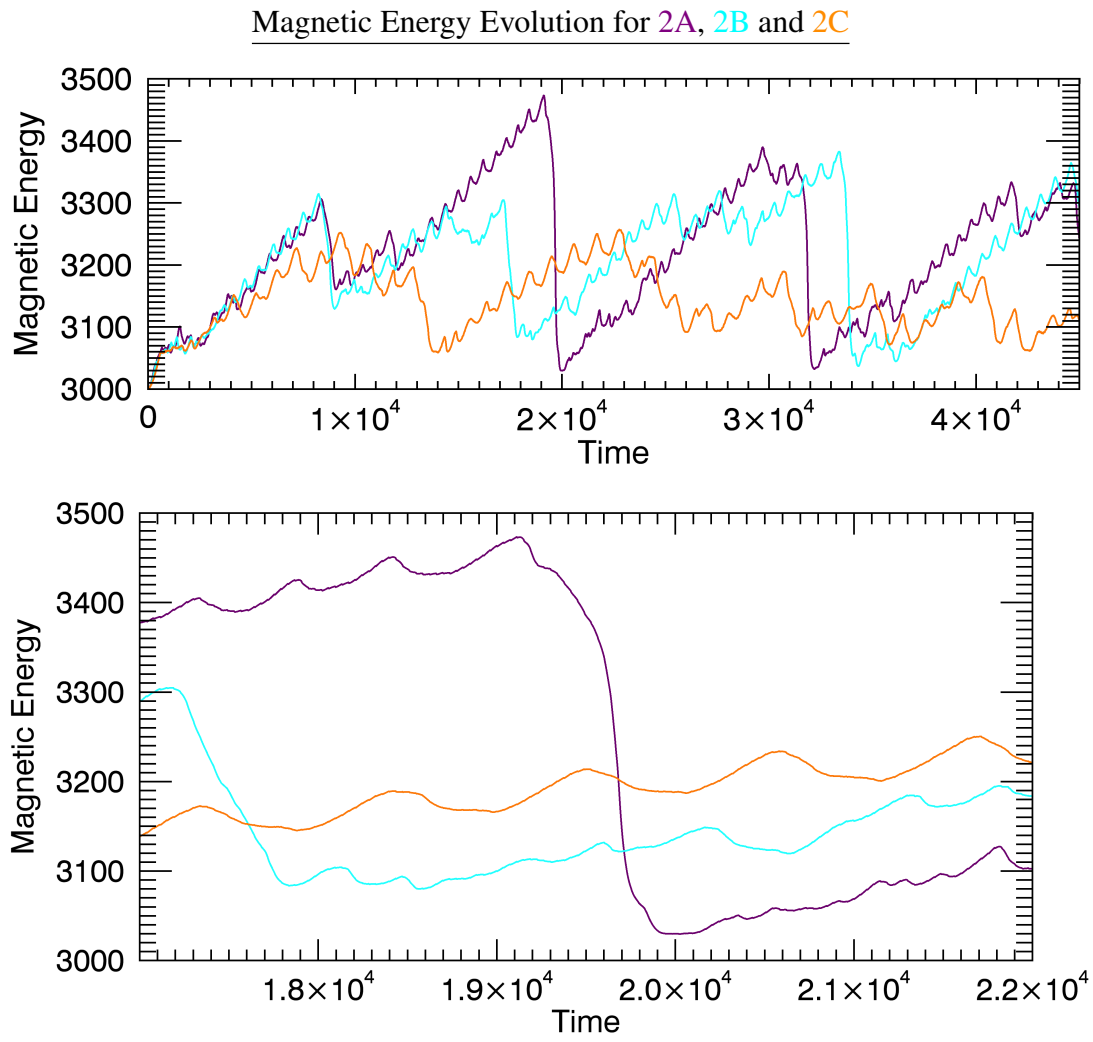

Figure 4.31: Magnetic energy plots for runs 2A, ($x_0 = 0.5$), 2B, ($x_0 = 1$), and 2C, ($x_0 = 2$) for blinking vortices of equal circulation. The run with the most coherent driver, 2A, has the highest levels of magnetic energy. Both this run and 2B, the second-most coherent run, feature build ups to large energy release events. The least coherent run has less dramatic details, but is still highly oscillatory.

Averages of Quantities for Case 2

Run	x_0	Excess	**Viscous/Ohmic**	**Heating Rate**	**Filling Factor**
2A	0.5	7.7%	0.30	0.342	$0.75\% > 1$ $0.043\% > 4$
2B	1.0	7.4%	0.28	0.332	$0.73\% > 1$ $0.044\% > 4$
2C	2.0	4.9%	0.30	0.235	$0.51\% > 1$ $0.019\% > 4$

Table 4.2: Case 2 average values: average magnetic energy in excess of potential value; average viscous to Ohmic heating; average heating rate (rate of change of Ohmic+viscous heating); average percentage volume filling

Kinetic Energy

The kinetic energy profiles differ less - we have very jagged plots, with small values relative to the magnetic energy behaviour. We see very prominent spikes protruding infrequently from a background level. These spikes correspond to times when the magnetic energy is going through large releases. This supports the idea of an instability - clearly there will be a surge in kinetic energy as the field reconnects, particles are accelerated and a flare is triggered. Looking more closely at the spike in run 2A seen in the last plot of Figure 4.32, we see that this lines up with the largest drop in magnetic energy between $19,000$ and $20,000$ time units. The rapid increase and decrease consisting this spike occurs towards the end, but well within, the time interval of the largest magnetic energy drop. This is not the whole story though. The magnitude of kinetic energy is still small in comparison with the magnetic energy, as in Case 1, however the fraction of magnetic energy converted into kinetic energy is even smaller than before. The average kinetic energies over the full simulations are $0.77, 0.74, 0.8$ for runs 2A, 2B and 2C respectively, with a ratio over the full simulation of average free magnetic energy to average kinetic energy for 1C of $298 : 1$. This even more significant dominance of magnetic energy over kinetic energy suggests that a smaller fraction of the free magnetic energy being built up is being dissipated through kinetic energy and

then viscous heating. A larger proportion of energy is being dissipated directly through Ohmic heating. Also, the kinetic energy appears to oscillate around a lower average level than in the complex cases, and there is less variation between these three lower complexity runs. Previously the most complex case had lowest kinetic energy.

We can, as before, look at the magnetic and kinetic energy averaged over half driver periods to extract larger scale variations. Doing this for the magnetic energy over the same time interval as in Figure 4.31 (b) showed very little change, indicating that in this case the driving has less of a directly translatable effect and oscillations are not correlated as directly to driving time. The kinetic energy averages in each half period, for the same time interval as in Figure 4.32 (b), are shown in Figure 4.33. The snapshot begins during the 477th half period, which corresponds to $18,000$ time units. It ends during the 583rd half period, at a simulation time of $22,000$. This does simplify the picture somewhat but we can still see variations above and beyond driver motion. The largest spike still dominates the picture, supporting the idea of a large reconnection event.

Internal Energy

The internal energy follows the same ordering as the magnetic energy, as seen in Figure 4.34. The lowest complexity but most coherent twisting run, 2A, displays highest internal energy, closely followed by the second most coherent run 2B. There are slight bumps visible in these two profiles, corresponding to large events in the magnetic energy profile. This is further evidence of reconnection. Note again however that the constantly increasing internal energy is due to the lack of radiation, and at some point in each of the simulations the plasma beta will become greater than one, switching from a low plasma beta regime to a larger one.

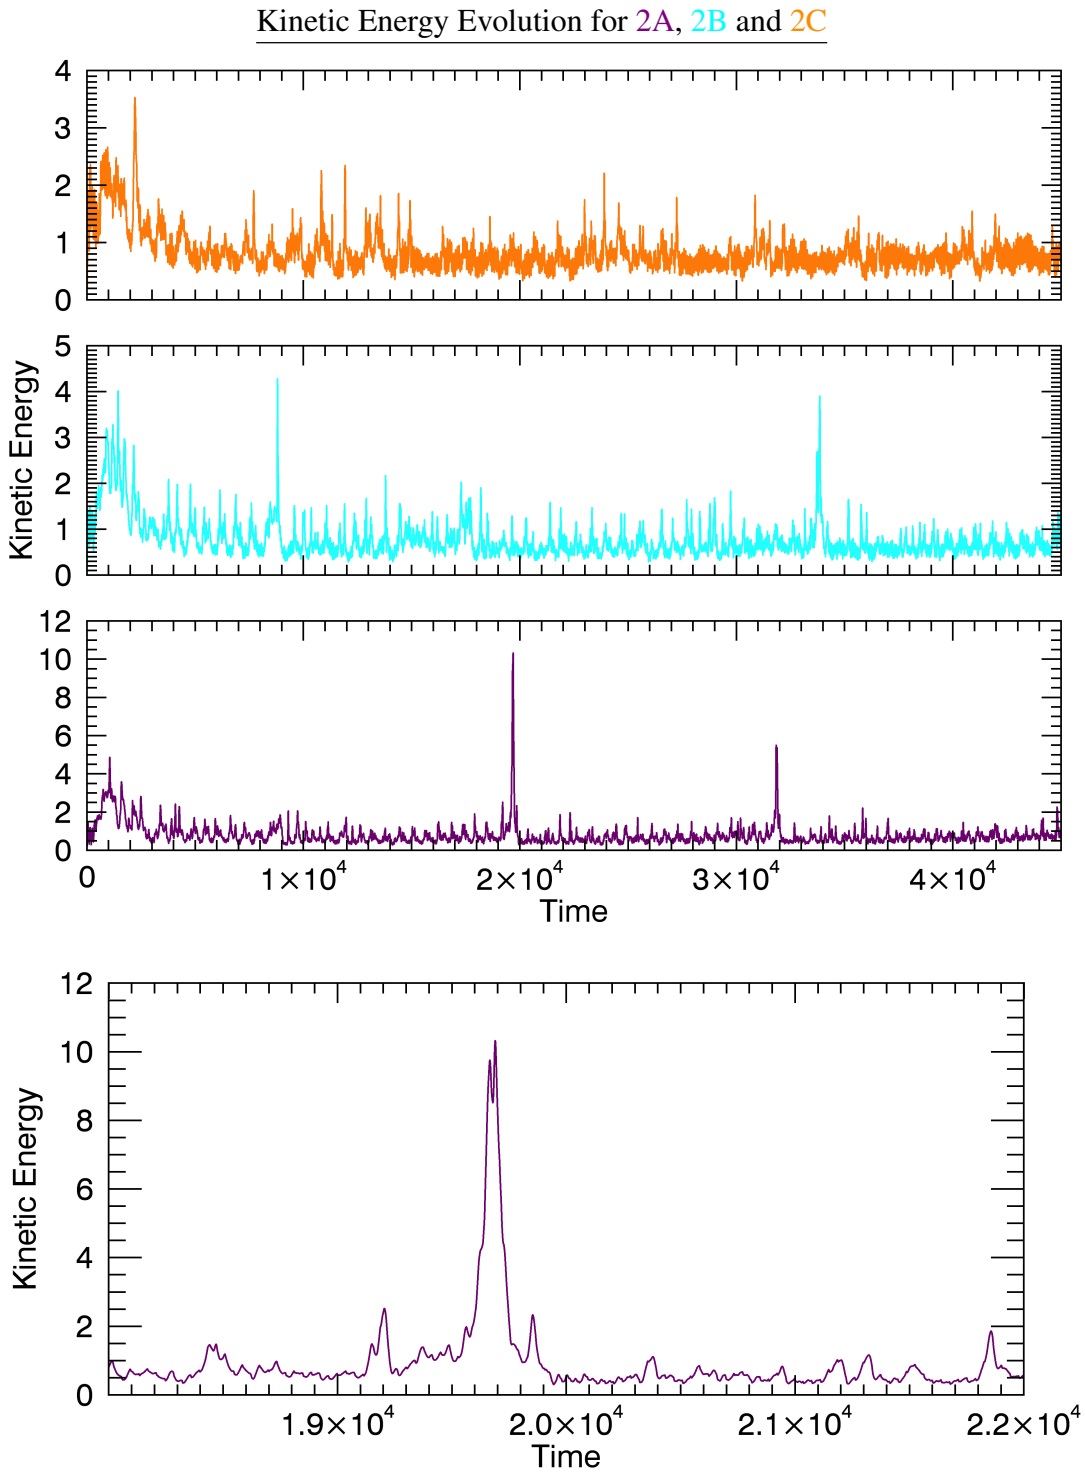

Figure 4.32: (a) Kinetic energy plots for runs 2A, ($x_0 = 0.5$), 2B, ($x_0 = 1$), and 2C, ($x_0 = 2$) for blinking vortices of equal circulation. (b) Zoomed in kinetic energy for run 2A including the spike corresponding to the largest magnetic energy release.

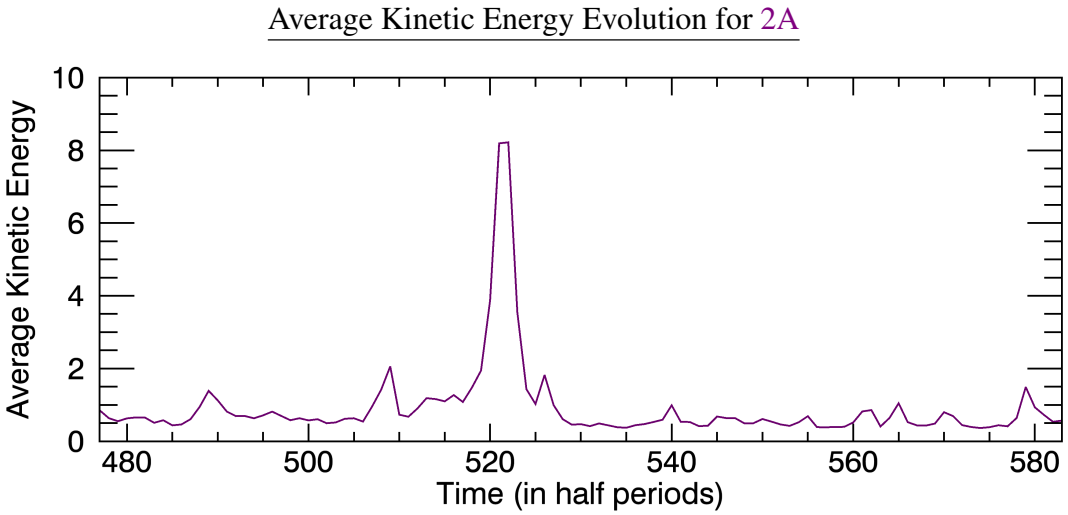

Figure 4.33: Average kinetic energy plots for the most tightly twisting run 2A, ($x_0 = 0.5$).

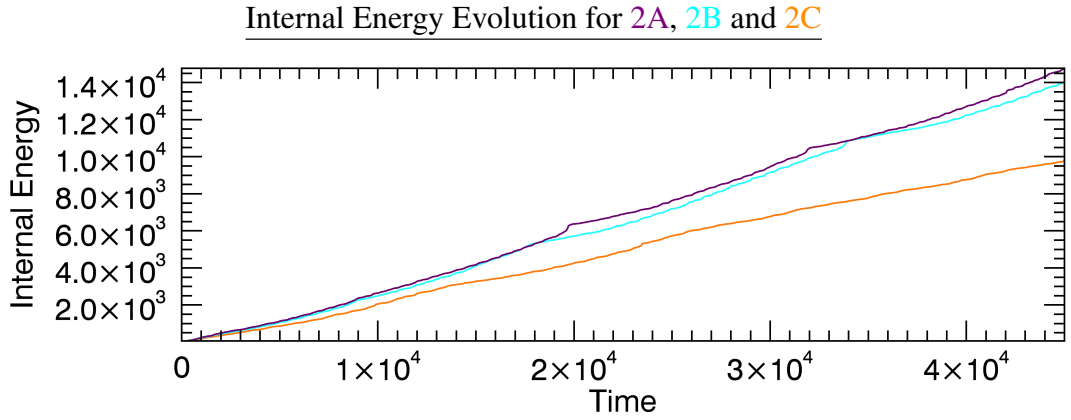

Figure 4.34: Internal energy plots for runs 2A, ($x_0 = 0.5$), 2B, ($x_0 = 1$), and 2C, ($x_0 = 2$) for blinking vortices of equal circulation. Run 2A and B both featured large magnetic energy releases and both have higher internal energy values than the highest complexity/lowest coherence in twisting run, 2C.

Heating

Looking at the heating properties, we see new ratios of average viscous to average Ohmic heating, which are presented in Table 4.2. The plots of viscous, Ohmic and total heating are presented in Figure 4.35. This time the ratio is roughly the same for each run. Ohmic heating dominates in each simulation to about the same degree. There is less variation between runs compared to the previous set of simulations. Viscous heating appears to increase in proportion with ohmic heating, leading to a similar ratio each time. Instead of the completely smooth heating profiles of the complex case, we see small jumps. As with internal energy the sudden small increases correspond to the large magnetic energy release events. The average total heating rates are also presented in Table 4.2. All are much higher than before. This tells us that even in lower complexity environments, higher helicity injection is able to lead to more energy release.

We look at the heating rate (rate of change of the viscous plus Ohmic heating) for these simulations in Figure 4.36. Once again there are prominent spikes in this quantity at the release event times. This points to flare like events occurring in these lower complexity but more tightly twisted runs.

4.1.2 Currents, Field Structure and Instability

Here we will compare the current levels of each run, and then focus in on the largest event occurring in Run 2A, a suspected instability.

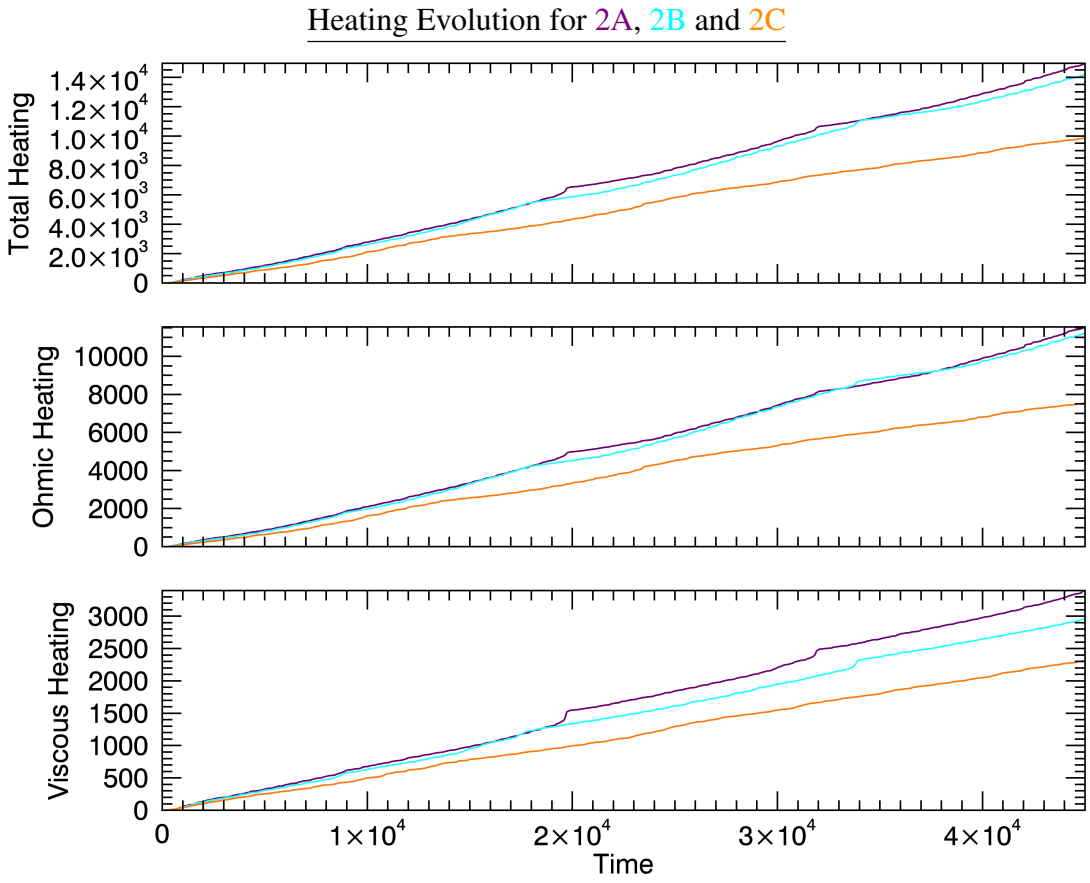

Figure 4.35: Heating plots for runs 2A, ($x_0 = 0.5$), 2B, ($x_0 = 1$), and 2C, ($x_0 = 2$) for blinking vortices of equal circulation. Again the most coherent run involves the largest heating levels.

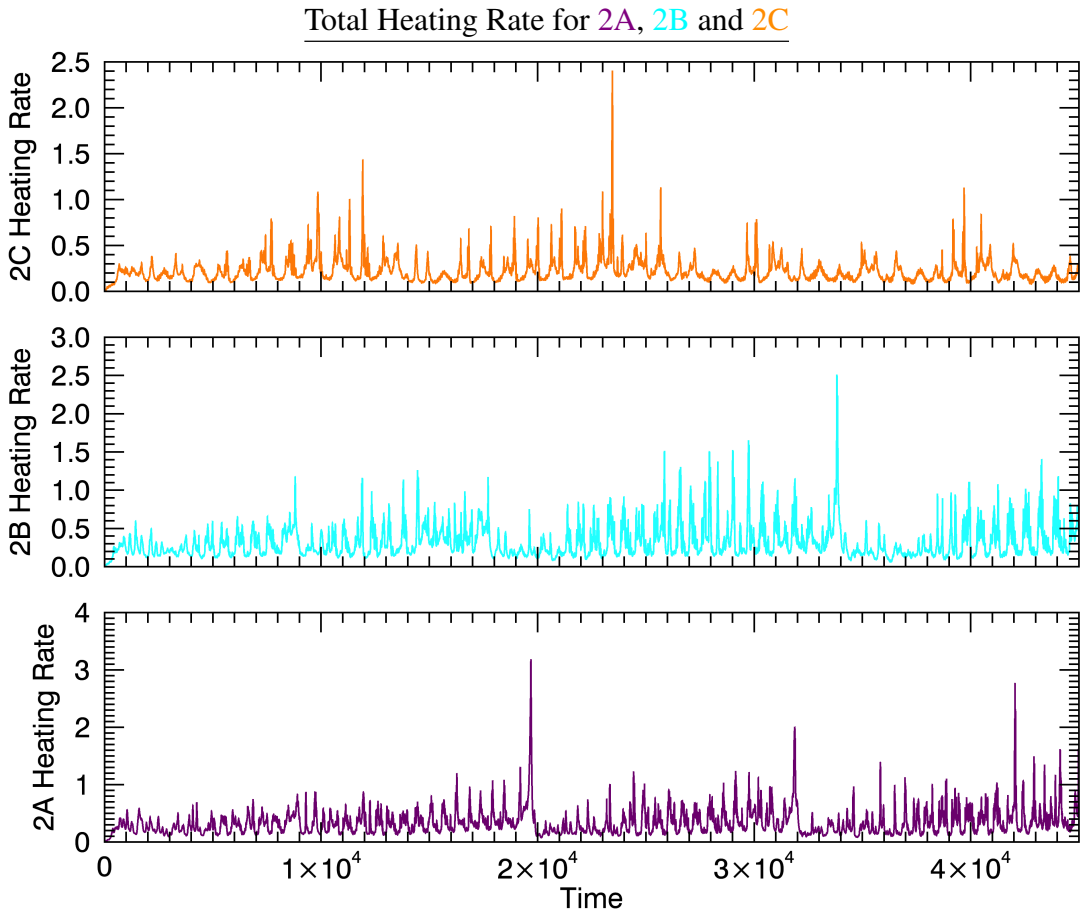

Figure 4.36: Heating rate plots for runs 2A, ($x_0 = 0.5$), 2B, ($x_0 = 1$), and 2C, ($x_0 = 2$) for blinking vortices of equal circulation.

Figure 4.37 displays the maximum current density present in the domain over time for each of the higher helicity scenarios. Clearly runs 2A and 2B see the largest currents. Run 2C varies over a smaller range of values with a smaller maximum. This fits with the narrative already established: the higher helicity levels in these runs is the dominant factor in behaviour, with instabilities and currents most significant for the most coherently twisted states.

Figures 4.38 and 4.39 illustrate how the percentage of the domain filled with current densities larger than 1 and 4 units, respectively, changes over time. Both show the runs with large magnetic energy release features containing more regions of higher currents. We continue to see an ordering with driver's abilities to twist the field into a tight spiral. This seems to be the determining factor in reconnection and heating opportunities for Case 2, rather than the complexity as in Case 1.

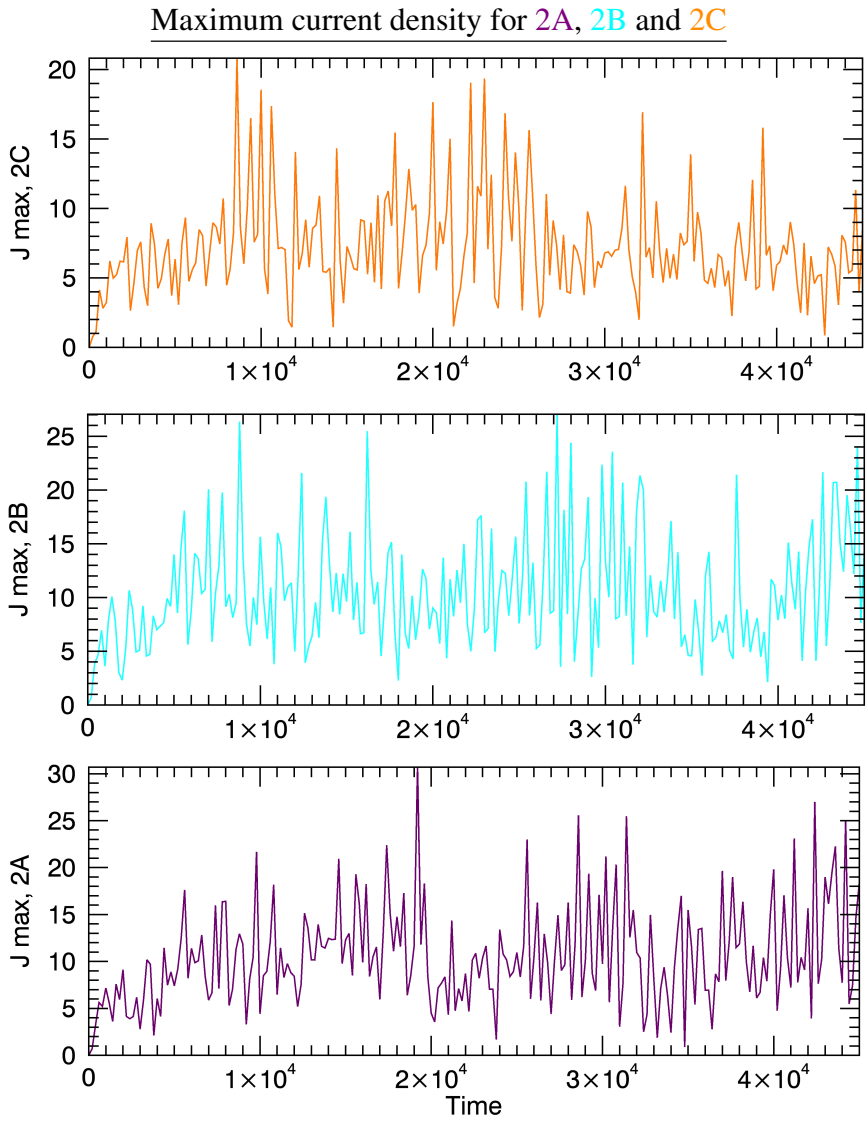

Figure 4.37: Comparison of the maximum current in the domain over the simulation, for runs 2A,B and C. The largest maximum currents are to be found in run 2A, closely followed by run 2B, echoing the trends seen so far in other variables.

Instability

Based on the vortical driver profile, we first considered the possibility of a kink instability occurring at the large release event. The work in [91] found a kink instability taking place as a result of the application of vortical motion on one photospheric boundary. Rather than blinking vortices, their method used one, time independent vortex only, driving for long time durations. Kink instabilities can develop in a system where some critical degree of twist is attained, and a kink is triggered along an axis of symmetry. In some cases the instability may lead to a solar flare as the system attempts to relax to a lower energy state, via reconnection and energy release.

Kink instabilities in MHD environments have been examined before. For example, the authors of [45] explores the phenomenon in an ideal environment, discussing how kink instabilities can lead to solar flares. They show how even with line tying (which our system has) of fieldlines at the photosphere, which should stabilise things, a flux tube still becomes unstable after sufficient twisting. Kink mode perturbations lead to reconnection and potential for flares. Resistive experiments include those by [37], who also used Lare3d. Some of these tests also found kink instabilities occurring in line-tied resistive loops under the influence of shearing. To ascertain if we have this type of instability here we will have to examine current and fieldline properties. In our setup we do not have a distinct axis of symmetry along which the magnetic field is twisted - in all of our simulations there are two vortices, placed apart. But note that the largest events happen in the run where the vortices are closest together. We do have continuous twisting of the overlying field in one direction. We see multiple occurrences in two of the three runs of an increase in magnetic energy, up to a point whereby a large amount of energy is relaxed, accompanied by a spike in kinetic energy (indicating acceleration of plasma) and heating. This fits with the profile of a flare resulting from some instability.

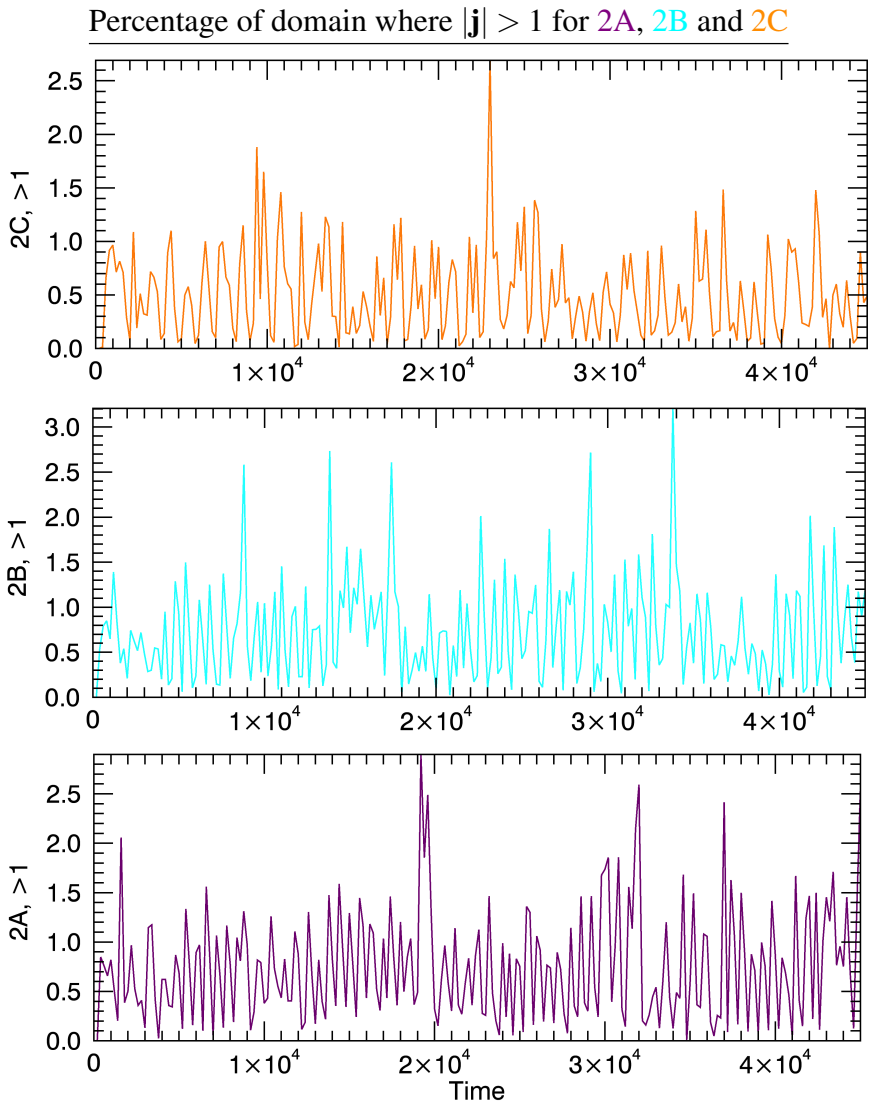

Figure 4.38: Percentages of each domain where the current density is greater than 1. Runs 2A and B are closely matched, whereas the least coherent run 2C generally oscillates about a lower average.

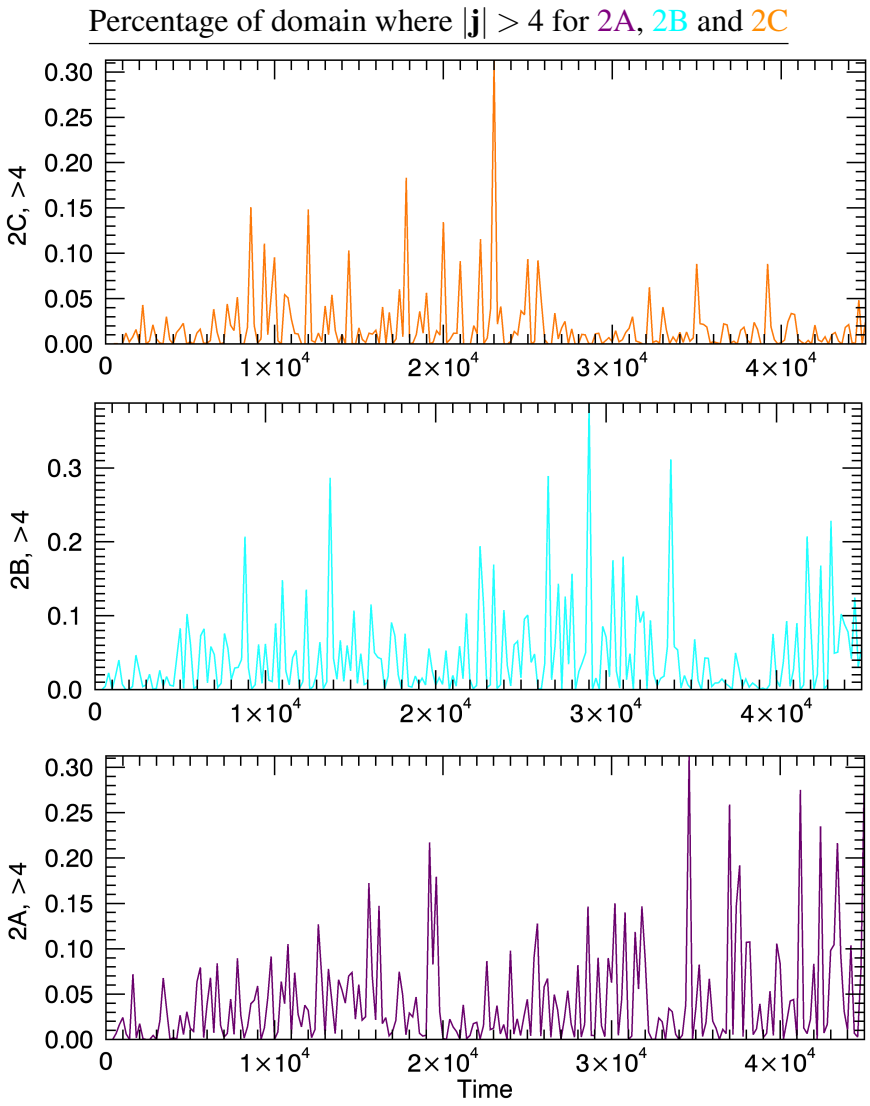

Figure 4.39: Profiles of the domain percentage filled with currents with value higher than 4 for 2A, B and C, where again we see the same ordering with coherence, rather than complexity.

Why do we observe evidence for instability in only 2 of the 3 runs? We can understand this by returning to the argument of limiting cases of vortex placement, which we discussed in Section 2.1.5. Were both vortices positioned at $x_0 = 0$, each vortex would twist the same volume of the field the same way. This would be the most coherent case of twisting we could implement and would be closer to the scenario mentioned in [91]. We would have a single, tightly wound structure. Taking x_0 to infinity, the action of one vortex would not affect the other. We would eventually see two tightly wound structures, constructed in twice the time. Between these two cases the twisting is less coherent and contained; we see twisting over a larger domain, but at a shallower angle. Run 2C is the furthest (of our 3 cases) from either of these limiting cases, and it is where we do not see large release events. In order to trigger something like a kink instability the field has to be sufficiently coherently wound up; the angle of twisting must reach a critical point. When $x_0 = 2$, perhaps the winding is not tight enough. This run has the highest complexity out of the Case 2 scenarios, and is most similar to the previous Case 1 scenarios. We see a magnetic energy profile closer to that of the previous runs, with many smaller events only. Run 2B has the vortices closer together and would lead to tighter winding, but not as tight as in run 2A, $x_0 = 0.5$. By our hypothesis run 2A should be most susceptible to instability; and indeed, this is where we see the largest reconnection events.

Firstly, we look to sample fieldlines and current density isosurface plots for Run 2A. Time $19,000$ to time $20,200$ covers the largest magnetic energy release, from the last of the build up, until the instability sets in and free magnetic energy is released, and the beginning of the next wind up of energy. Between these two times is where we see jumps in viscous and Ohmic heating and the internal energy, as well as a spike in kinetic energy. We focus on this interval.

The plots in Figure 4.40 give a sense of how tangled the fieldlines are. Looking down into the domain from above, fieldlines are coloured in shades corresponding to where on the $z = 0$ plane they began. Footpoints situated towards maximum x are darker, while those towards the minimum value are lighter. At $t = 19,000$, when the magnetic energy is still building, the field displays a localised twisted region in the middle of the box. Fieldlines outside this region maintain similar x coordinates but appear to reach across in y as they trace up the box, in opposing directions depending on whether they started at $x < 0$ or $x > 0$. In the next snapshot the picture is less ordered and by $t = 19,600$ we see a much messier structure. Light coloured fieldlines are seen to reach across in x to the darker half of the domain, indicating a great deal of tangling, and we see this to an even greater extent at the next snapshot at $t = 19,800$. By the last two times, the magnetic energy release has ended and the spikes in kinetic and internal energy and heating have occurred. Clearly a restructuring has taken place. At this point the magnetic energy has almost returned to its potential level, and the fieldlines start to re-form the swirling structure, particularly evident in the last plot at $t = 20,200$.

Looking now at current density isosurfaces, Figure 4.41 displays regions of constant current density over the event, starting from $|\mathbf{j}| = 1$, then 2, 4 and 6, at times $t = 19,000,\ 19,600$ and $19,800$. At $t = 19,000$ we have several ribbons of current flowing up through the domain at the lowest threshold. Increasing the level of $|\mathbf{j}|$ we very quickly see these ribbons all but disappear, with only one small current fragment evident at the highest thresholds. By $t = 19,600$, during the energy release, the domain is filled with many long, thin current layers at $|\mathbf{j}| = 1$. We see fewer as we increase the threshold and the extent of each ribbon becomes smaller, but even at $|\mathbf{j}| = 6$ there are still many small fragments visible. Towards the end of the energy release at $t = 19,8000$ we still see multiple current layers at the lowest current value but they have become less volume filling, narrower and more elongated. There are only some

Fieldlines for 2A

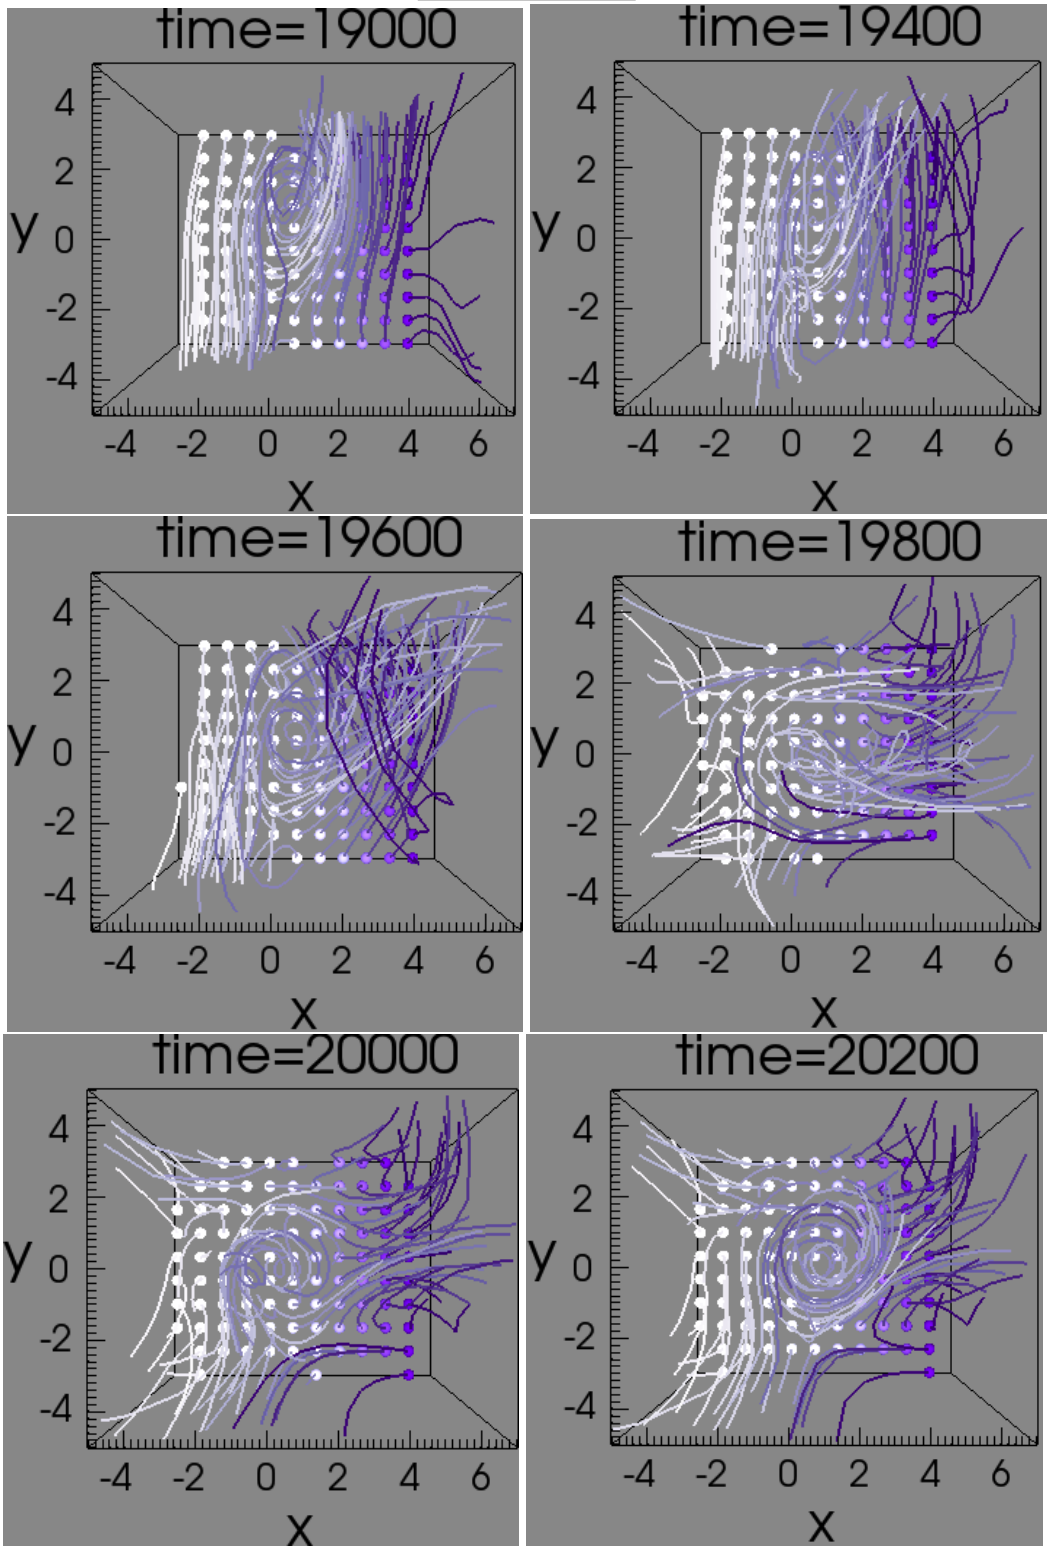

Figure 4.40: Fieldlines for run 2A, ($x_0 = 0.5$), seen from the top of the box over the course of the largest release event.

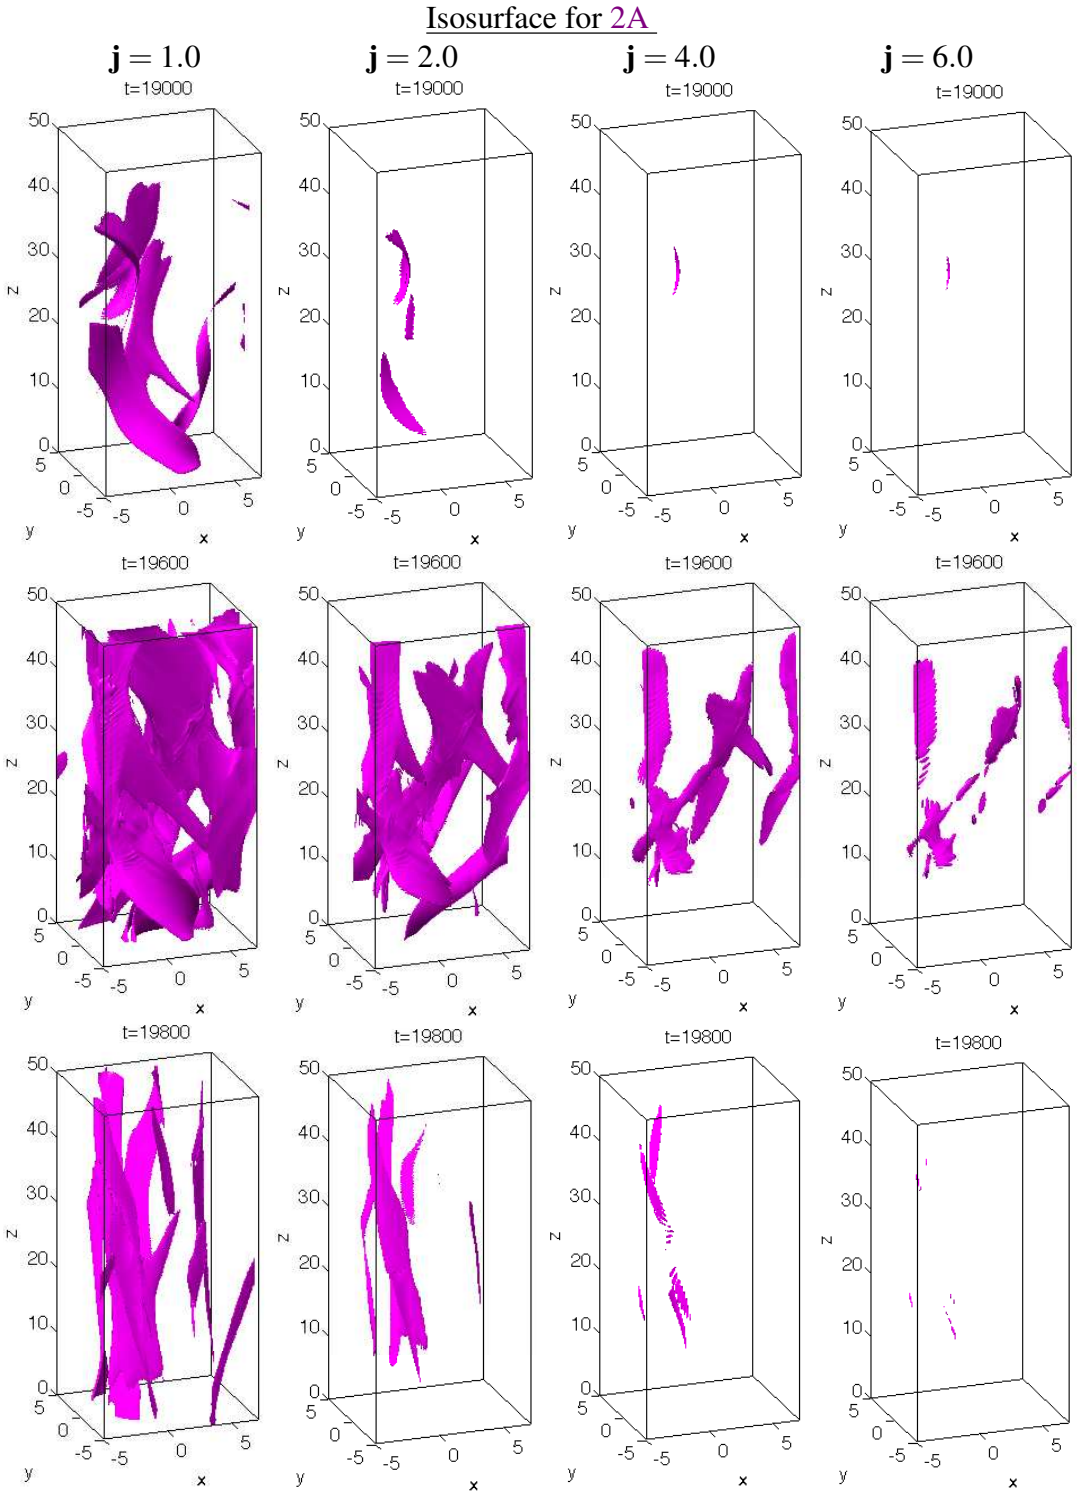

Figure 4.41: Isosurface plots for run 2A, ($x_0 = 0.5$), at values of $|\mathbf{j}| = (1, 2, 4, 6)$ during the large release event.

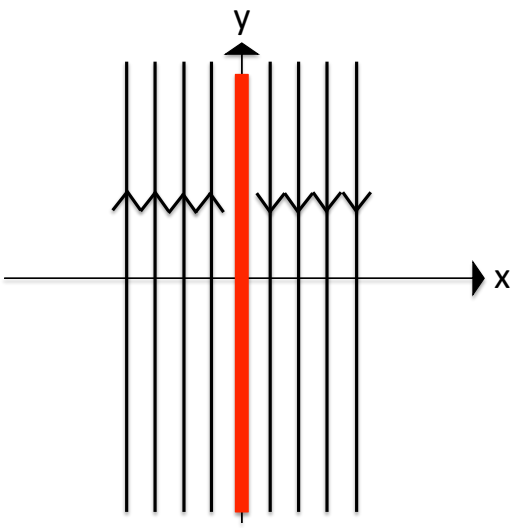

Figure 4.42: Sketch of a scenario in which a tearing mode instability may occur. Fieldlines are illustrated switching orientation across $x = 0$, where the gradient in the magnetic field causes the red current layer. This type of field is said to be unstable to tearing modes, such that reconnection can be triggered across the boundary.

tiny fragments still present at the highest level. These tell us two things. Firstly, we have further support for an instability being responsible for a large reconnection event. Across the time span the current structure changes from consisting of very volume filling, high valued, messy current layers (matching the messy appearance of the fieldlines) to narrower, smoother layers which do not attain the same levels. Secondly, while we see an instability, it does not appear to be a kink. There are no swirling helical structures in the isosurfaces, only broad vertical layers. This points instead to a tearing mode instability. There was evidence of this category of instability in the field line plots also: recall snapshot $t = 19,000$, where those fieldlines outside of the twisted region in the middle of the box appeared to shear across in y. When the fieldline was seeded between $x = [-5,0]$ the trajectory seemed to take it towards $y = -5$; many of those seeded beyond the swirling region at $x = [2,7]$ sheared across towards $y = 5$.

This type of field has been known to give rise to instability in the form of a tearing mode since the work of [34]. Tearing mode instabilities can occur when a field changes

direction across some interface. Figure 4.42 illustrates the nature of the structure which can give rise to this type of event. This simple 2D sketch shows that when the direction of the magnetic field changes across $x = 0$, a boundary is formed along the y-axis. The large gradient in the magnetic field caused by the sudden change in alignment can form a current layer, which if strong enough and for non-zero resistivity can trigger reconnection. This type of field set-up is therefore said to be unstable to tearing modes, almost as if the two opposing field directions tear the field apart. We mentioned before the work of [90], who observed tearing mode instabilities when applying shear flows on numerical photospheric boundaries. However, that work employed reduced MHD and saw only one reconnection event take place before the system settled into a statistically steady state. No further instabilities developed even though the simulations ran for 600 Alfvén times and were continuously driven. Note the difference to our fully 3D simulations, where two of our set-ups have triggered multiple events.

For further details and confirmation of the tearing instability, we turn to 2D fieldline plots and also consider the strength and direction of the field by taking slices through the domain at constant z. By plotting the vector field over fieldlines we can see exactly where and why boundaries are forming.

Figure 4.43 shows such slices at heights of $z = 10, 20, 30$ and 40 at $t = 19,000$. In the first plot at $z = 10$ we see a confirmation of the structure suggested by the fieldline plot. Between about $x = [-1..1.5]$ the streamlines and vector field outline a region where the field is twisting. It is not perfectly circular however, taking an elongated, elliptical form. On either side of this, the stronger, oppositely directed field regions are clear: on the left hand side the field points to negative y, while on the right hand side it points to positive y. At $z = 20$ we see a similar picture. In the illustrations at $z = 30$ and 40 also there is a clear switching in field direction between $x = 0$ and $x = 1$. There are some regions of stronger field along fieldlines outlining twisted structure too, for example in $\mathrm{x} \times \mathrm{y} = [-2, 2] \times [-5, -3]$, but the oppositely sheared regions dominate the picture.

$t = 19000$

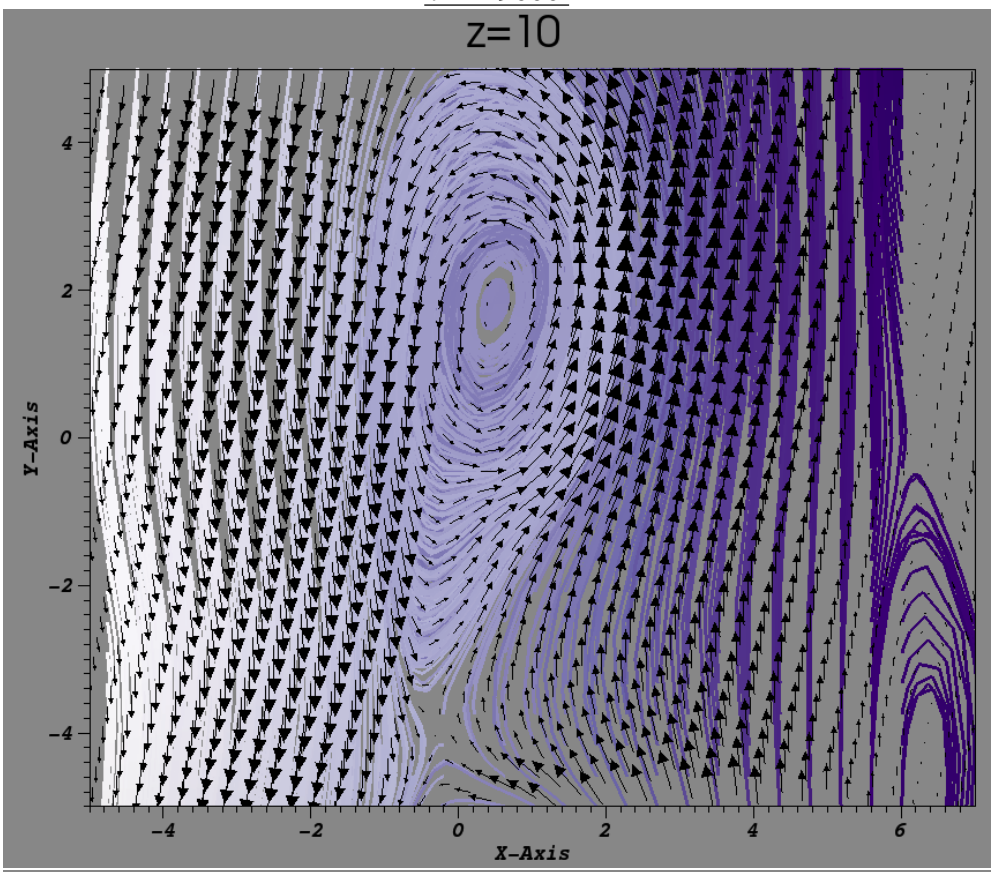

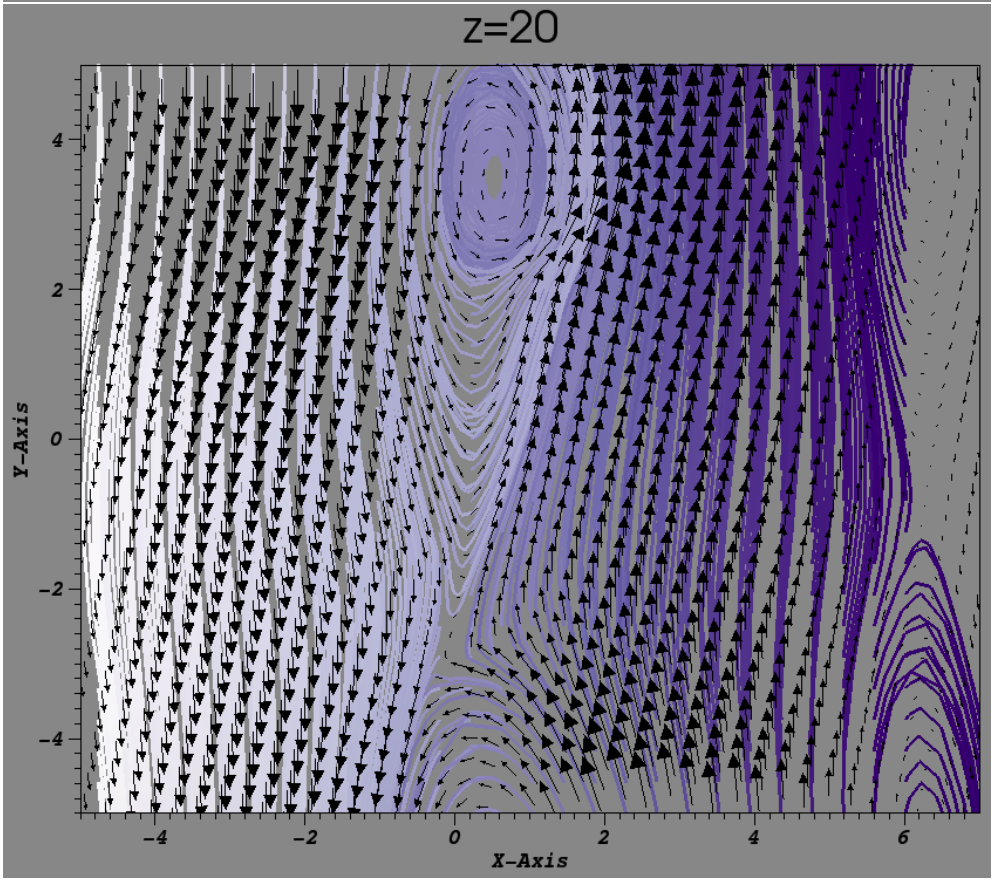

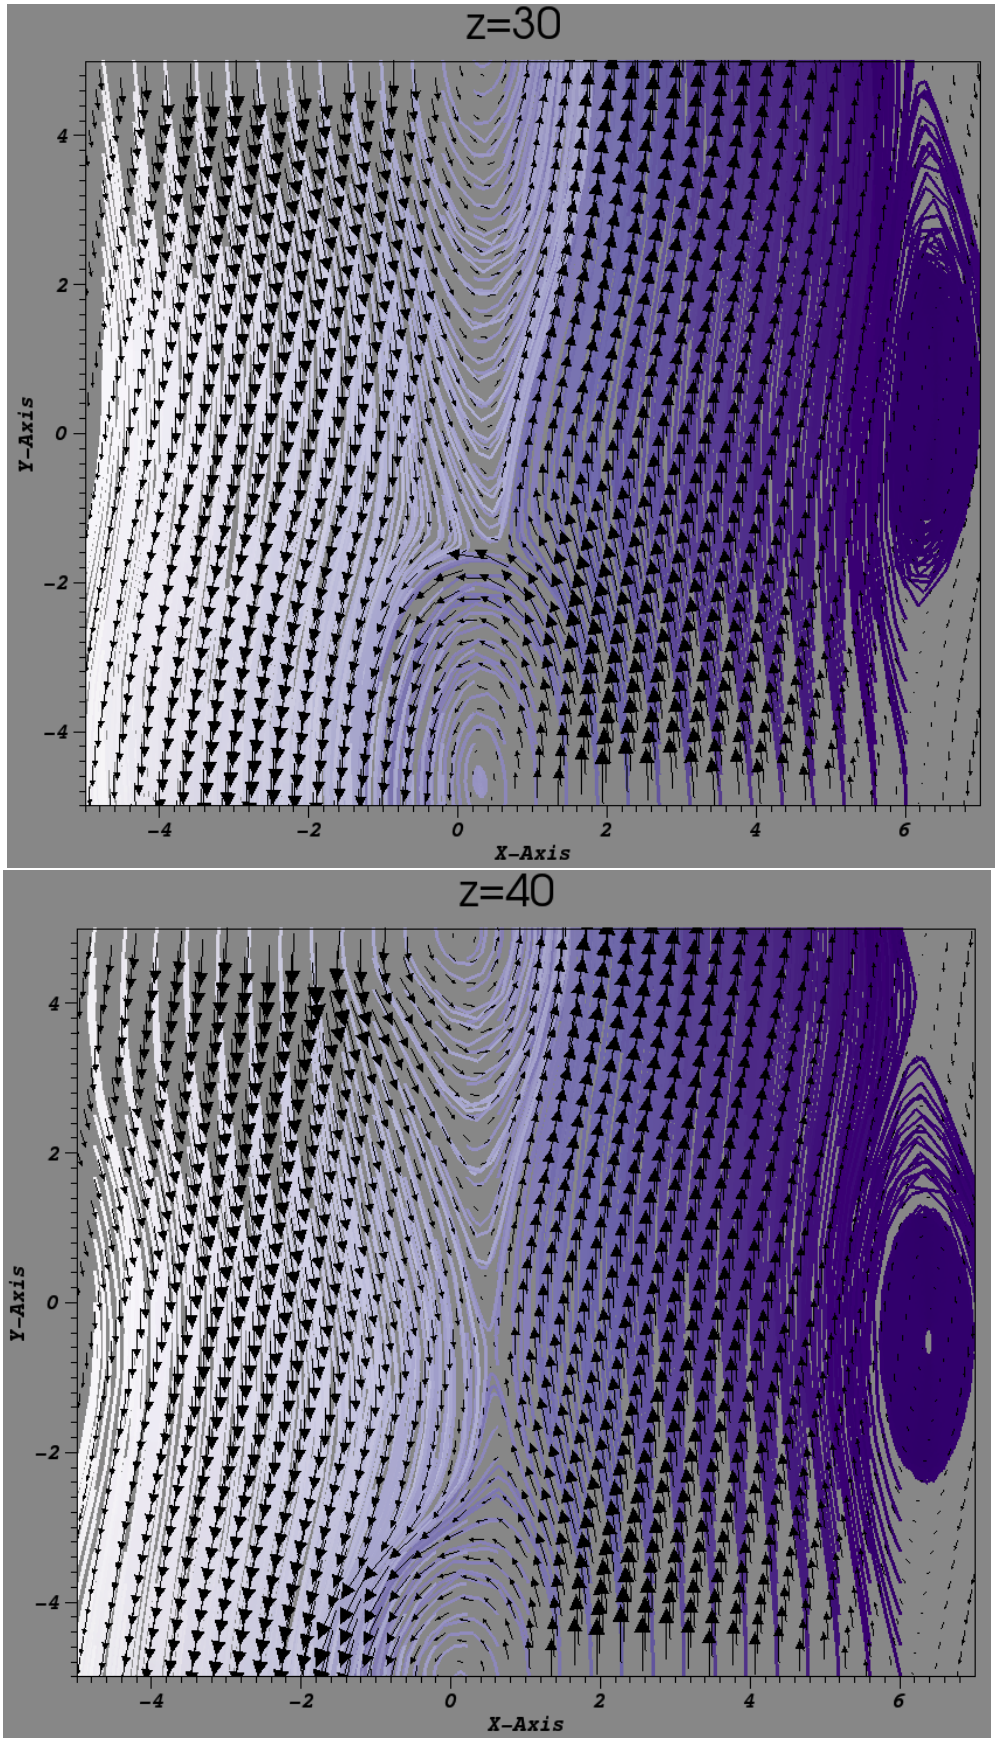

Figure 4.43: The magnetic field at slices of constant z at $t = 19,000$ for run 2A. The strength and direction of the field are represented by the black arrows, while the field-lines are plotted in purple-white.

We look also at the current density and Ohmic heating distribution at these same slices in Figures 4.44 and 4.45 respectively. In all slices, the regions of highest current and heating correspond primarily to the boundary layer between opposing field orientation and to a lesser extent, the twisted field regions.

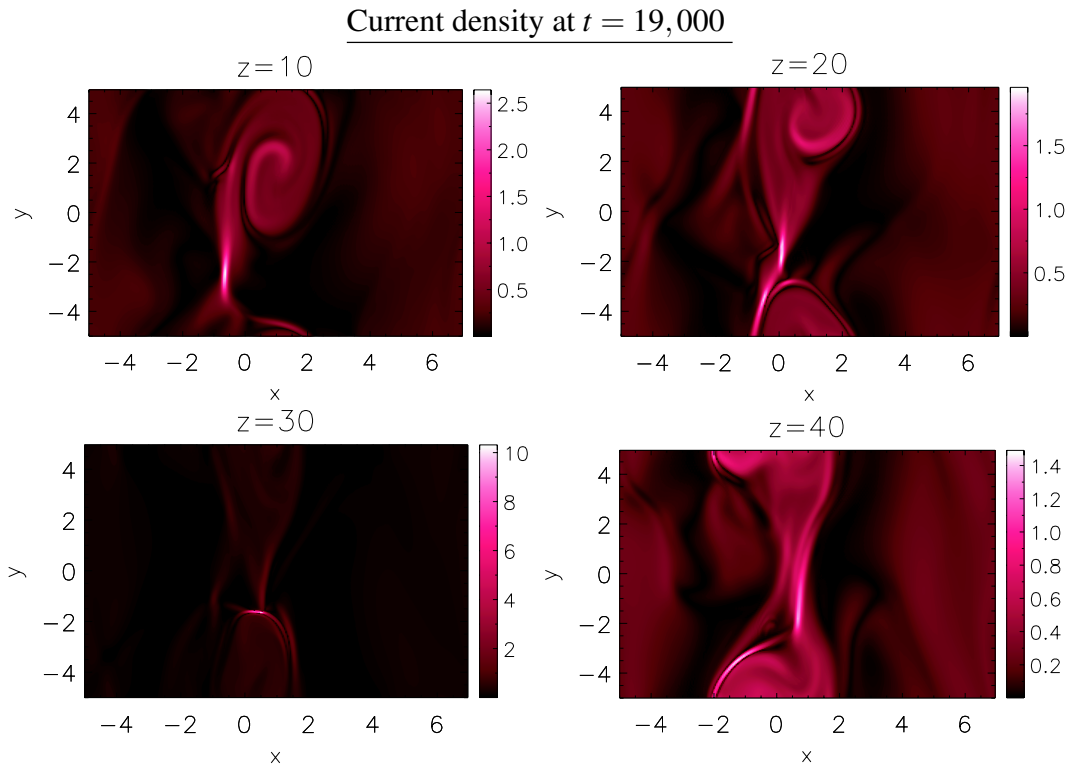

Figure 4.44: The value of $|j|$ in planes of constant z for run 2A at $t = 19,000$.

The fieldlines and orientation at the same heights for the snapshot at $t = 19400$ are shown in Figure 4.46. Again we see a pronounced division in field direction as we move across x, at all heights. There is twisted field also, but again this tends to be in the form of elongated, elliptical shapes lying along y between the regions of opposing field direction. Visually the streamlines are dominated by oppositely sheared field. As seen in Figures 4.47 and 4.48, the layers of current and heating are generally higher than before and are present at the boundaries between opposing field, supporting further the tearing mode theory.

Next we check time $t = 19600$, for slices at $z = 20$ and 40. Recall by this time the

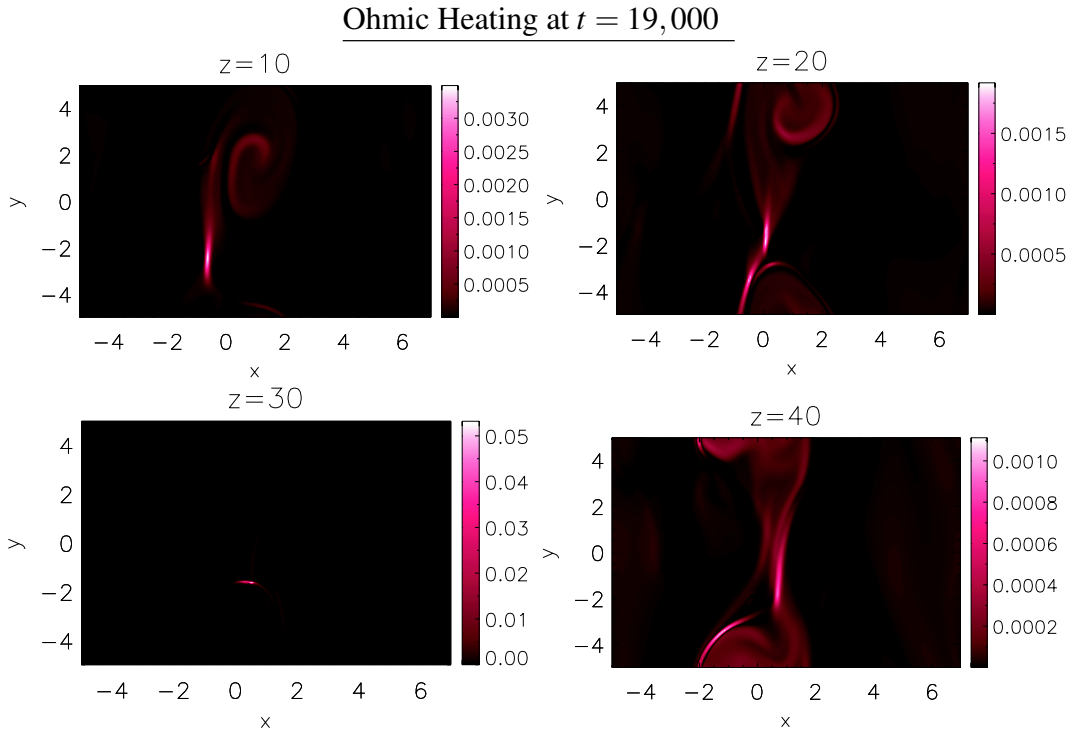

Figure 4.45: Distribution of the Ohmic heating at slices $z = 10, 20, 30, 40$ for run 2A, $t = 19,000$.

release event was well underway. The top plot in Figure 4.49, for $z = 20$, shows both a large vortical region in the upper half of the plot and a picture similar to our cartoon in Figure 4.42 in the lower half. The second plot at $z = 40$ shows the same features but this time the twisted region is in the bottom half, and we also see vortical geometry towards the extremes of the x axis. The topology is no longer so clearly dominated by sheared fieldlines: we see more circular paths forming. At this time, the isosurface plots showed lots of volume filling currents at low levels of $|\mathbf{j}|$ and the field line plots were very messy. The contours in Figures 4.50 and 4.51 show current and heating occurring in the boundaries between oppositely oriented field, but at this stage the reconnection is in full swing and the field clearly has many regions of interest. The strong opposing field is not the main factor. It is just after this time that the spike in kinetic energy and heating occurred.

Lastly, Figures 4.52 and 4.53 at times $t = 20,000$ and $t = 20,200$ suggest that the

$t = 19400$

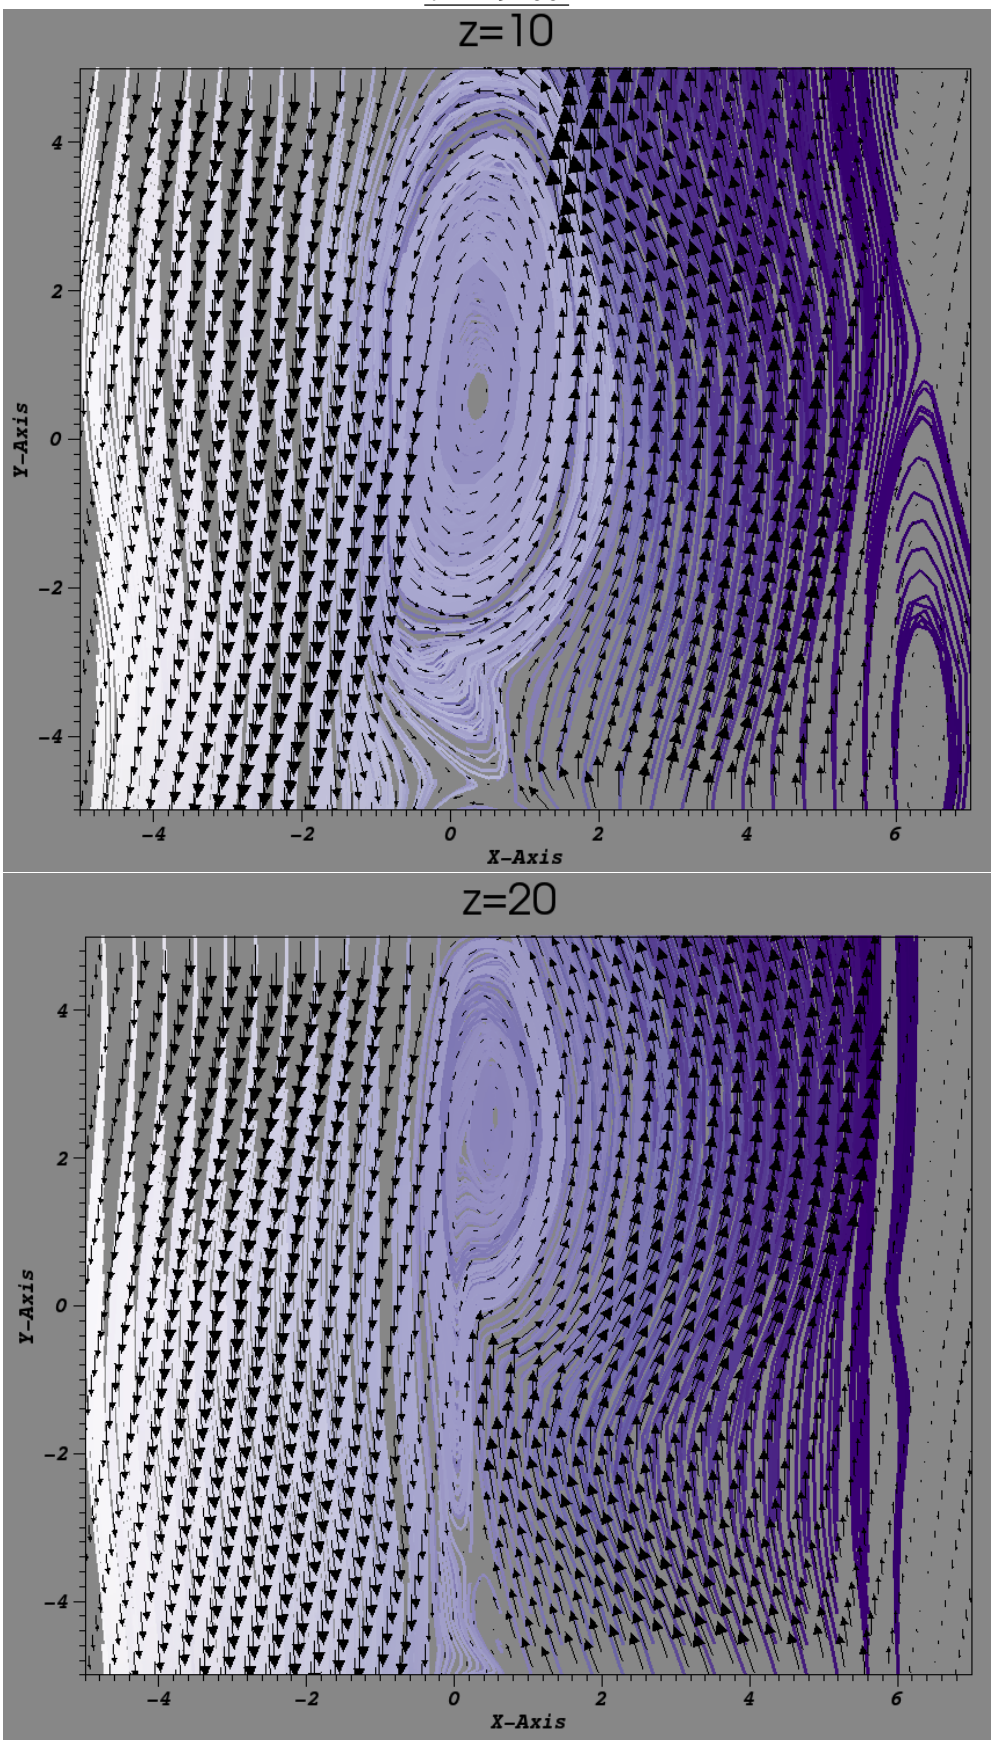

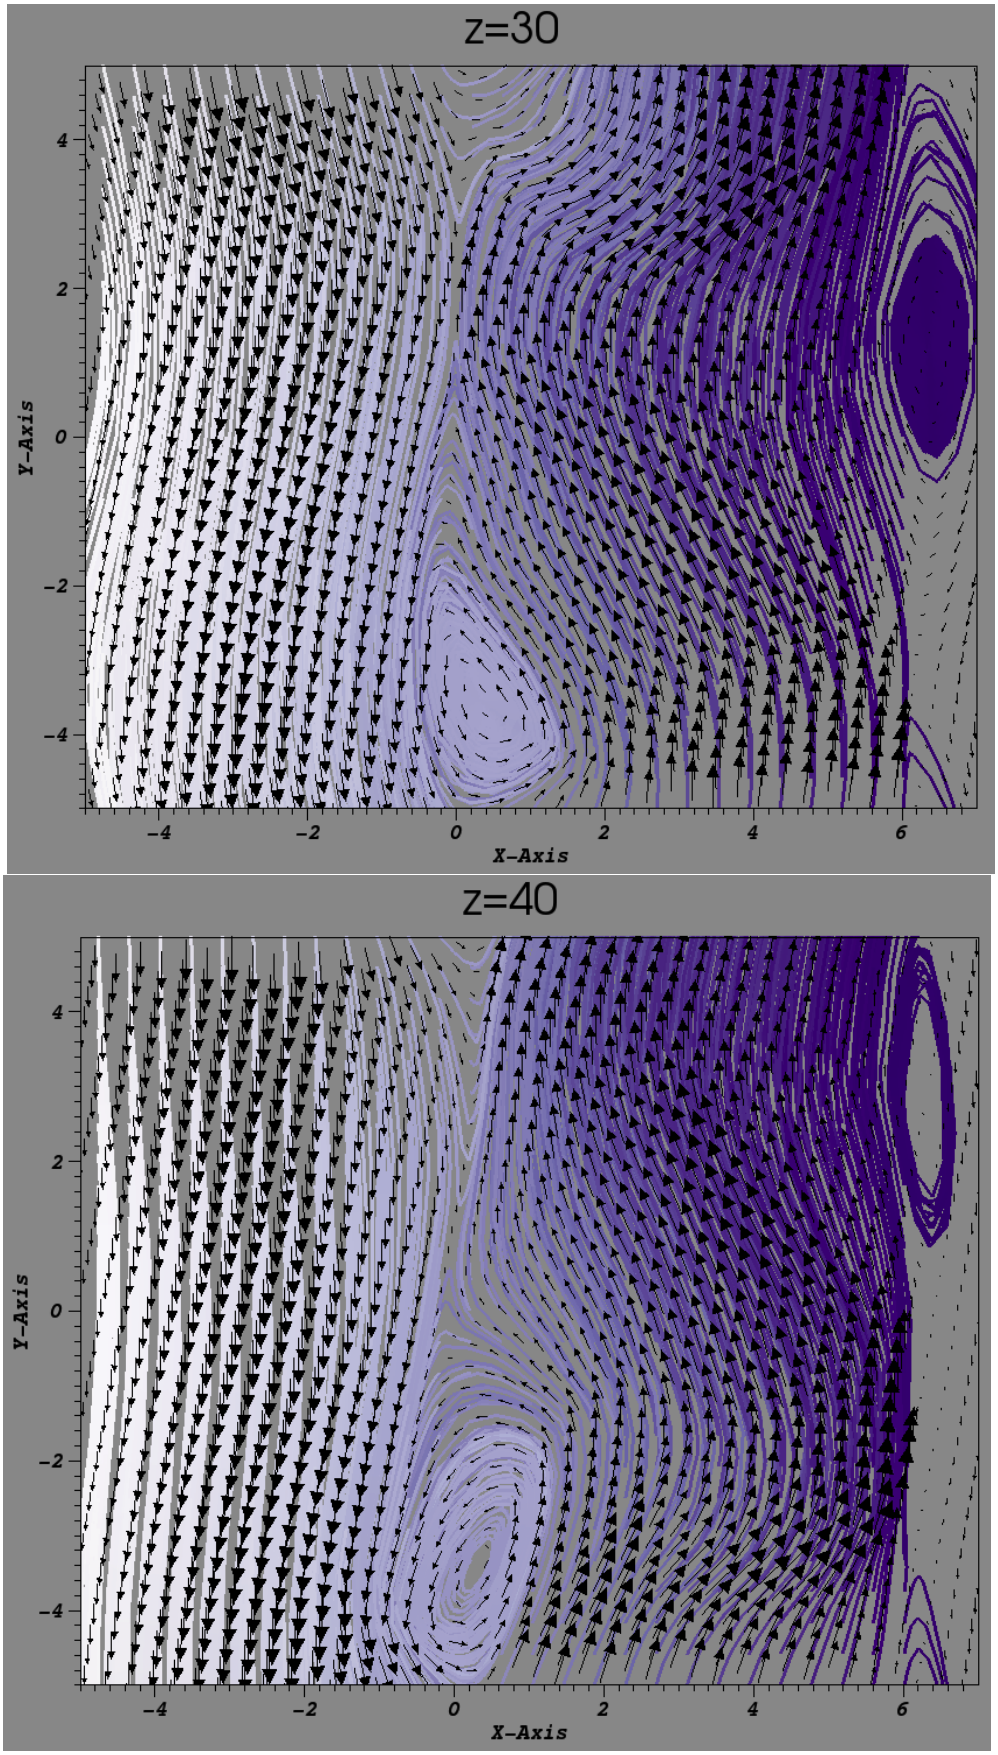

Figure 4.46: 2D fieldlines and vector field illustration for 2A, $t = 19400$.

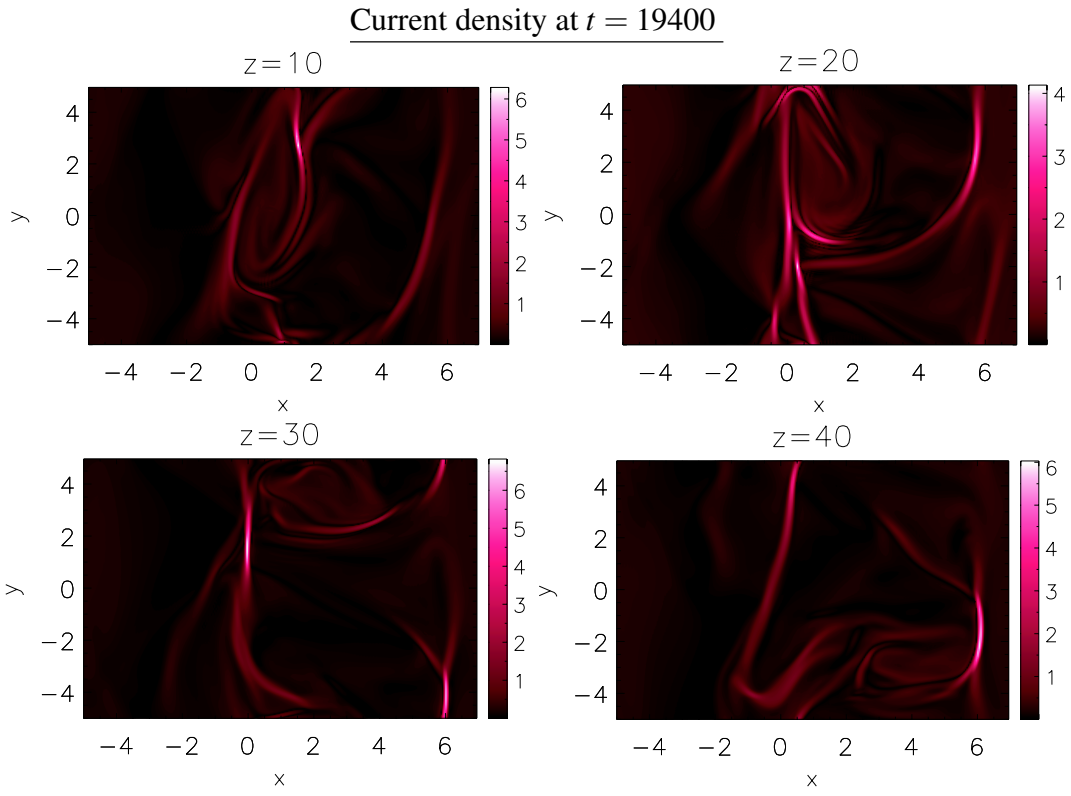

Figure 4.47: Current density in slices across the domain for 2A, $t = 19400$

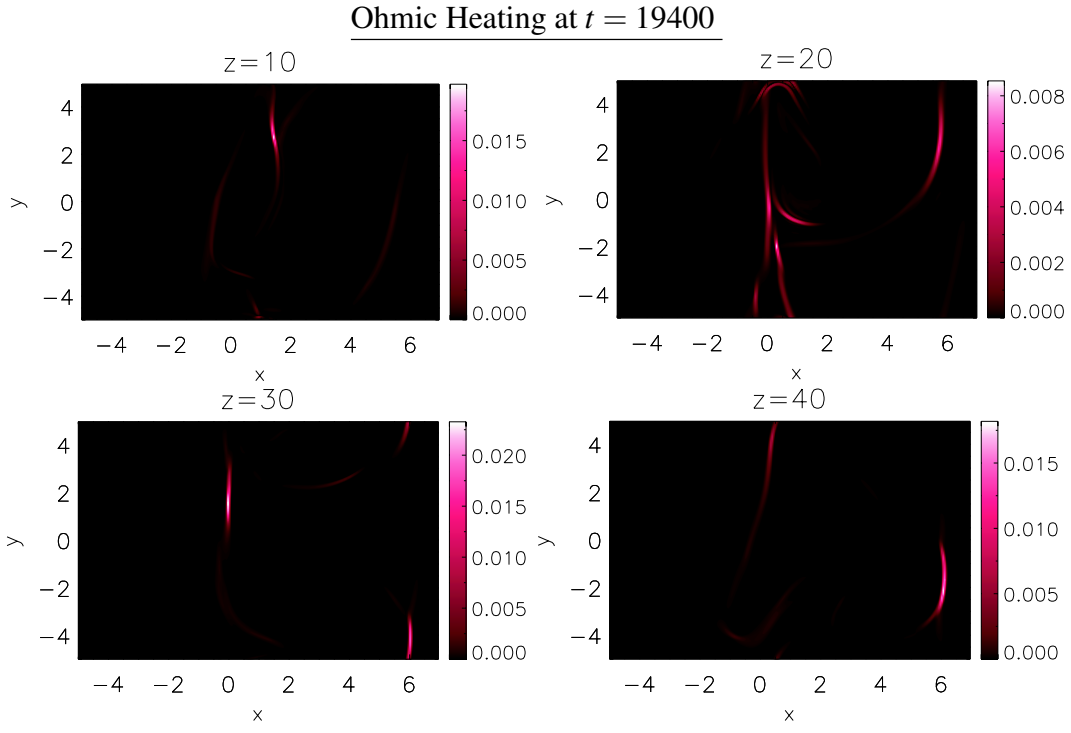

Figure 4.48: Ohmic heating distribution for 2A, $t = 19400$.

$t = 19600$

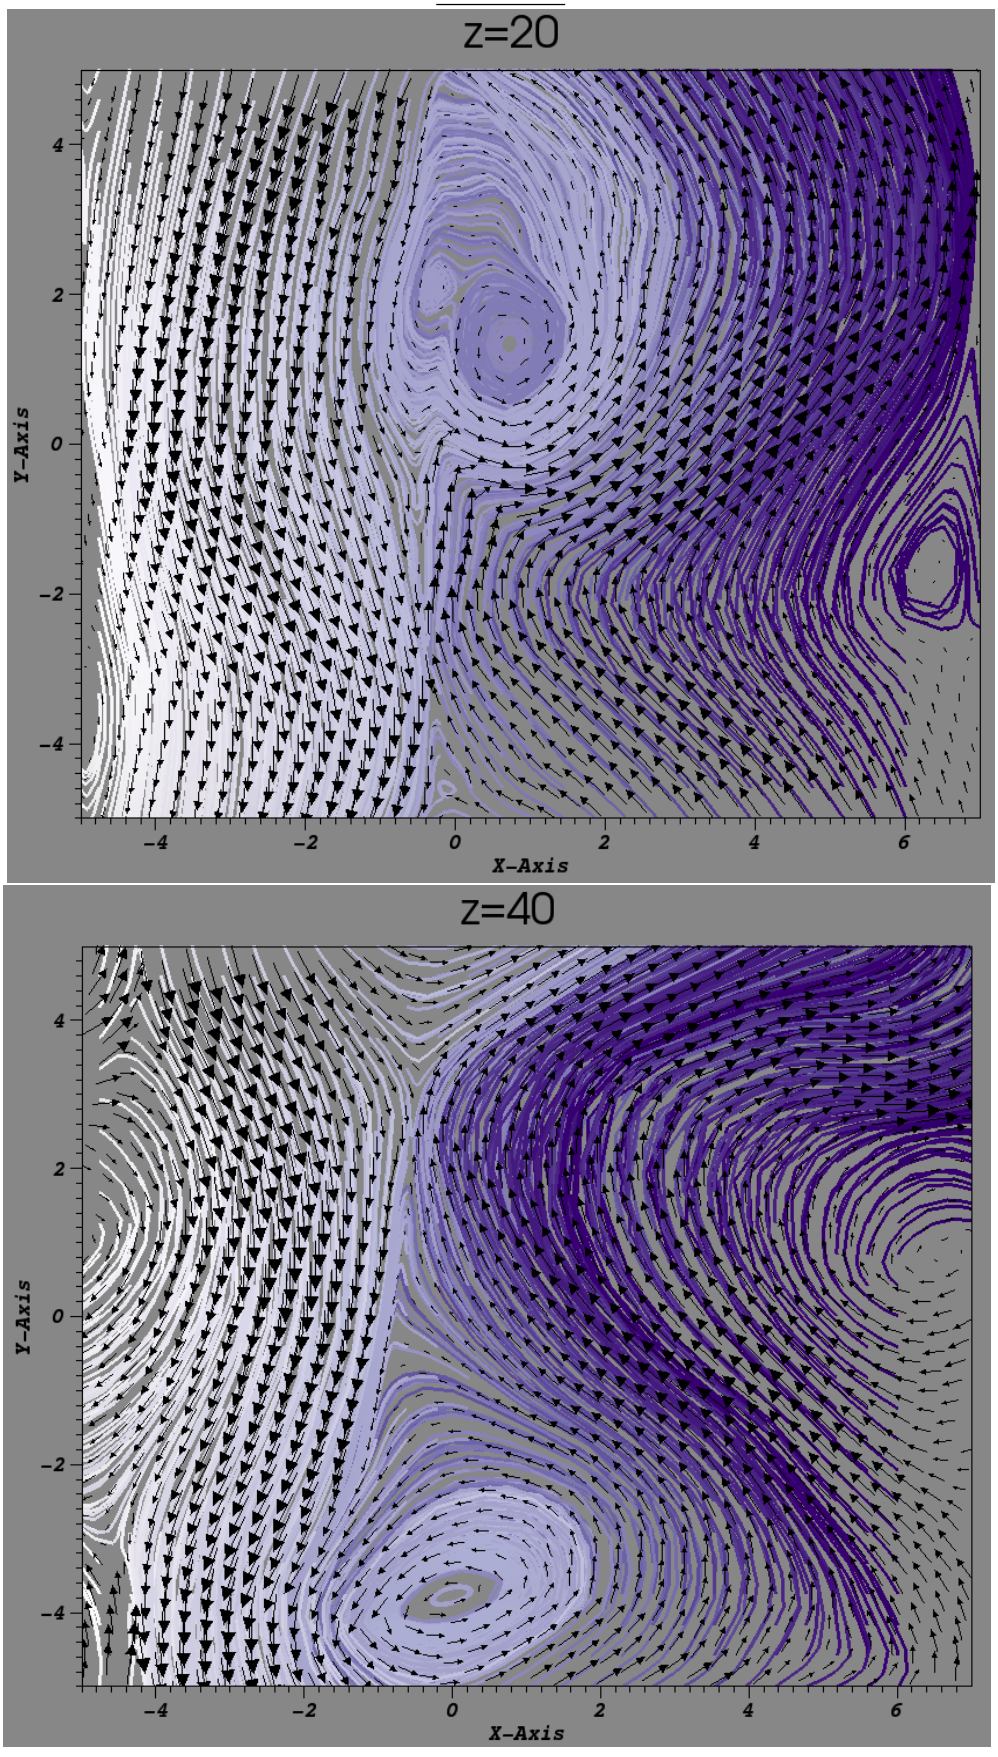

Figure 4.49: Fieldlines and the field direction and strength at $t = 19600$ for run 2A.

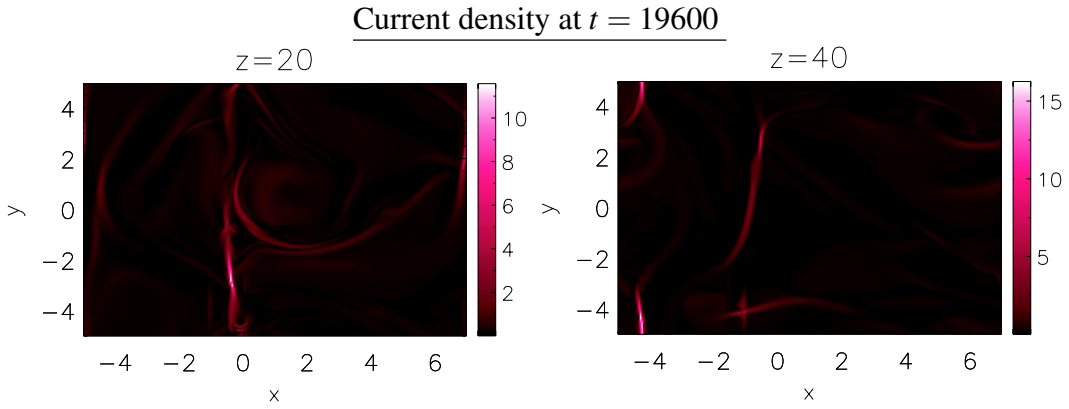

Figure 4.50: Current density layers for run 2A at $t = 19600$.

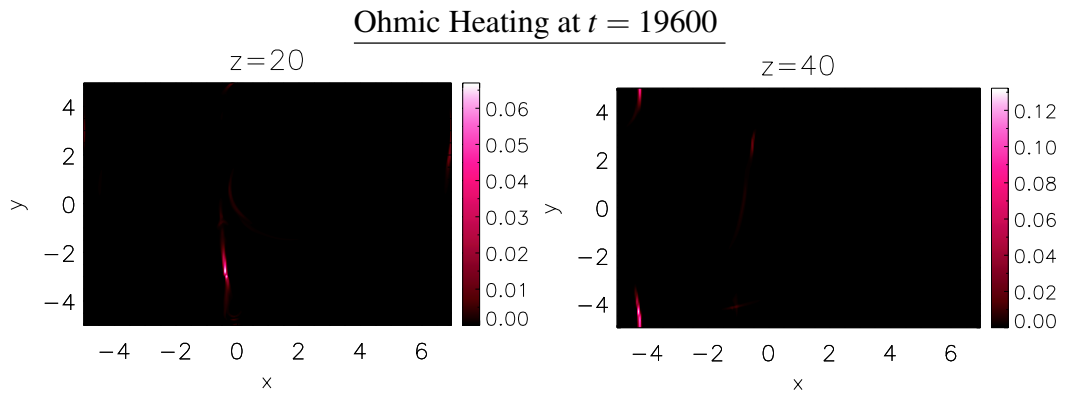

Figure 4.51: Ohmic heating distribution at $t = 19600$ for run 2A.

instability has been resolved. At heights of $z = 20$ and $z = 40$, we do not see the same clear division with the field oriented in opposing directions on either side. The picture is now created by lots of swirling structures. By these times the magnetic energy is ramping up again and the kinetic energy returns to its previous lower levels. The fieldlines did not show the domain split into two distinct regions any longer and display behaviour more akin to the field line slices we saw for Case 1.

Furthermore, Figures 4.54 and 4.55 illustrate these similar properties for two snapshots in Run 2B. Between $t = 33,000$ and $t = 34,000$ this simulation saw its largest drop in magnetic energy, accompanied by spikes in kinetic and internal energy and heating. At heights of $z = 20$ and $z = 40$ at times $t = 33,200$ and $t = 33,400$, we see again that the domain is split into two main regions, with the field pointing downwards on the left hand side of the plots and upwards on the right hand side, with elongated vortical structures lying along the boundary. Visually the plots are dominated by sheared field rather than vortical field.

As mentioned, the two large release events analysed here coincided with spikes in kinetic and internal energy and heating. Both runs 2A and 2B display two events like this. None of the smaller release events in run 2C produced such obvious indicators, but there is still a magnetic energy release event covering over 100 units. Slices at $z = 20$ are shown in Figure 4.56 for three times covering this event. We do see sheared fieldlines of opposing direction, but the division between the two regions is not as pronounced and the boundary layer is not as obvious. While there are regions of oppositely sheared field, there is more of a mix with swirling structures too. This picture is more like that of the plots of run 1C over one of the small drops present for that simulation. Figure 4.57 shows the current and Ohmic heating distribution at $z = 20$ at time $12,600$, and unlike as we saw in, for example, Run 2A, there are no strong currents lying along the boundary between sheared regions. Coupled with the lack of signatures in the other quantities, we surmise that in this case, the opposing shears in

$t = 20,000$

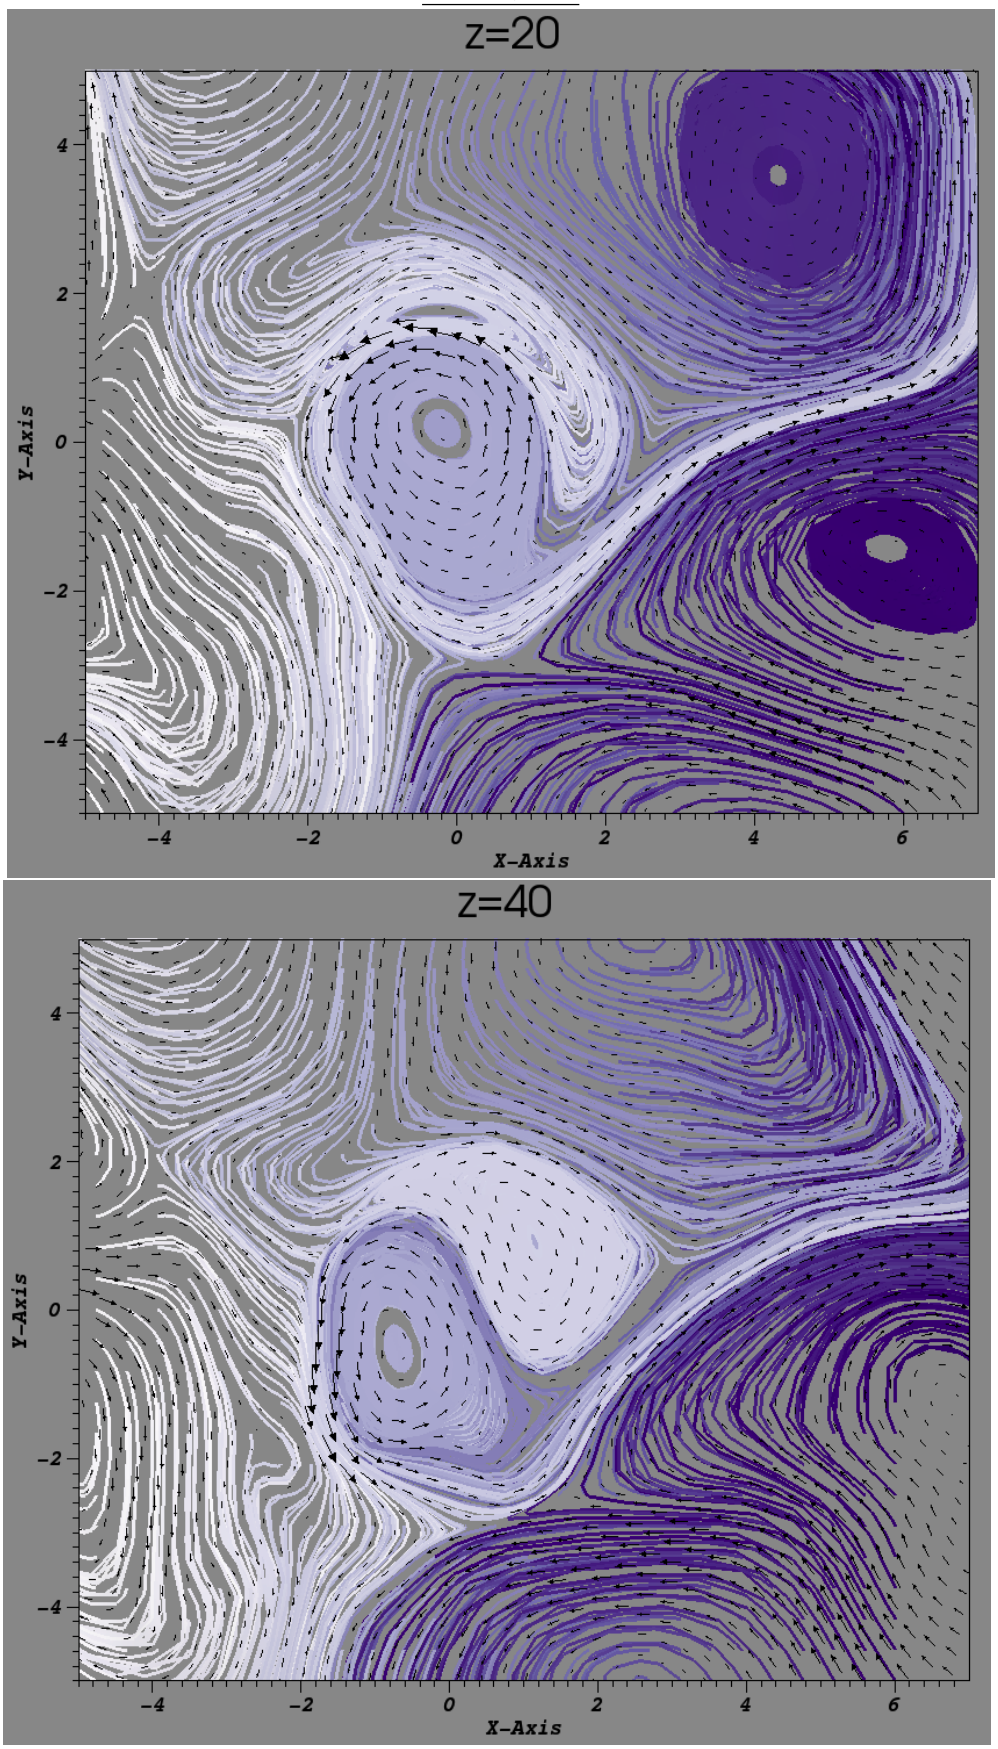

Figure 4.52: Field topology, strength and direction in 2D planes at $t = 20,000$, run 2A..

$t = 20,200$

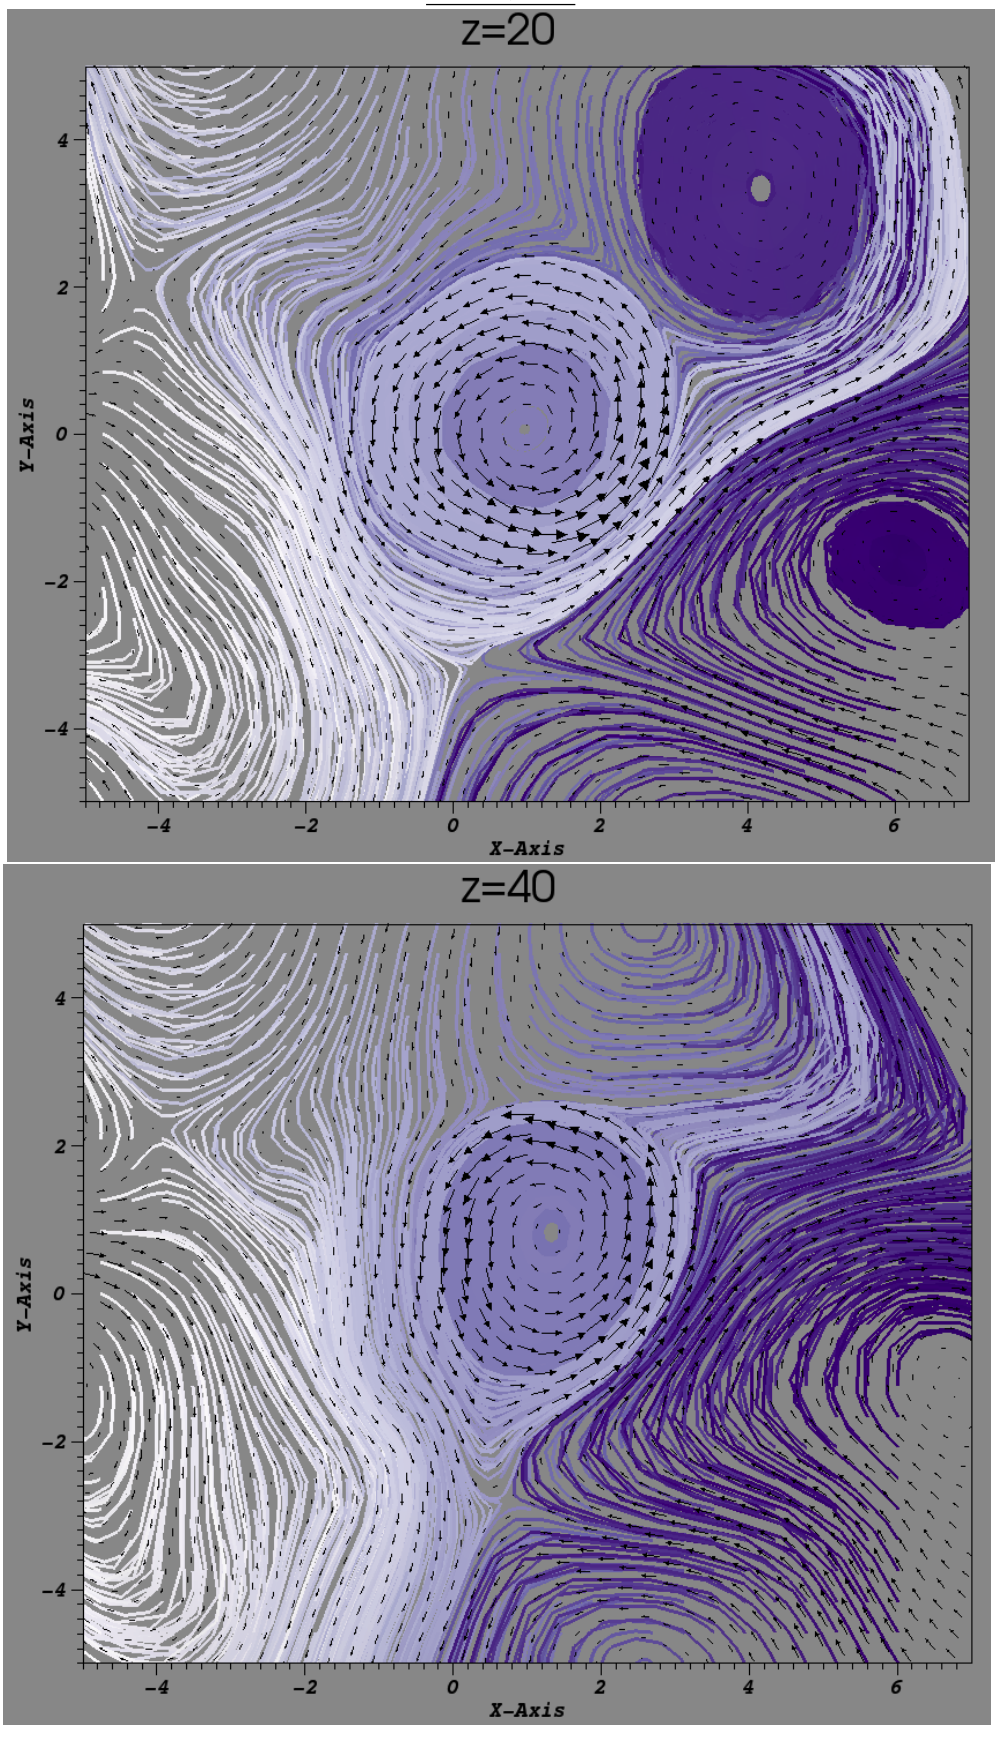

Figure 4.53: Fieldlines and vector field for run 2A at $t = 20200$.

2B, $t = 33,200$

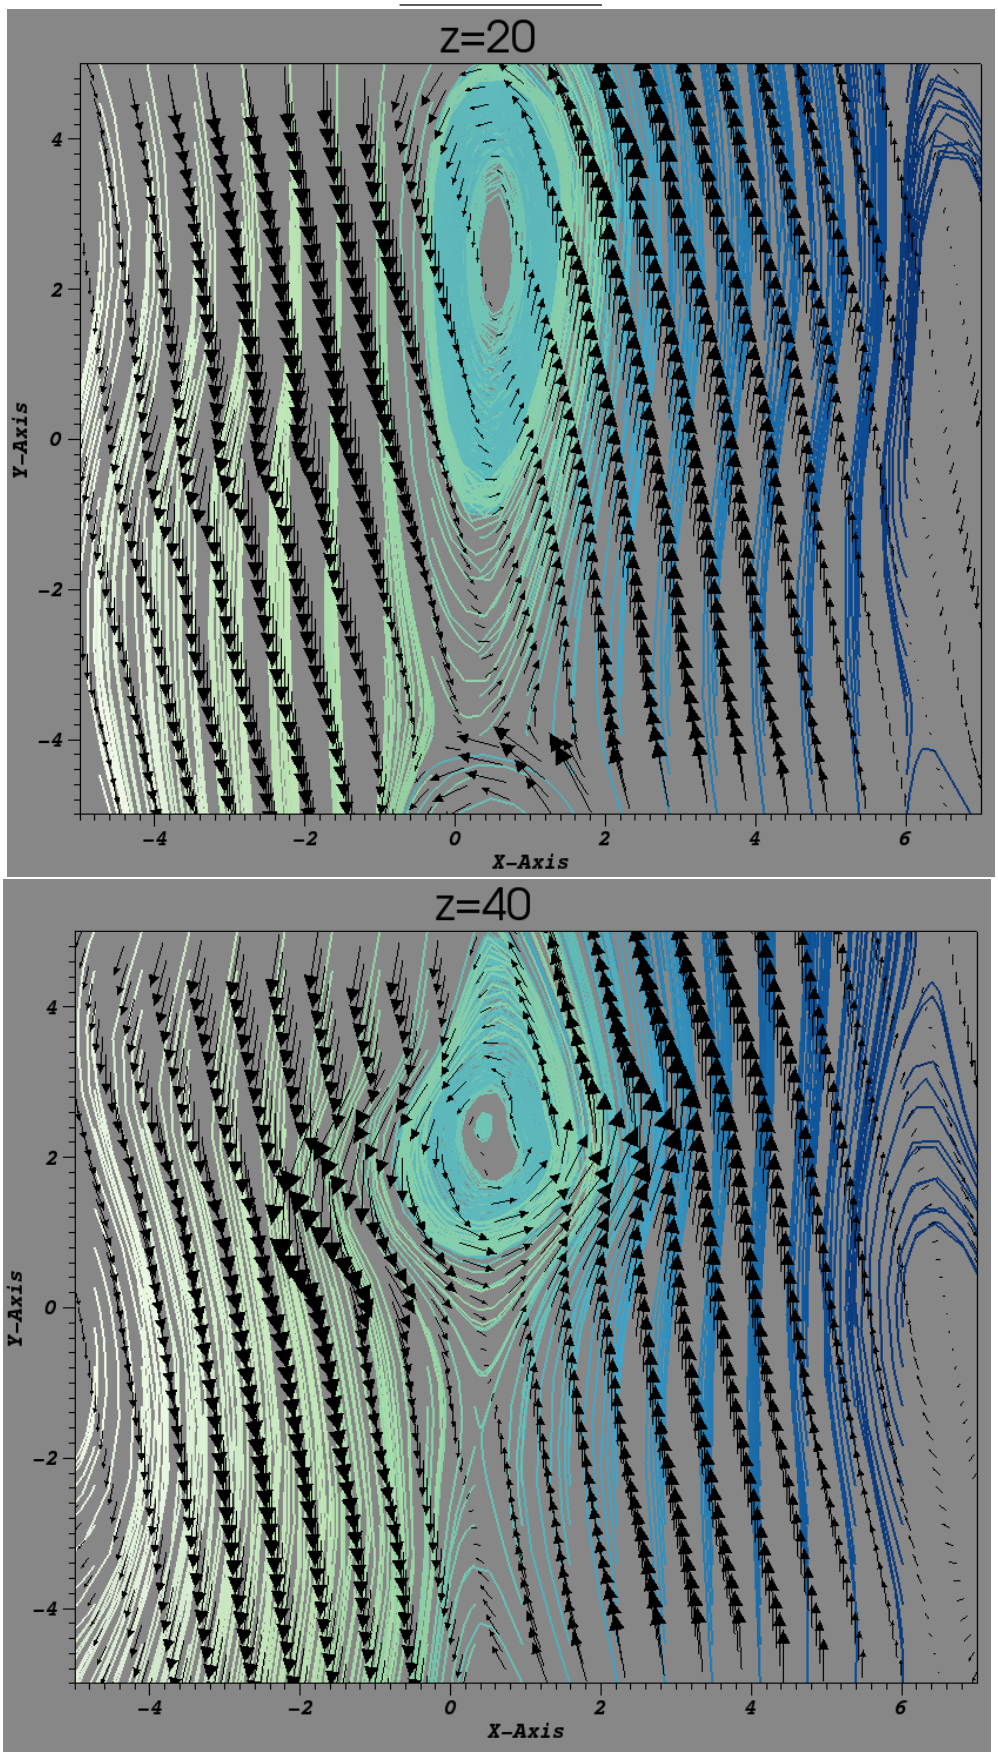

Figure 4.54: Field at $t = 33200$.

2B, $t = 33,400$

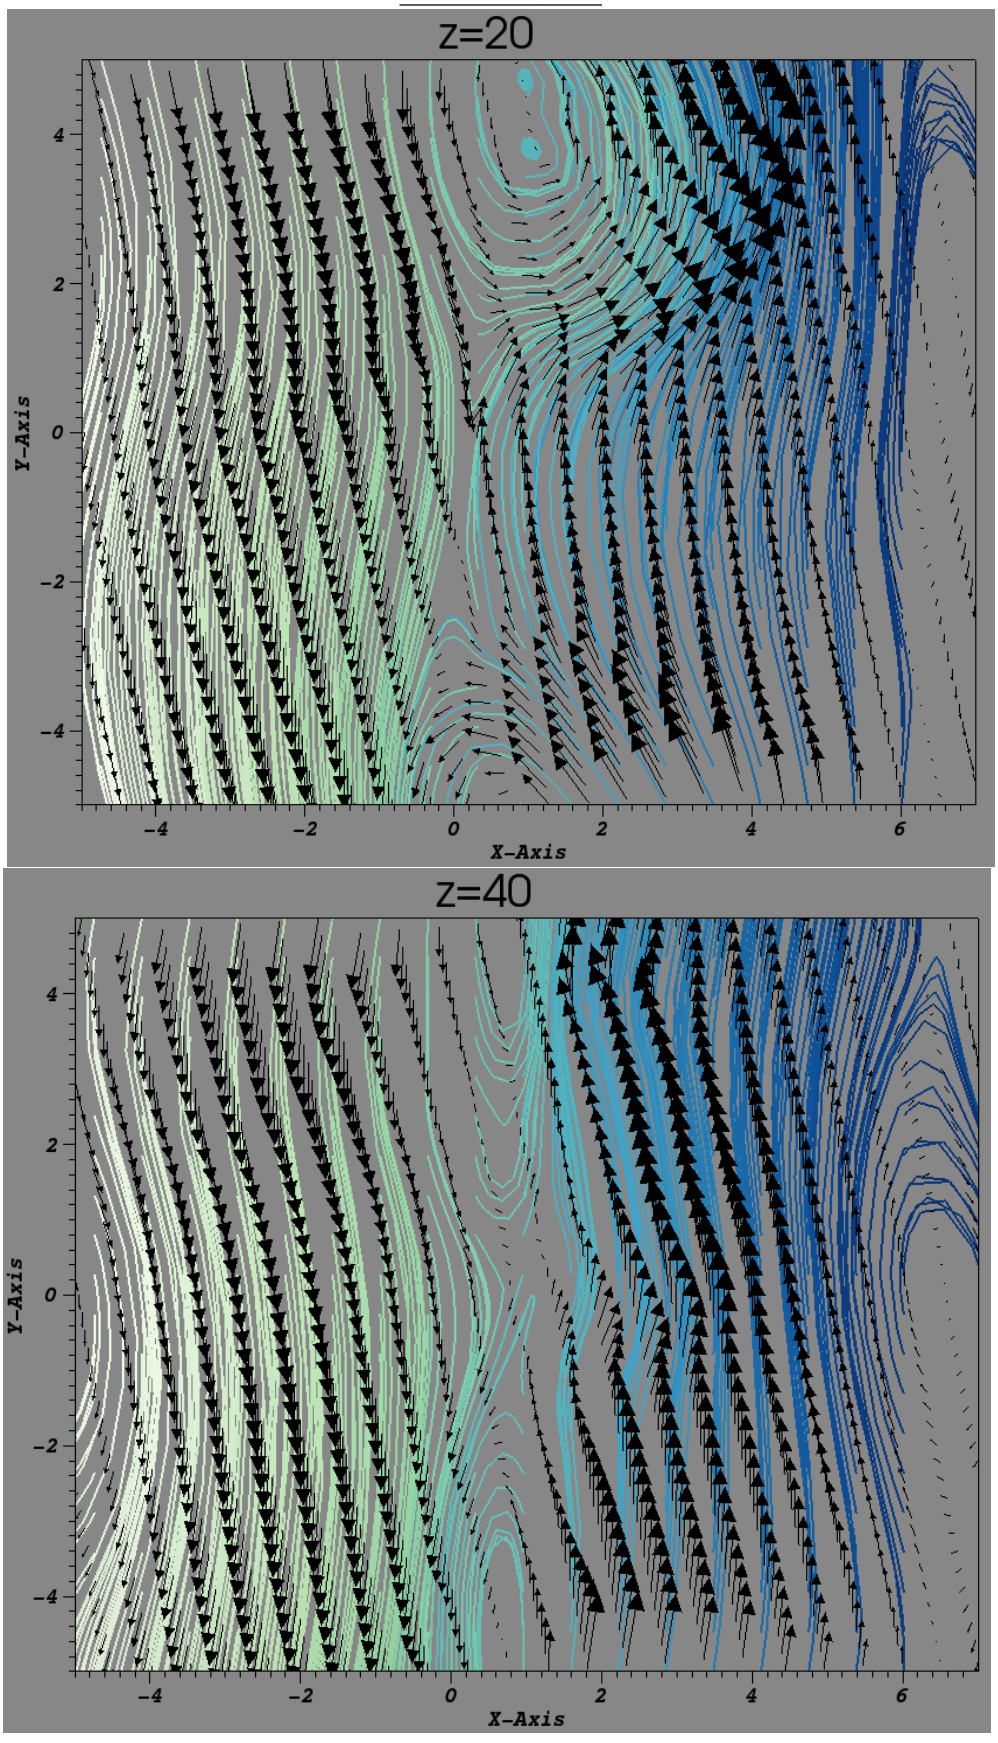

Figure 4.55: Field at $t = 33400$.

the field are not significant enough to result in a significant tearing mode instability, even if some reconnection does take place. It appears that with the vortices further apart, no long, narrow boundary layer is created, so we do not see such intense current layers and there is no large instability triggered.

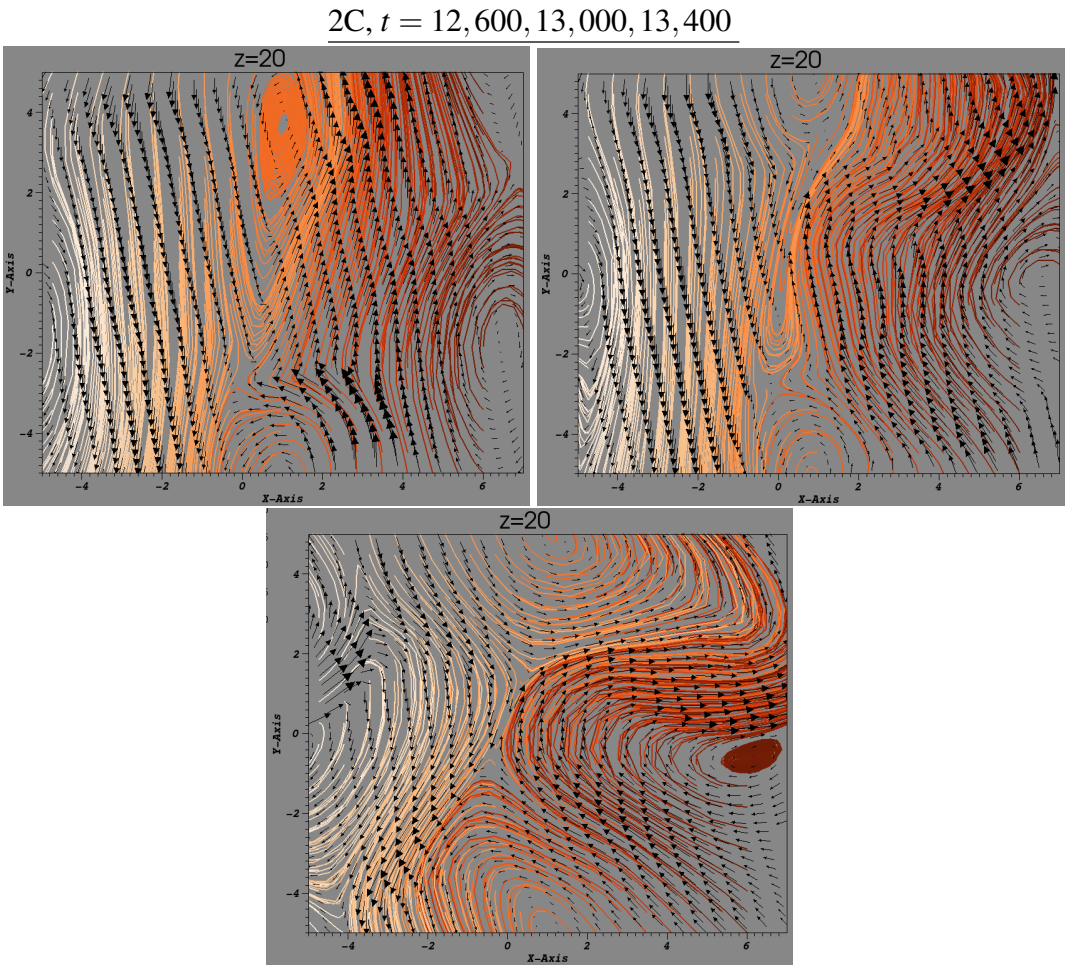

Figure 4.56: Fieldline and vector field slices at $z = 20$ for run 2C, at times covering one of the larger release events in this simulation.

We conclude then that we have a new result: the most coherent cases can trigger multiple tearing mode instabilities. The tangling of the field by the continuous driving causes many small events in all three coherent cases as in the Case 1 runs, but the most coherent are able to induce larger reconnection events by creating clear boundaries between regions of oppositely sheared field. The fact that our driver consists of two

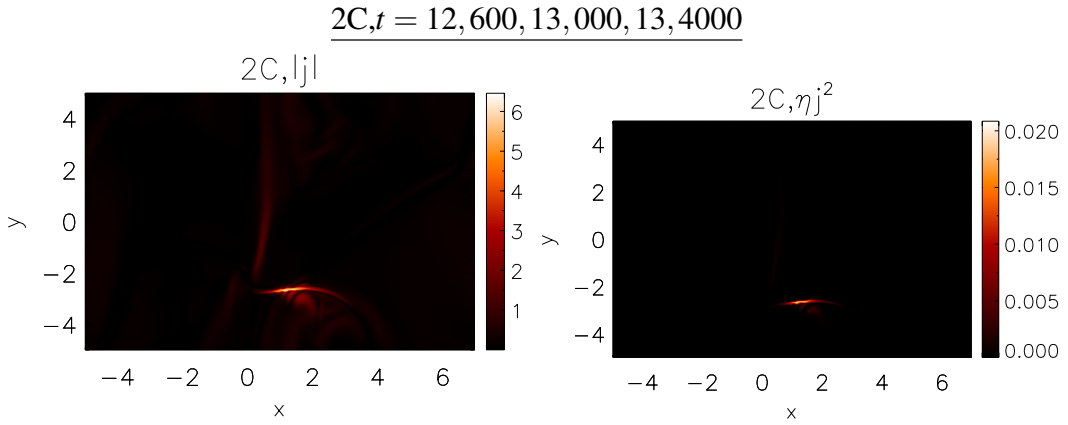

Figure 4.57: Current and Ohmic heating at $z = 20$ for run 2C at $t = 12,600$.

vortices placed at different points prevents the field from twisting uniformly along an axis to create a kink instability, as in papers such as [91]. Instead, the interaction of the vortical motions, while they are placed close enough together, cause distinct regions of oppositely sheared fieldlines, which are unstable to tearing modes and lead to reconnection. When $x_0 = 2$ it seems that the field is not sheared in opposite directions strongly enough and the distance between the vortices is too much to create clear narrow boundaries between the differing shears.

4.1.3 Power Spectrum Analysis

Our analyses to this point indicate that the different drivers can induce different types and levels of heating. We wish to go further now in quantifying the nature of heating and scales on which activity takes place. All simulations show small scale variations in the magnetic and kinetic energies, suggesting small scale reconnection events, which may possibly be categorised as nano or microflares. In addition, we have observed the large reconnection events in the high helicity runs. We wish to gain a sense of the distribution of energy over different scales and how this may differ for different drivers. Further, if we observe a cascade to small scales, i.e. a transfer of energy from the larger scales as we decrease in spatial resolution, we may have a turbulent system.

Turbulence can be a marker of particle acceleration and flaring, so this analysis can provide information about the nature of a system.

In Chapter 1, Section 1.1.1 we discussed the plotting flare energies versus frequency to examine the frequency of flares of different scales. It was concluded by the authors of the work that whichever scale you looked on, the smallest flares happened with larger frequency. We can use the same technique here. We will examine the power spectra of the magnetic and kinetic energy profiles to assess the distribution of energy through different scales and whether turbulence is taking place. A power spectrum considers an energy signal, where that signal has been broken down into its constituent sinusoids, and binned with the frequency k of each, and plotted against these frequencies, hence 'power' spectrum. The line formed over the frequency range may follow some power law, proportional to k^a, where a is called the power law index, and its value can tell us if, for example, the plasma is turbulent. We will think about this in further detail, but first let us discuss the method we use to calculate the power spectra.

4.1.3.1 Fourier Transform

Concept

The method we use to break a signal into sinusoidal elements is the discrete Fourier transform. There are two approaches we could take - we could look at energy variations in time over the full simulation duration, or at how energy density varies over different spatial scales within the domain at set times.

In the context of the energy signal in time, higher frequency components correspond to variations over smaller time scales. So a higher proportion of the signal coming from high frequencies would mean that the shortest lived events are more influential in the overall energy profile, and vice versa. This doesn't however provide any information

about the size of events and how reconnection sites are concentrated. Therefore, we will focus on the spatial variations.

The magnetic or kinetic energy density over the domain at a certain time can be seen as a three dimensional signal, so can be split up into its constituent sinusoids. Now the smallest 'frequencies' represent where the signal varies over the largest length scales and vice versa. We will refer to the inverse length scale as k. So physically, if we find a lot of energy coming from large inverse length scales then the energy in small sub-volumes is significant to the overall energy density in the domain. Then, things like nanoflares and reconnection over small scales may be more influential to heating.

Turbulence and Power Law Index

Turbulence is fluid phenomenon which occurs throughout the universe and has inspired a vast library of research. An overview of MHD turbulence can be found in [10], but the key point we are concerned with is this: a turbulent MHD system is a magnetised plasma containing chaotic flows. A turbulent plasma is full of swirling eddies and vortices creating complex, chaotic structure across a range of scales. Large-scale eddies induce smaller eddies, which in turn can induce even smaller swirls, and so on, transferring energy down the spatial scale - this is referred to as a turbulent energy cascade. Turbulent flows can be triggered by events such as reconnection, since reconnection can induce particle acceleration and complex plasma flows. Some of the largest and most powerful reconnection events are solar flares, with authors such as those in [32] using the concept of turbulence to help explain the mechanisms of a flare.

Power spectra can help us to identify turbulence if it is present. Some of the foundations of this area of research were laid in a famous paper by Kolmogorov, [50]. This work states that power spectra often show three distinct regions as we move across the spectrum of k. At the small k end of the scale energy variations are occurring on longer length scales, perhaps induced by external influences - for example our drivers

input kinetic energy, triggering motions on a relatively large scale within the domain. At very high k, variations are on very small scales and so the energy can be dissipated as heat via viscous and resistive effects. In between these two extremes is the inertial range - where the flow is dictated not by some driving force or viscous and resistive effects. It is within this inertial range where one can observe a turbulent cascade, if it exists. As mentioned briefly, the energy can be described as being proportional to k to some power a, so $E(k) \propto k^a$. The exact value of this power law index which indicates turbulence is debated and depends on the properties of the system being investigated. In paper [50], the author proposed that in a three dimensional (non-magnetised) turbulent fluid $a = -5/3$ - subsequently this is referred to as Kolmogorov turbulence. Later, Iroshnikov ([48]) and Kraichnan ([52]) considered MHD turbulence and calculated $a = -3/2$, hence this is referred to as Irosnikov-Kraichnan turbulence. However, both of these scenarios assumed turbulence occurred isotropically. Subsequent studies have shown that turbulent cascades occur anisotropically, in planes perpendicular to the field, when considering a plasma under the influence of a strong magnetic field in one dominant direction (for example [99]). Interestingly, however, the values of $-3/2$ and $-5/3$ still crop up in analysis of two dimensional planes in these more physically realistic experiments (the authors of [99] went on to obtain $-5/3$ in [39]), leading to much debate on the true value of the turbulent power law index. In the work of [58], the authors conducted a series of tests which involved drivers of various speeds acting on the velocity and/or the magnetic fields. They concluded that the power law index was $-3/2$, appearing in 2D planes, unless there was relatively fast driving of the velocity field only, when $a = -5/3$. Recall that we designed our drivers to be slower than the Alfvén speed in the initial field, so it will be interesting to discover if either of these indexes apply to any of our simulations.

Formulae and Fourier Transform

Before proceeding, let us make a quick note on the following notation and terminology.

Since we will be conducting spatial Fourier transforms, we will be breaking up functions into sinusoids with different lengthscales, rather than frequencies. The Fourier transformed energies will be integrated over k space, where k is an inverse length scale, but not technically wave number, since there is no factor 2π in the calculation of these values. The Fourier transforms themselves are calculated independently of spatial scale, in terms of κ, which is related to k by $k = \kappa/\Delta$ where Δ is the total length in the relevant direction. The Fourier constituents are then allocated to their relevant inverse length scale k in the integration.

The Transform: Taking a component $B_x(\mathbf{x})$, the three dimensional discrete Fourier transform $B_x(\kappa)$ is given by;

$$B_x(\kappa) = \sum_{j=0}^{N_1-1} \sum_{l=0}^{N_2-1} \sum_{m=0}^{N_3-1} B_x(\mathbf{x}) \exp^{-2\pi i(x_j\kappa_1/N_1 + y_l\kappa_2/N_2 + z_m\kappa_3/N_3)}, \tag{4.1}$$

where N_1, N_2 and N_3 are the number of elements in the x, y and z directions and $(\kappa_1, \kappa_2, \kappa_3)$ are the integer indices of each constituent sinusoid. Since the magnetic energy density is $|B|^2/2$, we simply calculate 4.1 for the B_y and B_z components also and calculate;

$$\frac{\mathbf{B}(\kappa)\cdot\hat{\mathbf{B}}(\kappa)}{2} = \frac{B_x(\kappa)\cdot\hat{B}_x(\kappa) + B_y(\kappa)\cdot\hat{B}_y(\kappa) + B_z(\kappa)\cdot\hat{B}_z(\kappa)}{2},$$

where hat variables represent the complex conjugates. Note that we are transforming each component individually and then calculating the energy in k-space, rather than transforming the magnetic energy directly. This is the approach taken in much of the literature, for example the method outlined in [15] performs the process in this way. If we transform the magnetic energy directly, then we simply obtain the transformed

energy. We want to look at how energy is stored at each spatial scale, which we obtain from transforming each component first.

We can also look at planar components of magnetic and kinetic energy and carry out two dimensional Fourier transforms in light of the research on how in magnetic environments with a strong guide field, for example one in which the vertical field is much stronger than the horizontal field, turbulent cascades occur in 2 planes perpendicular to the direction of the guide field. So, rather than calculating $E(\kappa)$ where $\kappa = (\kappa_1, \kappa_2, \kappa_3)$, we may calculate $E(\kappa_\perp)$ where $\kappa_\perp = (\kappa_1, \kappa_2)$, using 2D Fourier transforms. The volume integral over k-space then becomes a surface integral. In our case, the initial field is $(0,0,1)$ and checks generally show B_z larger than B_x and B_y through the simulation runtimes, so we do have a stronger vertical field although it is not significantly dominant. We will therefore investigate 3D and 2D transforms to gather as much information as possible. We use a Python function to compute the Fourier transforms of each field component and their complex conjugates. The k domain set-up and integration step requires further explanation.

K-space and Integration Over K-Shells

In the three dimensional scenario, we construct the three dimensional inverse length scale domain (k_1, k_2, k_3) using Python programs. These functions take the x, y and z Cartesian domains and calculate the corresponding k_1, k_2, k_3 in k-space. The x, y and z arrays have different interval spacing but the same size, so each k dimension also has the same size - any point $[x, y, z]$ in Cartesian space has a corresponding k-space vector $[k_1, k_2, k_3]$. We then create the three dimensional mesh in k-space using these vectors.

The next step is to sum the power in each 'shell'. These shells, or k-shells, can be thought of as spheres centred at the k-space origin, with radii of $i \cdot \mathrm{d}k$, where i is the number of the k-shell and $\mathrm{d}k$ is the radial distance between the boundaries of k-shells. For example - the zeroth shell has a radius of 0, while the first shell has radius $\mathrm{d}k$, and

so on. If the energy density has a value a at a point in k-space $[k_1, k_2, k_3]$, then this point lies at a magnitude from the k-space origin of $\sqrt{k_1^2 + k_2^2 + k_3^2}$. This point must sit within one of the k-shells, so its value a contributes to energy contained in that shell. Therefore, the only energy densities in k-space which contribute to the signal in the zeroth shell, for example, is that defined at the origin, $E(\mathbf{0})$. The first k-shell has a radius of dk and spans a radial distance between 0 and dk. The contribution to this shell will be the sum of the energy density at those points which lie at a radius r where $0 < r \leq \mathrm{d}k$. The next k-shell has a radius of 2dk and spans the region lying between a magnitude of dk and 2dk. So the energy in this shell is the sum of the energy densities defined at points which lie at a radius between dk and 2dk. This continues until the energy densities at all k-space points have been summed. We obtain an array of the total energy lying in each shell, where a shell corresponds to a volume in k-space. We plot this against the radius of the k-shell to obtain our power spectra.

In the two dimensional case, k-space is constructed from the perpendicular k vectors and the shell integration becomes a ring integration. We sum up the value of points at distances $r = \sqrt{k_1^2 + k_2^2}$ lying between circular regions instead. Figure 4.58 shows a schematic of the ring integration. The dots represent the k-space mesh, with some k-rings overlaid onto this mesh to illustrate how the domain is sectioned off. The orange double arrows represent the radial distance dk, the magnitude chosen as the separation between shells. The crosses are coloured to match the ring in which they lie and contribute their values to.

Choice of dk

We have some freedom to choose a value for dk, the radial distance between shells or rings. A sensible starting point is to take the minimum positive value from the k-arrays, k_i, thereby calculating the energy on the largest possible range of scales. Figure 4.59 (a) shows an example of a 3D dimensional power spectrum for the magnetic energy

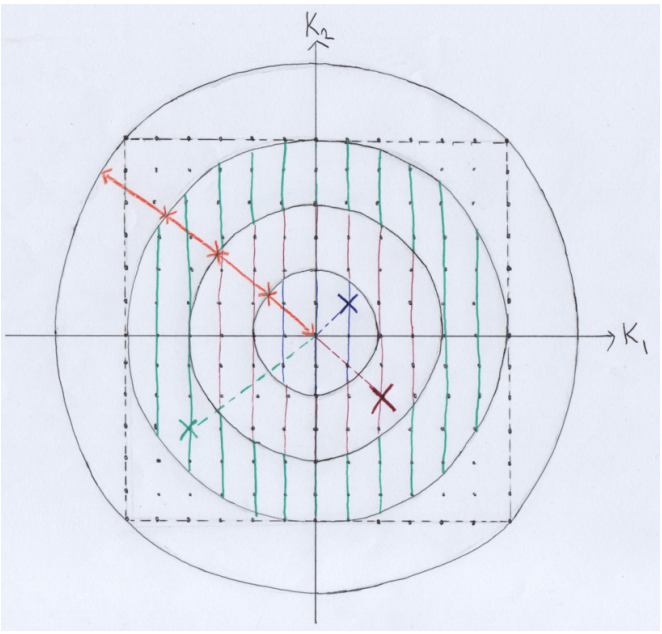

Figure 4.58: Diagram of k-shell mesh in 2D. The dots show how the inverse length scales translate to points defined in a grid. Each point has some energy density attached to it and a radius r from the origin which places it inside one of the correspondingly coloured rings. The orange double arrows highlight the distance dk marking the boundaries between rings. The concept is the same for the 3D case, but with spherical shells.

for run 2A at time 19400, where dk was taken to be the minimum of the minimum positive values of each of the three arrays. We observe many oscillations across the range of k, giving the profile a very jagged appearance. Figure 4.59 (b) shows the power spectrum of the same scenario but taking dk to be the maximum of the minima of the three arrays. This is much smoother, but still contains small spikes whose origin we must consider.

We can explain this by looking at slices through the original field and Fourier transforms, as in Figure 4.60. In plot (a) there appear to be small wave like features in the image of the real component of the Fourier transform of B_x, with regions of larger magnitude alternating with regions of smaller magnitude within small areas: this can account for the jagged spectrum, as there must be instances of summing a relatively large energy value in one shell and then a relatively small value in the next shell, and so

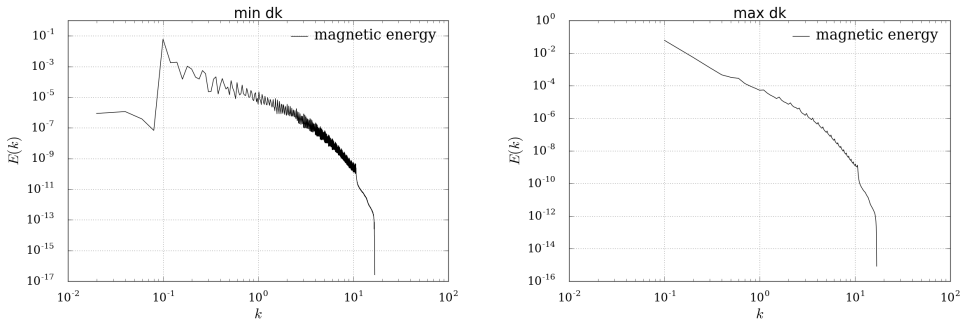

Figure 4.59: Three dimensional power spectra example with different values of dk. Plot (a) was calculated using the minimum of the minimum values of each of the k-arrays . Plot (b) used the maximum of the minimum values. We can see that summing over a slightly larger range at a time gives a smoother line.

on. This feature appears to be due to the nature of the original field. In Figure 4.60(b) a slice through B_x reveals structure in the field, where the gradient moving across the area changes quite significantly and quickly. Moving across the domain, regions of larger values switch quickly and dramatically to lower values. This results in ringing, or the so-called Gibbs phenomenon - an artefact of signals containing discontinuities or large gradients, whereby their Fourier transforms contain extra oscillations. We can confirm this by carrying out the same transform on a domain containing a spherical boundary with value two inside the sphere and one elsewhere - a type of step function given by $f = 2$ for $x^2 + y^2 + z^2 < 1$ and $f = 1$ otherwise. Figure 4.61 illustrates the same phenomenon occurring for this spherical test - the function value jumps at the boundary of the sphere, leading to waves appearing in the Fourier transform.

As we proceed we choose to take dk to be the maximum of the three k-array minima, to obtain smoother power spectra and a simpler picture of the resulting power law indexes.

The Spectra

We first examine three dimensional spectra at certain points, taking into account the

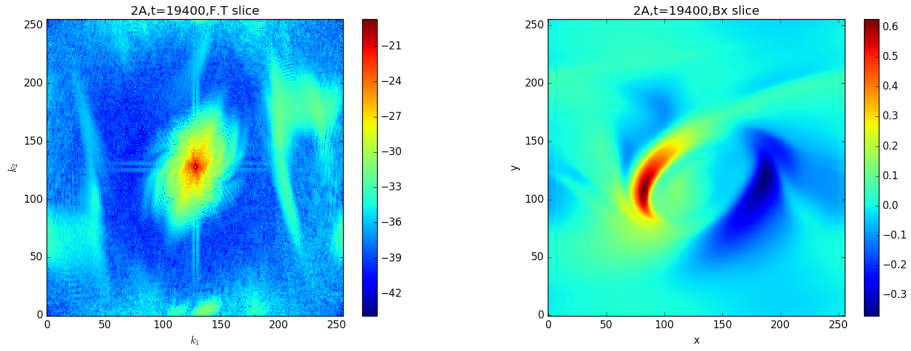

Figure 4.60: Plot (a) shows how the jumps in the power spectra result - small waves in the value of the quantity leads to shell energy values alternating rapidly between large and small values. The real space image slice in (b) shows the structure of B_x in the xy-plane. This has clear boundaries between regions of different values, which are to blame for the ringing in the Fourier transform.

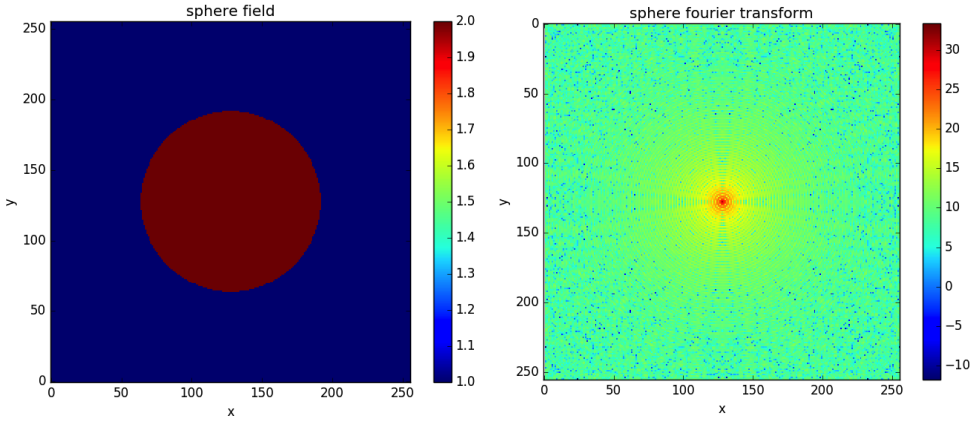

Figure 4.61: We see again the ringing effect in the Fourier transform of the sphere function. The discontinuity across the boundary of the sphere again leads to waves in the Fourier transform. Rings of alternating values propagate outwards across frequency space.

magnetic energy, kinetic energy, and magnetic plus kinetic energy. For Case 1 we start by looking at some of the data snapshots in the statistically steady state section, while for Case 2 we first look at the energy distribution at times before, during and after large drops in magnetic energy. Figures 4.62 to 4.67 display red solid lines representing the magnetic energy, blue solid lines the kinetic energy, and black solid lines the magnetic plus kinetic energy. We also plot $c_1 k^{-5/3}$ and $c_2 k^{-3/2}$ in green dotted and purple dotted lines respectively, in order to quickly compare the gradient of each power law

and identify potential turbulence. The constants c_1 and c_2 are chosen simply to position these plots at levels helpful to the comparison.

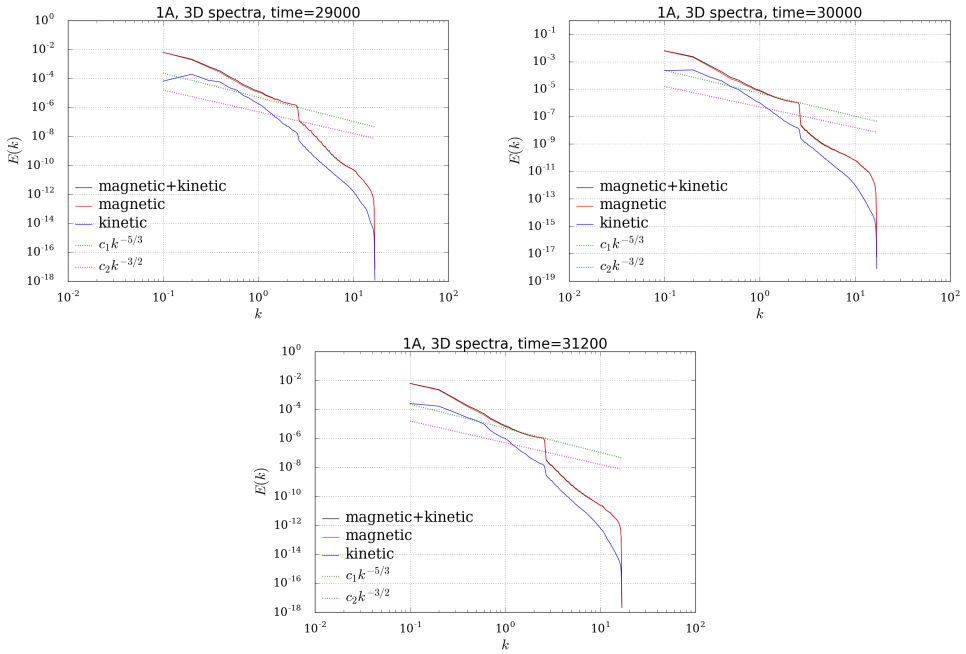

Figure 4.62: 3D spectra for run 1A, ($x_0 = 0.5$, opposite/zero helicity), at times $t = 29,000, 30,000, 31,200$.

One feature jumps out and is common to all runs: none of the runs exhibit slopes close to the Kolmogorov or Iroshnikov-Straichnan slopes. A few of the snapshots, for example Plot (a) in Figure 4.62, show a small range of k values in the inertial range (from around $k = 1$ to $k = 2.5$) where the slope is not too far off, however even then this brief correspondence is interrupted by the other main feature which is common to most of the snapshots - a sudden and large drop, generally in both the magnetic and kinetic energy, and always at a certain k-shell radius, $k = 2.5$.

Snapshots which either have a very minimal or no characteristic drop at this value of k are seen for runs 1C and 2A and B. These are the simulations with the largest energy values and largest energy releases. However generally there is at least a small drop at $k = 2.5$, regardless of the driver - the point of change is always the same. This suggests that driver characteristics are not responsible for the existence of the drop, but do affect

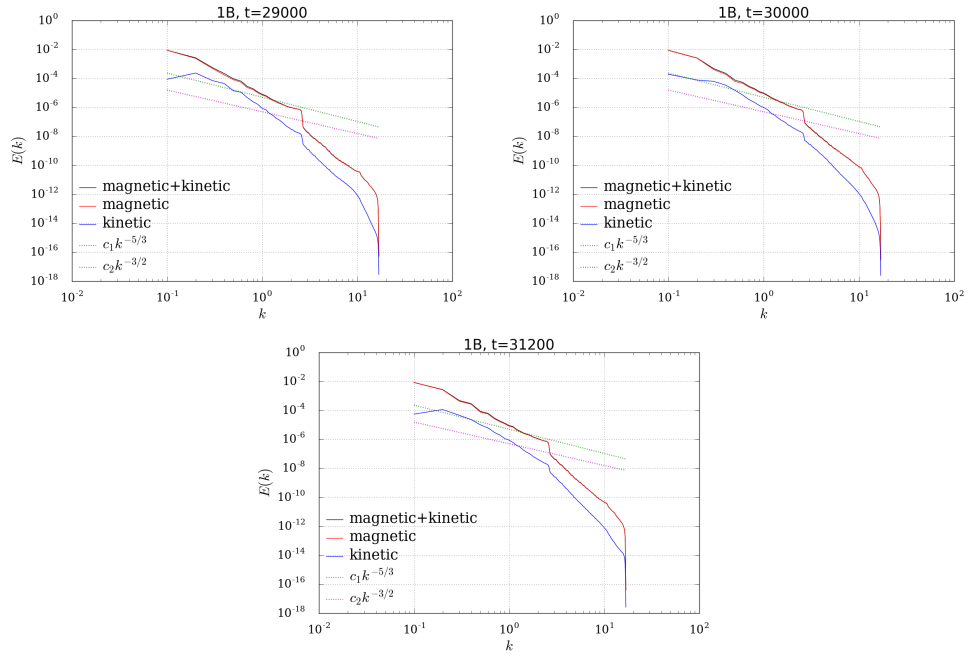

Figure 4.63: 3D spectra for run 1B, ($x_0 = 1.0$, opposite/zero helicity), at times $t = 29,000, 30,000, 31,200$.

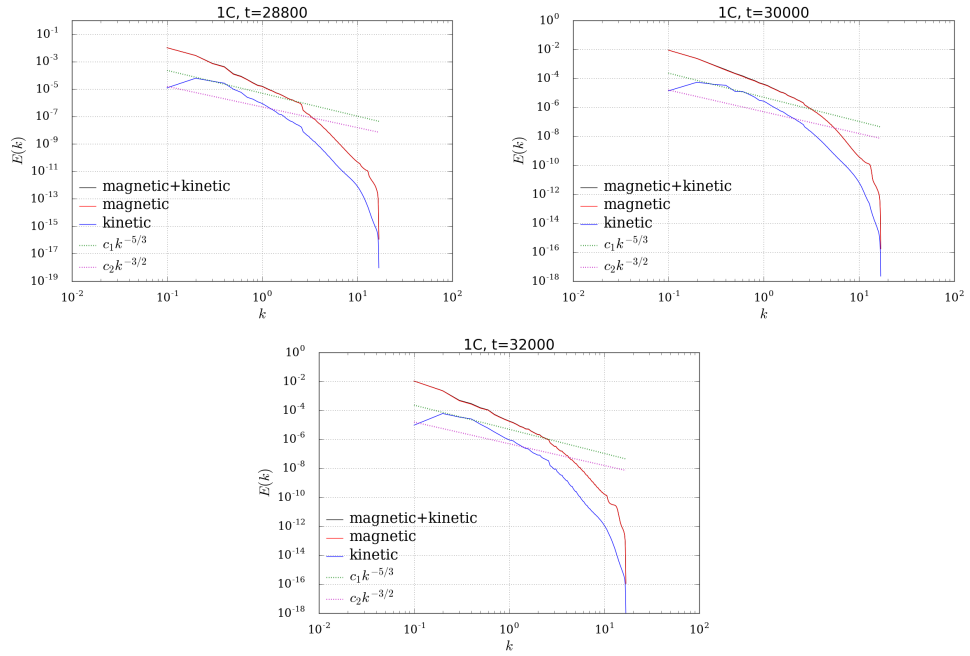

Figure 4.64: 3D spectra for run 1C, ($x_0 = 2.0$, opposite/zero helicity), at times $t = 28,800, 30,000, 32,000$.

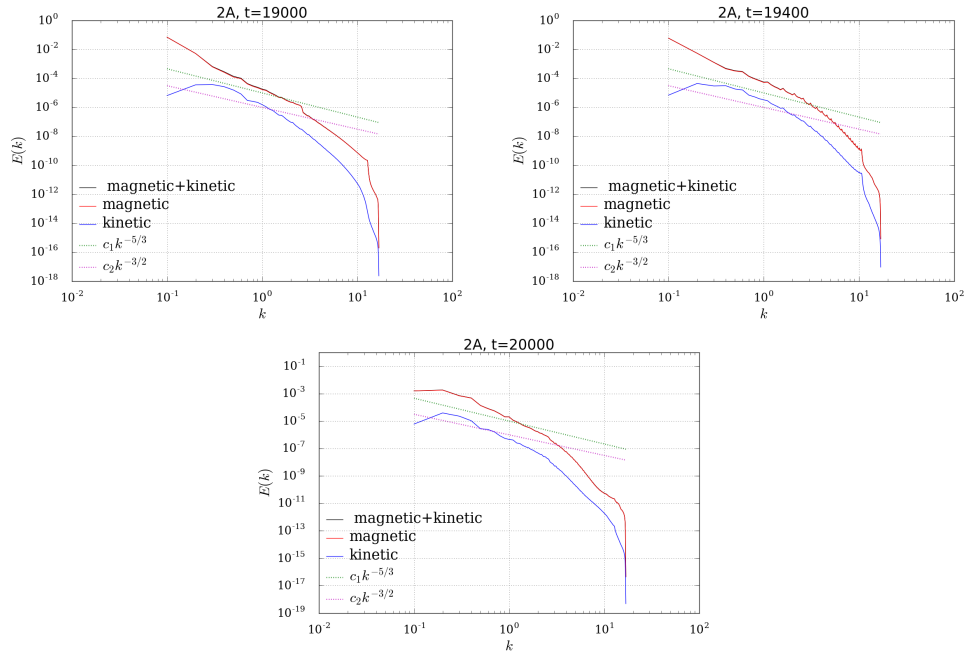

Figure 4.65: 3D spectra for run 2A, ($x_0 = 0.5$, equal/high helicity), at times $t = 19,000, 19,400, 20,000$.

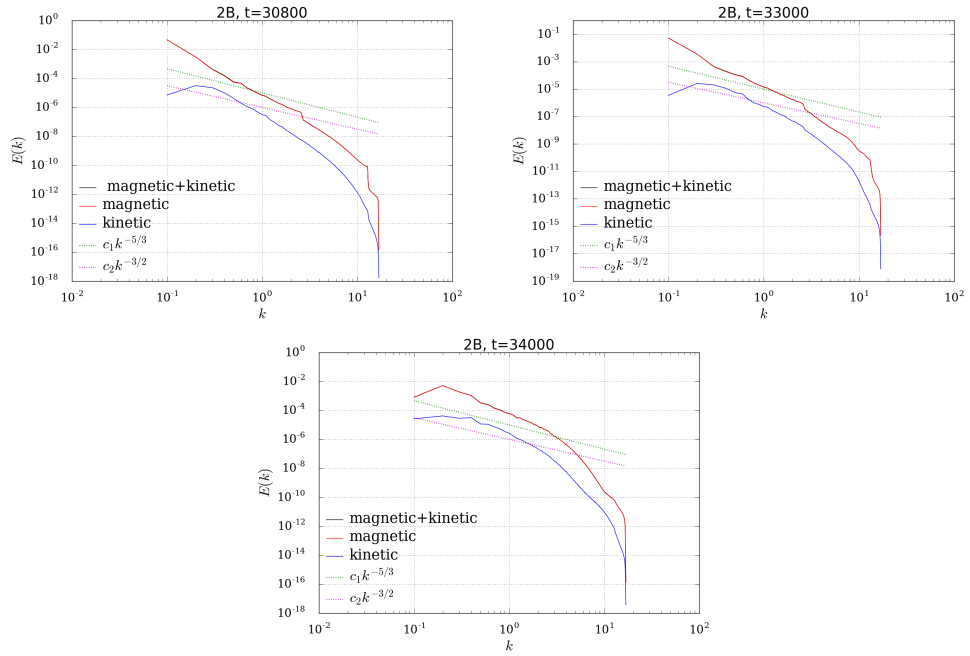

Figure 4.66: 3D spectra for run 2B, ($x_0 = 1.0$, equal/high helicity), at times $t = 30,800, 33,000, 34,000$.

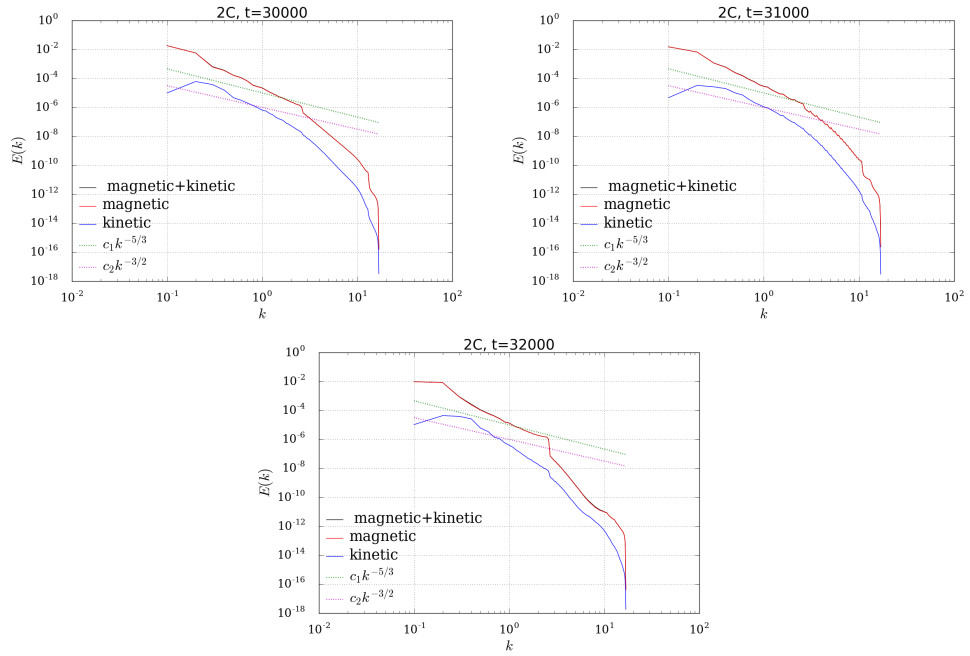

Figure 4.67: 3D spectra for run 2C, ($x_0 = 2.0$, equal/high helicity), at times $t = 30,000, 31,000, 32,000$.

the magnitude. The value $k = 2.5$ corresponds to a length scale in physical space of $\lambda = 1/2.5 = 0.4$. Therefore, for some reason, as we look at the energy at the largest length scales down, we reach a length scale of 0.4 and the energy value drops. The 3D spectrum does not offer any more information beyond this. We need to look at some 2D examples for further clues.

Figures 4.68 provide some examples of 2D spectra. We provide these plots for the initial snapshot analysed by 3D spectra for the run 1A, at $t = 29,000$. All other snapshots examined in 3D above showed the same picture in 2D. Taking the two dimensional transforms of the magnetic and velocity fields at for example, constant x, for all values of x, averaging, and then summing over 2D rings instead of the 3D shells, gives us a picture of the energy distribution in $k_2 - k_3/y - z$. Doing so for constant y and z also reveals, as in Figure 4.68, that it is only in $k_2 - k_3$ and $k_1 - k_3$ where the drop appears. The spectra in $k_1 - k_2$ are smooth, suggesting that it is in the z-direction that

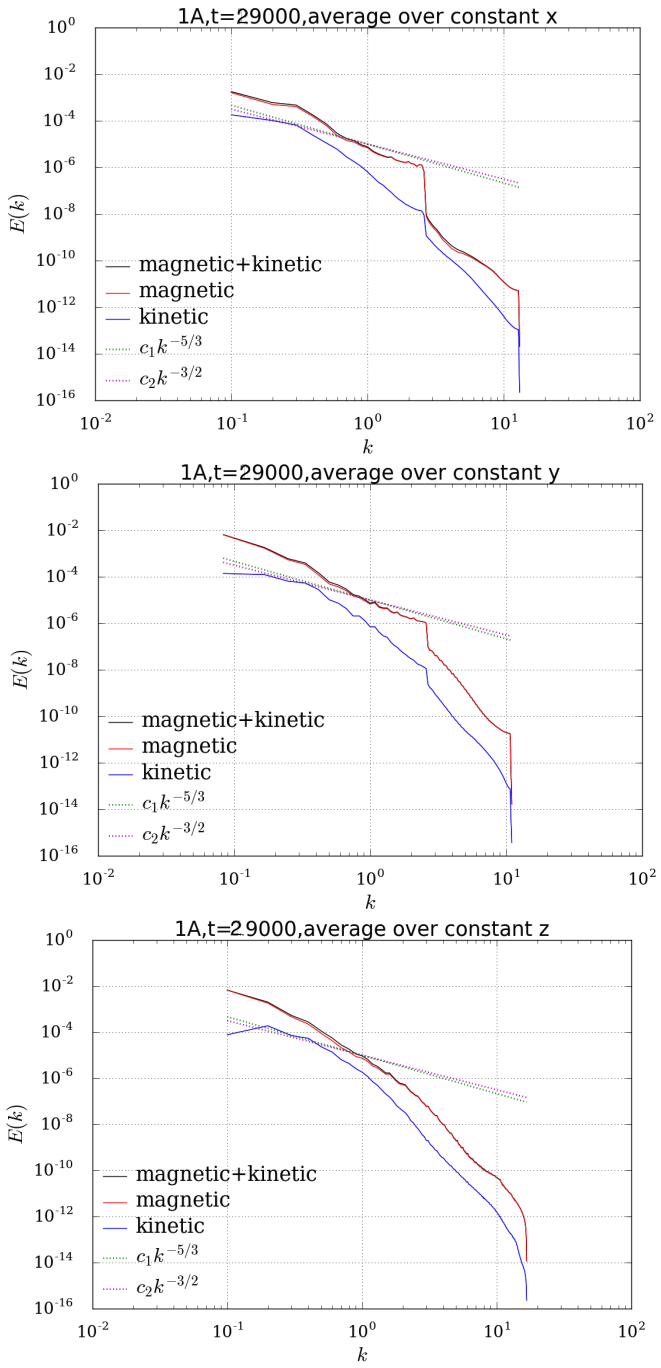

Figure 4.68: Two dimensional average power spectrum for run 1A at $t = 29,000$.

the characteristic length scale 0.4, which gives rise to the drop, lies. Also, the drop is larger in magnitude for constant x than for constant y (again this the same picture for all snapshots). To understand this we re-consider the numerical mesh constructed in

k-space.

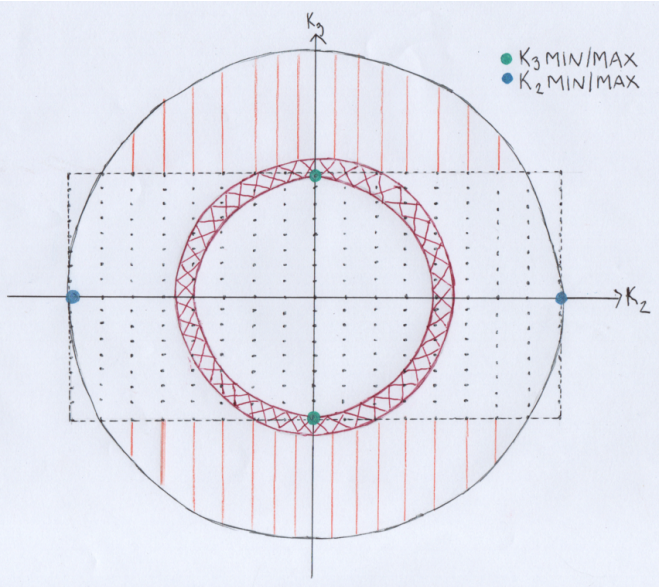

Figure 4.69: Shell drop off diagram.

The diagram in Figure 4.69 shows the k-space mesh for k_3 against k_2, with a couple of the rings in which values are summed superimposed. The k_3 axis corresponds to z, and k_2 corresponds to y. In physical space, the z-axis is the longest and the y-axis the shortest, therefore in k-space the corresponding axes k_3 and k_2 have the opposite attributes. The k values along these axes range from $k_2 = [-12.8, ..., 12.7]$, $k_3 = [-2.56, ..., 2.54]$, indicated in the diagram by blue and green markers respectively. Beyond the k-shell defined from a radius of 2.5 to 2.6 and shaded in red, the mesh has no more vertical extent, and larger and larger portions of each subsequent shell is empty - those areas shaded in orange. Beyond the red shell there are simply not as many points with values to sum, so the energy in larger shells takes a dip compared with those inside the red shell. Note also that the diameter of this boundary shell is 5, making the corresponding length-scale in physical space $\lambda = 0.2$, which is also around the value of dz in physical space ($50/256 = 0.195$). We simply do not have the resolution in z to keep counting energy values past this point. The resolution in x and y is much closer: $dx = 0.047$ and

$dy = 0.039$, so the mesh in the $x-y$ plane is not just as asymmetric. There is a ring past which there is no further extent in k_1, but it is much further down the length scales towards the limit of k, so we do not see the characteristic drops in the $x-y$ spectra.

The difference in the magnitude of the drop still requires an explanation. In Table 4.3 we show the average absolute value of each of the components of the magnetic and velocity fields in the domain at the first snapshot we have plotted 3D spectra for, for each simulation. We generally have B_z larger than B_x or B_y, so this is where the magnetic energy density receives its most significant amount. It is perhaps unsurprising then that when considering energy distribution involving the z axis that we can sometimes see such a large drop when we pass the smallest length scale in z - the signal has lost its largest factor. B_z is not defined at scales as small as B_x or B_y. We also note that the average $|B_y|$ is larger than the average $|B_x|$. Consider this: say each vortex's influence is only felt within a radius of 2 around the centre. Then the driver affects an area measuring about $-2...(2+x_0)$ units in x by 4 units in y. The action along x is distributed over a larger distance, while the action in y is confined between $y = -2...2$ for every driver. Recall that the most coherent driver twisting the field tighter lead to larger magnetic energies. If the energy is larger, the magnitude of the field components must be larger. The closer the vortices, the more confined the action, the tighter the twist, and the larger the field components. For each simulation the field is driven over a more confined distance in y, so this perhaps explains the field component in the y direction being generally larger than the field component in the x direction. Along x, the field is not twisted as tightly.

Further, note that in the 2D spectra the magnitude of the drop is larger in $k_2 - k_3$ than for $k1 - k_3$ - this is because dx is closer to dz than dy, so the $k_1 - k_3$ mesh is less asymmetric than the $k_2 - k_3$ mesh. Those runs with smaller or no drops - like 2A and 2B - have larger average B_x and B_y values (although still at least one order of magnitude smaller than B_z). These are also the runs with larger energy heating values and large

Average Magnetic and Velocity Components.

Run	B_x	B_y	B_z	v_x	v_y	v_z
1A	0.03	0.06	1.0	0.02	0.02	0.004
1B	0.04	0.12	1.0	0.02	0.02	0.005
1C	0.05	0.13	1.0	0.009	0.01	0.006
2A	0.06	0.341	1.0	0.005	0.008	0.005
2B	0.06	0.28	1.0	0.004	0.008	0.005
2C	0.07	0.18	1.0	0.007	0.010	0.005

Table 4.3: Average values of magnetic field and velocity field components at the first snapshot illustrated by the 3D spectra plots for each simulation.

reconnection events. The drop in the kinetic energy profile tends to be smaller than in the magnetic energy profile despite being defined on the same mesh - this can be explained by the fact that v_z is generally smaller than v_x and v_y, so when we pass $k = 2.5$ the loss of values further up the k_3 axis is not as important and leads to a smaller drop.

To confirm the theory that the difference in resolution in z is the key to this feature, we can run some other tests. We will use the driver from run 1A, which appears to result in a drop which, we hypothesise, is due to the resolution in z being much smaller than in x or y. In our simulations, recall that the physical domain spans $[-5,7]\times[-5,5]\times[0,50]$. We will compare the spectrum of this original simulation to spectra of test domains:

(1) $[-5,7]\times[-5,5]\times[0,75]$. Based on our our theory, we would expect to see some drop at a specific value of k based on the relative resolutions of each direction. As before dx and dy are similar, but now $dz = 75/256 = 0.293$. The diameter of the corresponding k-shell would be $1/0.293 = 3.413$. We expect then to see some sort of decrease in the energy distribution around $k = 3.413/2 = 1.7$.

(2) $[-5,7]\times[-6,6]\times[0,12]$. Now the resolution is the same in each direction we predict that we will not see an obvious drop in energy as we pass along the k-shells.

In Figure 4.70 we compare 3D spectra at $t = 200$ for the original 1A run, Test 1 and Test 2. As predicted, we see a distinct decrease in energy at $k = 1.7$ in Test 1. For Test 2 we do not see a significant fall in energy until we pass $k = 10$, which corresponds to a shell of diameter 20, or a length scale of 0.05 - which is the value of dx, dy and dz. Also interesting is the fact that in this case, there is a better similarity between our magnetic energy gradient and the turbulence power laws. Between about $k = 1..10$, the slopes are much closer. Considering that all other factors in these tests where the same, this raises an interesting question about the role of resolution in affecting results.

We should also note that the magnetic energy spectrum dominated the kinetic energy spectrum (other than in the test snapshots which were early in the simulations) and the magnetic plus kinetic energy looks much the same as the magnetic energy alone. This is simply consistent with what we have observed before.

Overall then, it is difficult to say whether turbulence is present in any of the simulations. None appear to induce a classical turbulent cascade with a power law index $-5/3$ or $-3/2$. Generally the slopes in 2D and 3D for both kinetic, magnetic and magnetic plus kinetic energy are much steeper. This suggests that energy is not being transferred down spatial scales, but is dissipated at larger scales. One of the more promising candidates for turbulence is Run 1A in Figure 4.62 - the magnetic energy in particular shows more of an agreement with the turbulence slopes in range $k = 1..2.5$, but then the profile is interrupted by the characteristic drop. The Test 2 profile for this same driver suggests that had the asymmetric mesh resolution not been an issue, we may have seen more evidence of turbulence. Future work could involved integrating in such a way as to avoid this issue.

The spectra with no characteristic drop are those with (relatively) larger B_x and B_y values. These are also the simulations which were driven by lower complexity, high helicity motions. These motions must be able to generate structures varying more from

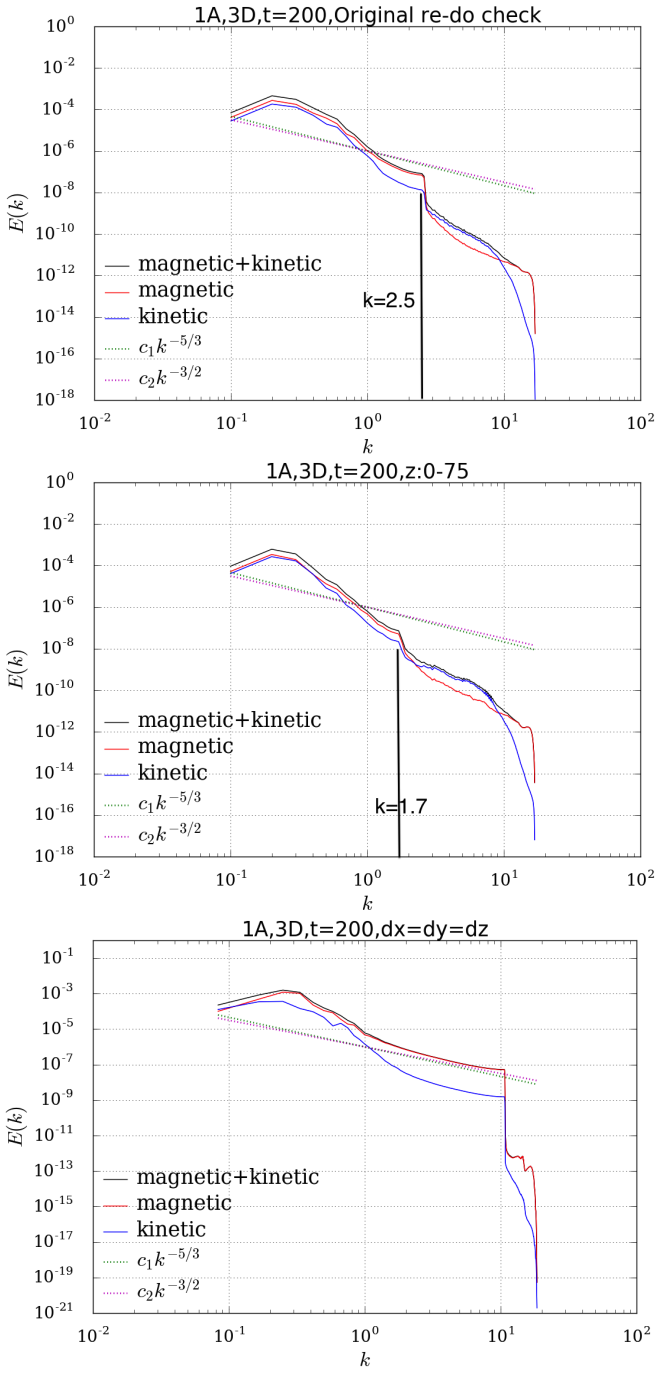

Figure 4.70: Comparison at time= 200 for run 1A, Test 1 and Test 2.

the initial field $(0,0,1)$. The studies of the energy distribution over the large magnetic energy releases also show more of the small oscillations we discussed as being ringing artefacts. In particular run 2A at time $19,400$ (plot 2 of Figure 4.65), during the largest

magnetic energy discharge, shows several small oscillations. This itself tells us that the field at this point must have large gradients, which fits with the theory that current layers are present and reconnecting in some large event. The smoother profile of the next plot in this figure suggests that current layers which had build up have been dissipated and the field now has a smoother structure once again.

Chapter 5

Magnetic Carpet Experiments

5.1 Initial Attempt

In this Chapter we will discuss the multiple problems that we ran into when testing our drivers on the magnetic carpet environment. We will look further into the code details to explain why these problems appear and how we tried to solve them.

Our first attempt at these simulations used the parameters presented in Table 5.1. In this scenario, we are investigating the carpet field, with magnetic features at the base and straight field rising up through the corona, rather than a straightened loop. Therefore, we initially set the upper z boundary to open - i.e. the velocity on the upper z boundary is not set to zero. All other parameters were the same as previous runs in Chapter 4.

Property	Value/Nature
x, y boundaries	periodic
Upper z boundary	open
Lower z boundary	driven, $\|\mathbf{v}\| \leq 0.1 v_A$
$\mathbf{B_0}$	$\mathbf{e_z}$ + CARPET
$\mathbf{v_0}$	$\mathbf{0}$
Resolution	$256 \times 256 \times 256$
Domain	$[-5,7] \times [-5,5] \times [0,50]$
η	0.0005, uniform
viscosity	shock capturing form
Duration	45000 time units (596 driver periods, 900 Alfvén times)

Table 5.1: Simulation details for carpet field tests.

5.1.1 First Instance of Problems

We began to experiment with the initial carpet field, applying each of our established drivers in turn; instantly we had the problem that all six simulations progressed incredibly slowly. We illustrate this in plot (a) of Figure 5.1. Here is plotted the time against the timestep (the number of predictor-corrector + remap steps taken) for the simulation under driver 2A, but all other new simulations showed the same problem. We can see that it took around 4500 time steps to reach 10 time units, but then a further $25,500$ time steps to evolve another 5 time units to reach 15 time units - a rapid drop in the size of successive time steps and efficiency of the code. This was far too slow to be practical and affected all other simulations in the same way, regardless of driver type.

As a first test then, since the driver type seemed to be innocent, we considered the boundary conditions. We had the x and y boundaries set to be periodic as before. Unlike the simple uniform field $(0,0,1)$ which was automatically periodic across these boundaries, the magnetic carpet field is not. This means that the carefully constructed magnetic carpet, designed to be divergence free and potential, is being forced into periodicity by the code, giving the field different values at some points than the initial

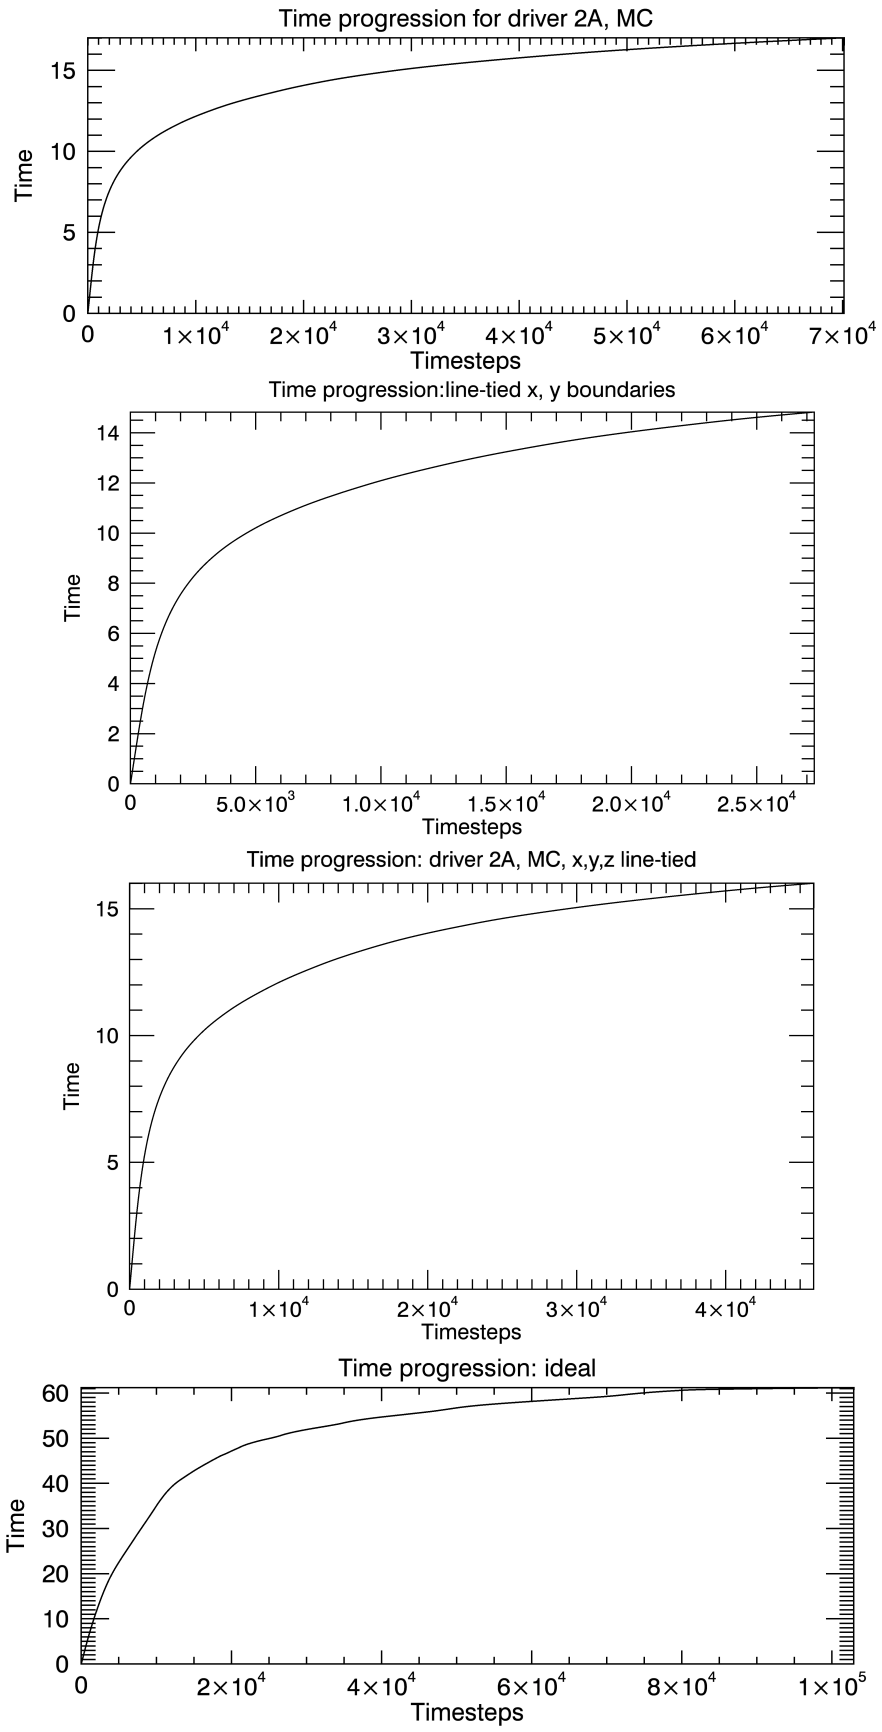

Figure 5.1: Time progressions for driver 2A acting on the magnetic carpet field.

design - could this affect the set up in a way that could cause the time progression to stall? We set the x,y boundaries to be closed too, and checked that the field values at random points on the numerical grid now gave the same values as for the same points substituted into the original algebraic expression. This change was to no avail, as seen in plot (b) of Figure 5.1. This simulation reached around 10 time units in a little under 5000 time steps, and had almost made it to 15 time units by a further $22,500$ time steps. This is comparable to this first attempt, so the periodic boundaries were not the problem, even though the field was more accurate. Lastly we tried making the upper z boundary line tied too, since that was the only other difference with the previous experiments. Plot (c) of Figure 5.1 shows the time progression for this test and again we see no improvement - it made 10 time units in just under 5000 time steps and 15 time units in around a further $22,500$ time steps, so the difference in boundary conditions are not to blame for this issue. Therefore, clearly the approach taken previously is not appropriate for this new scenario. We now set out to understand why, and what we should do differently.

5.1.2 Timestep Factors: CFL Condition and Simple Fix Attempts

Now we focused specifically on factors which could affect the size of the time step. The interval in time covered by each timestep depends on a number of variables and therefore does tend to vary slightly, but ideally (for example in our previous simple field simulations) each timestep will advance the simulation by roughly the same time. Our problem here is that each subsequent timestep is evolving the system over a shorter and shorter period of time until we effectively have a stalled simulation, with such a small timestep that time is practically no longer advancing.

The timestep in a numerical simulation may be able to change from step to step, but there are restrictions placed on it due to the CFL condition: that is, the Courant,

Friedrichs and Lewy condition (see [23] for the original work). The solution to a partial differential equation at time $t > 0$ depends on the data at some continuous set of points at $t = 0$ - this set of points is called the mathematical domain of dependence. The finite difference scheme chosen to approximate the differential equation is defined on a mesh of discrete points. Data from a discrete set of points will be needed in order to update each variable at each step, and so the scheme has a numerical domain of dependence. Courant, Friedrichs and Lewy showed that a condition for convergence of a finite difference scheme is that the mathematical domain of dependence must be contained within the numerical domain of dependence. In other words, finite difference regimes update the value of a variable at some point in space and time using the value of variables at neighbouring positions in space and time. Therefore, for a scheme to be successful there must be a way to guarantee that the code will have access to all potentially necessary points at each step, so the timestep must be limited such that we will not run into the situation where the value at a point is relying on an undefined value.

This can be quantified algebraically. The variable C is referred to as the Courant number and is a ratio involving timestep, velocity and grid spacing as in Equation 5.1. Grid spacing is chosen by the user of the code and the velocity is dependent on system dynamics. Therefore the timestep Δt is set at each iteration in order to keep the Courant number at every point less than or equal to some number C_{max}. Physically, the Courant number quantifies how information travels at velocity u over a distance measuring Δx in time Δt. For a given grid spacing, the velocities generated will dictate how large we can take Δt. We want this time interval to be as large as possible in order to complete the simulation in as few time steps as possible, completing the simulation efficiently, but we cannot make it so large that C exceeds C_{max} somewhere on the grid. If C becomes too large, the physical speed of information in the system is faster than the speed $\Delta x/\Delta t$ in which the code can update the variables. The value of C_{max} also

varies depending on the finite differencing scheme utilised - Lare3d uses $C_{max} = 0.8$. So if, for example, the grid spacing becomes smaller, the ratio of the velocities to grid spacing may become larger, therefore Δt may have to decrease in order to maintain $C \leq 0.8$. Similarly, an increase in velocities may lead to Δt being decreased. This condition must be met for every point on the computational grid, so at each iteration Δt must be set to the value which will satisfy Equation 5.1 for every single point. Velocities can depend on the current density and plasma density amongst other things. Therefore at each iteration the code considers a number of properties and settles on a feasible timestep length.

$$C = \frac{u\Delta t}{\Delta x} \leq C_{max} \tag{5.1}$$

The next step we took attempted to prevent this timestep magnitude falling. We tried simple alterations - some made minimal improvements - i.e. the time progression did not stall as early or as steeply - but none removed the problem or even sped things up enough to be practical solutions. These attempts included the following:

1. We tried things which would have an immediate and direct effect on the CFL limit - by equation 5.1, increasing the grid spacing may allow the timestep to increase. We therefore increased the domain range in x, y and z by various different test amounts, to no avail - it was not enough to make a difference.

2. By Equation 5.1, if we can decrease the velocities we may be able to bring up the timestep. Reconnection can lead to plasma flows, therefore reducing η may reduce reconnection opportunities and the velocities around potential reconnection sites. Setting η to various smaller values, some an order of magnitude smaller, led to a small improvement, but again not enough to be a viable solution. We also turned off resistive effects entirely, but the time progression for this test is

shown in Plot (d) of Figure 5.1. It illustrates how the point of stalling was only delayed a little.

3. We tried both a slower boundary motion and turning off driving. The fact that this did not help, again helps exclude the driving process from any suspicion.

4. We also tried making the grid spacing equal in all dimensions, so that $dx = dy = dz$. Irregularity in cell dimensions can sometimes lead to problems. However adjusting this made no significant difference here.

5. We tried both turning off shock viscosity and using a different form, as per the Lare3d user guide suggestion when a code seems to hang. The problem still did not go away.

We now turned again to the initial condition and other simulation variables which can affect the timestep more indirectly. For example the current density depends on the magnetic field, whose evolution depends on velocity; the velocities are also affected by plasma pressure and density. We found three key issues:

1. the divergence of the numerical field was non-zero,

2. the initial currents and therefore initial Lorentz forces were non-zero, and

3. the minimum plasma density in the domain drops further and further as the simulation slows and is close to zero when it halts almost entirely.

Clearly items 1 and 2 should not be the case at all based on the design of the field, and item 3 must somehow be caused by one or both of these inaccuracies, which in turn is causing the timestep to drop. We sought now to pinpoint the details.

5.1.3 Discretisation Error

There is naturally a discretisation error introduced when one translates a field, given by an algebraic expression, into a variable prescribed only on the finite number of points of a numerical grid. This error vanishes in the limit of vanishing grid point spacing, so the higher the resolution employed, the closer to the analytic form the numerical field will be. The results obtained will be more accurate with a smaller error. We are also obviously aware that the higher the resolution, the more computationally expensive a simulation will be. We found that a grid of 256^3 was a good resolution in terms of accuracy and practicality in previous simulations, and continued with that for our first attempt at the magnetic carpet experiments. The other factor affecting the accuracy of a numerical scheme is the order of the finite differences. One could increase the order of the derivatives calculated and gain precision without increasing resolution. Lare3d is second order, while codes such as the Pencil code (we will return to this in Section 5.3) can use differences of up to order six. The nulls and topology of our carpet field result in highly variable fields within small regions, the behaviour of which perhaps second order differencing cannot capture as accurately as one might hope.

Let us return to the error which we may try to address within these particular experiments: the discretisation error. One consequence of this is that an initial field whose algebraic expression is, as in our magnetic carpet case, divergence and current free, may have a non-zero divergence and current in the code. It appears that the non-zero current leads to non-zero Lorentz forces acting on the plasma and influencing the plasma flows, and therefore plasma density distribution, from the very beginning of these simulations.

Let us consider the divergence first of all. Both our very first attempt using the set up described in Table 5.1 and the test setting the x and y boundaries to be line-tied had stalling. Intuitively we expect there to be more of a discretisation error in the section

of the domain containing the parasitic polarities and their local field. As we move up the computational box to higher values of z we move into a uniform field region like that of the previous chapter's initial condition where the numerical divergence should be close to zero. This leads us to think of the maximum absolute divergence in planes of constant z, as in Figure 5.2. We have plotted the maximum absolute value of the divergence at each value of z for different boundary conditions, and we see that it does indeed decrease with z. The maximum divergence occurs near $z = 0$ for simulations with x and y boundaries periodic and line-tied - but the value of the maximum varies significantly. For x and y periodic, where the field is forced into periodicity, the maximum magnitude of divergence is 25.26 - this is wildly inaccurate. For x and y boundaries closed, the maximum magnitude of the divergence is 0.56 - two orders of magnitude smaller than the first set-up and clearly the better of the two, but still much larger than we want. The tests for all boundaries closed for resistive and ideal circumstances tell the same story. While the divergence issue is very worrying, the fact that the difference in divergence does not appear to make a direct difference in the timestep issue as seen from the time progression plots led us to first consider the Lorentz forces in more detail.

To investigate the details of the Lorentz forces we focus on a less complicated carpet field, to make things clearer.

Simple Carpet Tests

The simplified test uses an initial condition with a source placed at $(-1, 0, -0.85)$, so that our null is positioned at around $(-1, 0, 0.737)$ - simply put we have one of the dome structures we illustrated in the basic unit, centred at $x = -1$. This field is sketched in Figure 5.3 - in the first plot we are lying along the y-axis through the middle of the separatrix dome created, with the same picture lying along $x = -1$ in the second plot.

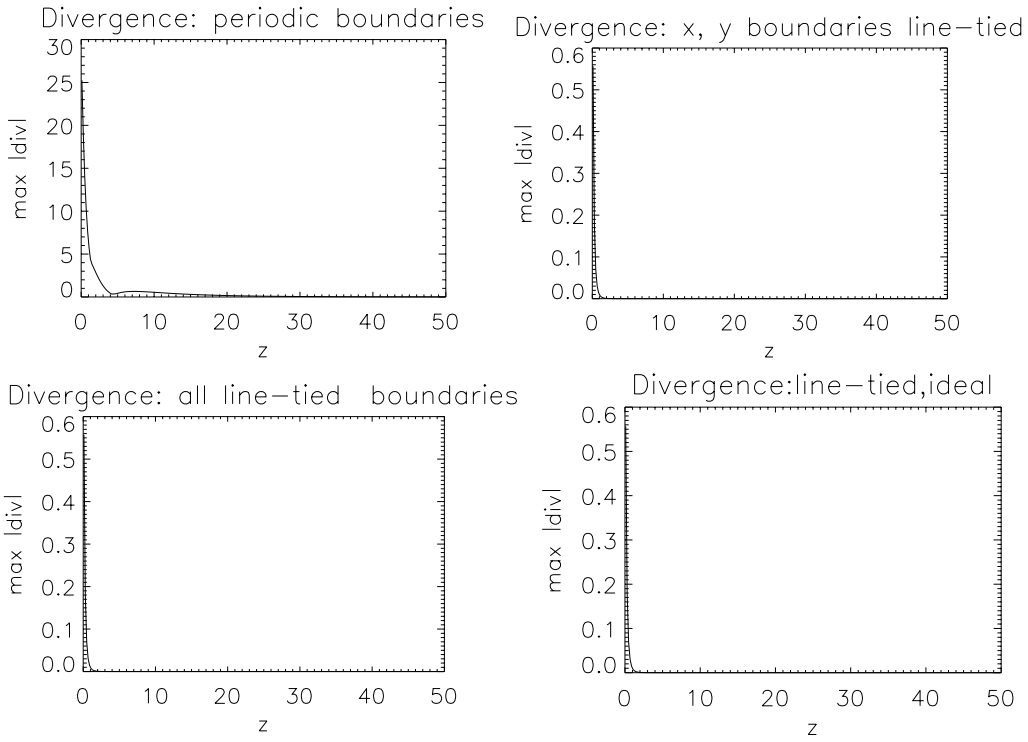

Figure 5.2: Maximum absolute value of $\nabla \cdot \mathbf{B}$ in planes of constant z. Clockwise from top left, the first plot is for a simulation with periodic x and y boundaries, open upper z boundary and line tied lower z boundary. Next is the scenario with line tied x,y and lower z boundaries and an open upper z boundary. Then we have the same quantity when setting all boundaries to be line tied and taking an ideal environment. Lastly, we return to a resistive simulation and keep all boundaries line tied.

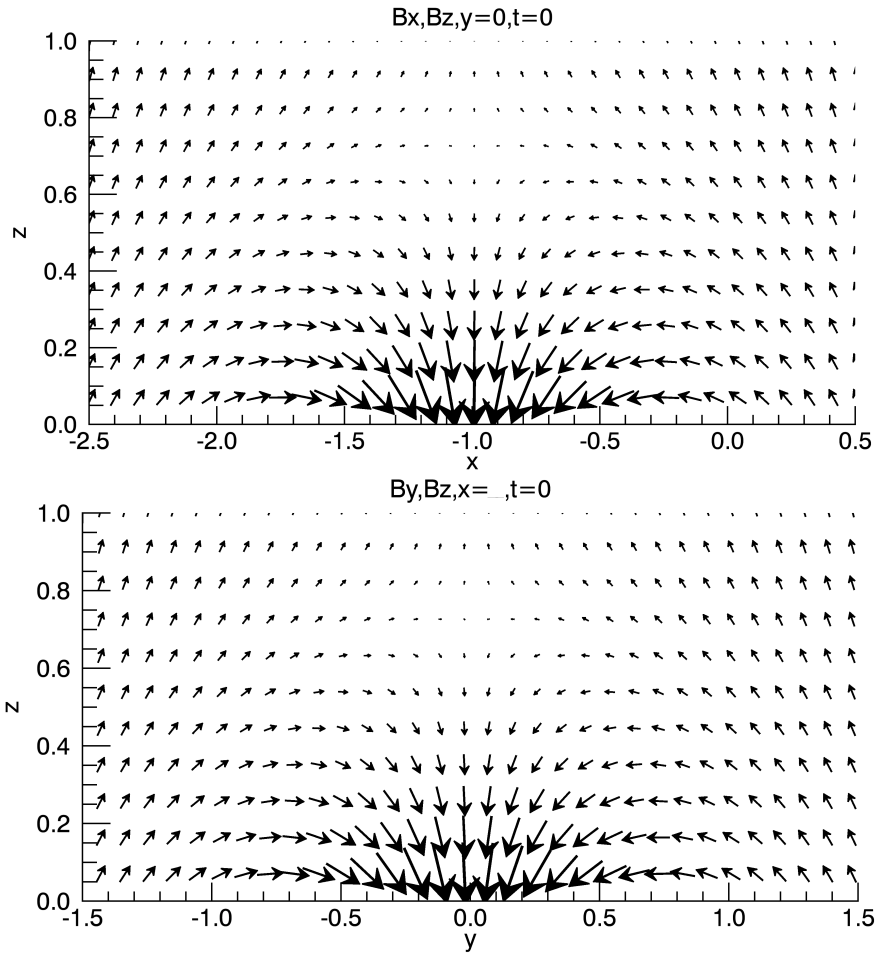

Figure 5.3: Original field structure around our single parasitic polarity field. The single null lies at $[-1, 0, 0.737]$, with fieldlines curving over and up from the null above the separatrix dome and up and back down again inside the dome.

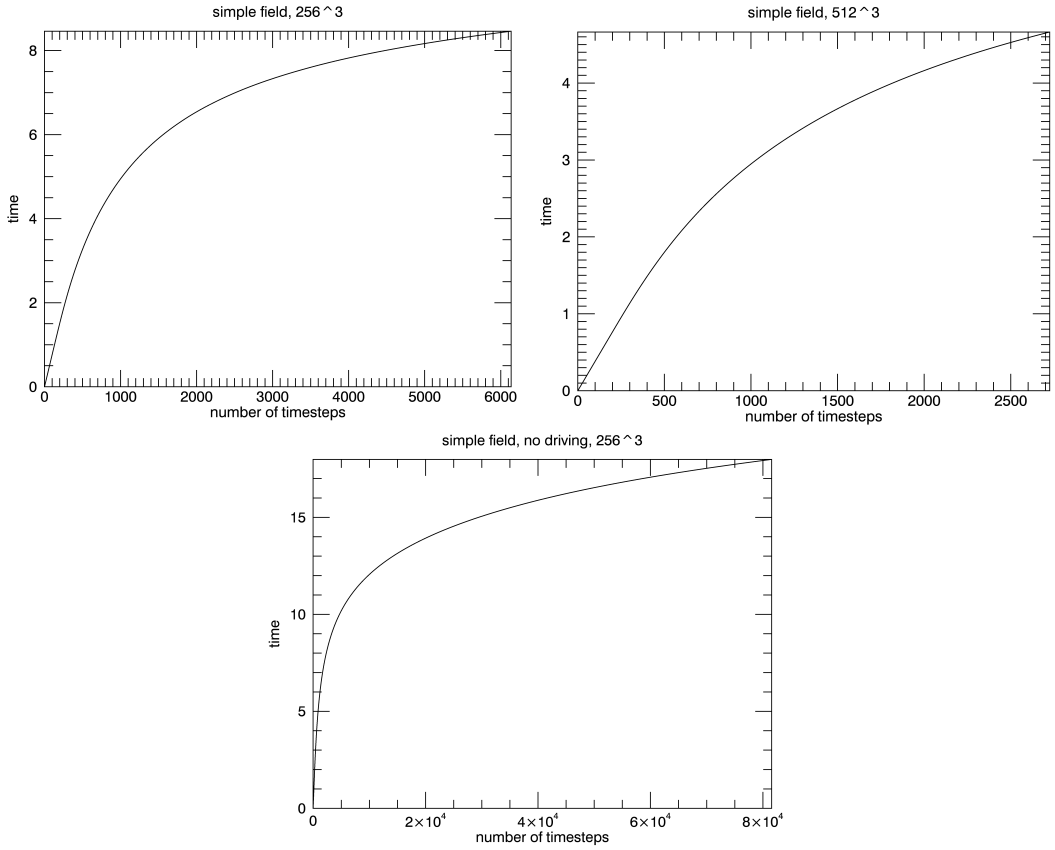

Figure 5.4: Time progressions of the simple field tests for (a) 256^3, with driving, (b) 512^3, with driving and (c) 256^3, no driving.

We ran these tests on a smaller domain consisting $[-5,7] \times [-6,6] \times [0,12]$ to keep everything as symmetric and simple as possible and avoid additional complications. We kept the x and y boundaries line-tied and returned to an open upper z boundary. Initially, to check if this simplification itself gave rise to any improvement, we tried running tests on this setup with driving (driver 2A) for resolutions of 256^3 and 512^3, and then without driving for 256^3. The time progression plots for these tests are shown in Figure 5.4. Again none of the changes have removed the problem, illustrating that this problem lies in the nature of the translation of the basic parasitic polarity structure to the numerical grid.

Proceeding on a 256^3 grid, the investigations showed that the unintended non-zero

Lorentz forces act from the beginning within the dome region to force plasma towards the separatrix surface, letting the density within fall close to zero. We illustrate this in detail in the following plots, and refer to the Lorentz force as $\mathbf{L} = (L_x, L_y, L_z)$.

Figure 5.5 illustrates the magnitude of these forces. We see that that L_x, L_y, L_z are small (indeed close to zero) everywhere but inside the region where the dome structure sits. The black regions have zero magnitude where there is uniform field, but under the separatrix layer the Lorentz magnitude reaches up to around 21 units. Figure 5.6 (a) shows a cross-section of the 2D vector field (L_x, L_z) at $y = 0$, and (b) shows (L_y, L_z) at $x = -1$. Clearly Lorentz forces are acting from the centre of the dome down and out. Initially the density is set to be constant throughout the domain and velocity is zero, however right from the very beginning we have magnetic forces acting within the dome. We will call the region with these anomalous Lorentz forces region L.

We reassess the state of the field once the simulation has started. For this set-up we saw the time progression slowing significantly by 8 time units. At this point in the simulation we are not even yet halfway through the first driving period. The first vortex is spinning from $t = 0..12\pi = 0..37.7$, so it is still early in its first cycle. The density distribution has already been altered greatly, as seen in Figure 5.7, where we observe plasma density along $(a)y = 0$ and $(b)x = -1$. Outside the dome the density is still fairly uniform but plasma has migrated from the region L up to the cusp of the null region and out towards the separatrix surface, leaving a region with very low plasma density - region L has effectively been evacuated by plasma. Figure 5.8 shows the plasma pressure and also how material is collecting around the location of the original null, leaving low pressure regions within.

Velocity fields are illustrated in Figures 5.9 and 5.10 and show exactly how the plasma is being moved, the arrows of the field echoing the outline of the plasma suggested by the density and pressure plots. The magnitude of the velocities range around levels

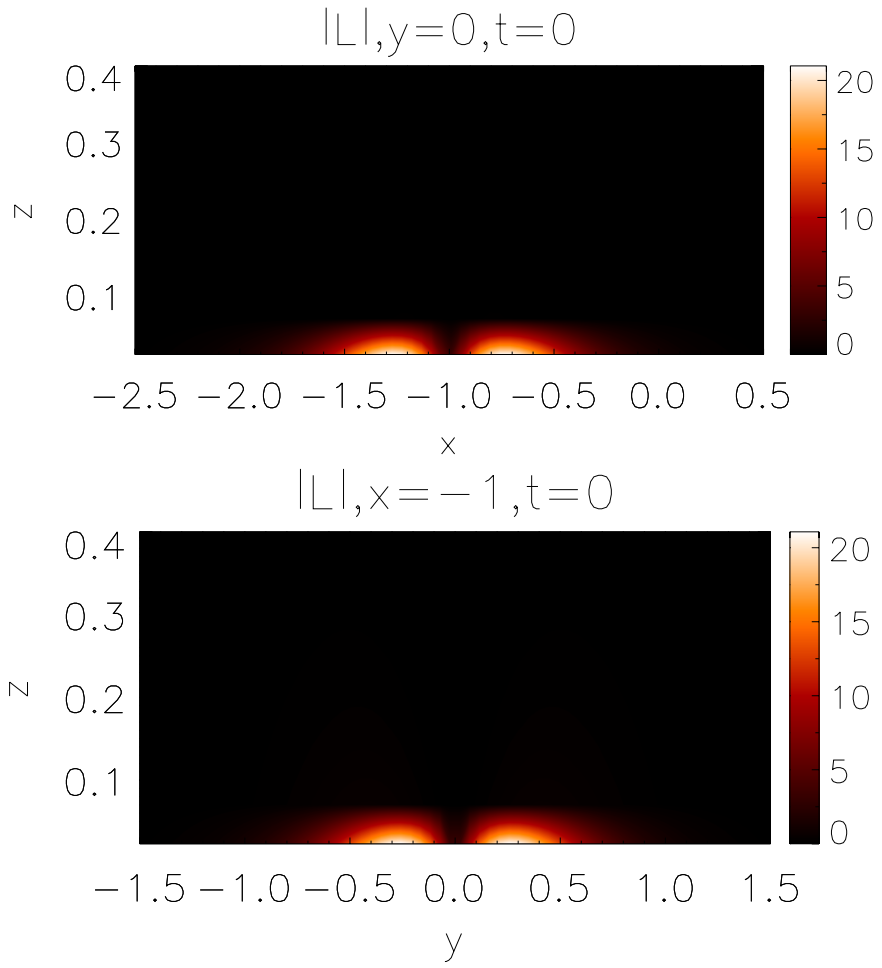

Figure 5.5: Magnitude of the Lorentz forces at t=0. Highly localised non-zero forces are at work.

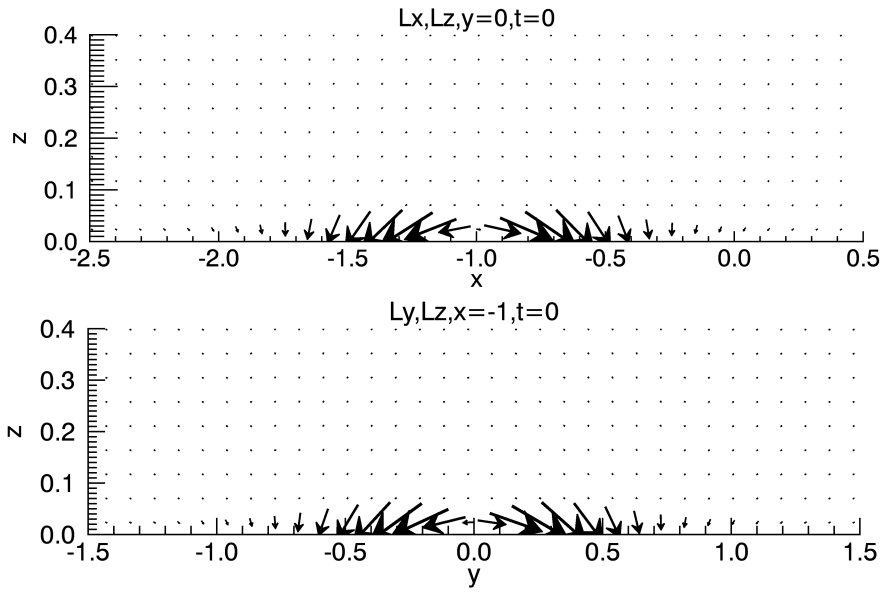

Figure 5.6: Initial Lorentz forces - these should be zero everywhere, but we see non-zero forces acting in the region of the magnetic . These act away from the middle of the dome from the very beginning of the simulation.

comparable to that of the driver - i.e. slower than Alfvén speed but not insignificant.

Again, this is still very early in the simulation, and there are still significant Lorentz forces acting as seen in Figures 5.11 and 5.12. They appear to continue to be acting away from the already depleted region of low density, and indeed the problem escalates further.

Figure 5.13 shows the structure of the magnetic field at this point - the centre of the action has shifted a little but the overall topology is still very similar. Therefore, we end up in a situation that is doomed from the beginning, with our original region L becoming increasingly depleted, the timestep decreasing continuously, until the simulation effectively stops.

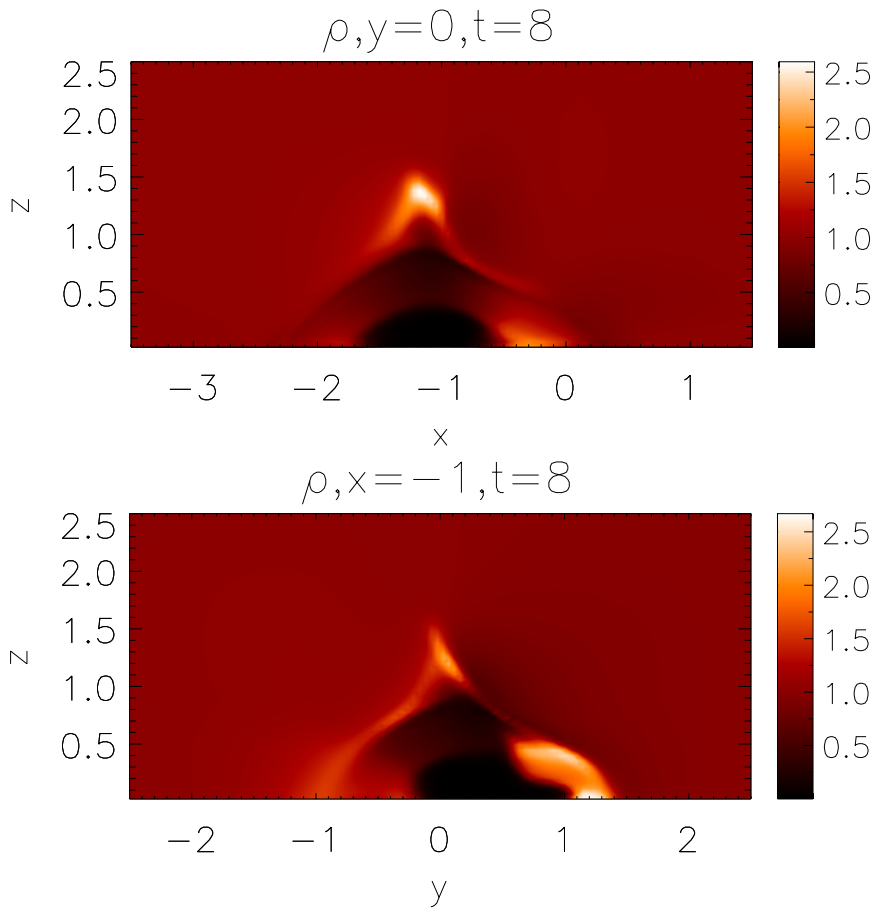

Figure 5.7: Plasma density at t=8. Plasma has moved from region L outwards and upwards, leaving behind a volume of very low density.

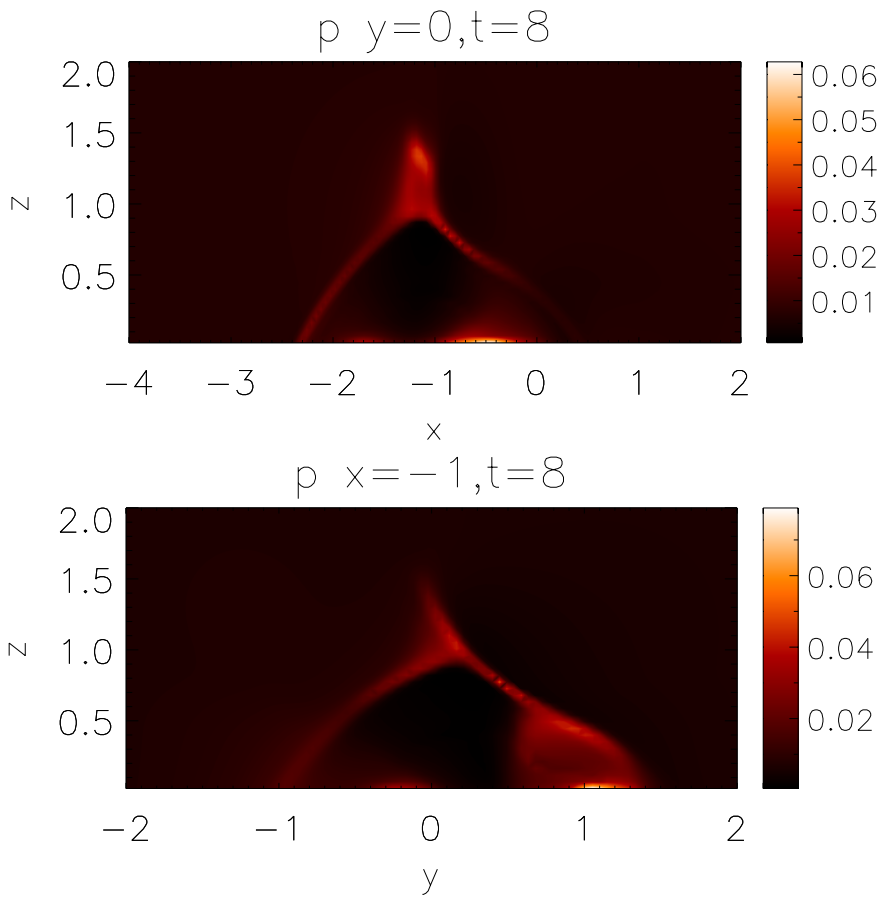

Figure 5.8: Plasma pressure at t=8. Plasma appears to be collecting in layers around the original null location leaving region L depleted in its wake.

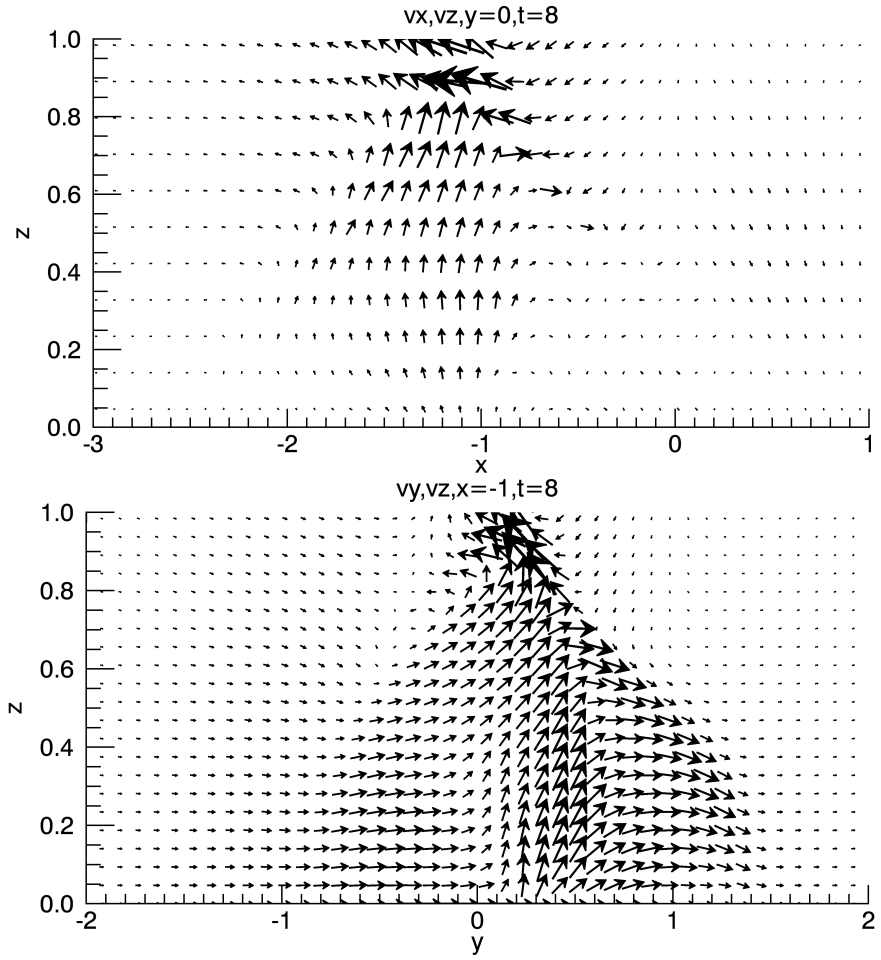

Figure 5.9: Velocity profile at time=8. In cross section of the *xz* and *yz* planes we see how the plasma is moving up and away from region L.

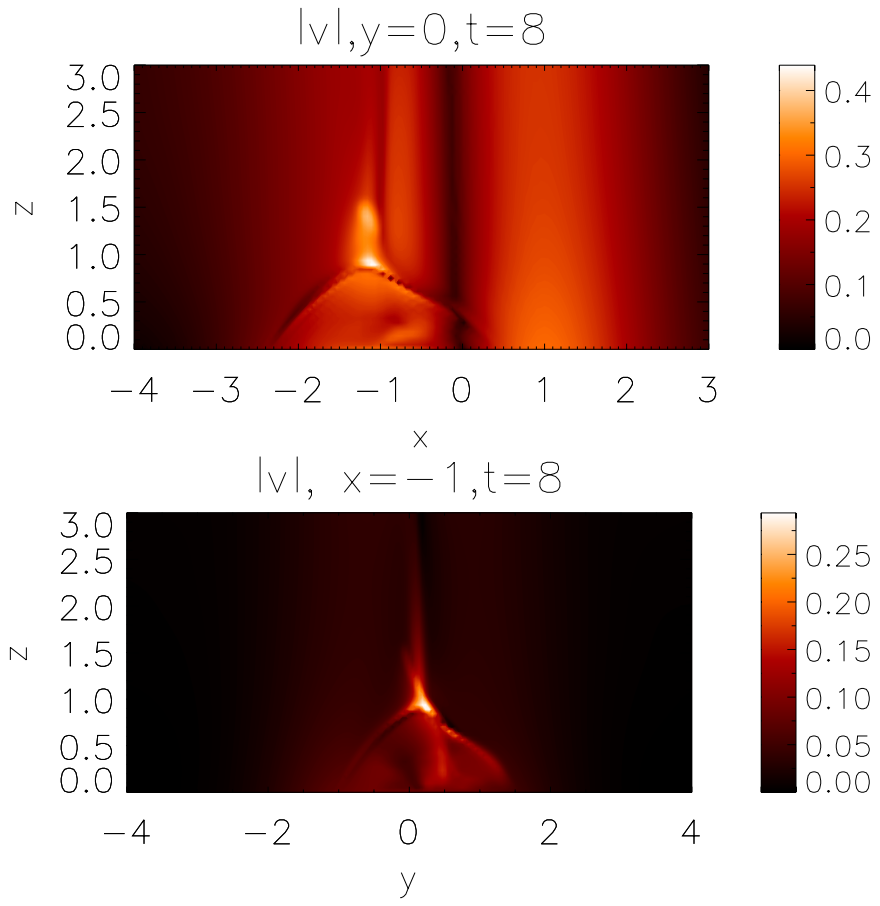

Figure 5.10: Velocity magnitudes at 8 time units in the magnetic feature region.

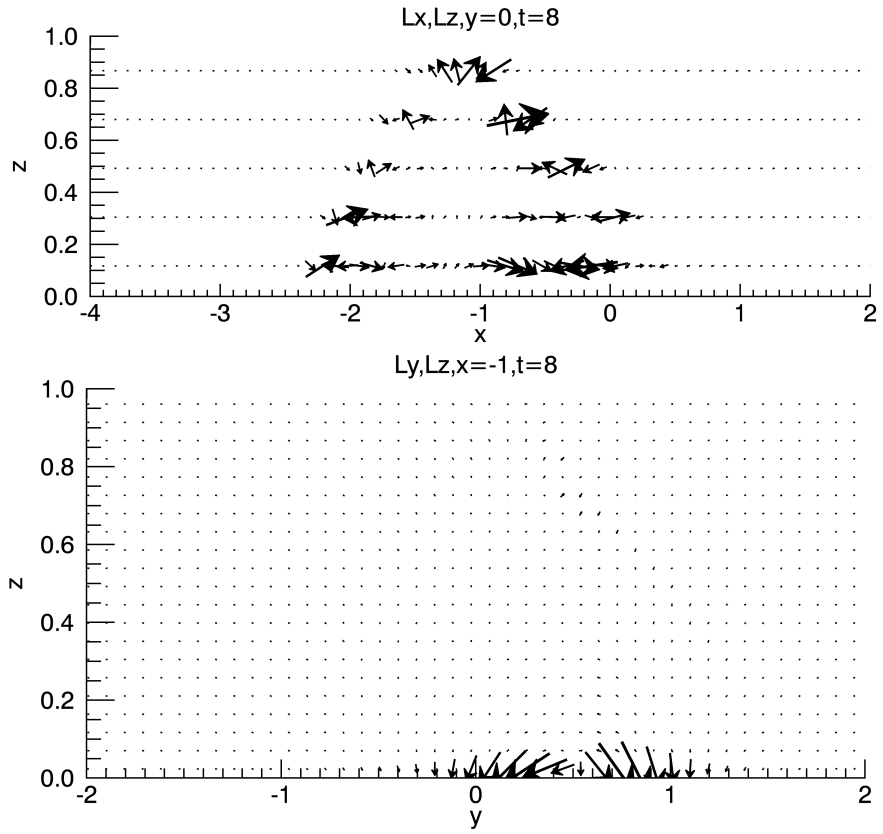

Figure 5.11: Lorentz forces at t=8.

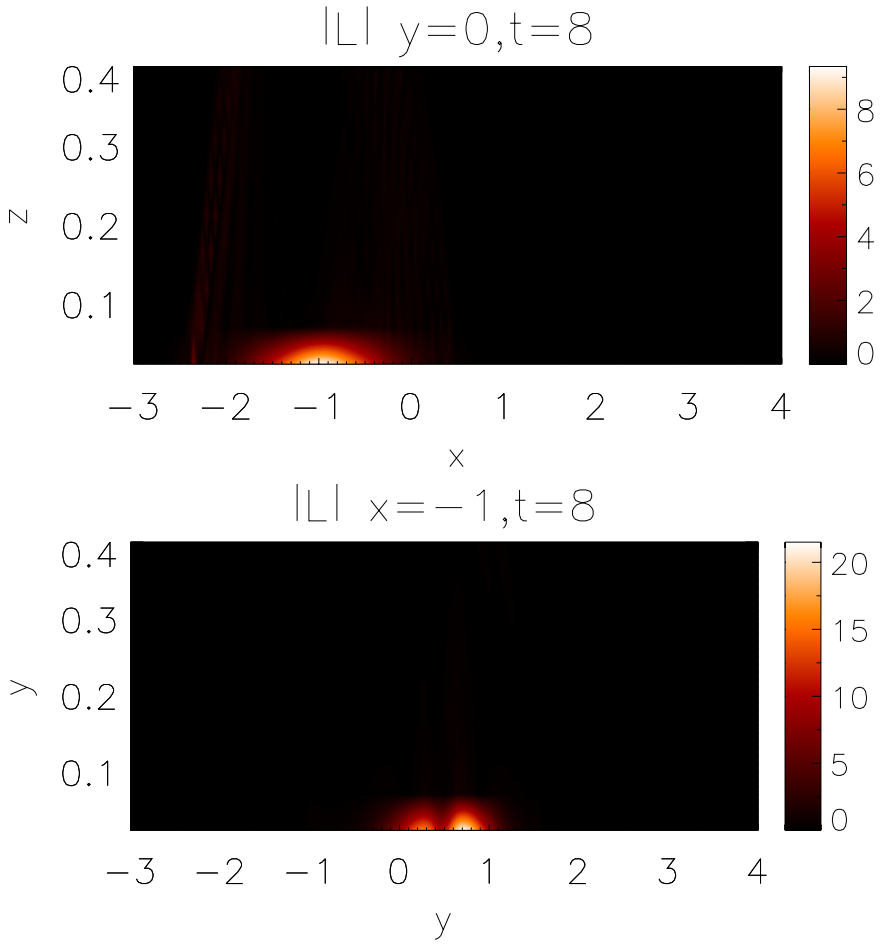

Figure 5.12: Lorentz force magnitude at t=8.

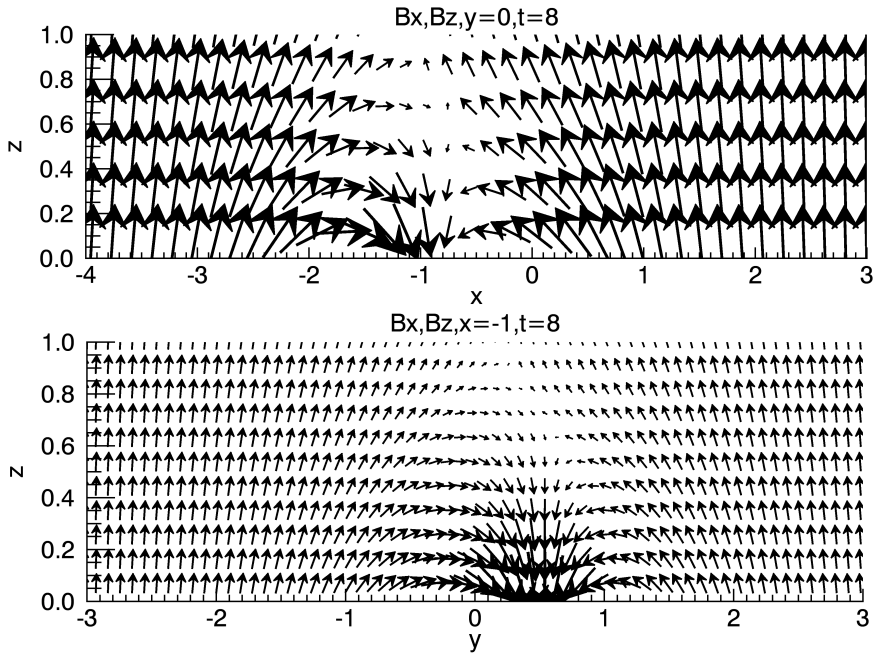

Figure 5.13: Magnetic field at t=8. This is still very early in the simulation and the field is mostly unchanged, and the plasma migration continues, leading to stalling of the simulation.

The details of key values and locations are summarised in Table 5.2. The maximum divergence is located just under the null, but the real problem appears closer to the *xy* plane. By 8 time units the minimum density is 0.000177, at a point inside the dome. By 10 time units, this minimum value has dropped to 0.00000172 and the simulation is not making any progress. Again this occurs within the dome, close to the *xy* plane. The simulation has failed before it even really got going.

	Simulation Details: Null Location: $[-1,0,0.737]$	
Variable	Value	Location
$\nabla \cdot \mathbf{B}$	max= 0.0859	$[-0.99,0.0234,0.02]$
ρ	$t=8$: min= 1.77×10^{-4} $t=10$: min= 1.72×10^{-6}	$t=8$: $[-0.945,0.26,0.02]$ $t=10$: $[-0.945,0.44,0.02]$

Table 5.2: Details of problem variables. By the time we reach 10 time units, the minimum density has dropped to the order of 10^{-6}, within region L, and the simulation has stalled.

We saw two mains options to deal with this issue:

1. We could re-write the initial condition in some way, so that the numerical curl was identically zero. This is highly non-trivial task, however.

2. We could accept the non-zero forces and attempt to 'limit the damage'. Plasma flow is dictated by the balance between the Lorentz force and the plasma pressure. In this case the Lorentz force appears to be too dominant - could we reduce its influence instead?

We could reduce it, rather crudely, simply by dividing the magnetic field by some constant. This would also affect the normalised Aflvén speed, so we would have to adjust our driving speed and period to keep everything in the same relative proportions as before. The normalising constants could be adjusted to keep everything the same in real terms. However, this is rather complicated and unsatisfactory.

Rather than actually reducing the Lorentz force then, we can try to reduce its level of dominance over the problem plasma flows by increasing the plasma pressure, thereby increasing the plasma beta, making the two factors more equally matched. Initial pressure is not explicitly set in Lare3d, but by changing the initial specific internal energy density ε, we change the plasma pressure. This would not have any other knock-on effects either, as nothing else is defined in terms of ε. We did have $\varepsilon = 0.01$. Trial and error tests of different values led us to set $\varepsilon = 0.25$ - this gave good improvement on simulation speed without changing things any further than necessary. Most of the field is uniform and therefore the previous normalised plasma beta in most of the domain in terms of ε was

$$\beta = \frac{2(\gamma-1)\rho\varepsilon}{B^2} = 0.0133.$$

Therefore, setting $\varepsilon = 0.25$ leads to multiplying this by 25, giving $\beta = 0.332$ in most of the domain.

Figure 5.14 shows the time progression for this particular test up to 8 time units in plot (a) and 190 time units in plot (b). For comparison, (by eye estimates from the plots), the run with $\varepsilon = 0.01$ was slow at around 8 time units, and it took about 6000 time steps to reach that. We can see from plot (a) that the simulation with $\varepsilon = 0.25$ reaches around 8 time units after about 1500 time steps. Plot (b), while far from a straight line, shows that time is not stalling as it was before. The timestep does not remain as consistent for the duration as it was up until 8 time units. For comparison, the uniform field under the influence of this driver had consistent time progression and reach 200 time units in around 10% of the time steps, as shown in Figure 5.15. However, the rate of progression of the simulation is now much better than it was and the time step does not plummet any more. This is a definite improvement, and seemed a good compromise.

We also questioned whether this approach would simply slow down how long it took for the plasma to migrate and end up leading to the same problem further down the line. Looking at the Lorentz profiles in Figure 5.16, in the same planes as before, at $t = 190$ for this new experiment, shows that these forces are changed from their initial states - the original field has evolved past the point where we met problems before, and will hopefully not run into the stalling problem later in the simulation.

5.2 Increased Plasma Pressure Runs

We proceeded with this new approach and applied each driver to the initial carpet field, this time with $\varepsilon = 0.25$, line-tied x and y boundaries and an open upper z boundary - everything else remaining as originally intended. We did not see any further significant

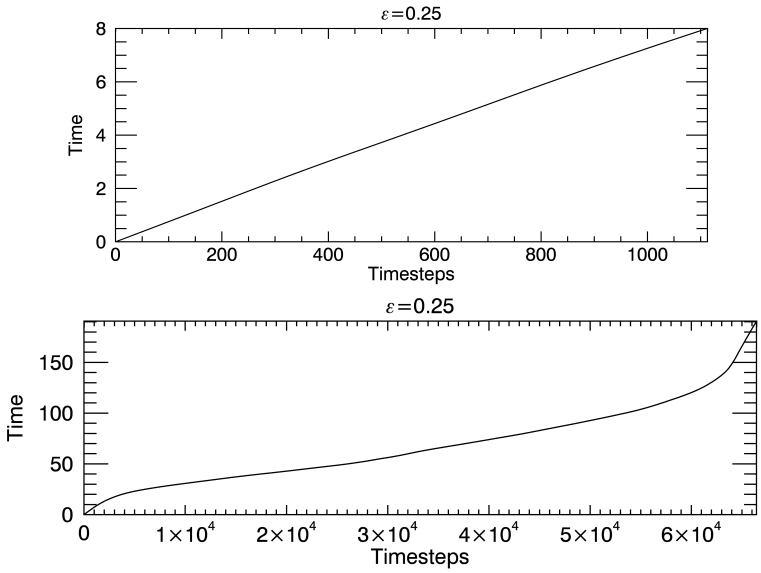

Figure 5.14: Time progression for $\varepsilon = 0.25$ up to $t = 8$ for plot (a) and $t = 190$ for plot (b). Previously the time step had fallen significantly by 8 time units and the simulation was slowing. Here we see a fairly uniform timestep and no loss of progress. Plot (a) shows how the simulation makes it to 8 time units in just under 1500 time steps - a vast improvement on the almost 6000 time steps it took to reach this stage previously. Plot (b) shows the time progression up to $t = 190$ time units. The timestep has not stayed consistent but the simulation has not stalled.

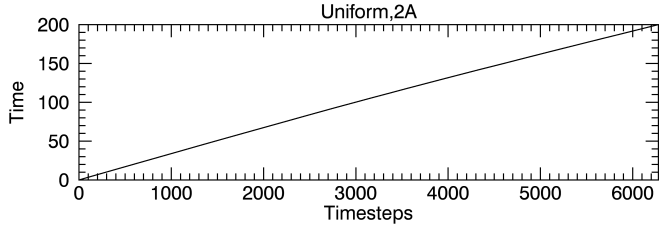

Figure 5.15: Time progression up to $t = 200$ for the uniform field under the influence of driver 2A.

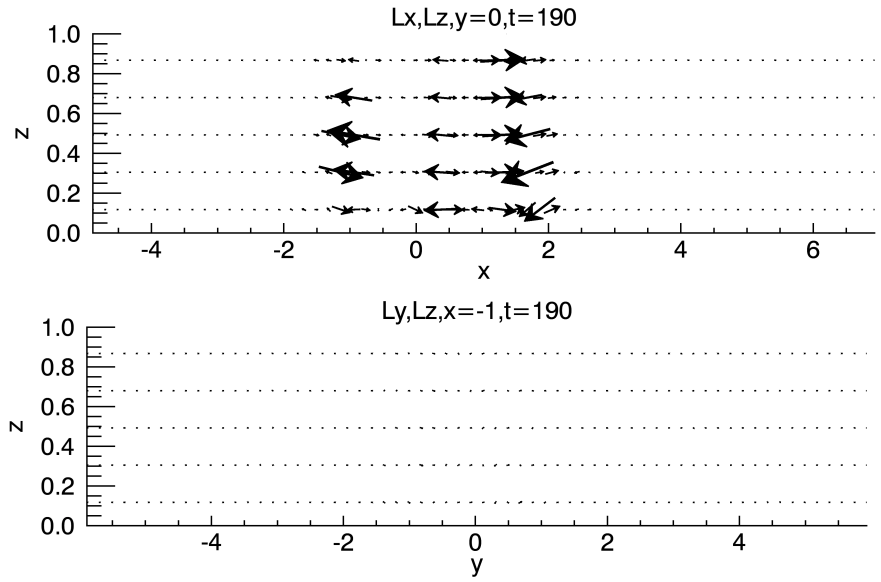

Figure 5.16: Lorentz forces at $y = 0$ (plot (a)), and $x = -1$ (plot (b)).

time progression issues - some started more slowly than others, but we had no further problems with very slow progress .

Figure 5.17 shows energy diagnostics for the system under the influence of driver 2A. These plots run to almost $24,000$ time units - but there is not much happening. The profiles appear thick due to many small oscillations, but overall the quantities have very smooth profiles. The magnetic energy in particular seems unusually lacking in features. After a brief decrease at the beginning of the simulation it begins to rise slowly. The internal energy has a spike at the beginning and then stays roughly constant, while the Ohmic and viscous heating rise steadily. The magnitude of the heating is small compared with the levels seen for the uniform case over this length of time, but is larger than the change in internal energy. This can be explained due to the fact that the Ohmic and viscous heating are calculated cumulatively and in this case we have an open upper z boundary. These plots and the following analyses point to theory that the initial structure is quickly destroyed by the driving at the very beginning of the simulation, so we see a drop in magnetic energy, a spike in internal energy and rise

in Ohmic heating at the very start. We will see that the field then remains quite uniform, with the only activity coming from the driving action which triggers a relatively small amount of heating via Ohmic and viscous effects. The heating variables are stored cumulatively, so they continue to rise. It seems that a stable system is reached whereby the total internal energy is not changing particularly because the heating from the driving and that lost through the top boundary appear to be balanced. Checking this by carrying out a second simulation with the top boundary closed and everything else the same confirms this, since in this test the internal energy continued to rise, rather than staying along a roughly constant level. All other variables exhibited the same kind of behaviour, indicating that the boundary difference is not responsible for their characteristics.

Figure 5.18 displays the magnetic energy profiles for each of the other five runs - they all look suspiciously similar. We began some further checks to explain this.

Figure 5.19 displays energy diagnostics for a test with no driving, all boundaries closed and the larger ε value. The kinetic and internal energies look sensible - after some initial activity, they settle to roughly constant values (kinetic to zero as expected). The heating quantities also level out after some initial changes. The magnetic energy however, seems to settle to a lower level, as we would expect for a system under no external forcing which approaches equilibrium, but then starts increasing slowly but steadily again. There is no physical reason why this should happen.

This initially seemed very worrying. Take driver 2A again. The initial magnetic field has a magnetic energy of 2969.2 units. It falls to a minimum of 2955.1 units at $t = 1,586$. By $23,599.6$ time units, i.e. $22,013.6$ time units later, the magnetic energy value has risen back to 2968.1 units. So in the space of around 292 driver periods, the magnetic energy has risen by 0.44% of its minimum value, and is almost back to the same value as the initial state. Comparing to the non-driven initial field, it reached its

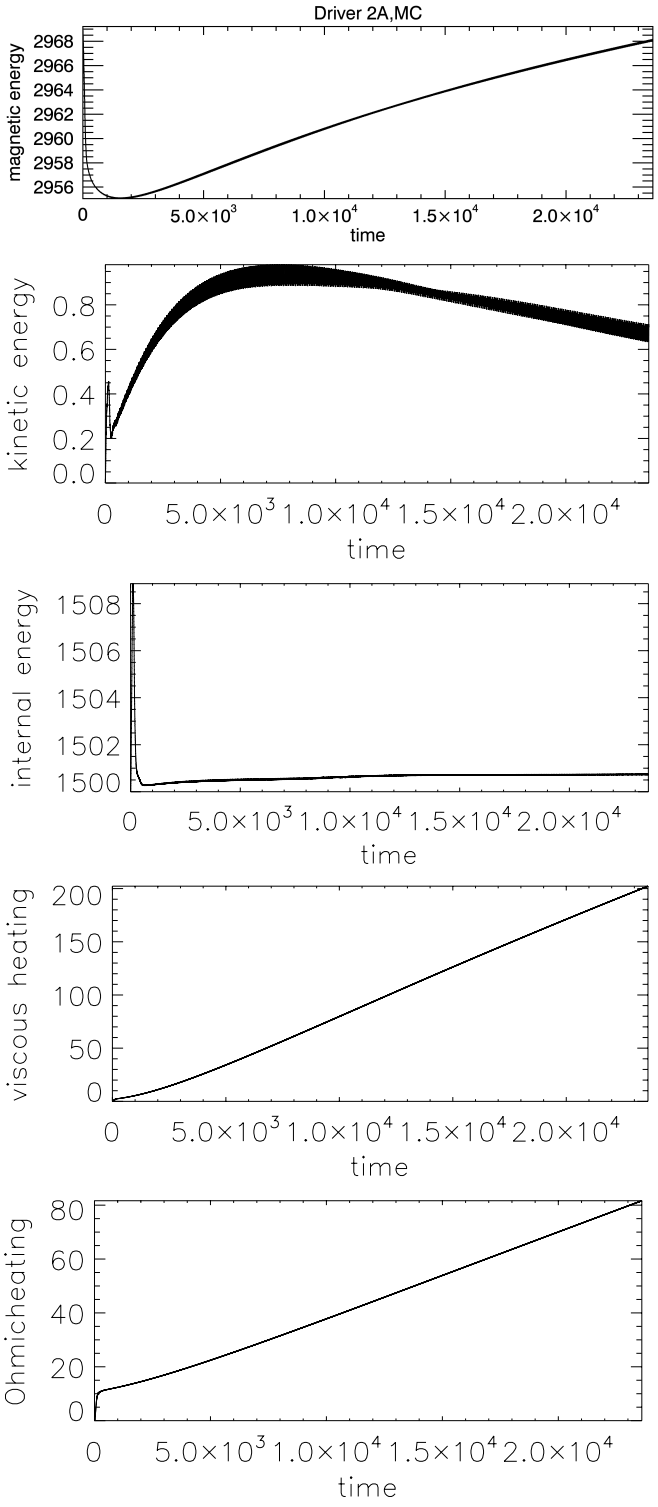

Figure 5.17: Driver 2*A*, magnetic carpet field original attempt with $\varepsilon = 0.25$.

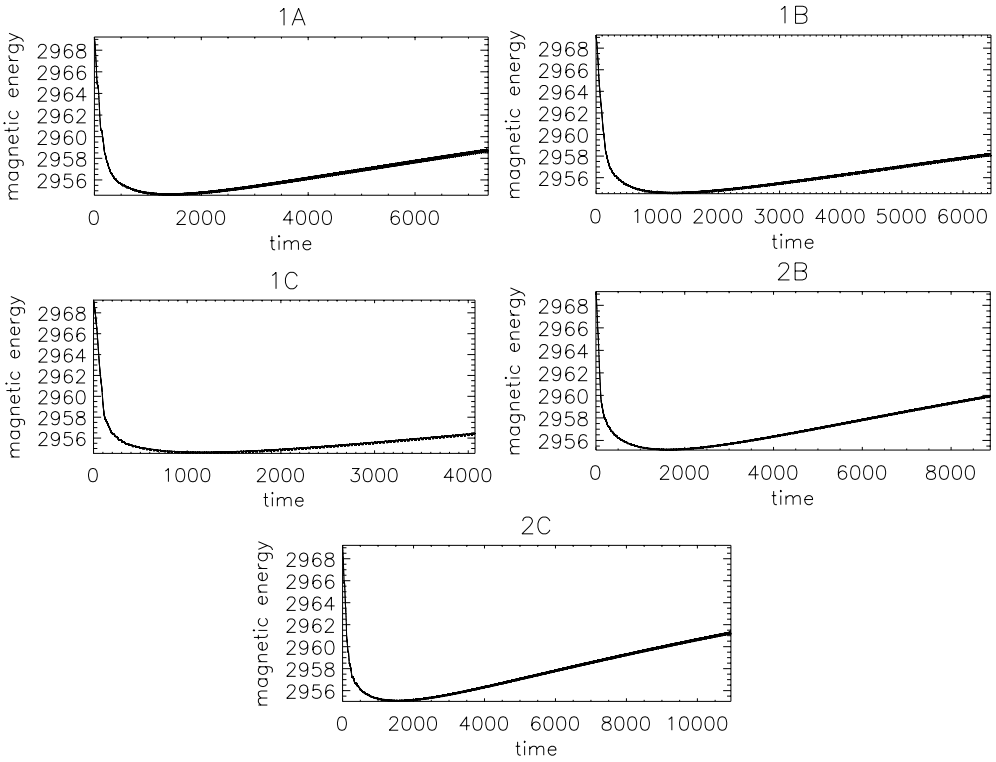

Figure 5.18: Magnetic energy evolutions for the other runs.

minimum value of 2954.9 and then started rising again to a value of 2957.1, 4,2298 time units (or 57 driver periods) later. This is an increase of 0.0744%. To establish an explanation for this consider the following facts. For a direct comparison of these two cases, note that our non-driven case ran until 6631 time units. By this same time, the driven case had a magnetic energy value of 2958.4 units - a 0.11% increase from the minimum. Furthermore, it took the driven case 951,684 time steps to reach this point, whereas the non-driven case took 551,930 timsteps. So, the driven case has taken 399,754 more time steps to reach the same simulation time as the non-driven case and when it gets to that time the increase in magnetic energy from the minimum values is larger. These properties suggest that the increase is due to an error - the non-driven case simply should not display this behaviour, but if there are errors in current calculation in parts of the code from the very beginning (and these errors are passed on in the magnetic field evolution through the induction equation), then this effect will

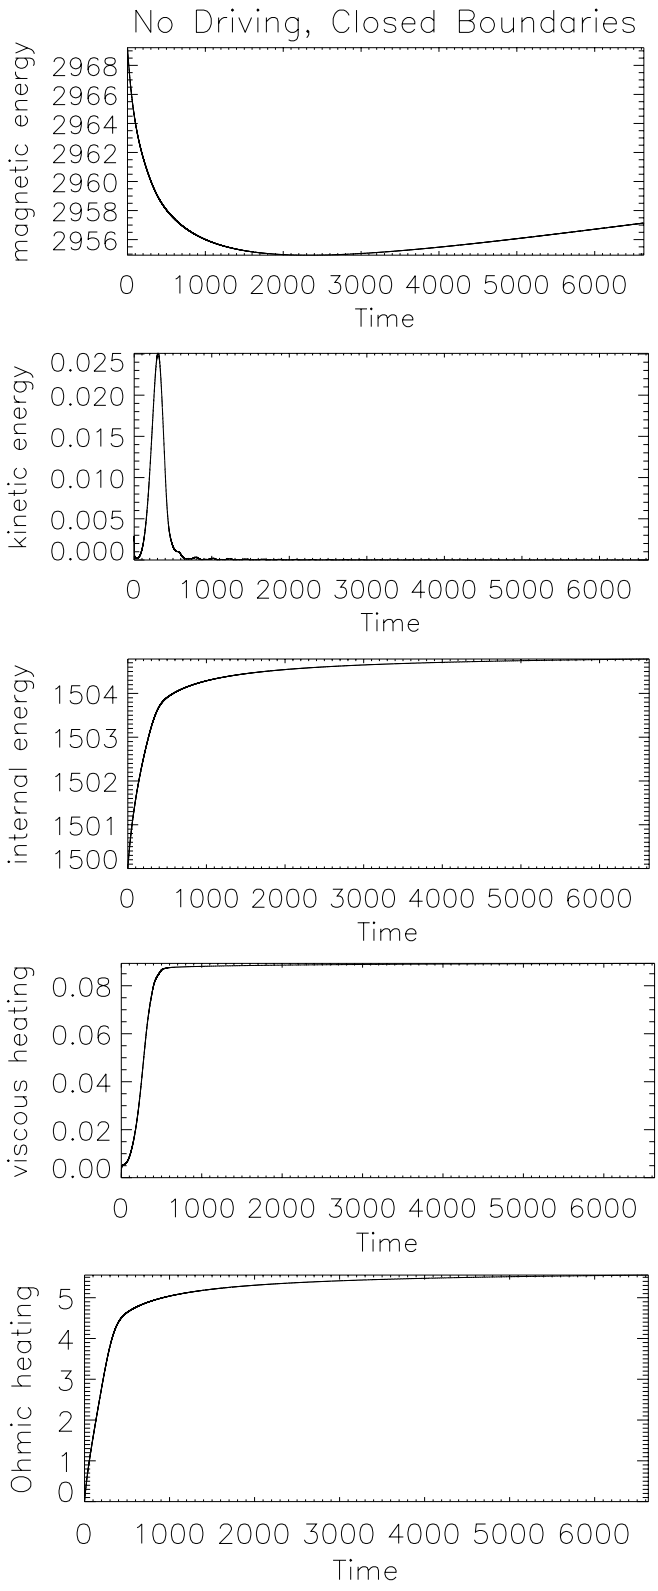

Figure 5.19: No driving, magnetic carpet field with $\varepsilon = 0.25$ and all boundaries closed.

Driver	% Increase from min	Timesteps from min
1A	0.030	115,587
1B	0.034	127,544
1C	0.037	133,453
2A	0.020	106,005
2B	0.020	100,527
2C	0.024	107,325

Table 5.3: Increase in magnetic energy by 3000 time units in terms of percentage of the minimum value, and the number of time steps it takes to reach this point, for each driver.

build up with each successive iteration. The more iterations, the more the error should build, and this is indeed confirmed by the fact that the driven case takes more timesteps to reach the same time and in turn shows a larger increase in magnetic energy.

To check if this concept is consistent we can check the other runs too. Table 5.3 compares for each different driver what was the percentage increase in magnetic energy from the minimum value by 3000 time units. It also quotes how many time steps it took to reach this point from the time of the minimum value. In general we can see that the size of the increase does appear to increase with number of timesteps. We can also see a correlation with the characteristics of the drivers - the Case 1 drivers were the higher complexity runs, with 1C being the most complex. It is also driver 1C which takes the most time steps and leads to the highest increase after the minimum, with 1B and 1C the next highest timestep resolutions. Driver 2C is the most complex of the other drivers and is the next most labour intensive run. So, higher complexity runs, seem to induce conditions whereby timestep is lowered further and numerical errors can build up more.

Above and beyond this, none of the simulations show the energy diagnostics changing in significant ways, relative to the previous chapter runs, and we are not able to learn much from them. Focusing again on the driver 2A run, which we eventually ran to 45,000 time units, we see that there continues to be very little activity compared with

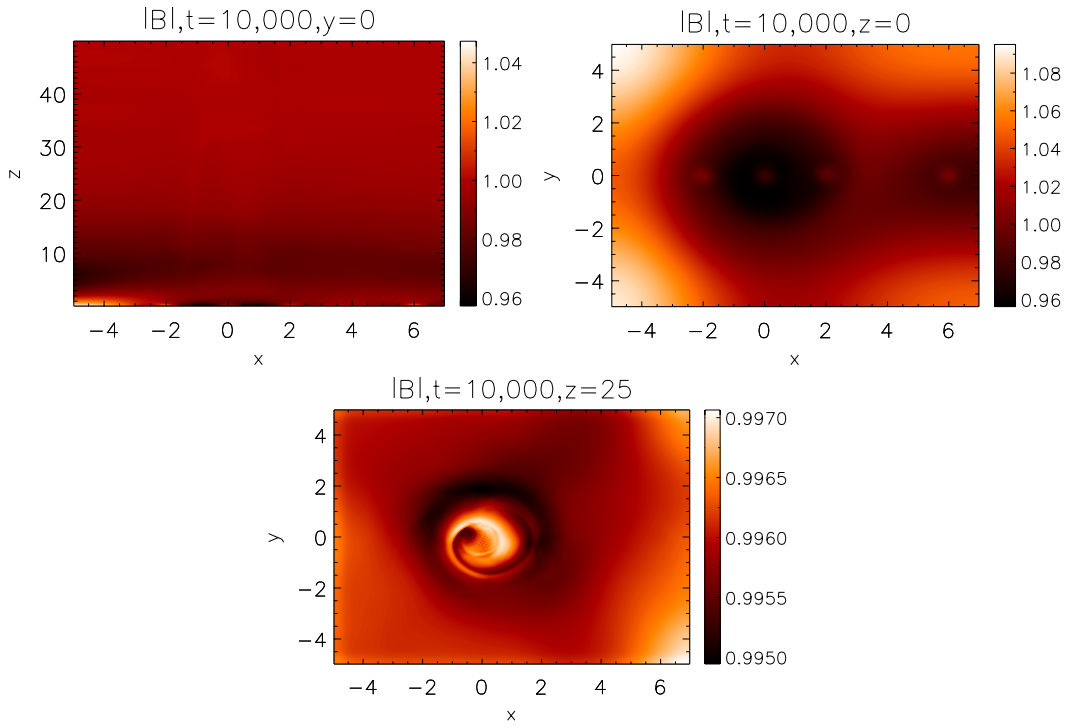

Figure 5.20: Contours of $|B|$ at $t = 10,000$. Planes lying at $z = 0$, $z = 25$ and $y = 0$ show that the magnetic field strength is fairly homogenous compared to the initial state. The field has been smoothed out.

the previous uniform field tests. Figure 5.20 illustrates the field strength at $t = 10,000$ in cross sections of the domain. This is during the 133^{rd} driving period. The driver has therefore had lots of time to act on the initial structures and we expected to see the field beginning to to twist up in the same way as the initially uniform field. However it seems that the field is quite smooth, with the field strength actually varying very little in all of the planes shown. We can see this in Figure 5.21, where we plot the magnetic field.

Figure 5.22 shows contours of density at this same time, for planes lying at $z = 25$ (halfway up the box), $z = 50$ (at the top of the box) and $y = 0$. Outside of the regions with initial polarity features and directly above the driven area, the density is largely unchanged and remains at the initial value of 1.

Lastly, Figure 5.23 shows contours of the magnitude of the velocity at $z = 25$, halfway

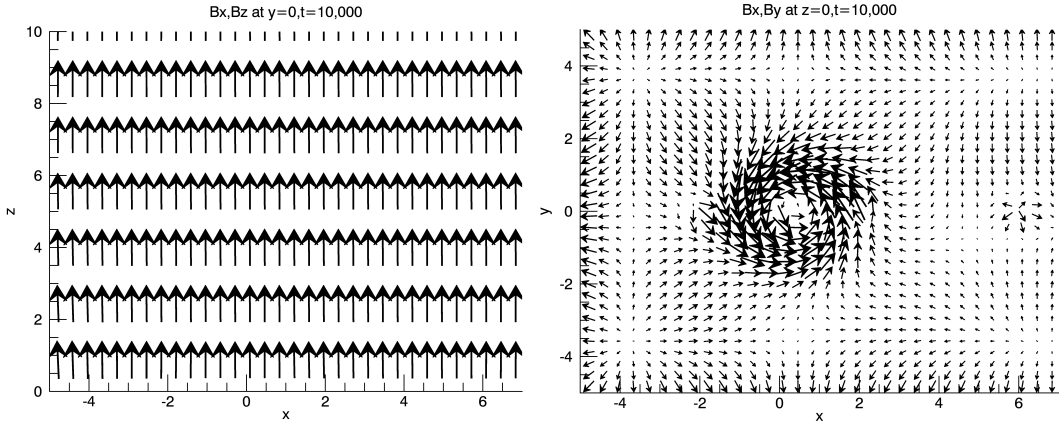

Figure 5.21: (a) B_x, B_z field through y mid plane at $t = 10,000$. (b) B_x, B_y field on $z = 0$ at $t = 10,000$.

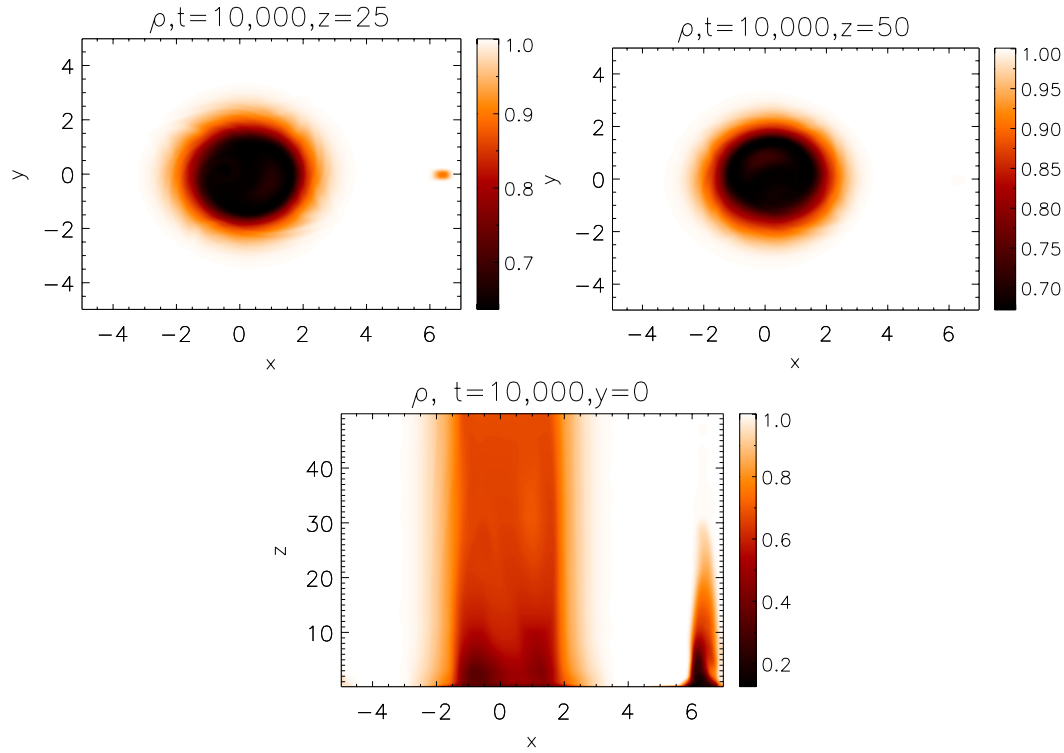

Figure 5.22: Density plots for $t = 10,000$. Plot (a) is the xy plane lying halfway up the numerical box, (b) is at the top of the box, and (c) is along $y = 0$.

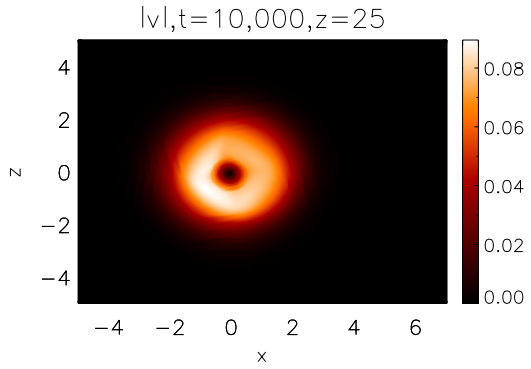

Figure 5.23: Velocity magnitude for $t = 10,000$. There is very little happening outside of the driver influence.

up the box. Again we can see that beyond the direct action of the driven region, velocity is unchanged from its initial value of zero. It seems that really, not much is happening in the system relative to the activity we saw before.

We ran this simulation to the intended end point of $45,000$ time units and found that it continued in the same way, with very little activity. Figure 5.24 shows the density and magnetic field strength in cuts through the domain at the end of the runtime. The density profile has not changed much from that at $t = 10,000$, and the magnetic field strength is looking fairly uniform.

It seems that here the increased plasma pressure experiments have stopped timestep failure, but we simply do not see much activity. We considered that maybe the technique made competing forces too well balanced. After a relaxation of the initial field the only changes we see are from numerical error build up, which we may not even have noticed if other things were happening in the plasma. On the other, more likely hand, there remained the fact of the non-zero divergence and its impact; this is in fact probably the fundamental flaw.

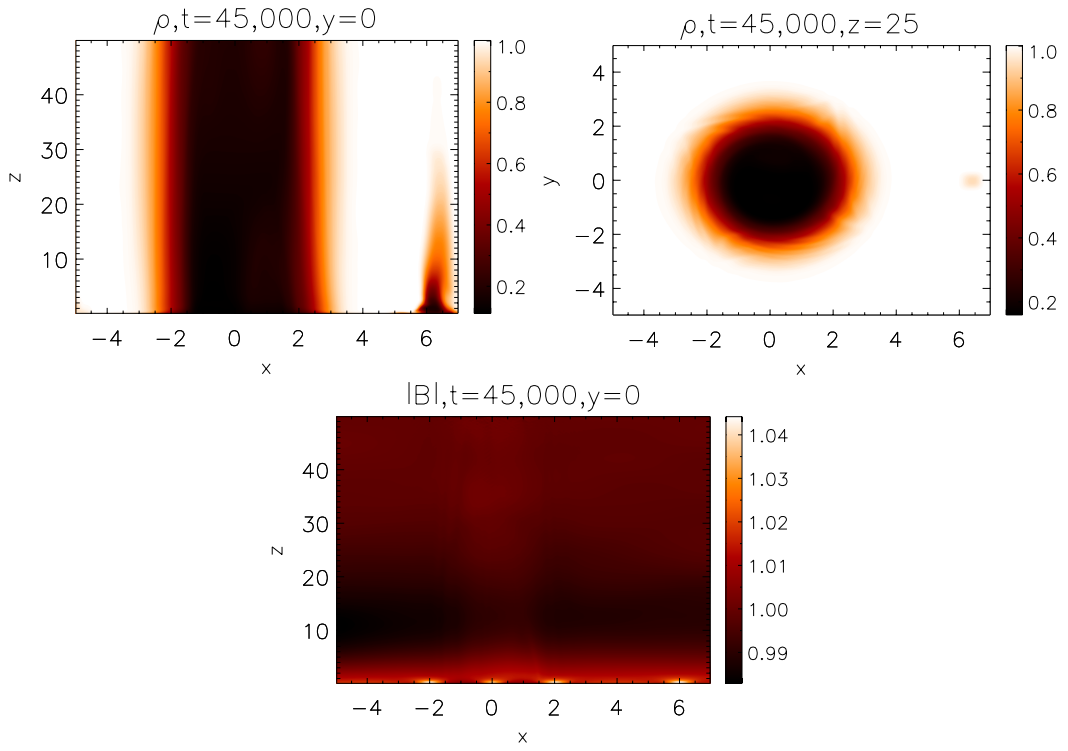

Figure 5.24: Density, (plots (a) and (b)) and $|B|$ (plot (c)) for $t = 45,000$. The magnetic field is quite smooth and the density profile looks at lot like it did at $t = 10,000$.

5.3 Magnetic Carpet by Scalar Potential

First Attempt

To remove any doubt of the exact role of the finite divergence, we return to the second option - giving the code an initial field which does automatically have zero current density. This should remove the initial Lorentz forces to stop timestep collapse, while avoiding the change in plasma pressure too. If we can run a simulation with steady progress and still see very little activity, then the non-zero divergence must destroy the null structures and stop any further formations arising.

The problem we have is that the numerical curl of the discretised analytic field we give to Lare3d is not zero. Other codes, such as the Pencil code, appear to avoid this problem. The Pencil code (available at `pencil-code.nordita.org`, another MPI MHD code) accepts a vector potential as the initial condition and calculates the numerical curl of this to obtain $\mathbf{B}$. This guarantees a divergence free field. Further, this code's calculations using the numerically calculated field rather than a discretised analytical field appear to be more accurate and free of the unintentional currents. The authors of Paper [19] used this code for their own investigations, of a field containing the same parasitic polarities we have, and were successful, free of the unexpected complications we have encountered with Lare3d.

Based on the form of the calculation of $\mathbf{j}$ in Lare3d's predictor-corrector step (note that the staggered nature of the magnetic field components mean there is more than one method to calculate current!), a scalar potential ψ was created such that the curl of $\mathbf{B} = -\nabla\psi$ is, by definition, zero, and we should have no troublesome initial currents or Lorentz forces to worry about.

Unfortunately, in practice this did not work either; the timestep crash returned.

Digging deeper into the code, we saw that this was due to the fact that we can calculate current in different ways and define its components at different places on each cell. For our illustrations of Lorentz forces we calculated the current density and interpolated both it and the magnetic field such that all components were defined at the centre of each grid cell, to make the calculation of $\mathbf{j} \times \mathbf{B}$ simple, and this was fine in order to get a big picture, qualitative view of what was happening. It turns out however that Lare3d not only calculates the current in a different way during the predictor-corrector step (i.e. the method for which ψ was created), but also calculates a further, different version of the current in other parts of the code which take into account resistive effects. In other words, the predictor-corrector step updates variables ideally, calculating one current density - let us call this $\mathbf{j}_i$. Then, if resistive MHD is set, further routines are called where another version of the current is calculated and used in other variable updates. Let us call this current $\mathbf{j}_r$. Using the carpet-by-ψ method, zero currents were satisfied in ideal parts of the code but not in the resistive parts. Therefore the problem remained.

Second Attempt

Since the code calculates different methods, there is no way to design a scalar potential which satisfies both $\mathbf{j}_i = 0$ and $\mathbf{j}_r = 0$. Instead, we try to manipulate the code to create a new starting point, using the carpet-by-ψ field as part of a relaxation process.

Starting again from the carpet-by-ψ field which we know satisfies $\mathbf{j}_i = 0$, we ran a non-driven resistive simulation. We overwrote the density at each timestep to reset the value at all points to their initial starting value of 1, and did the same for the specific internal energy density ε, which we set back to 0.01. These steps keep the timestep up and remove heating and the increasing magnetic energy error. This process was allowed to run until things seemed stable, with the topology and magnetic energy still very close to the original parasitic polarity field. This took around $6,200$ time units. We then tried using this as the initial state for new driven runs.

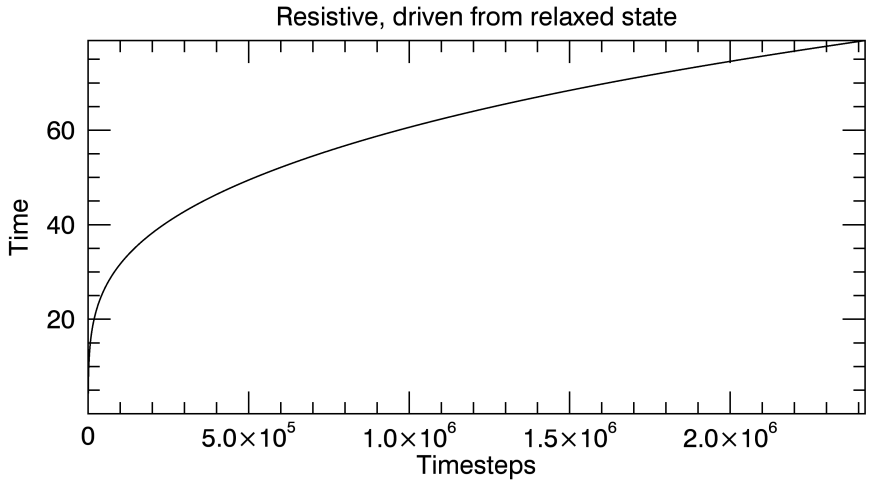

Figure 5.25: Time progression for the relaxed state attempt. This driven, resistive simulation under driver 2A encounters the same problem as before - a continuous decrease in timestep.

Unfortunately, the old problem of timestep fall returned. It did not stall completely, but slowed to the point that it took over $2,000,000$ time steps to cover around 75 time units as shown in Figure 5.25. Clearly this is still impractical. The relaxation process avoided a timestep fall by overwriting the density and pressure at each step, but this did not necessarily mean the original problem of $\mathbf{j}_R$ went away and furthermore, while $\mathbf{j}_i$ was zero in the carpet-by-ψ field, the relaxation process may lead to it becoming non-zero. Therefore as soon as we have stopped overwriting in the latest test, there is nothing to stop the currents having an effect and the density has been allowed to fall in places again, as can be see in Figure 5.26. Again the regions inside the dome structures have lost plasma density.

At this point, we must conclude that while there may be other ways to work around this problem, when undertaking further research with this type of scenario one must consider priorities. Once one becomes acquainted with a code it is of course useful to be able to continue using the same one where appropriate. Furthermore, in many cases when problems do arise it may be easier to find solutions if one is already quite familiar with the workings of the code. There may be other times however, when it is more

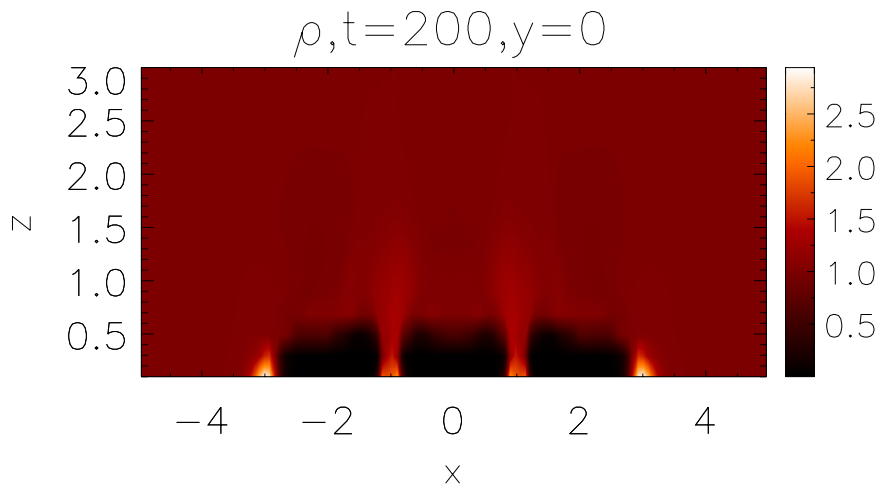

Figure 5.26: Density distribution along $y = 0$. The regions inside the original dome structures have become depleted, resulting in a slump in the time progression.

useful to investigate the possibility of a different code. For example if the problem in question is far from trivial, and the user's priority is obtaining results rather than developing robust numerical schemes. In our case, we had two problems - non-zero divergence and non-zero initial Lorentz forces. Giving the code a vector potential for the magnetic field to guarantee zero divergence does not guarantee zero initial currents. The non-zero initial currents lead to non-zero Lorentz forces and immediate runaway loss of density in some regions, causing the timestep to drop and the simulation to stall. Giving the code a scalar potential for the magnetic field does not result in zero initial currents in all parts of the code and so the simulation still stalls. Increasing plasma pressure to reduce the effect of the initial currents allows the simulation to progress since plasma density is no longer decimated in certain pockets, but the non-zero divergence now has the chance to fully come into play and destroy any potential activity. Our approach going forward would be to switch to a different code entirely.

Chapter 6

Conclusions and Future Work

6.1 Motivation and Starting Point

The purpose of this work was to investigate how, if at all, the nature of coronal loop evolution and heating depend on the type of motions taking place on the photosphere below. In Chapter 1 we discussed how the initial inspiration for this question began with Parker and the theory of topological dissipation (see [68]). This work posed that for a field anchored at two perfectly conducting parallel plates representing the photosphere, sufficient tangling by footpoint motions would lead to tangential discontinuities in the form singular current sheets between braided strands. Parker hypothesised there would be no smooth force free equilibria to relax to while maintaining ideal conditions. In a resistive environment, stresses built up by such braiding motions would lead to sufficiently small scale currents for magnetic reconnection to occur. This would provide heating in the form of free magnetic energy release as the field reconnects to a simpler, smooth state. Since photospheric motions are ubiquitous on the solar surface, this could be happening throughout the corona on a range of scales and could therefore

be responsible for a significant amount of heating. If so, this heating could be a contributing factor to the unexpectedly high temperatures observed in the solar corona, a mystery solar scientists have been trying to solve for decades.

Subsequent studies by many authors in the years since this theory was formed provide support for and against it. Most experiments, analytical and numerical, suggest that true tangential discontinuities do not form and there may exist smooth force-free equilibria for an ideal field to relax to. However, many find that while truly singular currents may not form, sufficiently strong finite current layers do, leaving open the possibility that field line braiding could trigger reconnection events and provide heating in a resistive environment. Various numerical simulations have applied shear and rotational boundary motions to magnetic field set-ups in ideal and resistive environments. Some have taken initially braided fields and analysed the subsequent behaviour. Key research was carried out in papers such as [112], [115] and [82], where the authors examined a field containing pairs of twists of opposite sense, such that the straightened coronal loop being considered had braiding but no net twist. In a series of ideal and resistive experiments, while no truly singular currents formed, current layers did and reconnection occurred for non-zero η. Furthermore, increasingly well resolved observations are beginning to provide better visual evidence of tangled field, for example [21]. The idea behind this work was to design a driving function which would induce braiding in the fashion of the initial condition used in the aforementioned papers. The function would have adjustable parameters which would allow us to create multiple drivers with different, measurable characteristics. We would then apply each driver to a magnetic field and analyse the particular effect of each, in order to ascertain how different types of motions could bring about reconnection and heating.

The driving motion would be based on a pair of blinking vortices - one vortical motion would occur at a time, alternating between the two. They would have an exponential form and be slow relative to the Alfvén speed. Initially we considered each vortex

spinning with opposing circulation acting on the lower z boundary of a uniform field representing a straightened coronal loop. So, one vortex would turn anticlockwise, twisting the field above it, then stop, and then the second vortex would spin in a clockwise direction, contributing the same amount of opposite twist and creating a field with braiding but no net twist. We then created two key categories from this basic idea.

6.2 Drivers and Categories

The first question was where to place the blinking vortices. The closer together we place them, the more the action of one will interfere with that of the other. A thought experiment tells us that different degrees of braiding will occur depending on this. If two oppositely twisting vortical motions are centred at the same place, then in an idealised environment where no relaxation at all can occur, the action of the second vortex would completely undo the braiding brought about by the first. The net result would be an unchanged field and no tangling of any degree. Similarly, if we were to place each at an infinite distance from each other, the action of one would not affect the action of the other and we would not see complex tangling between fieldlines affected by the different vortices. Between these two cases the braiding will be more complex with fieldlines crossing over and under each other and the braiding done and undone to different degrees, creating varyingly complex tangling. We needed a way to quantify this, and chose the concept of topological entropy, as described in Chapter 2.

Topological entropy is a measure of complexity. Mathematically it approximates the rate at which a material line length grows in some flow due to the ability of that flow to mix. In practice it is not always possible to calculate the true entropy, but instead estimate it using braids. So for example, a stirring protocol in a 2D fluid will move

seed points, initially at different coordinates, along different trajectories, and those trajectories will create braids when considered in time. The braid representing these motions can be applied to a random material line, and the rate of stretching of this line under the braid action tells us how complex the motions are. The more initial seed points we take the closer we get to the true entropy value. We carried out these calculations using the MATLAB code 'braidlab' (see [103] and [104]).

To create different driver functions, one vortex was held fixed centred on the origin $(0,0)$ and the second was placed at $(x_0,0)$, where we varied x_0 between 0 and 4 in intervals of 0.25. This analysis showed us that the complexity versus distance between vortex centres followed a roughly Gaussian shape, as predicted.

We also decided to examine a second set of drivers whereby the vortices would twist in the same direction. This type of motion would be less complex, as fieldlines are twisting with the same sense and forming a coherent structure. It too reached a peak complexity as x_0 increased before dropping off for larger values, but was generally of lower entropy than the first case.

The introduction of the second set of drivers also created another way to categorise them: the concept of magnetic helicity. Helicity is the measure of twist in a magnetic system. The oppositely twisting cases would inject equal and opposite amounts, and in an ideal case, have no net helicity. Where reconnection can take place, the twist in the overlying field may not cancel entirely due to changes in topology over the course of a driving period, but the driver action considered in isolation would contribute zero net helicity. The drivers with equally spinning vortices will always be injecting more helicity as the field is twisted the same way. Varying x_0 changes the tightness, or coherence, of the twist. Therefore the lower complexity case is also referred to as the high helicity case or the coherent case. In summary then, having carried out this analysis in Chapter 3, we chose three different values of x_0 for both the equal and

Run	Circulation	Helicity	x_0	Entropy (95% C.I)	Category
1A	Opposite	Zero	0.5	0.161 ± 0.004	Complex
1B	Opposite	Zero	1	0.786 ± 0.017	Complex
1C	Opposite	Zero	2	1.321 ± 0.028	Complex
2A	Equal	65.14	0.5	0.265 ± 0.104	Coherent
2B	Equal	65.14	1	0.444 ± 0.046	Coherent
2C	Equal	65.14	2	0.888 ± 0.039	Coherent

Table 6.1: Entropy calculations with 95% confidence intervals for each driver and their characteristics.

opposite cases, giving us six different drivers falling into categories as detailed in Table 6.1

At this stage we also described the types of magnetic field we planned to set the drivers to act on. The first set of experiments took a uniform field, $\mathbf{e}_z$, representing a straightened coronal loop. For the second, we wanted to investigate the effect on a magnetic carpet field - a field with magnetic features like those observed on the solar surface. The magnetic carpet comprises many different structures on different scales, creating a complex topology near the photosphere. Null points and their associated separatrix surfaces and separator fieldlines are thought to be common features. Our magnetic carpet setup featured parasitic polarities, consisting of multiple occurrences of a null point sitting within a small (relative to the height of the domain) dome-like structure which intersects the surface, in a background field of $\mathbf{e}_z$.

6.3 Uniform Field Experiments

Once the driver functions were finalised we incorporated them into the MPI code, Lare3d [2]. This is a staggered grid code which solves the full resistive 3D MHD equations and is able to deal with shocks through artificial viscosity.

The following simulations are analysed in Chapter 4. We applied the drivers to the

lower z boundary continuously, throughout the runtime of the simulations. We wanted to investigate the long term effects of such activity and so ran each simulation for $45,000$ time units. The domain extended over $[-5,7]\times[-5,5]\times[0,50]$. The resistivity is included by way of inverse Lundquist number, which we set to 0.0005. The upper z boundary was closed and the x and y were boundaries periodic.

6.3.1 Case 1

Case 1 includes three simulations with the zero helicity, high complexity, oppositely twisting vortex drivers. We saw that many properties were ordered with complexity. Run 1C, where $x_0 = 2$, had the highest complexity and turned out to inject the highest levels of magnetic energy, followed by the second highest complexity run 1B, and lastly the lowest complexity run 1A. Runs 1B and C in particular appear to settle into a statistically steady sates after $30,000$ time units. All three profiles showed a range of energy release events throughout each simulation, corresponding in real terms to picoflares and nanoflares, lending support to the idea that the continuous braiding of a field could contribute to a background level of heating in the solar atmosphere. The levels of kinetic energy are ordered in the opposite direction, with the highest complexity run settling into an average level around half that of the other two. The subsequent examination of current levels and structures also showed that the highest complexity run induced more and higher currents lying in finite layers through the domain, and the corresponding levels of both viscous and Ohmic heating show more heating coming from this run also. Looking at isosurfaces, fieldlines and contours of current before and after release events shows evidence of reconnection, as current levels decrease and fill less volume across the drop. This also all comes together to point to the conclusion that the more complex the driving motion, the more tangled the overlying field becomes, injecting more free magnetic energy and leading to more build up of current layers

and Ohmic heating as a result of reconnection events. We see a correlation of heated layers with layers of localised current. It seems also that in this case more kinetic energy is dissipated as viscous heating, leading to this run having lower kinetic energy and highest viscous heating. The most complex driver leads to smaller scales in the magnetic and velocity fields. The main point then is that complexity of photospheric motions may be a key factor in the formation of finite current layers and opportunities for heating.

6.3.2 Case 2

We then turned to the low complexity, high helicity, equally twisting vortex drivers. Keeping everything else as before we applied these to the lower z boundary of a uniform field. We again saw a specific ordering of properties, but this time helicity and coherence of twist appeared to be the dominant factors. The case with the most tightly, coherently twisted field, run 2A (which also had the lowest complexity), induces the highest levels of free magnetic energy, and also triggers the most dramatic features in the form of large release events. Run 2B also exhibits these features, but they are generally smaller. Run 2C, with the least tightly twisting driver, does not result in discharge events on these larger scales. Each of the three runs again show many smaller magnetic energy release events indicating nanoflare type occurrences. The largest event, taking place during Run 2A, releases 93.9% of the free magnetic energy, corresponding to something verging on a microflare. It seems that in the runs with vortices closest together, where the overlying field is twisted most tightly, instability has the chance to take hold. We pose that the large events are due to the triggering of a tearing mode instability after sufficient twisting of the field. Vector plots illustrate sheared fields running in opposite directions on either side of $x \approx 0$, a configuration prone to tearing modes. We also observe spikes and small bumps in the kinetic and internal energies

and viscous and Ohmic heating at times when these instabilities take place, supporting the theory. Again we see an ordering of these other quantities, but this time it is the closest placed vortex runs which exhibit the highest levels across the board. We examined current values and structures again, seeing that the highest magnetic energy run also contained currents which were the largest and most volume filling. Fieldline plots over the course of the largest drop in run 2A show the degree of tangling reducing as time goes on. Fieldlines evolve from criss-crossing through each other and the box to lying in simpler, helical patterns. Contours and isosurface plots of currents around these same times illustrate how ribbons of higher valued currents flowing through the domain at the beginning of the event have decreased in value and number by the end. Heating is shown to correspond to regions with current layers. In this lower complexity experiment it is the driver's ability to trigger instabilities which is the defining characteristic.

Comparing between the two Cases, it appears that coherent motions are effective at inducing both higher general levels of Ohmic heating through reconnection and larger release events. The lowest energy and heating run of the coherent cases, 2C, still has higher levels of magnetic energy, heating, current levels and volume filling levels than run 1A, the complex case run with lowest values. Therefore, in scenarios with zero net twist, complexity of driving is the key factor, while in high helicity environments it is the degree of twist itself that dictates reconnection and heating level.

6.3.3 Power Spectrum Analysis

Next in Chapter 4 we considered the power spectra of each simulation. This technique uses Fourier transforms to break signals, or in our case energy data, into the sum of their constituent sinusoidal parts. We transformed the magnetic and kinetic energies in space, therefore the process binned these data into the different spatial scales present

in the domain. Different fractions of the overall energy lie in different spatial scales, ranging from the grid cell volume up to the volume of the whole domain. Plotting this with inverse length scale, which we call k, gives a 'power spectrum'. Examining power spectra provides a picture of the way energy is distributed through a system and tells us if there is turbulent motion. Turbulence describes a plasma whereby eddies and vortical motions induce smaller eddies and vortical motions, transferring energy down into smaller regions until the spatial scales become small enough for viscous effects to come into force and dissipate the energy. Power spectra can often be divided into three sections: the smallest k values, corresponding to the largest length scales, where external forcing, such as our driver motions, dictate behaviour; the largest k values, corresponding to the smallest spatial scales where viscous effects are dominant; and an intermediate range, the inertial range, where a turbulent flow will follow a power law ck^a with a particular value of index a. This value has been calculated by many different authors. Values of $-5/3$ (the Kolmogorov value) and $-3/2$ (the Iroshnikov-Kraichnan value) have been posed, appearing in transforms of the magnetic and kinetic energies. Often $-3/2$ is calculated in 2D transforms of a field dominated by one component, in the plane perpendicular to this component. We calculated 3D transforms of magnetic and kinetic energy and some 2D transforms also.

Once we had transformed the magnetic and kinetic energies into k-space, we integrated the energy densities over k-space using shell integration (for 3D transforms) and ring integration (for 2D transforms). We construct k-space using the k-vectors corresponding to the Cartesian space vectors. This space is then divided up into shells or rings spanning some radial distance $i \cdot \mathrm{d}k$, where i is the number of shells or rings needed to take in all the data. Each transformed energy value lies at some radius from the k-space origin and falls within one particular shell or ring, and so counts towards the total energy contained on that particular scale.

We focused on times during the statistically steady state section for the complex runs

and around large drops in the case of the coherent runs. None of the spectra exhibited inertial range gradients of either Kolmogorov or IK values. However, each spectrum was interrupted during the inertial range by a sudden drop at a characteristic k-value of around $k = 2.5$. The magnetic energy profile for driver 1A looked promising for turbulence within the range $k = 1..2.5$, before being disrupted by the drop. Looking into the reasons for this we concluded that this drop is due to the non-uniform nature of the grid resolutions. The 2D spectra showed that the drop only occurred when transforms along the z direction were involved. All Cartesian space arrays had the same number of elements but different extents, such that $\mathrm{d}z > \mathrm{d}x > \mathrm{d}y$, therefore $\mathrm{d}k_3 < \mathrm{d}k_1 < \mathrm{d}k_2$. Although $\mathrm{d}x$ and $\mathrm{d}k_1$ are still close to $\mathrm{d}y$ and $\mathrm{d}k_2$, $\mathrm{d}z$ and $\mathrm{d}k_3$ are sufficiently different such that when k-shells or k-rings are constructed, many of them extend pass the maximum k_3 value, at which point these shells have large, empty regions. The value k_3 varies between -2.56 and 2.54, while k_2 extends from -12.8 to 12.7. Once we move into shells with radii greater than about 2.5 we have no contributions at all from sections of the shell. This theory is supported by the fact that 2D transforms in x and y where the k-grid is much more uniform do not suffer from this interruption. Furthermore, a test with a uniform grid produced smooth profiles, while a test with an even smaller extent in k_3 did result in a power spectrum with a drop, at the predicted value of $k = 1.7$.

This all means that it is currently difficult to say if turbulence is present or not. Future work could include using a method of integration which does not involve placing uniform shells or circles over a highly non-uniform grid. If we could sum up energy contributions in such a way that many of the k-space bins are not largely empty by design, we may get a clearer picture of whether turbulence occurs in these experiments. If carrying out similar investigations in the future we would seek to have similar resolutions in each dimension to avoid the problem altogether.

6.4 Magnetic Carpet Experiments

In Chapter 5 we applied the same six driving motions to an initial condition designed to have features of magnetic carpet field. We hoped to establish whether the different types of footpoint motions could, as with the uniform field, induce specific types of behaviours when acting on a non-trivial, more realistic atmospheric field.

These simulations used the same physical domain as before, with the same initial parameter values as before. Initially we set the upper z boundary to be open, but also experimented with a closed upper boundary and closed/periodic x and y boundaries, but these changes were shown not to make any difference to the subsequent problems.

Each of the six simulations suffered from a fatal decrease in timestep, to the point where progress slowed effectively to a halt. The timestep at each iteration of the code depends on the CFL condition, which places restrictions on the timestep in order to ensure that the code has access to all necessary information at each iteration. Therefore initial attempts to solve this issue involved changes which would directly affect the CFL condition, for example slowing driver speeds, changing the resolution, changing η, turning off driving and resistivity, and changing the artificial viscosity. None of these resulted in a big enough improvement.

We moved on to considering the discretisation error involved in translating the analytic expression for $\mathbf{B}$ to the numerical grid. We discovered that the numerical divergence was non-zero. Although unsatisfactory, this still did not seem to be the problem with the timestep. We also had a non-zero current density, for a field which was designed to be potential. This led to non-zero Lorentz forces, acting in particular within the separatrix dome structure, down and out underneath the null point from the very beginning of the simulations, as illustrated when we took a simpler, one parasitic polarity field to investigate further. This led to flows of plasma away from this region, depleting it

and leaving a very low density volume. As the minimum density becomes smaller, the timestep drops accordingly, and the simulation stalls.

The simplest way to work around this was to reduce the influence of these unintended Lorentz forces by increasing the plasma β, by increasing the initial specific internal energy density. Trial and error gave us a new value of ε which prevented the density drop by counteracting the influence on the plasma of the Lorentz force. The one polarity simple test field was shown as a result to be able to evolve past the point where the initial Lorentz forces were acting uniformly away from the problem area, and the simulation could continue hopefully without encountering the same issue.

Having set up all six simulations again as before with the new ε, all six made good progress. However, looking at the data showed that there seemed to be very little activity of any kind. For each of the runs, after an initial burst of activity where the parasitic polarity structures were destroyed, the energy diagnostics remained very unchanging. Small oscillations were seen in the kinetic energy in particular, corresponding to the driver. The magnetic energy however, decreased by about 15 units and then very slowly but consistently increased for the rest of the simulation. Switching driving off revealed an unexpectedly similar picture. The magnetic energy once again decreased and then began increasing slowly, even though in this non-driven scenario there should be nothing to cause this.

The increase in ε, while preventing the timestep fall, did not capture any interesting behaviour. We contemplated that this could be due to the fact of balance in the forces themselves, or more likely the still non-zero divergence annihilating the structure and preventing any more features forming. The increase in magnetic energy observed even without driving is due to numerical error, which shows up because there is not really any other activity. The more timesteps a simulation took to reach a certain time, the larger the increase in magnetic energy by that time, indicating that this is an error

which builds up with each iteration. It may be something that would not even have been picked up on if there was other activity taking place to mask it.

We made an attempt to re-write the initial condition such that the numerical error in the current responsible for the timestep fall and magnetic energy increase did not occur. A scalar potential, ψ, was constructed and the magnetic field calculated from this within the code, so that $\mathbf{j} = \nabla \times \mathbf{B} = -\nabla \times (\nabla \psi)$ is identically zero. If we could run a simulation in this way and still see no activity after the initial structures were destroyed, then this would be a final confirmation that the non-zero divergence is a fundamental problem when using this code in this scenario. Unfortunately the timestep problem returned.

It was then discovered that this is due to the fact that Lare3d calculates currents in different ways in different parts of the code. The function ψ was designed such that currents in the ideal parts of the code, $\mathbf{j}_i$, are zero. However, the code calculates a second current density, $\mathbf{j}_R$, in the resistive subroutines, and so error arises in this way still and we see the timestep drop again. Since it is not possible to design a scalar potential which will give rise to zero currents by both methods, we attempted another approach.

If we could somehow let the initial currents dissipate and the field relax to a stable state with the topology mostly preserved, we may be able to use this an initial condition instead. Using the scalar potential ψ so $\mathbf{j}_i$ is zero, we ran a non-driven resistive simulation whereby we over-wrote the density and the specific internal energy density to their initial values at each step. This removed any heating due to $\mathbf{j}_R$ and reinstated a uniform plasma density. We ran this relaxation until we were sure no magnetic energy increase or timestep decrease due to current errors was going to occur. We took this relaxed state as the initial condition in a new driven simulation with ε set back to the

lower value we used before. Unfortunately, it appeared that the initial Lorentz forces continued to be large enough to cause problems, and once we stopped overwriting, there was nothing to undo the damage of the initial $\mathbf{j}_R$. The same problems returned.

Other work, for example that of [19], investigated systems containing the same type of parasitic polarity under the influence of driving motions, but using the Pencil code, an alternative MHD code. This code handles variables in such way that this work did not run into these issues. Future work could continue to attempt to find a solution to the problems encountered using Lare3d, or consider a different code. In general, in future it would be instructive to think carefully about which codes are most suited to the problem one is investigating, since it may be that some will handle certain things more easily than others. Furthermore, in the event that one does run into significant issues with a code, one should be prepared to consider an alternative approach in terms of the numerical scheme employed.

6.5 Summary

We have examined the effect of different types of photospheric motion on a coronal loop and concluded that both motion complexity and ability to inject helicity have an impact on the evolution of the overlying field. When the driving does not inject any net helicity, it is the complexity which dominates the energy and heating levels observed. The most complex blinking vortex driver set-up was shown to inject the highest levels of free magnetic energy, create larger valued and more volume filling currents, and result in the highest levels of Ohmic heating. It also led to smaller scales in velocities, creating further heating in the form of viscous effects. In the case of the equally

twisting vortices with lower complexities and positive helicity injection, the level of coherence was the key factor. The simulation with the closest placed vortex centres, while the lowest complexity of the three, twisted fieldlines most tightly and led to the highest levels of magnetic energy, current formation and heating, as well as inducing the largest energy release events. Both this and the next most coherent run were found to trigger two instances of tearing mode instability. The least coherent and most complex was not able to bring about events of this size, which were also registered as spikes in kinetic and thermal energy and heating. We suggest that when considering similar driven simulations one should consider the properties of complexity and helicity injection and keep in mind that results can fundamentally differ depending on what driving profile is employed. Furthermore, we judge that it is important to use full MHD in this type of research as we have found activity, such as repeated instabilities, which have not previously been captured using reduced MHD.

We were unable to reach a conclusion on whether or not a turbulent cascade occurred in any of the runs: the power spectra were interrupted by a characteristic drop, seen for all drivers, and appearing to be due to the non-uniform gird resolution meaning that B_x and B_y are defined beyond the spatial scales of B_z and so energy shells past a certain scale are partially empty. In future we would seek to set grid spacing in each dimension to be as close as possible in order to avoid this problem again.

Finally, our experiments with a magnetic carpet field containing null points and separatrices were hampered by the discretisation error and low order difference scheme of the code causing non-zero magnetic divergence and and currents in the initial condition. The initial currents generate Lorentz forces which diminish plasma density in some regions to a point where the code stalls. Lare3d calculates current in two different ways in the code and so we could not design a scalar potential for the initial magnetic field such that currents were identically zero for both methods. Reducing the influence of the Lorentz forces by increasing plasma pressure allows the code to run

longer, however the non-zero divergence destroys the field structure and prevents any further feature formation occurring. The best solution to this type of problem in future work would be to use a vector potential initial condition in a higher order differencing scheme code, which would prevent numerical error entirely dominating the evolution, as encountered in this work.

Bibliography

[1] R. L. Adler, A. G. Konheim, and M. H. McAndrew. Topological Entropy. *Transactions of the American Mathematical Society*, 114:309–319, 1965.

[2] T. D. Arber, A. W. Longbottom, C. L. Gerrard, and A. M. Milne. A Staggered Grid, Lagrangian-Eulerian Remap Code for 3-D MHD Simulations. *Journal of Computational Physics*, 171:151–181, July 2001.

[3] M. J. Aschwanden, T. D. Tarbell, R. W. Nightingale, C. J. Schrijver, A. Title, C. C. Kankelborg, P. Martens, and H. P. Warren. Time Variability of the "Quiet" Sun Observed with TRACE. II. Physical Parameters, Temperature Evolution, and Energetics of Extreme-Ultraviolet Nanoflares. *The Astrophysical Journal*, 535:1047–1065, June 2000.

[4] G. Aulanier, P. Démoulin, C. J. Schrijver, M. Janvier, E. Pariat, and B. Schmieder. The standard flare model in three dimensions. II. Upper limit on solar flare energy. *Astronomy and Astrophysics*, 549:A66, January 2013.

[5] L. Balmaceda, S. Vargas Domínguez, J. Palacios, I. Cabello, and V. Domingo. Evidence of small-scale magnetic concentrations dragged by vortex motion of solar photospheric plasma. *Astronomy and Astrophysics*, 513:L6, April 2010.

[6] Luis Barreira. *Metric Entropy and Topological Entropy*, pages 107–146. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[7] M. A. Berger. Introduction to magnetic helicity. *Plasma Physics and Controlled Fusion*, 41:B167–B175, December 1999.

[8] M. A. Berger and G. B. Field. The topological properties of magnetic helicity. *Journal of Fluid Mechanics*, 147:133–148, October 1984.

[9] D. Biskamp. *Nonlinear magnetohydrodynamics*. Cambridge University Press, 1993.

[10] D. Biskamp. *Magnetohydrodynamic Turbulence*. Cambridge University Press, July 2003.

[11] J. A. Bonet, I. Márquez, J. Sánchez Almeida, I. Cabello, and V. Domingo. Convectively Driven Vortex Flows in the Sun. *The Astrophysical Journal Letters*, 687:L131, November 2008.

[12] P. Boyland. Exponential growth in two-dimensional topological fluid dynamics. *Procedia IUTAM*, 7:107–116, October 2013.

[13] P. Boyland, M. Stremler, and H. Aref. Topological fluid mechanics of point vortex motions. *Physica D Nonlinear Phenomena*, 175:69–95, January 2003.

[14] P. L. Boyland, H. Aref, and M. A. Stremler. Topological fluid mechanics of stirring. *Journal of Fluid Mechanics*, 403:277–304, January 2000.

[15] A. Brandenburg and K. Subramanian. Astrophysical magnetic fields and nonlinear dynamo theory. *Physics Reports*, 417:1–209, October 2005.

[16] P. N. Brandt, G. B. Scharmer, S. Ferguson, R. A. Shine, and T. D. Tarbell. Vortex flow in the solar photosphere. *Nature*, 335:238–240, September 1988.

[17] D. S. Brown, R. W. Nightingale, D. Alexander, C. J. Schrijver, T. R. Metcalf, R. A. Shine, A. M. Title, and C. J. Wolfson. Observations of Rotating Sunspots from TRACE. *Solar Physics*, 216:79–108, September 2003.

[18] S. Candelaresi, D. I. Pontin, and G. Hornig. Magnetic Field Relaxation and Current Sheets in an Ideal Plasma. *The Astrophysical Journal*, 808:134, August 2015.

[19] S. Candelaresi, D. I. Pontin, and G. Hornig. Effects of Field-line Topology on Energy Propagation in the Corona. *The Astrophysical Journal*, 832:150, December 2016.

[20] P. Charbonneau. Dynamo Models of the Solar Cycle. *Living Reviews in Solar Physics*, 7, September 2010.

[21] J. W. Cirtain, L. Golub, A. R. Winebarger, B. de Pontieu, K. Kobayashi, R. L. Moore, R. W. Walsh, K. E. Korreck, M. Weber, P. McCauley, A. Title, S. Kuzin, and C. E. Deforest. Energy release in the solar corona from spatially resolved magnetic braids. *Nature*, 493:501–503, January 2013.

[22] R. M. Close, C. E. Parnell, D. W. Longcope, and E. R. Priest. Recycling of the Solar Corona's Magnetic Field. *The Astrophysical Journal Letters*, 612:L81–L84, September 2004.

[23] R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen Differenzengleichungen der mathematischen Physik. *Mathematische Annalen*, 100:32–74, 1928.

[24] T. G. Cowling. *Solar Electrodynamics*, page 532. University of Chicago Press, 1953.

[25] I. J. D. Craig and A. D. Sneyd. The Parker Problem and the Theory of Coronal Heating. *Solar Physics*, 232:41–62, November 2005.

[26] C. Day. SOHO observations implicate 'magnetic carpet' as source of coronal heating in quiet Sun. *Physics Today*, 51:19–21, March 1998.

[27] B. De Pontieu, S. W. McIntosh, M. Carlsson, V. H. Hansteen, T. D. Tarbell, C. J. Schrijver, A. M. Title, R. A. Shine, S. Tsuneta, Y. Katsukawa, K. Ichimoto, Y. Suematsu, T. Shimizu, and S. Nagata. Chromospheric Alfvénic Waves Strong Enough to Power the Solar Wind. *Science*, 318:1574, December 2007.

[28] J. W. Dungey. Interplanetary Magnetic Field and the Auroral Zones. *Physical Review Letters*, 6:47–48, January 1961.

[29] I. Dynnikov and B. Wiest. On the complexity of braids. *Journal of the European Mathematical Society*, 009(4):801–840, 2007.

[30] I. A. Dynnikov. COMMUNICATIONS OF THE MOSCOW MATHEMATICAL SOCIETY: On a Yang-Baxter map and the Dehornoy ordering. *Russian Mathematical Surveys*, 57:592–594, June 2002.

[31] B. Edlén. An Attempt to Identify the Emission Lines in the Spectrum of the Solar Corona. *Arkiv för Matematik, Astronomi och Fysik*, 28B(1), 1941.

[32] G. Eyink, E. Vishniac, C. Lalescu, H. Aluie, K. Kanov, K. Bürger, R. Burns, C. Meneveau, and A. Szalay. Flux-freezing breakdown in high-conductivity magnetohydrodynamic turbulence. *Nature*, 497:466–469, May 2013.

[33] M. D. Finn and J.-L. Thiffeault. Topological Entropy of Braids on the Torus. *SIAM Journal on Applied Dynamical Systems*, 6:79–98, January 2007.

[34] H. P. Furth, J. Killeen, and M. N. Rosenbluth. Finite-Resistivity Instabilities of a Sheet Pinch. *Physics of Fluids*, 6:459–484, April 1963.

[35] K. Galsgaard and . Nordlund. Heating and activity of the solar corona 1. Boundary shearing of an initially homogeneous magnetic field. *Journal of Geophysical Research*, 101:13445–13460, June 1996.

[36] K. Galsgaard and . Nordlund. Heating and activity of the solar corona 1. Boundary shearing of an initially homogeneous magnetic field. *Journal of Geophysical Research*, 101:13445–13460, June 1996.

[37] C. L. Gerrard, T. D. Arber, A. W. Hood, and R. A. M. Van der Linden. Numerical simulations of kink instability in line-tied coronal Loops. *Astronomy and Astrophysics*, 373:1089–1098, July 2001.

[38] R. G. Giovanelli. A Theory of Chromospheric Flares. *Nature*, 158:81–82, July 1946.

[39] P. Goldreich and S. Sridhar. Toward a theory of interstellar turbulence. 2: Strong alfvenic turbulence. *The Astrophysical Journal*, 438:763–775, January 1995.

[40] D. O. Gomez, P. A. Dmitruk, and L. J. Milano. Recent theoretical results on coronal heating. *Solar Physics*, 195:299–318, August 2000.

[41] W. Grotrian. Zur Frage der Deutung der Linien im Spektrum der Sonnenkorona. *Naturwissenschaften*, 27:214–214, March 1939.

[42] I. G. Hannah, H. S. Hudson, M. Battaglia, S. Christe, J. Kašparová, S. Krucker, M. R. Kundu, and A. Veronig. Microflares and the Statistics of X-ray Flares. *Space Science Reviews*, 159:263–300, September 2011.

[43] D. L. Hendrix and G. van Hoven. Magnetohydrodynamic Turbulence and Implications for Solar Coronal Heating. *The Astrophysical Journal*, 467:887, August 1996.

[44] M. Hesse and K. Schindler. A theoretical foundation of general magnetic reconnection. *Journal of Geophysical Research*, 93:5559–5567, June 1988.

[45] A. W. Hood and E. R. Priest. Kink instability of solar coronal loops as the cause of solar flares. *Solar Physics*, 64:303–321, December 1979.

[46] F. Hoyle. *Some recent researches in solar physics.* Cambridge University Press, 1949.

[47] H. S. Hudson. Solar flares, microflares, nanoflares, and coronal heating. *Solar Physics*, 133:357–369, June 1991.

[48] P. S. Iroshnikov. Turbulence of a Conducting Fluid in a Strong Magnetic Field. *Astronomichesskii Zhurnal*, 40:742, 1963.

[49] . M. Janse and B. C. Low. The Topological Changes of Solar Coronal Magnetic Fields. I. Spontaneous Current Sheets in Three-Dimensional Fields. *Astrophysical Journal*, 690:1089–1104, January 2009.

[50] A. Kolmogorov. The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers. *Akademiia Nauk SSSR Doklady*, 30:301–305, 1941.

[51] A. G. Kosovichev. Advances in Global and Local Helioseismology: An Introductory Review. In J.-P. Rozelot and C. Neiner, editors, *Lecture Notes in Physics, Berlin Springer Verlag*, volume 832 of *Lecture Notes in Physics, Berlin Springer Verlag*, page 3, 2011.

[52] R. H. Kraichnan. Inertial-Range Spectrum of Hydromagnetic Turbulence. *Physics of Fluids*, 8:1385–1387, July 1965.

[53] R. B. Leighton. The Solar Granulation. *Annual Review of Astronomy and Astrophysics*, 1:19, 1963.

[54] A. W. Longbottom, G. J. Rickard, I. J. D. Craig, and A. D. Sneyd. Magnetic Flux Braiding: Force-free Equilibria and Current Sheets. *Astrophysical Journal*, 500:471–482, June 1998.

[55] D. W. Longcope and R. N. Sudan. Evolution and statistics of current sheets in coronal magnetic loops. *The Astrophysical Journal*, 437:491–504, December 1994.

[56] M. C. López Fuentes, J. A. Klimchuk, and C. H. Mandrini. The Temporal Evolution of Coronal Loops Observed by GOES SXI. *The Astrophysical Journal*, 657:1127–1136, March 2007.

[57] B. C. Low. The Parker Magnetostatic Theorem. *The Astrophysical Journal*, 718:717–723, August 2010.

[58] J. Mason, F. Cattaneo, and S. Boldyrev. Numerical measurements of the spectrum in magnetohydrodynamic turbulence. *Physics Review*, 77(3):036403, March 2008.

[59] S. W. McIntosh, B. de Pontieu, M. Carlsson, V. Hansteen, P. Boerner, and M. Goossens. Alfvénic waves with sufficient energy to power the quiet solar corona and fast solar wind. *Nature*, 475:477–480, July 2011.

[60] Z. Mikic, D. D. Schnack, and G. van Hoven. Creation of current filaments in the solar corona. *Astrophysical Journal*, 338:1148–1157, March 1989.

[61] L. J. Milano, D. O. Gómez, and P. C. H. Martens. Solar Coronal Heating: AC versus DC. *Astrophysical Journal*, 490:442–451, November 1997.

[62] J.-O. Moussafir. On computing the entropy of braids. *Functional Analysis and Other Mathematics*, 1:37–46, July 2006.

[63] S. Newhouse and T. Pignataro. On the estimation of topological entropy. *Journal of Statistical Physics*, 72:1331–1351, September 1993.

[64] C. S. Ng, L. Lin, and A. Bhattacharjee. High-Lundquist Number Scaling in

Three-dimensional Simulations of Parker's Model of Coronal Heating. *The Astrophysical Journal*, 747:109, March 2012.

[65] . Nordlund, R. F. Stein, and M. Asplund. Solar Surface Convection. *Living Reviews in Solar Physics*, 6:2, April 2009.

[66] J. Palacios, S. Vargas Domínguez, L. A. Balmaceda, I. Cabello, and V. Domingo. Multi-wavelength Observations of Photospheric Vortex Flows in the Photosphere Using Ground-based and Space-borne Telescopes. In I. Dorotovic, C. E. Fischer, and M. Temmer, editors, *Coimbra Solar Physics Meeting: Ground-based Solar Observations in the Space Instrumentation Era*, volume 504 of *Astronomical Society of the Pacific Conference Series*, page 139, April 2016.

[67] E. N. Parker. Sweet's Mechanism for Merging Magnetic Fields in Conducting Fluids. *Journal of Geophysical Research*, 62:509–520, December 1957.

[68] E. N. Parker. Topological Dissipation and the Small-Scale Fields in Turbulent Gases. *The Astrophysical Journal*, 174:499, June 1972.

[69] E. N. Parker. Equilibrium of magnetic fields with arbitrary interweaving of the lines of force. I - Discontinuities in the torsion. *Geophysical and Astrophysical Fluid Dynamics*, 34:243–264, 1986.

[70] E. N. Parker. Equilibrium of magnetic fields with arbitrary interweaving of the lines of force. II - Discontinuities in the field. *Geophysical and Astrophysical Fluid Dynamics*, 35:277–301, 1986.

[71] E. N. Parker. Nanoflares and the solar X-ray corona. *The Astrophysical Journal*, 330:474–479, July 1988.

[72] E. N. Parker. Spontaneous current sheets in magnetic fields : with applications

to stellar x-rays. *Spontaneous current sheets in magnetic fields : with applications to stellar x-rays. International Series in Astronomy and Astrophysics, Vol. 1. New York : Oxford University Press, 1994.*, 1, 1994.

[73] C. E. Parnell. Nature of the magnetic carpet - I. Distribution of magnetic fluxes. *Monthly Notices of the Royal Astronomical Society*, 335:389–398, September 2002.

[74] C. E. Parnell and I. De Moortel. A contemporary view of coronal heating. *Philosophical Transactions of the Royal Society of London Series A*, 370:3217–3240, July 2012.

[75] C. E. Parnell, C. E. DeForest, H. J. Hagenaar, B. A. Johnston, D. A. Lamb, and B. T. Welsch. A Power-Law Distribution of Solar Magnetic Fields Over More Than Five Decades in Flux. *The Astrophysical Journal*, 698:75–82, June 2009.

[76] C. E. Parnell, A. L. Haynes, and K. Galsgaard. Recursive Reconnection and Magnetic Skeletons. *The Astrophysical Journal*, 675:1656–1665, March 2008.

[77] C. E. Parnell and P. E. Jupp. Statistical Analysis of the Energy Distribution of Nanoflares in the Quiet Sun. *The Astrophysical Journal*, 529:554–569, January 2000.

[78] H. E. Petschek. Magnetic Field Annihilation. *NASA Special Publication*, 50:425, 1964.

[79] D. I. Pontin and G. Hornig. The Structure of Current Layers and Degree of Field-line Braiding in Coronal Loops. *The Astrophysical Journal*, 805:47, May 2015.

[80] D. I. Pontin, M. Janvier, S. K. Tiwari, K. Galsgaard, A. R. Winebarger, and J. W. Cirtain. Observable Signatures of Energy Release in Braided Coronal Loops. *The Astrophysical Journal*, 837:108, March 2017.

[81] D. I. Pontin, E. R. Priest, and K. Galsgaard. On the Nature of Reconnection at a Solar Coronal Null Point above a Separatrix Dome. *The Astrophysical Journal*, 774:154, September 2013.

[82] D. I. Pontin, A. L. Wilmot-Smith, G. Hornig, and K. Galsgaard. Dynamics of braided coronal loops. II. Cascade to multiple small-scale reconnection events. *Astronomy and Astrophysics*, 525:A57, January 2011.

[83] E. Priest. *Magnetohydrodynamics of the Sun*. Cambridge University Press, May 2014.

[84] E. R. Priest, C. R. Foley, J. Heyvaerts, T. D. Arber, J. L. Culhane, and L. W. Acton. Nature of the heating mechanism for the diffuse solar corona. *Nature*, 393:545–547, June 1998.

[85] E. R. Priest, J. F. Heyvaerts, and A. M. Title. A Flux-Tube Tectonics Model for Solar Coronal Heating Driven by the Magnetic Carpet. *The Astrophysical Journal*, 576:533–551, September 2002.

[86] E. R. Priest, D. W. Longcope, and J. Heyvaerts. Coronal Heating at Separators and Separatrices. *The Astrophysical Journal*, 624:1057–1071, May 2005.

[87] E. R. Priest and D. I. Pontin. Three-dimensional null point reconnection regimes. *Physics of Plasmas*, 16(12):122101–122101, December 2009.

[88] M. R. E. Proctor. Magnetoconvection: Solar convection and magnetic fields. *Astronomy and Geophysics*, 45(4):4.14–4.20, August 2004.

[89] R. Ramesh, K. Sasikumar Raja, C. Kathiravan, and A. S. Narayanan. Low-frequency Radio Observations of Picoflare Category Energy Releases in the Solar Atmosphere. *The Astrophysical Journal*, 762:89, January 2013.

[90] A. F. Rappazzo, M. Velli, and G. Einaudi. Shear Photospheric Forcing and the Origin of Turbulence in Coronal Loops. *The Astrophysical Journal*, 722:65–78, October 2010.

[91] A. F. Rappazzo, M. Velli, and G. Einaudi. Field Lines Twisting in a Noisy Corona: Implications for Energy Storage and Release, and Initiation of Solar Eruptions. *The Astrophysical Journal*, 771:76, July 2013.

[92] A. F. Rappazzo, M. Velli, G. Einaudi, and R. B. Dahlburg. Coronal Heating, Weak MHD Turbulence, and Scaling Laws. *The Astrophysical Journal Letters*, 657:L47–L51, March 2007.

[93] A. F. Rappazzo, M. Velli, G. Einaudi, and R. B. Dahlburg. Nonlinear Dynamics of the Parker Scenario for Coronal Heating. *The Astrophysical Journal*, 677:1348–1366, April 2008.

[94] F. Reale. Coronal Loops: Observations and Modeling of Confined Plasma. *Living Reviews in Solar Physics*, 11:4, December 2014.

[95] M. L. Ritchie, A. L. Wilmot-Smith, and G. Hornig. The Dependence of Coronal Loop Heating on the Characteristics of Slow Photospheric Motions. *The Astrophysical Journal*, 824:19, June 2016.

[96] K. Schindler, M. Hesse, and J. Birn. General magnetic reconnection, parallel electric fields, and helicity. *Journal of Geophysical Research*, 93:5547–5557, June 1988.

[97] J. T. Schmelz, S. H. Saar, K. Nasraoui, V. L. Kashyap, M. A. Weber, E. E. DeLuca, and L. Golub. Multi-stranded and Multi-thermal Solar Coronal Loops: Evidence from Hinode X-ray Telescope and EUV Imaging Spectrometer Data. *The Astrophysical Journal*, 723:1180–1187, November 2010.

[98] C. J. Schrijver, A. M. Title, K. L. Harvey, N. R. Sheeley, Y.-M. Wang, G. H. J. van den Oord, R. A. Shine, T. D. Tarbell, and N. E. Hurlburt. Large-scale coronal heating by the small-scale magnetic field of the Sun. *Nature*, 394:152–154, July 1998.

[99] S. Sridhar and P. Goldreich. Toward a theory of interstellar turbulence. 1: Weak Alfvenic turbulence. *The Astrophysical Journal*, 432:612–621, September 1994.

[100] R. F. Stein and . Nordlund. Simulations of Solar Granulation. I. General Properties. *The Astrophysical Journal*, 499:914–933, May 1998.

[101] P. A. Sweet. The Neutral Point Theory of Solar Flares. In B. Lehnert, editor, *Electromagnetic Phenomena in Cosmical Physics*, volume 6 of *IAU Symposium*, page 123, 1958.

[102] P. Testa, B. De Pontieu, J. Martínez-Sykora, E. DeLuca, V. Hansteen, J. Cirtain, A. Winebarger, L. Golub, K. Kobayashi, K. Korreck, S. Kuzin, R. Walsh, C. DeForest, A. Title, and M. Weber. Observing Coronal Nanoflares in Active Region Moss. *The Astrophysical Journal*, 770:L1, June 2013.

[103] J.-L. Thiffeault. Braids of entangled particle trajectories. *Chaos*, 20(1):017516, March 2010.

[104] Jean-Luc Thiffeault and Marko Budišić. Braidlab: A software package for braids and loops, 2013–2017. Version 1.0.4 or earlier.

[105] S. K. Tiwari, C. E. Alexander, A. R. Winebarger, and R. L. Moore. Trigger Mechanism of Solar Subflares in a Braided Coronal Magnetic Structure. *The Astrophysical Journal Letters*, 795:L24, November 2014.

[106] G. S. Vaiana, A. S. Krieger, and A. F. Timothy. Identification and Analysis of Structures in the Corona from X-Ray Photography. *Solar Physics*, 32:81–116, September 1973.

[107] A. A. van Ballegooijen. Force free fields and coronal heating part I. The formation of current sheets. *Geophysical and Astrophysical Fluid Dynamics*, 41:181–211, 1988.

[108] A. A. van Ballegooijen. Magnetic fine structure of solar coronal loops. In R. C. Altrock, editor, *Solar and Stellar Coronal Structure and Dynamics*, pages 115–124, 1988.

[109] A. A. van Ballegooijen, M. Asgari-Targhi, and M. A. Berger. On the Relationship Between Photospheric Footpoint Motions and Coronal Heating in Solar Active Regions. *The Astrophysical Journal*, 787:87, May 2014.

[110] G. E. Vekstein and R. Jain. Signatures of a nanoflare heated solar corona. *Plasma Physics and Controlled Fusion*, 45:535–545, May 2003.

[111] S. Wedemeyer-Böhm, E. Scullion, O. Steiner, L. Rouppe van der Voort, J. de La Cruz Rodriguez, V. Fedun, and R. Erdélyi. Magnetic tornadoes as energy channels into the solar corona. *Nature*, 486:505–508, June 2012.

[112] A. L. Wilmot-Smith, G. Hornig, and D. I. Pontin. Magnetic Braiding and Parallel Electric Fields. *The Astrophysical Journal*, 696:1339–1347, May 2009.

[113] A. L. Wilmot-Smith, G. Hornig, and D. I. Pontin. Magnetic Braiding and Quasi-Separatrix Layers. *The Astrophysical Journal*, 704:1288–1295, October 2009.

[114] A. L. Wilmot-Smith, D. I. Pontin, and G. Hornig. Dynamics of braided coronal loops. I. Onset of magnetic reconnection. *Astronomy and Astrophysics*, 516:A5, June 2010.

[115] A. L. Wilmot-Smith, D. I. Pontin, A. R. Yeates, and G. Hornig. Heating of braided coronal loops. *Astronomy and Astrophysics*, 536:A67, December 2011.

[116] G. L. Withbroe. The temperature structure, mass, and energy flow in the corona and inner solar wind. *The Astrophysical Journal*, 325:442–467, February 1988.

[117] G. L. Withbroe and R. W. Noyes. Mass and energy flow in the solar chromosphere and corona. *Annual Review of Astronomy and Astrophysics*, 15:363–387, 1977.

[118] A. R. Yeates, F. Bianchi, B. T. Welsch, and P. J. Bushby. The coronal energy input from magnetic braiding. *Astronomy and Astrophysics*, 564:A131, April 2014.

[119] L.-S. Young. Mathematical theory of Lyapunov exponents. *Journal of Physics A Mathematical General*, 46(25):254001, June 2013.

[120] E. G. Zweibel and H.-S. Li. The formation of current sheets in the solar atmosphere. *The Astrophysical Journal*, 312:423–430, January 1987.