

Chapter 5

Numerical simulations

5.1 Introduction

Having looked at a number of simplified models, we are now ready to look at full three-dimensional numerical simulations. Although still idealized in many ways, these simulations will be a step closer to reality than the kind of models we have considered in previous chapters. Of course, this extra realism is balanced by the long run times required by three-dimensional computations, meaning that only a few isolated points in the parameter space can be examined, rather than being able to plot complete bifurcation diagrams, as we have done previously.

We will begin, in section 5.4, by looking at simulations with simple uniform magnetic fields. This will allow us to study the basic phenomena of magnetoconvection in inclined fields, and to look at what changes occur as the inclination angle is varied. These simulations will correspond to the type of situation considered in Chapters 2 and 3. Our aims will be to confirm the results from those chapters, and to examine the new behaviour that occurs in more strongly nonlinear regimes.

In a (relatively strong) vertical field, previous calculations have shown that a stationary hexagonal convection pattern will be found. (See section 1.5 from the Introduction for more information and references.) We have discussed in previous chapters the changes that would be expected to this pattern once the field becomes tilted. Of these, the main ones were that we would expect the cells to become elongated along the direction of tilt of the field for moderate tilt angles, and we would expect the pattern to travel, either towards or against the tilt. For larger tilt angles, a pattern of rolls would be expected.

The calculation of Hurlburt et al. (2000) appears to confirm these expectations. Their simulations show a travelling pattern of cells for $\phi = 22^\circ$, and a more roll-like pattern for $\phi = 67^\circ$. However, their simulations used relatively small boxes, and only three different runs were shown. In this chapter a larger number of runs is presented, and this reveals a clearer picture of which patterns occur in which areas of parameter space.

In sections 5.5 and 5.6 we move on to more complicated situations in which the angle of inclination of the field varies with position. This is motivated by the structure of sunspots, and in particular the sharp distinction observed between the umbra and the penumbra. We conjecture that this is ultimately caused by a change in the nature of convection that occurs once the field becomes sufficiently inclined to the vertical. In order to investigate this conjecture, therefore, we would like to answer the following question: Is it possible to produce a model in which there is a sharp distinction between two different forms of convection in two different regions of the computational domain, one where the field is near-vertical, and one where the field is near-horizontal?

5.2 Definition of the problem

In this chapter, we will use a standard model of compressible magnetoconvection, as described (for example) by Matthews et al. (1995); Bushby and Houghton (2005). This model assumes that the following quantities are all constant: the heat capacities c_p and c_v , the thermal conductivity K , the dynamic viscosity μ , the magnetic diffusivity η , the permeability μ_0 and the acceleration due to gravity g . The calculation is performed within a Cartesian box of depth d and with periodic boundary conditions in the horizontal directions. The lengths of the box in the x and y directions will be $\lambda_x d$ and $\lambda_y d$ respectively, where λ_x and λ_y are the (dimensionless) *aspect ratios*. The box is filled with conducting plasma which is initially static, with a known temperature T_0 and density ρ_0 at the top of the box.

The equations of motion will be expressed in non-dimensionalized form with lengths being scaled by d (the depth of the box), densities by ρ_0 , temperatures by T_0 and magnetic fields by the initial magnetic field strength, B_0 (the details of the initial magnetic field setup will be discussed further below). Time will be measured in units of the isothermal sound travel time at the top of the layer, $d/(R_* T_0)^{1/2}$, where R_* is the gas constant (equal to $c_p - c_v$).

The non-dimensionalized governing equations are as follows: the equation of mass conservation

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{u}), \quad (5.1)$$

the momentum equation

$$\frac{\partial}{\partial t}(\rho \mathbf{u}) = -\nabla \left(p + \frac{F \mathbf{B}^2}{2} \right) + \nabla \cdot (F \mathbf{B} \mathbf{B} - \rho \mathbf{u} \mathbf{u} + \sigma \kappa \boldsymbol{\tau}) + \theta(m+1)\rho \mathbf{e}_z, \quad (5.2)$$

the induction equation

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \wedge (\mathbf{u} \wedge \mathbf{B}) + \zeta_0 \kappa \nabla^2 \mathbf{B} \quad (5.3)$$

and the heat equation

$$\frac{\partial T}{\partial t} = -\mathbf{u} \cdot \nabla T - (\gamma - 1)T \nabla \cdot \mathbf{u} + \frac{\gamma \kappa}{\rho} \nabla^2 T + \frac{\kappa(\gamma - 1)}{\rho} \left(\frac{\sigma \boldsymbol{\tau}^2}{2} + F \zeta_0 \mathbf{j}^2 \right), \quad (5.4)$$

together with the equation of state

$$p = \rho T \quad (5.5)$$

and Maxwell's equation

$$\nabla \cdot \mathbf{B} = 0. \quad (5.6)$$

(Note that we use a coordinate system in which z increases downwards.) The symbol \mathbf{j} represents the current density ($\mathbf{j} = \nabla \wedge \mathbf{B}$) while $\boldsymbol{\tau}$ is the viscous stress tensor, defined as follows:

$$\tau_{ij} = \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3} \delta_{ij} \frac{\partial u_k}{\partial x_k}. \quad (5.7)$$

The non-dimensionalization introduces a number of dimensionless parameters, and these are summarized in Table 5.1.

In the absence of convection, these equations admit a static trivial solution where $\mathbf{u} = 0$ and ρ and T take the form of a polytrope of index m :

$$T = 1 + \theta z \quad (5.8)$$

$$\rho = (1 + \theta z)^m \quad (5.9)$$

$$p = (1 + \theta z)^{m+1}. \quad (5.10)$$

We will use this state as the initial condition for our simulations (together with a small random perturbation added to the temperature field). We have not yet specified the initial condition for \mathbf{B} , but note for now that we may take \mathbf{B} to be any field satisfying $\nabla \cdot \mathbf{B} = 0$ and $\nabla \wedge \mathbf{B} = 0$ (this ensures that there is no Lorentz force to disturb the static state, and that $\partial \mathbf{B} / \partial t = 0$ in the induction equation).

<i>Parameter</i>	<i>Symbol</i>	<i>Definition</i>	<i>Value</i>
Ratio of specific heats	γ	c_p/c_v	5/3
Prandtl number	σ	$\mu c_p/K$	1
Ratio of magnetic to thermal diffusivity (at $z = 0$)	ζ_0	$\eta c_p \rho_0/K$	0.2
Polytropic index	m	$gd/R_*\Delta T - 1$	1
Dimensionless temperature difference	θ	$\Delta T/T_0$	10
Dimensionless thermal conductivity	κ	$K/d\rho_0 c_p (R_* T_0)^{1/2}$	variable
Dimensionless magnetic field strength	F	$B_0^2/R_* T_0 \rho_0 \mu_0$	variable

Table 5.1: *Dimensionless parameters and their meanings. Note that ΔT is the imposed (dimensional) temperature difference between the top and the bottom of the layer.*

We now discuss the values chosen for the parameters in Table 5.1. Note first of all that due to the large number of parameters, it would be impossible to conduct a full survey of the parameter space. Therefore, we have chosen fixed values for most of the parameters, and these values are shown in the final column of the Table. For example, γ is chosen to be the ratio of specific heats for an ideal monatomic gas (as appropriate for the Sun). The other parameters are largely arbitrary; we use values consistent with the work of Rucklidge et al. (2000). We note that the choice of $\theta = 10$ implies a fairly deep layer, in which the temperature varies by a factor of 11 from top to bottom. Note also that the effective value of ζ is a function of depth; the choice $\zeta_0 = 0.2$ corresponds to a value of 0.2 at the top of the layer, 1.2 at mid-layer and 2.2 at the bottom. Previous calculations (e.g. Weiss et al., 1990; Rucklidge et al., 2000) have shown that in this situation convection is steady, rather than oscillatory, at onset (although oscillatory behaviour can appear above onset via secondary bifurcations).

The remaining parameters – κ and F – are allowed to vary between different runs. It is more convenient to express these in terms of the Rayleigh number R and Chandrasekhar number Q , defined as follows:

$$R = \frac{(m+1)\theta^2}{\sigma\kappa^2\gamma}(m+1-m\gamma)(1+\theta/2)^{2m-1}, \quad (5.11)$$

$$Q = \frac{F}{\zeta_0\sigma\kappa^2}. \quad (5.12)$$

(Given values for R and Q , these equations can be inverted to find the appropriate

values for κ and F .)

At the top and bottom surfaces, we use standard impenetrable, stress-free boundary conditions, with $u_z = \partial u_x / \partial z = \partial u_y / \partial z = 0$. The temperature will be fixed ($T = 1$ at $z = 0$, and $T = 1 + \theta$ at $z = 1$). For the magnetic field, we will match to a potential field above the layer, while ‘tying’ the field lines at the lower boundary (i.e. the top boundary is insulating, and the lower boundary is conducting). As was discussed in Chapter 2, this seems an appropriate choice of boundary conditions for the Sun. (The equations for these magnetic boundary conditions are derived in Appendix A.)

5.3 Numerical methods

In this section we describe the methods used to solve (5.1)–(5.4) numerically.

Any numerical code for magnetohydrodynamics must have some method for ensuring that $\nabla \cdot \mathbf{B} = 0$. We can see from the induction equation that provided $\nabla \cdot \mathbf{B}$ is zero initially, it will, mathematically speaking, remain zero for all time. However, when implementing the equation on a computer, we will find, if we are not careful, that numerical errors in the simulation build up and lead to a non-zero $\nabla \cdot \mathbf{B}$ being produced after a certain time.

One way to combat this problem is to write the induction equation in conservative form, as $\partial \mathbf{B} / \partial t = -\nabla \wedge \mathbf{E}$. In theory, solving the equation this way means that $\nabla \cdot \mathbf{B}$ will be conserved to machine precision (because $\nabla \cdot (\nabla \wedge \mathbf{E}) = 0$). However, while the method works well within the interior of the layer, there can be problems at the top and bottom boundaries. These problems arise when numerically evaluating the required vertical derivatives of \mathbf{E} ($\partial E_x / \partial z$ and $\partial E_y / \partial z$) – if this is not done carefully, and in a manner consistent with the boundary conditions, then errors will be introduced and $\nabla \cdot \mathbf{B}$ will grow. For simple vertical field boundary conditions ($B_x = B_y = 0$), this will not be a problem (since we know $\partial E_x / \partial z = \partial E_y / \partial z = 0$ in such cases), but in general we will not be able to write down explicit boundary conditions for $\partial E_x / \partial z$ and $\partial E_y / \partial z$.

A more robust approach is to re-write the equations in poloidal-toroidal form. This basically means that \mathbf{B} is written as

$$\mathbf{B} = \nabla \wedge (B_T \mathbf{e}_z) + \nabla \wedge (\nabla \wedge (B_P \mathbf{e}_z)) \quad (5.13)$$

where B_T and B_P are the toroidal and poloidal potentials respectively. We can now derive induction equations describing the evolution of B_T and B_P , instead of B_x , B_y

and B_z (we also have to add a ‘mean field’ term, $\bar{\mathbf{B}}(z)$; see Appendix B for the details). The advantage of doing this is that it is now *impossible* for $\nabla \cdot \mathbf{B}$ to become non-zero – whatever values B_T and B_P take, the form of (5.13) is such that we will have $\nabla \cdot \mathbf{B} = 0$. The only disadvantage of this method is that the simulation will run slightly more slowly, since a few extra spatial derivatives need to be computed at each timestep.

Note that the other equations (the continuity, momentum and heat equations) do not change when using the poloidal-toroidal formulation. Formally, we should substitute (5.13) into (5.1)–(5.4), but in fact it is more efficient to compute \mathbf{B} once, using equation (5.13), and then simply to use that value wherever \mathbf{B} is needed in the other equations.

Once the equations have been rewritten in poloidal-toroidal form, they need to be discretized. For the spatial discretization we use a mixed finite difference/pseudospectral code. The term ‘pseudospectral’ refers, in our case, to the horizontal directions, where, since periodic boundary conditions are being used, all variables can be written as Fourier series. (The Fourier coefficients can be computed efficiently by using the Fast Fourier Transform algorithm, as described for example by Press et al. 1992.) Differentiation then becomes a simple matter of multiplying by ik (where k is the wavenumber) in Fourier space. The advantage of this method is that the spatial derivatives can be computed much more accurately than with a finite difference method, if the functions being differentiated are smooth (see e.g. Boyd 2001). However, this method cannot be used for the vertical derivatives, because the vertical boundary conditions are not periodic. Instead, a fourth order finite difference formula is used to compute the derivatives in this direction.

For the time discretization, the explicit third-order Adams-Bashforth method is used. The timestep is limited by stability constraints, derived both from the diffusive timescale and from the wave travel time across a mesh interval (i.e. the Courant condition; see Press et al. 1992). The computed maximum timestep is multiplied by a safety factor of 0.2. The scheme is stable provided that the diffusion coefficients are sufficiently large.

At the end of each timestep, boundary conditions are imposed. This is done simply by setting each variable to an appropriate value on the boundary. For example, the boundary values of u_z are set to zero, while the boundary values of u_x and u_y are set so as to ensure that $\partial u_x / \partial z$ and $\partial u_y / \partial z$ evaluate to zero at the boundary (when calculated using a one-sided finite difference formula). The boundary conditions are also utilized when calculating vertical derivatives near the boundary (where this is possible).

Also at the end of each timestep, we ‘dealias’ the variables, by applying the usual ‘2/3

rule’. This means that we filter out any modes with wavenumber (with respect to the horizontal directions) greater than two thirds of the Nyquist wavenumber (sometimes called ‘Nyquist frequency’, although in this case it is really a wavenumber rather than a frequency because it applies to space rather than time). This is easily implemented in Fourier space: the relevant Fourier amplitudes are simply set to zero.

The dealiasing is necessary because, although our computational grid can only represent modes of wavenumbers up to the Nyquist wavenumber (equal to π/λ , where λ is the grid spacing), modes of higher wavenumbers will in fact be generated by the nonlinear terms in the equations of motion. These modes do not simply disappear; instead their energy reappears at lower wavenumbers, in a process known as ‘aliasing’, and this can sometimes cause ‘aliasing instability’ (see e.g. Boyd 2001 for further explanation). The problem is avoided by filtering out the highest-frequency modes.

Note that the ‘2/3 rule’ applies to systems in which the nonlinearities are quadratic. This is true for most of our nonlinear terms, although there are some terms that involve divisions by ρ , for which the ‘2/3 rule’ is not strictly speaking appropriate. Nevertheless, we have found that in practice, this rule has been sufficient to prevent the aliasing instability from occurring.

One further problem with the code, as described so far, is that it is sometimes prone to numerical instabilities leading to fluctuations in the density. This can be a particular problem near the top boundary, where the lowest values of density can be found, and sometimes the density can even go negative, with disastrous results. We have found that this can often be solved by increasing the numerical resolution in the vertical direction. There are also two other ad hoc methods that we have used in order to improve the stability.

The first is to impose a boundary condition on the density. Strictly speaking, the density does not satisfy a boundary condition, but we can derive one by observing that $u_z = 0$ on the boundaries. This means that the z -component of the right-hand side of the momentum equation (5.2) will be zero on the boundary. The resulting equation can be interpreted as a boundary condition, as follows:

$$0 = -\frac{\partial}{\partial z} \left(\rho T + \frac{F \mathbf{B}^2}{2} \right) + \frac{\partial}{\partial x_j} (F B_j B_z + \sigma \kappa \tau_{jz}) + \theta(m+1)\rho. \quad (5.14)$$

In theory this equation should be satisfied automatically (provided that we are solving the momentum equation correctly), but in practice, numerical errors can cause the right-hand side of (5.14) to take on a small non-zero value. We correct this by explicitly

applying (5.14) as a boundary condition, i.e. by adjusting ρ on the top and bottom boundaries at the end of each timestep such that (5.14) is satisfied exactly. Doing this helps to ensure that the momentum equation is fully consistent with the boundary conditions, and this appears to improve the stability of the code in some cases.

Another method that can improve the stability of the density near the top boundary is to rewrite the continuity equation. Writing it as

$$\frac{\partial \rho}{\partial t} = -\rho \nabla \cdot \mathbf{u} - \mathbf{u} \cdot \nabla \rho, \quad (5.15)$$

instead of the conservative form (5.1), appears to improve stability (P. Bushby, private communication). One advantage of this formulation is that we no longer need to evaluate the vertical derivative of the density on the boundary (because $u_z = 0$ there). One disadvantage of not using the conservative form is that the method may no longer conserve mass exactly; however, in our simulations we typically find that mass is conserved to better than 1% accuracy, so this is not a serious problem.

All of these methods are implemented in a code which was originally written by Cattaneo et al. (1991) for the purely hydrodynamic case. Magnetic fields were added by Matthews et al. (1995), and the code has been worked on by various people since then (including P. Bushby, S. Houghton and N. Roxburgh). The changes made by the present author were the addition of the poloidal-toroidal field representation, and the addition of new magnetic boundary conditions. The modified code was tested by comparing it against linear theory.

Runs with resolutions of up to $256 \times 128 \times 144$ points were carried out. The code is parallelized (using MPI), so that several processors can be used for a single run. Smaller runs (of resolutions of around 64^3) were carried out on clusters of standard PCs, but for the larger runs we made use of the Cambridge-Cranfield High Performance Computing Facility, using up to 72 processors for each run.

5.4 Runs with uniform fields

In this section we take as our initial condition for \mathbf{B} a straight, uniform field, inclined at an angle ϕ to the vertical, as shown in Figure 5.1. We can then investigate the different types of behaviour that occur as the tilt angle is varied.

We choose in this section to keep B_z constant (as opposed to keeping $|\mathbf{B}|$ constant)

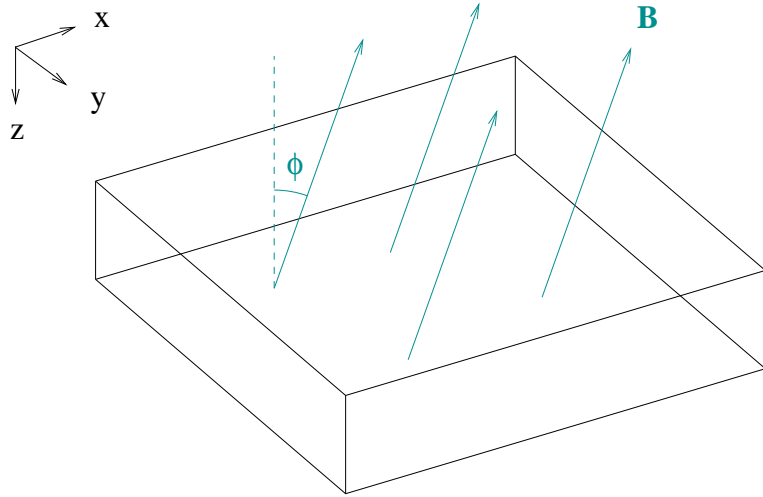


Figure 5.1: Diagram showing the initial field configuration for the ‘uniform field’ runs. Note the directions of the three coordinate axes. (The field is tilted towards the positive x -direction.)

as ϕ is varied. In other words, we take

$$\mathbf{B}_{\text{initial}} = \begin{pmatrix} -\tan \phi \\ 0 \\ 1 \end{pmatrix}. \quad (5.16)$$

This is different from the convention used in previous chapters (where we fixed $|\mathbf{B}|$). The motivation for this change is that, at least near onset, the amplitude of convection is controlled by the size of B_z , not $|\mathbf{B}|$. In other words, if we fix the Rayleigh number (and all other parameters) and vary ϕ , then for fixed B_z , the amplitude of convection stays roughly constant, while for fixed $|\mathbf{B}|$ the amplitude increases with increasing ϕ . The former seems more convenient, since we can control the amplitude of convection (and degree of nonlinearity) by varying R , independently of the changes in tilt angle.

For the numerical runs themselves, we chose to vary R and ϕ , while keeping Q fixed, at 2500 (for a vertical field, this would correspond to the ‘strong field’ regime of Weiss et al. 2002). We used aspect ratios of $\lambda_x = \lambda_y = 4$, with resolutions of 64 grid points in each horizontal direction, and between 40 and 80 grid points in the vertical.

The results are summarized in Figure 5.2. Each point on this graph represents the result of one simulation run, and the different symbols represent the different types of convection pattern that were found. Two basic patterns were found, namely rolls and cells. The cells also divided into three main types, which we have denoted ‘normal’, ‘wavy’ and ‘turbulent’. Finally, at the points marked ‘subcritical’, the trivial static

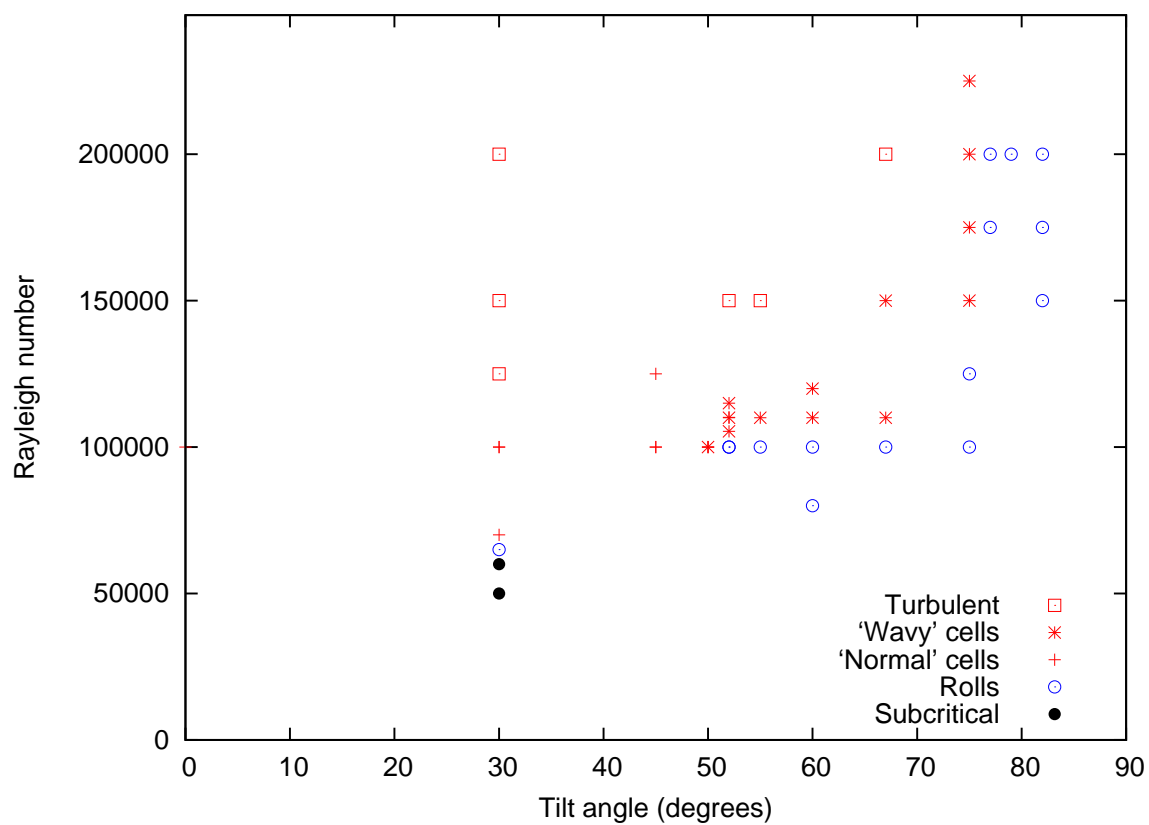


Figure 5.2: Diagram showing the solution types found at different values of R and ϕ , for $Q = 2500$.

solution was found to be stable, and no convection took place (since the simulations were started from scratch).¹

Some of the cellular solutions are illustrated in Figure 5.3. These runs all converged to uniformly drifting, but otherwise steady patterns (i.e., patterns which would be steady if viewed in an appropriately moving frame). The top two pictures in the Figure show what we have called ‘normal’ cellular solutions. Here the cells line up in a regular, more or less rectangular pattern. This is in contrast to the bottom two pictures, which show the ‘wavy’ cells. Here the convection cells form the same basic rectangular pattern, but this time it is subject to a zig-zag-like modulation.

There are a number of points to note about these solutions. Firstly, the basic pattern formed is one of rectangles, but the expected pattern in compressible convection is hexagonal. However, as pointed out by Matthews (1998), it is not possible to have a perfect hexagonal pattern in a square computational domain. In large domains we can obtain near-hexagonal patterns, but here we have used a relatively small aspect ratio (of 4) and this may be the reason that rectangular patterns are found instead.

Another observation is that the form of the solutions changes as the field becomes more inclined. For example, the cells become more and more elongated along the direction of tilt as ϕ increases. Moreover, the cells start to ‘run together’ forming more roll-like structures, particularly at the bottom of the box. At $\phi = 30^\circ$ (top left picture in Figure 5.3), there is a clear pattern of cells at both the top and the bottom of the box, but at $\phi = 45^\circ$ (top right picture), the cells are less clearly separated at the bottom of the box. By $\phi = 52^\circ$ (bottom left picture), the pattern at the bottom boundary is roll-like, although there are still clearly separated cells at the upper boundary. By $\phi = 75^\circ$ the cells at the top of the box have started to blend together as well.

The two other main types of solution are illustrated in Figure 5.4. The left-hand panel shows a roll solution; these are found mostly for very large ϕ , but they can also be found for smaller ϕ if R is only just above critical. For very large R we have found more turbulent solutions; a snapshot from one of these can be seen in the right-hand panel of Figure 5.4.

All of our solutions (except the straight parallel rolls) were found to travel. The

¹This means that these points are *linearly* subcritical. We did not attempt to look for *nonlinear* subcritical solutions here. For example, one might expect hexagons to be stable subcritically (see also Chapter 3); to see these, we would have had to run simulations starting from a (perturbed) hexagon solution, instead of starting from scratch.

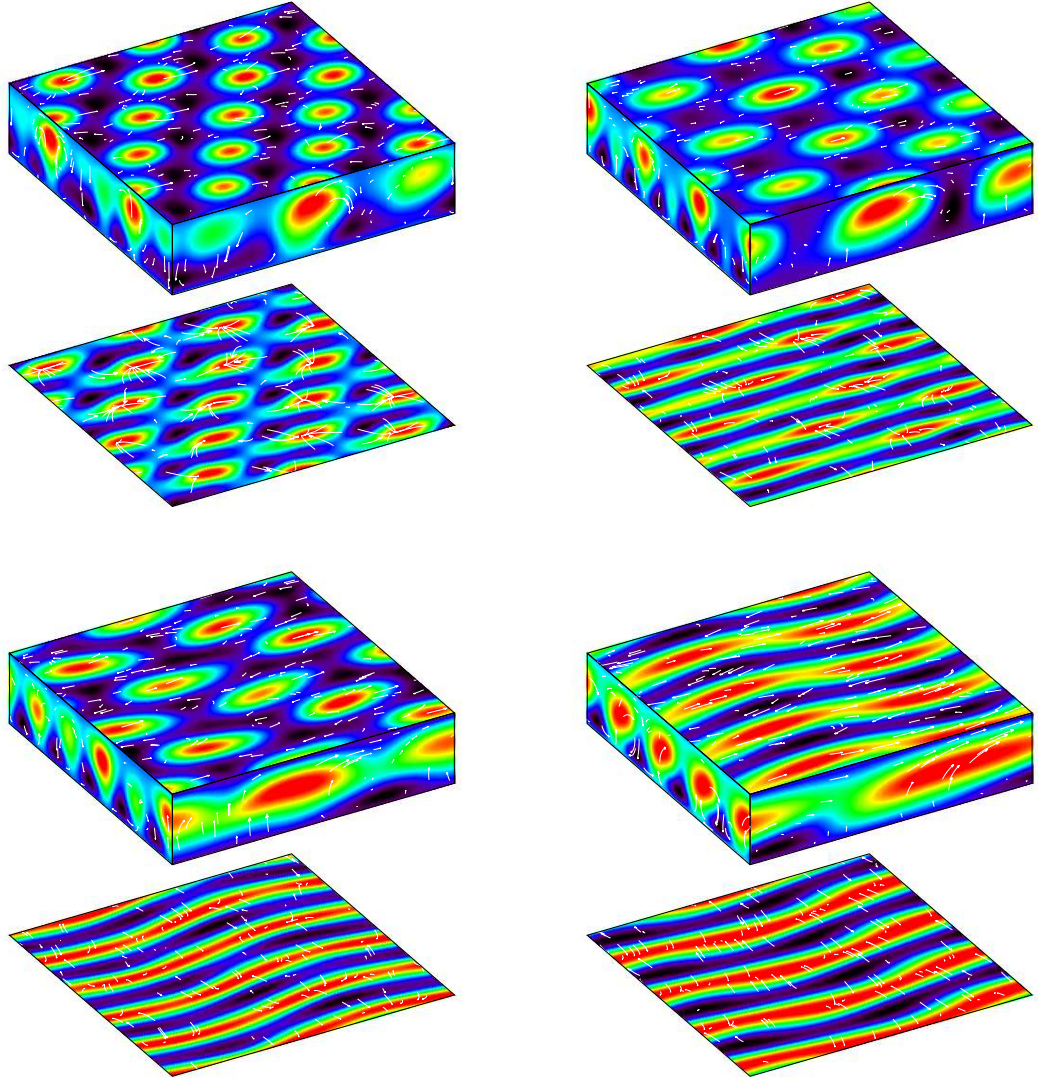


Figure 5.3: Solutions with uniform fields (all with $Q = 2500$). Top left: $R = 10^5$, $\phi = 30^\circ$; top right: $R = 10^5$, $\phi = 45^\circ$; bottom left: $R = 1.05 \times 10^5$, $\phi = 52^\circ$; bottom right: $R = 1.75 \times 10^5$, $\phi = 75^\circ$. The colours indicate temperature perturbation (with black, blue representing cool fluid, and green, yellow, red representing warm fluid), and the small arrows show the direction and magnitude of the fluid velocity (when projected onto each of the surfaces of the box). Note that each of these patterns drifts in the negative x -direction (against the direction of tilt).

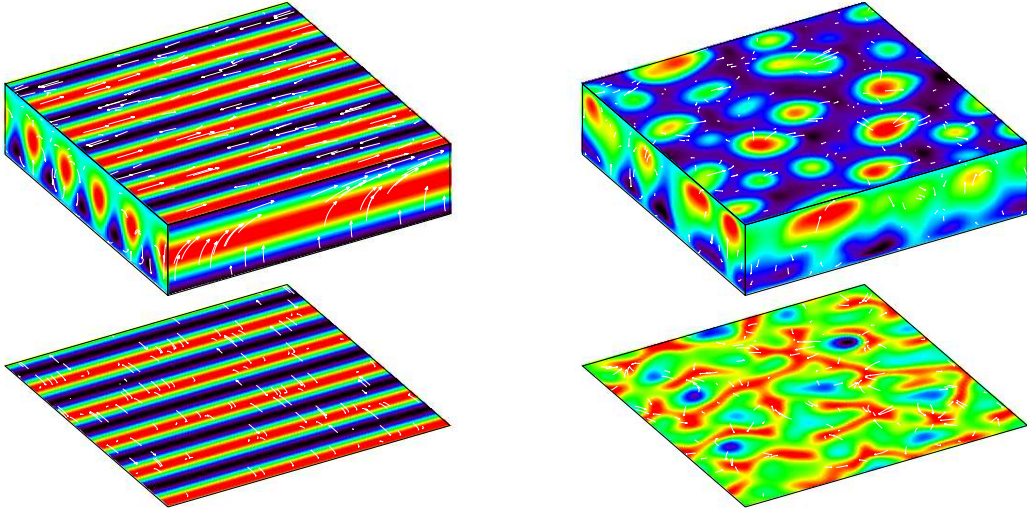


Figure 5.4: Additional solutions in a uniform field (with $Q = 2500$). Left-hand picture: $R = 10^5$, $\phi = 60^\circ$. This is a steady solution. Right-hand picture: $R = 2 \times 10^5$, $\phi = 30^\circ$. This solution is turbulent – the pattern is constantly evolving – but there is also a noticeable leftward drifting motion.

direction of travel was always against the tilt (i.e. in the negative x -direction) – there was no reversal of the direction of travel with increasing ϕ , although we do see a transition from drifting solutions to steady rolls as ϕ becomes large. We also find, in agreement with Hurlburt et al. (1996, 2000), that the mean *flow* at the surface is always in the opposite direction to this phase velocity (i.e. the pattern drifts to the left, but the mean flow at the surface is to the right).

5.5 Runs with non-uniform fields: ‘Arch’ initial condition

We now turn to runs with in which the initial condition involves a non-uniform magnetic field. As mentioned in the introduction to this chapter, these calculations are motivated by the difference in appearance between the umbra and the penumbra of a sunspot, and the aim will be to try to produce a model in which there is a sharp change between two different forms of convection. We begin with a simple ‘arch’-shaped field, and then move on (in the next section) to slightly more sophisticated initial configurations.

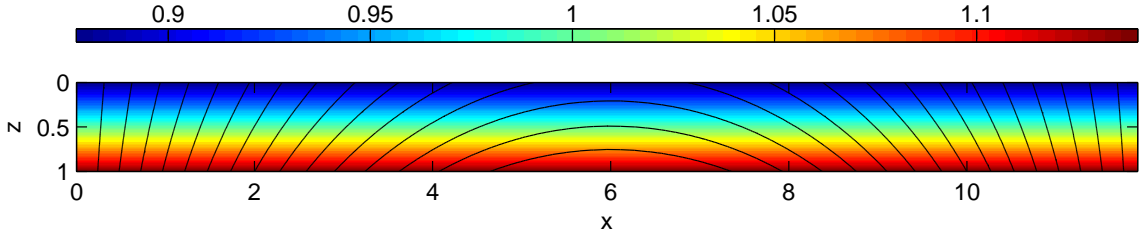


Figure 5.5: The ‘arch’ initial field (for $\lambda_x = 24$). The black curves are field lines, and the colours show the field strength $|\mathbf{B}|$.

We describe the ‘arch’-shaped field first. Here, the initial magnetic field configuration will be a two-dimensional field, given by

$$B_x(x, z) = -\frac{\partial\chi}{\partial z}, \quad B_y = 0, \quad B_z = \frac{\partial\chi}{\partial x} \quad (5.17)$$

where

$$\chi(x, z) = \frac{\lambda_x}{2\pi} \sin\left(\frac{2\pi x}{\lambda_x}\right) \exp\left(\frac{2\pi(z - 1/2)}{\lambda_x}\right). \quad (5.18)$$

Note that this field satisfies $\nabla \wedge \mathbf{B} = 0$, so that it is force-free and current-free, and therefore our trivial, motionless solution (5.8)–(5.10) is still an equilibrium state. Note also that

$$|\mathbf{B}| = \exp\left(\frac{2\pi(z - 1/2)}{\lambda_x}\right), \quad (5.19)$$

so that the field strength increases with depth; the normalization has been chosen such that $|\mathbf{B}| = 1$ at mid-layer ($z = 1/2$).

The field is illustrated in Figure 5.5. The Figure only shows the field for $x \in [0, \lambda_x/2]$; the second half (for $x \in [\lambda_x/2, \lambda_x]$) is identical to the first, but with a sign change. In fact, the entire field can be generated from the first quarter (for $x \in [0, \lambda_x/4]$), by using the following symmetries:

$$\mathbf{B}(x + \lambda_x/2, z) = -\mathbf{B}(x, z), \quad (5.20)$$

$$B_x(-x, z) = -B_x(x, z), \quad (5.21)$$

$$B_z(-x, z) = B_z(x, z). \quad (5.22)$$

Therefore, it is in fact possible to restrict attention to one quarter of a full period (i.e., $x \in [0, \lambda_x/4]$). By running simulations in only this subset of the box, we can in effect reduce the run time by a factor of 4. (This can be implemented by using so-called ‘quarter-wave’ Fourier transforms instead of ordinary ones.) We will use the symbol λ'_x to represent the width of the box actually used in the simulation (i.e. $\lambda'_x = \lambda_x/4$).

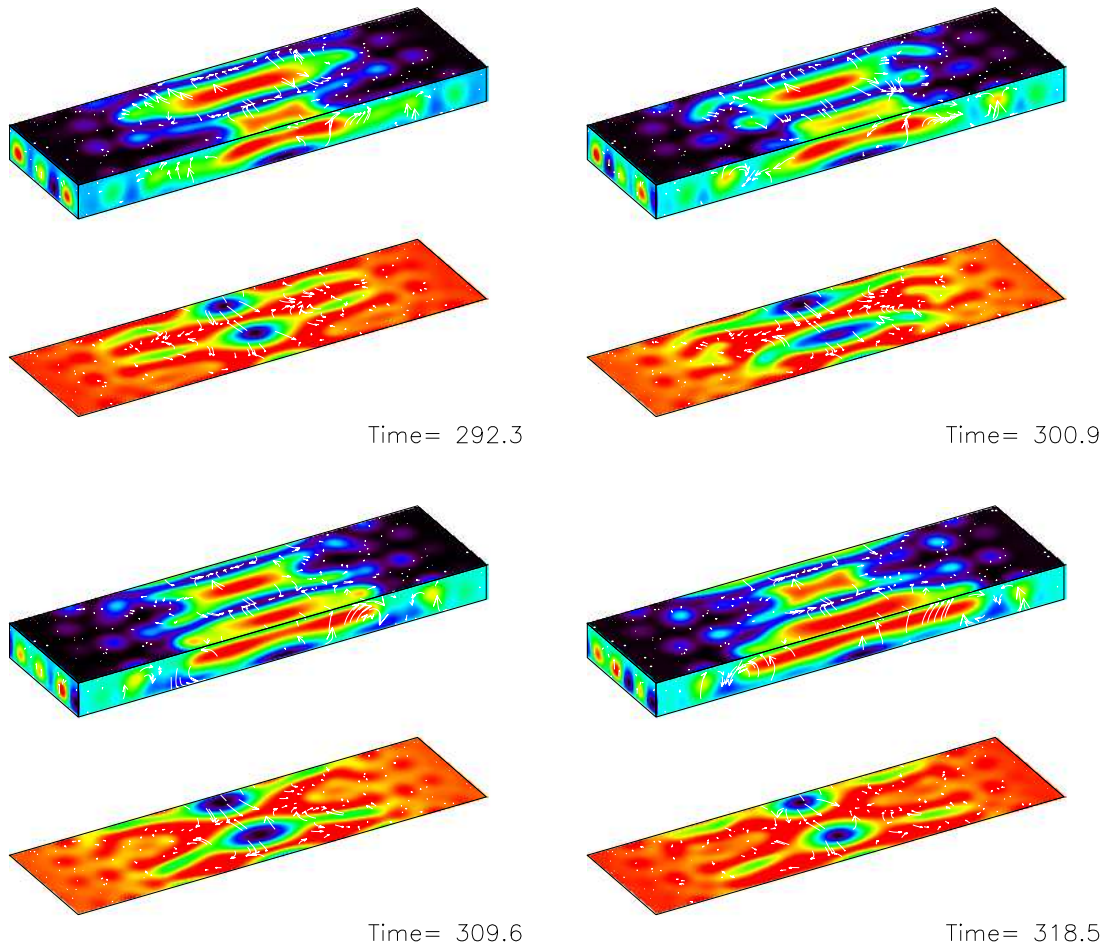


Figure 5.6: Snapshots at four different times during the ‘arch’ run with $Q = 1000$ and $R = 2.5 \times 10^4$, showing how the adjacent rolls wax and wane in amplitude. The colours show temperature perturbation while the arrows show the projection of the fluid velocity onto each surface of the box.

The reason for choosing the initial condition shown in Figure 5.5 was as a simple way of modelling a sunspot. The ‘sunspot’ is represented by the region with a near-vertical field – in fact, we have two ‘sunspots’ in our domain, one at each end of the box. This could also be seen as an infinite chain of ‘sunspots’, in view of the periodic boundary conditions. Note that the field is relatively strong everywhere – there is no region of weaker field to represent the ‘outside’ of a sunspot – so it would perhaps be more accurate to think of two ‘umbrae’ (the near-vertical fields near the edges of the box) connected by a shared ‘penumbra’ (the region of near-horizontal field in the centre).

We ran a simulation starting from this initial condition and using the parameters $R = 25000$, $Q = 1000$, $\lambda'_x = 6$, $\lambda_y = 3.5$, and a resolution of $65 \times 128 \times 48$. Note that a greater resolution is required in the y -direction than in the x -direction; this is

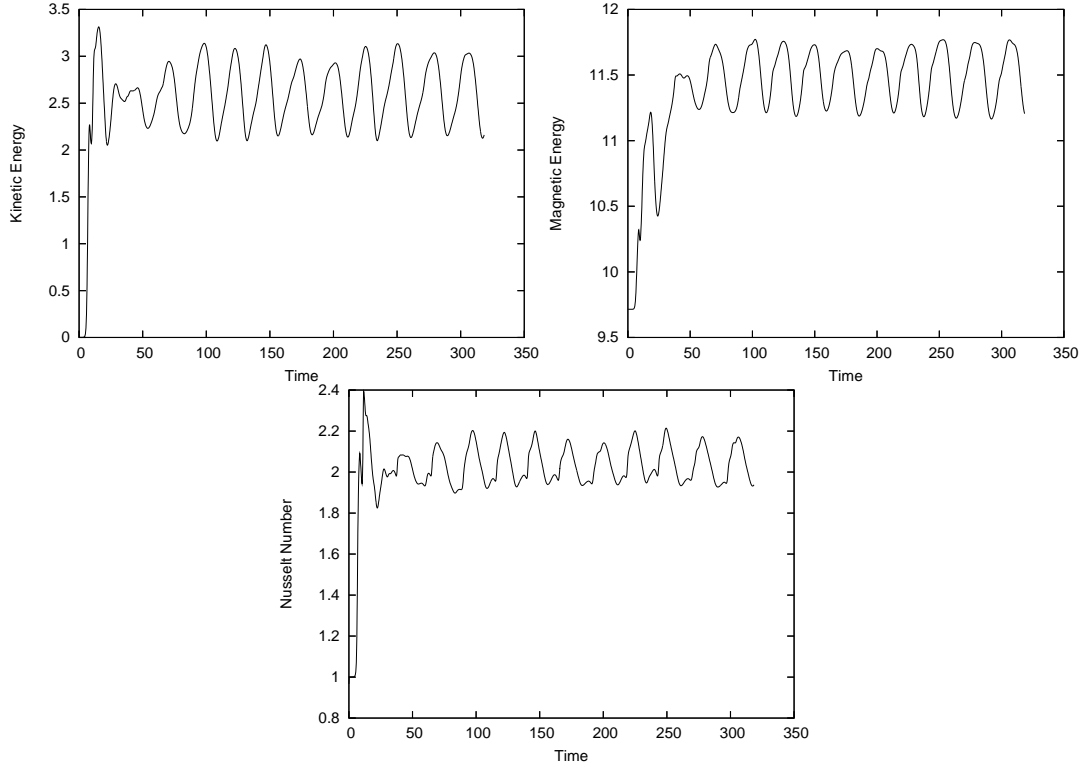


Figure 5.7: Plots of mean kinetic energy, mean magnetic energy and Nusselt number against time for the ‘arch’ run with $Q = 1000$ and $R = 2.5 \times 10^4$.

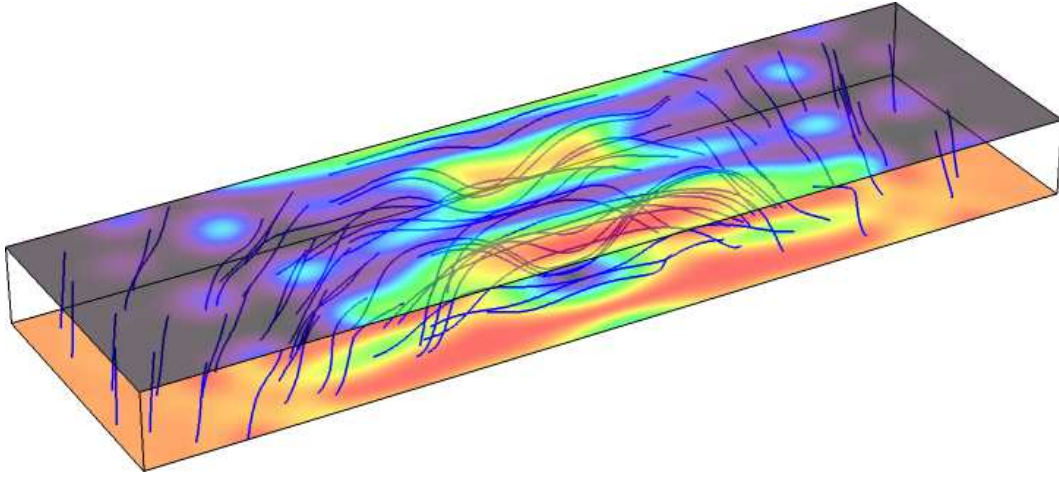


Figure 5.8: *Three-dimensional visualization of the magnetic field lines within the computational box (at the end of the calculation). For reference, the temperature perturbations on the top and bottom surfaces are also shown, where black, purple and blue indicate cooler fluid, and yellow, orange and red indicate warmer fluid.*

because the simulation produces rolls aligned with the magnetic field, which tend to have thin downflow sheets aligned in the x - z plane, and these downflows need to be properly resolved. (Note also that the vertical resolution is a little lower than what we have used for some of our other runs.)

Figure 5.6 shows snapshots of this run at four different times near the end of the calculation (after a quasi-steady state has been reached). The results clearly show a region of hexagon-like convection towards the outer edge of the domain, where the field is more vertical, with a transition to roll-like features in the centre, where the field is close to horizontal. Note the time dependence; the ‘rolls’ are undergoing a sort of oscillation, with adjacent rolls waxing and waning in amplitude. Figure 5.7 shows time series for kinetic and magnetic energy and the Nusselt number, which illustrate these oscillations more clearly. (The Nusselt number is a dimensionless measure of the superadiabatic temperature gradient, defined as follows:

$$N = \frac{d\bar{T}/dz - \theta(m+1)(1-1/\gamma)}{\theta - \theta(m+1)(1-1/\gamma)}, \quad (5.23)$$

where $d\bar{T}/dz$ is the mean temperature gradient at the bottom of the layer. Recall that $d\bar{T}/dz = \theta$ in the absence of convection; thus $N = 1$ if convection is not present, and $N - 1$ measures the convective contribution.)

Figures 5.8 and 5.9 show three-dimensional visualizations of the magnetic field struc-

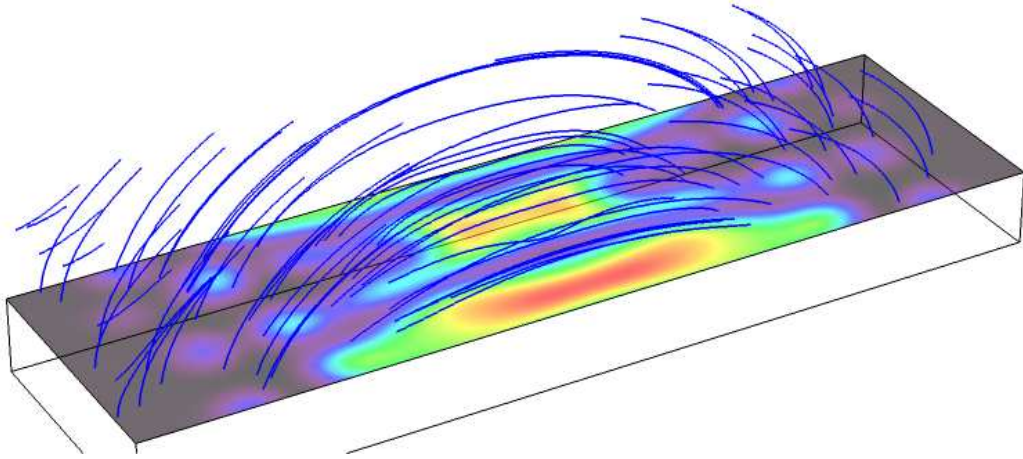


Figure 5.9: Three-dimensional visualization of a potential-field extrapolation of the field lines above the computational box. The temperature perturbation on the top of the box is also shown.

ture. The former shows the actual magnetic field structure within the box, while the latter shows the potential field extrapolation of the magnetic field into the region above the box.

In Figure 5.8, notice that at the edge of the box (in the region where hexagonal convection is found, and the field lines are near-vertical), the magnetic field structure is largely unaltered from its initial configuration. However, towards the centre of the box, where the near-horizontal fields can be found, the field lines are visibly pushed downwards compared to the initial condition. This can perhaps be seen more clearly in Figure 5.10 which shows the y -averaged magnetic field. Here, the depression of the field lines towards the right-hand end, and the increased field strength towards the bottom of the layer, are visible.

This seems to be evidence that *flux pumping* (Tobias et al., 2001; Thomas et al., 2002a; Weiss et al., 2004) is taking place within our simulation. This is a phenomenon in which magnetic field is transported downwards by the convection, due to the asymmetry – found only in compressible convection – between the fast-flowing downward plumes and the slow-moving upflows. In order to investigate this further, we have plotted the *centre of magnetic mass* as a function of both x and t . The centre of magnetic mass, z_{mm} , is in some sense a measure of the ‘average depth’ of the (horizontal) field, and we

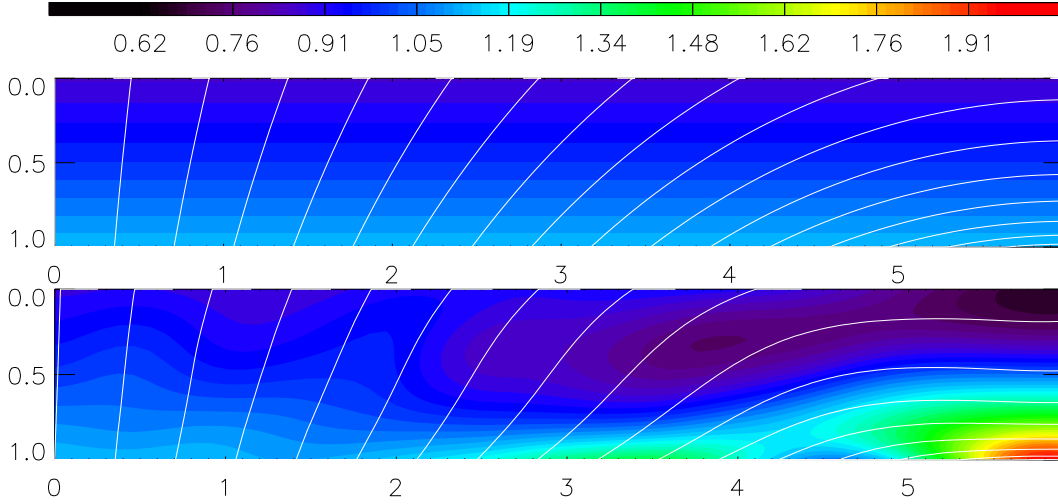


Figure 5.10: The magnetic field within the computational box, after averaging over the y -direction (thus obtaining a two-dimensional field). The colours indicate $|\mathbf{B}|$, and field lines are also plotted. The upper panel shows the initial condition, while the lower panel shows the final state (time-averaged over one oscillation period).

define it as follows:

$$z_{\text{mm}}(x, t) = \frac{\int_0^{\lambda_y} \int_0^1 z B_x \, dz dy}{\int_0^{\lambda_y} \int_0^1 B_x \, dz dy}. \quad (5.24)$$

(Note that flux pumping acts only on the horizontal flux, so we are only interested in the ‘centre of mass’ of the horizontal, not the vertical, field.)

We also define a version averaged over x as well as y :

$$\bar{z}_{\text{mm}}(t) = \frac{\int_0^{\lambda'_x} \int_0^{\lambda_y} \int_0^1 z B_x \, dz dy dx}{\int_0^{\lambda'_x} \int_0^{\lambda_y} \int_0^1 B_x \, dz dy dx}, \quad (5.25)$$

The first version is useful for describing how the effectiveness of pumping depends on position within the box, while the second version gives an average measure of the pumping over the entire box. (The second version corresponds to the definition adopted by Tobias et al. 2001, except that here we use B_x instead of B_y .)

Figure 5.11 shows z_{mm} as a function of x for our run with $R = 25000$ and $Q = 1000$. The dotted line shows the initial value, while the solid line shows the final value of z_{mm} ; the latter has been time-averaged over one oscillation period. The graph shows that at the right-hand end of the box, where the field is near-horizontal, the centre of magnetic mass has been pumped downwards significantly, but the pumping is less effective in regions where the field is more vertical (indeed, z_{mm} actually reduces in value near the left-hand end of the box).

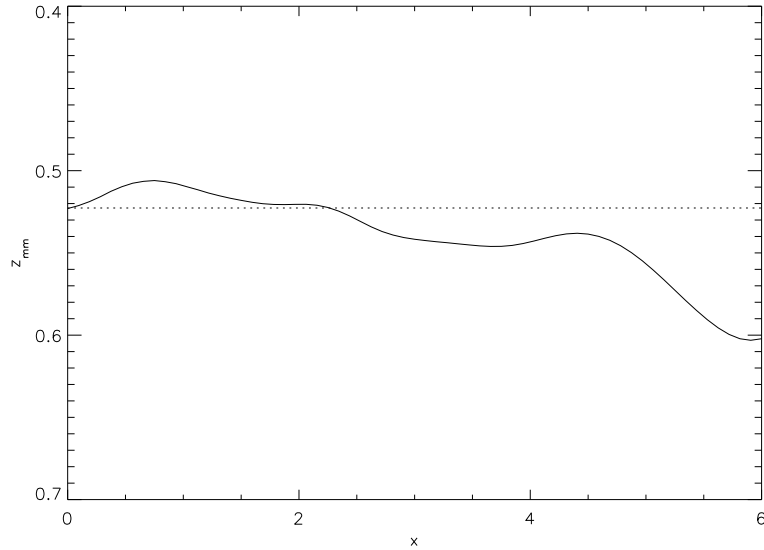


Figure 5.11: *The centre of magnetic mass $z_{\text{mm}}(x)$, time-averaged over one oscillation period. (The dotted line shows the value of z_{mm} at $t = 0$, for comparison.)*

Even at the right-hand end of the box, the change in z_{mm} is relatively small – only about 10% of the layer depth. In the calculation of Tobias et al. (2001), the horizontal field was pumped downward by a somewhat larger amount. The difference is that in their calculation, the computational domain was extended by adding a stable layer underneath the normal convecting region; the field could therefore be pumped down into this stable layer. In our calculation, we do not have this extra layer underneath the box, and so the field is blocked by the impenetrable lower boundary. This presumably limits the amount by which the field can be pumped downwards.

Figure 5.12 shows the depth of the centre of magnetic mass as a function of time. In agreement with Tobias et al. (2001), we find that there is a very rapid pumping phase at the beginning of the calculation, after which the centre of magnetic mass remains at a more or less constant level.

5.6 Runs with non-uniform fields: ‘Flux sheet’ initial condition

The above results, using the ‘arch’ initial condition, showed an interesting transition between hexagonal and roll-like convection. However, the results were a little unsatisfactory, because the length of the region of roll-like convection was rather small (compared

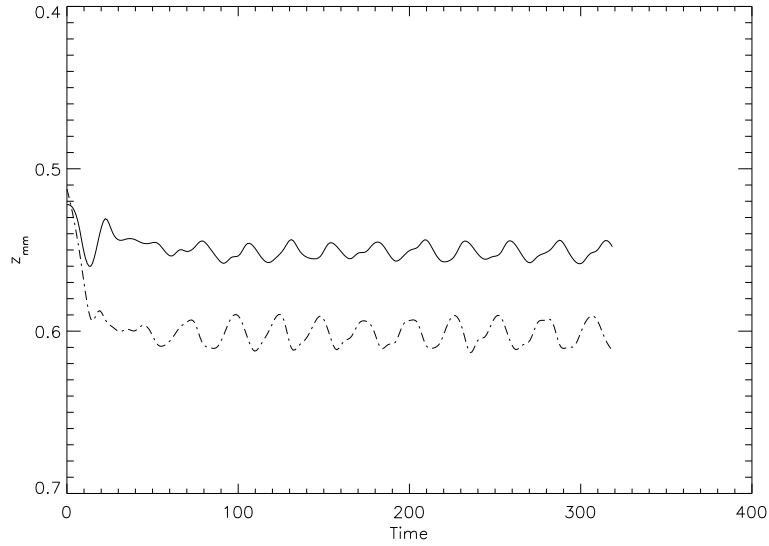


Figure 5.12: *Centre of magnetic mass against time (for the ‘arch’ run). The solid line shows \bar{z}_{mm} (i.e. the centre of magnetic mass averaged over all x), while the dash-dotted line shows z_{mm} at the right-hand end of the box (where pumping is found to be most effective).*

to its width), making it difficult to see precisely what was happening in this region.

We have therefore refined the choice of initial condition in order to allow greater control over the relative proportions of ‘umbra’ and ‘penumbra’ in the model (and hopefully to allow a wider region of roll-like convection in the results). The central observation is that we can specify B_z on the lower boundary; this, together with the requirement that the field within the layer be a potential field, is enough to uniquely specify the initial field.

We choose the following profile for B_z on the lower boundary ($z = 1$):

$$B_z = \begin{cases} \lambda'_x/x_0 & (0 \leq x \leq x_0) \\ 0 & (x_0 \leq x \leq \lambda'_x) \end{cases} \quad (5.26)$$

(This choice of scaling ensures that the average of B_z over the surface $z = 1$ is 1, or in other words, that the total flux of \mathbf{B} through the lower boundary is independent of the choice of x_0 .)

This field is sketched in Figure 5.13. The field below the box is assumed to continue downwards indefinitely, so that we have a column or tube of flux extending downwards in the region $x < x_0$, $z > 1$. (This choice was inspired by the flux tubes that are believed to be present underneath sunspots – although what we have here is more of a flux sheet than a flux tube, given that it extends indefinitely in the y -direction.) The potential

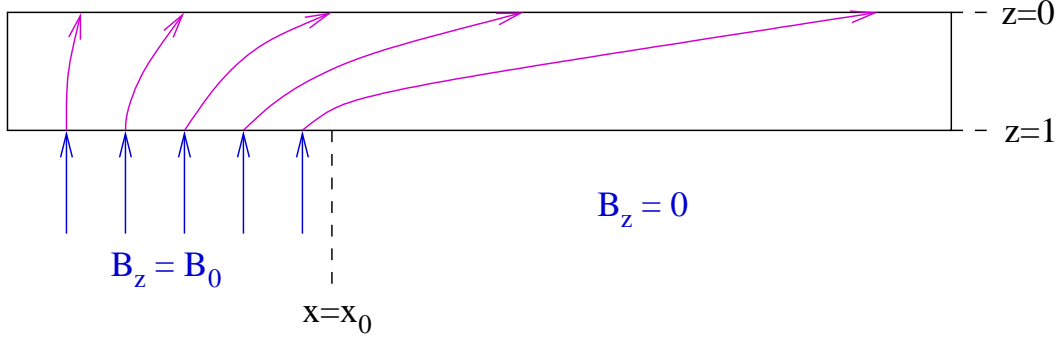


Figure 5.13: Sketch of the ‘flux sheet’ initial condition. Below the box, there is imagined to be a uniform sheet of flux in the region $x < x_0$ (blue field lines), with zero field for $x > x_0$. We match this onto a potential field (purple field lines), which forms the initial condition for the magnetic field within the computational box ($0 < z < 1$).

field above the surface $z = 1$ is what will form our actual initial condition.

Note that by varying x_0 , we can adjust the width of the flux sheet, and thus control the relative width of the ‘umbral’ and ‘penumbral’ regions within the simulation.

We begin constructing this initial condition by defining the *flux function* χ , as follows:

$$B_x = -\frac{\partial \chi}{\partial z}, \quad (5.27)$$

$$B_z = \frac{\partial \chi}{\partial x}. \quad (5.28)$$

The flux function defines a two-dimensional magnetic field satisfying $\nabla \cdot \mathbf{B} = 0$. If we also want χ to represent a potential field, then we must have $\nabla^2 \chi = 0$. We may now write χ as a Fourier series as follows:

$$\chi(x, z) = \sum_{k \text{ odd}, k \geq 1} A_k \sin\left(\frac{\pi k x}{2\lambda'_x}\right) \exp\left(\frac{\pi k(z-1)}{2\lambda'_x}\right). \quad (5.29)$$

(As before, $\lambda'_x = \lambda_x/4$, i.e. λ'_x denotes the actual width of the box used in the simulation, as opposed to one full period.) Note that we have restricted the series to sin terms, with the cos terms being omitted, and we have also included only the odd-numbered modes. These restrictions arise because of the ‘quarter-wave’ symmetry that we want our field to obey – see equations (5.20)–(5.22) above.

In view of our choice (5.26) for B_z , we now have the following expression for χ at $z = 1$:

$$\chi(x, 1) = \begin{cases} x\lambda'_x/x_0 & (0 < x < x_0) \\ \lambda'_x & (x_0 < x < \lambda'_x) \end{cases} \quad (5.30)$$

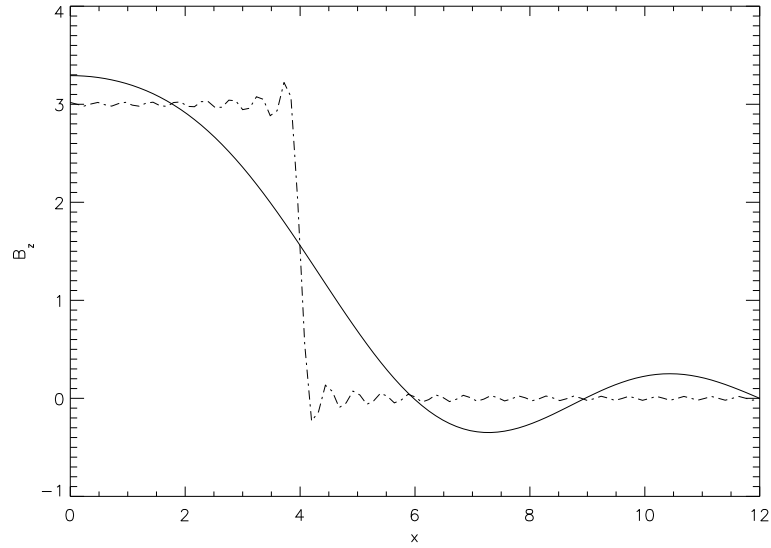


Figure 5.14: B_z as a function of x at the lower boundary, as given by the series (5.29), with A_k as defined in (5.34). Dot-dashed curve: series truncated at the Nyquist frequency. Solid curve: series truncated after 4 terms.

We must now compute values for A_k that make (5.29) and (5.30) consistent with one another. To do so, we proceed as follows:

$$\int_0^{\lambda'_x} \chi(x, 1) \sin\left(\frac{\pi k x}{2\lambda'_x}\right) dx \quad (5.31)$$

$$= \int_0^{\lambda'_x} \sum_{j \text{ odd}, j \geq 1} A_j \sin\left(\frac{\pi j x}{2\lambda'_x}\right) \sin\left(\frac{\pi k x}{2\lambda'_x}\right) dx \quad (5.32)$$

$$= \frac{A_k \lambda'_x}{2}. \quad (5.33)$$

To compute A_k , therefore, we need to evaluate (5.31), substituting for χ from (5.30), and multiplying the result by $2/\lambda'_x$. After a little algebra we find that

$$A_k = \frac{8\lambda_x'^2}{\pi^2 k^2 x_0} \sin\left(\frac{\pi k x_0}{2\lambda'_x}\right). \quad (5.34)$$

One unforeseen problem with this particular choice of initial condition is that it produces a number of undesirable oscillations in the value of B_z at $z = 1$. To see this, consider the representation (5.29). In this equation, we should in theory sum over all positive odd integers k ; however, in practice, only a finite number of Fourier modes can be represented in the calculation, and we must truncate the sum at some point. It is this truncation that generates the oscillatory behaviour; see Figure 5.14, dash-dotted curve, in which the series was truncated at the Nyquist frequency. The oscillations arise because Fourier series cannot represent discontinuities well (they are an example of the Gibbs phenomenon).

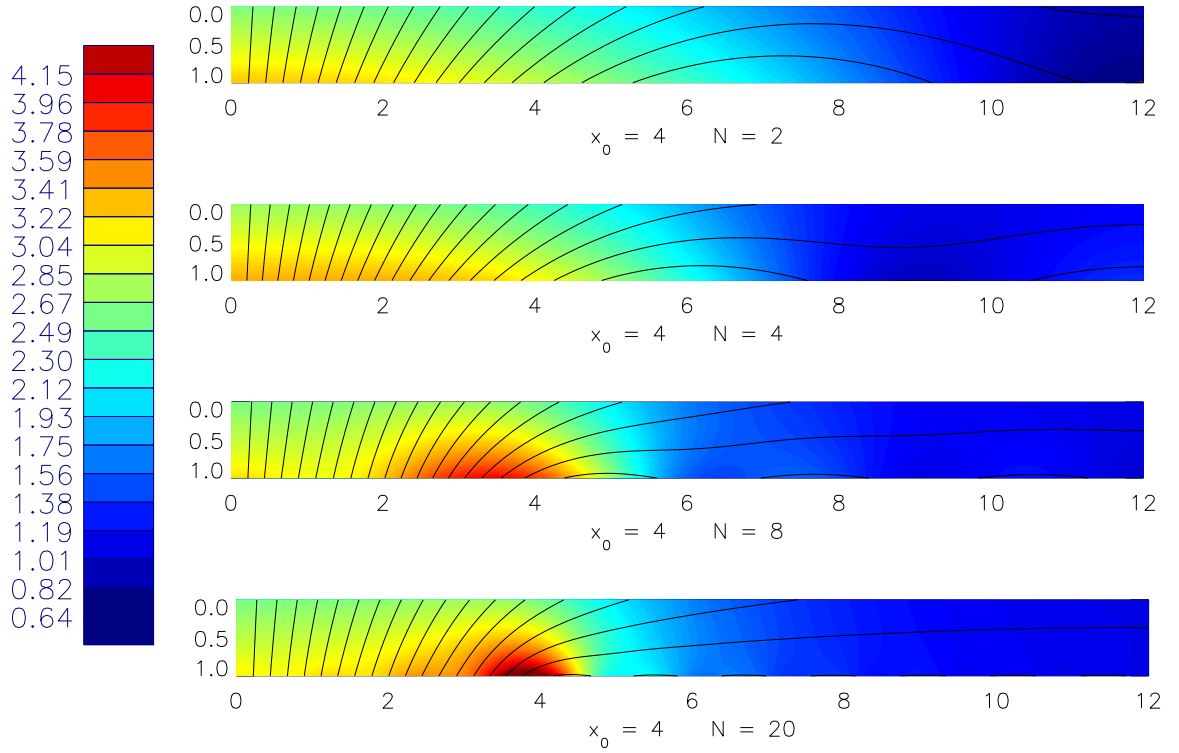


Figure 5.15: Illustration of the initial condition given by (5.29) and (5.34), plotted as a function of x and z . The colours indicate $|\mathbf{B}|$ and the curves are field lines.

The oscillations are undesirable because they lead to numerical instabilities. When using a ‘tied field’ boundary condition, the problem is worse, because the value of B_z is fixed on such a boundary; therefore, the oscillations can never smoothed away by diffusion, and they remain present for the entire simulation.

After discovering this problem, we resolved it in a fairly simple way: by truncating the sum at a lower value of k .² We will use N to represent the total number of modes used (i.e., the odd modes from $k = 1$ to $k = 2N - 1$ are included). The result when choosing $N = 4$ is shown on Figure 5.14 (solid curve); the sharp oscillations have disappeared, and have been replaced by a smooth curve.

Plots of the magnetic field for various N are shown in Figure 5.15. Here we have chosen $x_0 = 4$, since this gives a relatively large ‘penumbral’ region to the right-hand side, whilst still retaining a reasonable amount of ‘umbra’ at the left-hand end. Notice

²An alternative solution would be to choose a different profile for B_z in the first place, i.e. to choose a continuous function in place of the discontinuous (5.26). We will look at this possibility as well, later on in this chapter.

that choosing N too low results in a field which is not much different from the ‘arch’ field (which would correspond to $N = 1$), while choosing N too high produces a region of quite strong field at the base of the layer near $x = 4$ (which seems unrealistic). In our simulations, we have chosen to use the value $N = 4$, as a compromise between these two extremes.

5.6.1 Initial results

We began by running a simulation in a 12×2 box (i.e. $\lambda'_x = 12$ and $\lambda_y = 2$). The resolution was $257 \times 64 \times 144$, and we set $Q = 45$ and $R = 25000$. The simulation converged to the final state shown in Figure 5.16, with a region of hexagon-like cellular convection to the left, where the field is mainly vertical, and a region with long roll-like structures to the right, where the field is near-horizontal. The pattern is time-dependent: the hexagons travel to the left, while the rolls undulate and continually evolve. The pattern in the central part of the roll region also appears to travel leftwards, while the pattern near the right-hand end of the box does not seem to travel in any particular direction.

We must be careful in interpreting this result, since the box is rather narrow in the y -direction – with room for only a single convection roll – and the motion is therefore highly constrained. To reduce this effect we performed a second calculation using the same parameters but in an enlarged box with $\lambda_y = 4$ instead of 2. (The resolution in the y -direction was also correspondingly increased.) As an initial condition we used the result from the run in the narrower box, duplicated twice in the y -direction, and with a small perturbation added to the temperature field.

Initially, this new calculation proceeded similarly to the original run, but once enough time had passed, the original pattern of rolls decayed and broke up. This is shown in Figure 5.17; the top picture shows the situation when the rolls first begin to break up, while the bottom picture shows the situation some time later when almost all trace of the original roll pattern has disappeared. (The break-up begins from the right-hand end of the box, and proceeds leftwards.)

To understand the cause of this phenomenon we must look at the magnetic field structure. This is illustrated in Figure 5.18, which plots magnetic field lines both before and after the break-up of the roll pattern. In the earlier picture, although the field lines snake from side to side a little, they are still predominantly horizontal and aligned with

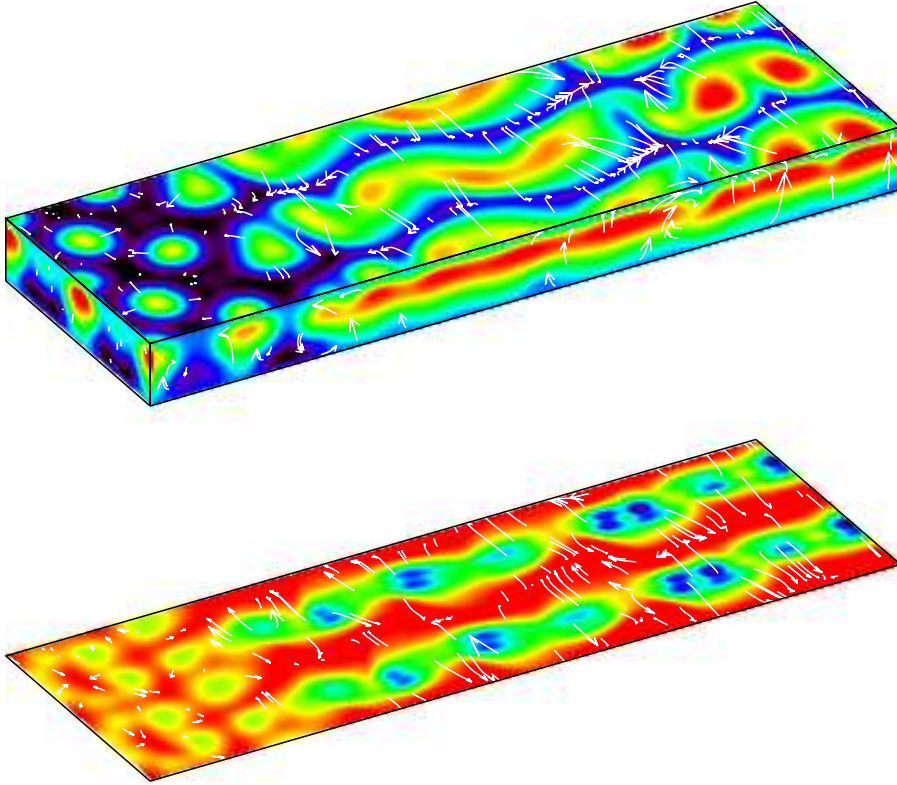


Figure 5.16: Temperature plot for the solution with the ‘flux sheet’ initial condition, in the narrow box ($\lambda_y = 2$). Note that in order to better show the periodicity, the diagram actually shows two copies of the box, placed side by side along the y -direction (so that the actual computational domain is half of what is shown above).

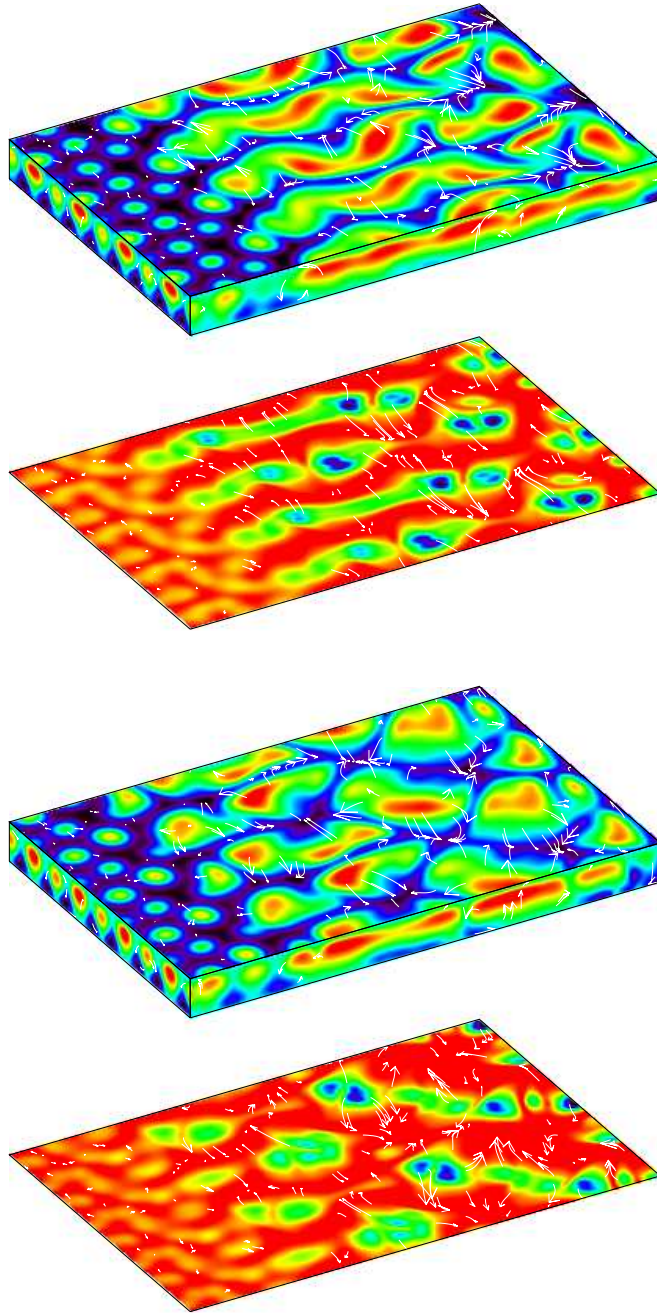


Figure 5.17: As Figure 5.16, but extended to a wider box (with $\lambda_y = 4$ instead of 2). Once again, two copies of the box are shown in the y -direction. The upper picture shows the situation mid-way through the calculation, and the lower diagram shows the final state.

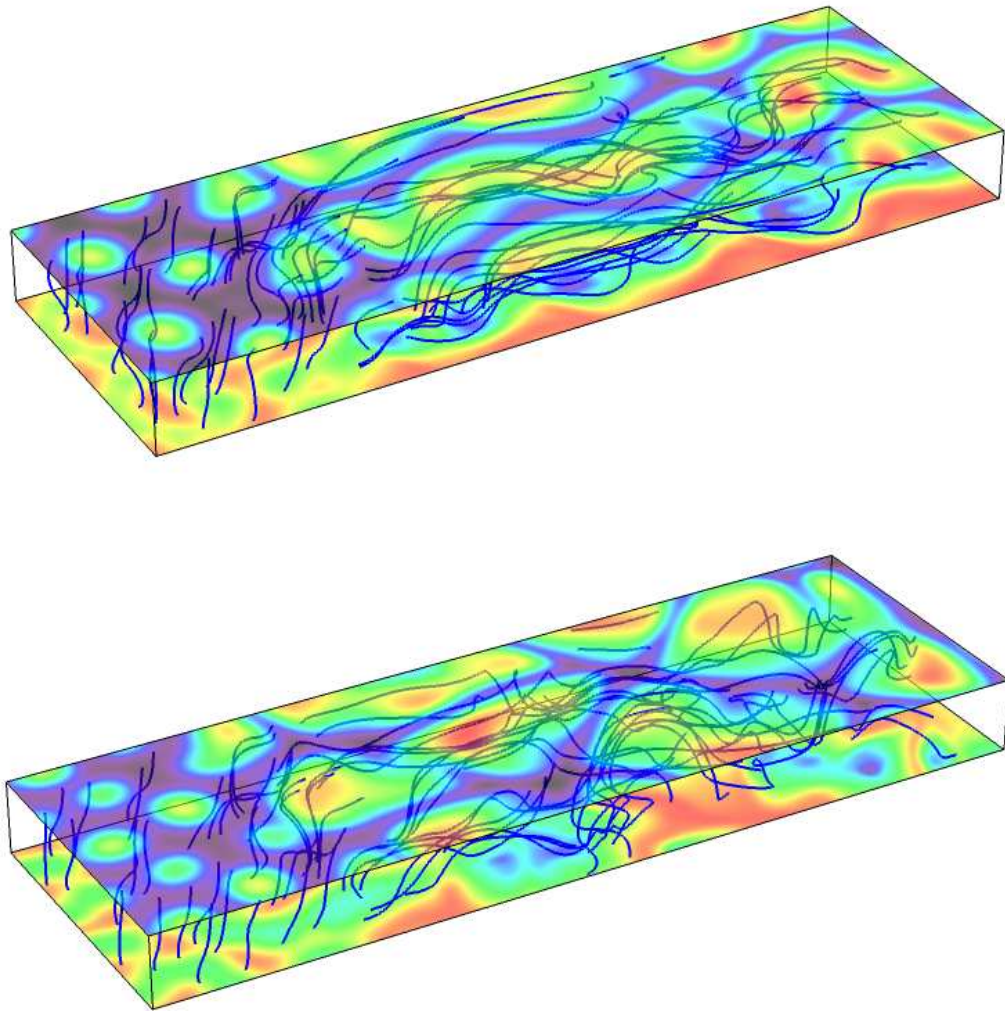


Figure 5.18: *Plots of magnetic field lines within the (wider) box, with temperature shown in colour on the top and bottom surfaces. (Only one copy of the box is shown in this diagram.) The top picture shows the situation just after the width of the box was doubled and the calculation restarted. The bottom picture shows the final state.*

the roll axes. In the later picture, the field lines are much more tangled. What appears to have happened is that part of the field has risen upwards (presumably by some sort of magnetic buoyancy process) and this has disrupted the roll pattern.

This can be seen clearly if one plots $\langle B_x \rangle$, where the angle brackets denote a horizontal average, against z and t . (This measures the amount of magnetic flux in the x -direction at each depth within the box.) See Figure 5.19, where $\langle B_x \rangle$ is plotted on the x -axis, with depth on the y -axis, and with different times represented by differently-coloured curves (lighter colours representing earlier times). The first graph represents the first calculation (the one with the narrower box). Here, the initially almost-uniform flux profile is quickly redistributed and pumped downwards. The total amount of flux in the box also increases during the calculation, indicating that flux is coming down into the domain through the top boundary. (Recall that although the lower boundary is a perfect conductor, the upper boundary is not, and therefore flux can enter or leave the computational domain via the top boundary.)

The second panel of Figure 5.19 shows what happens in the second calculation (after the width of the box has been increased). (The final, i.e. darkest, curve in the top picture therefore corresponds to the initial, i.e. lightest, curve in the bottom picture.) Initially, the profile remains fairly constant, but then the flux starts to move upwards (with $\langle B_x \rangle$ increasing for $z \gtrsim 0.4$ and decreasing for $z \lesssim 0.4$). In the final stages the flux seems to decrease at all depths, indicating that flux is being lost through the upper boundary. The corresponding reduction in field strength, together with the rearrangement of the field from a more-or-less coherent horizontal structure to a more tangled pattern, seems to be enough to prevent the long field-aligned convection rolls from forming.

This can also be seen by looking at plots of the centre of magnetic mass \bar{z}_{mm} (equation 5.25) against time, as shown in Figure 5.20. The left-hand picture shows the plot for the initial run, in the narrower box; here, the centre of magnetic mass descends quite rapidly at the beginning of the simulation, showing that flux is pumped downwards. The right-hand picture shows the plot for the second run, in the wider box, and here we see that the initial flux pumping is reversed (to a certain extent) with flux rising back out of the top of the box again.

In summary, therefore, we see that in the narrow box, flux is brought down from above, and this strengthens the field within the penumbral region. In the wider box, the reverse happens and flux escapes from the domain, causing a somewhat weaker field in the penumbral region. In the former case a pattern of rolls is seen, while in

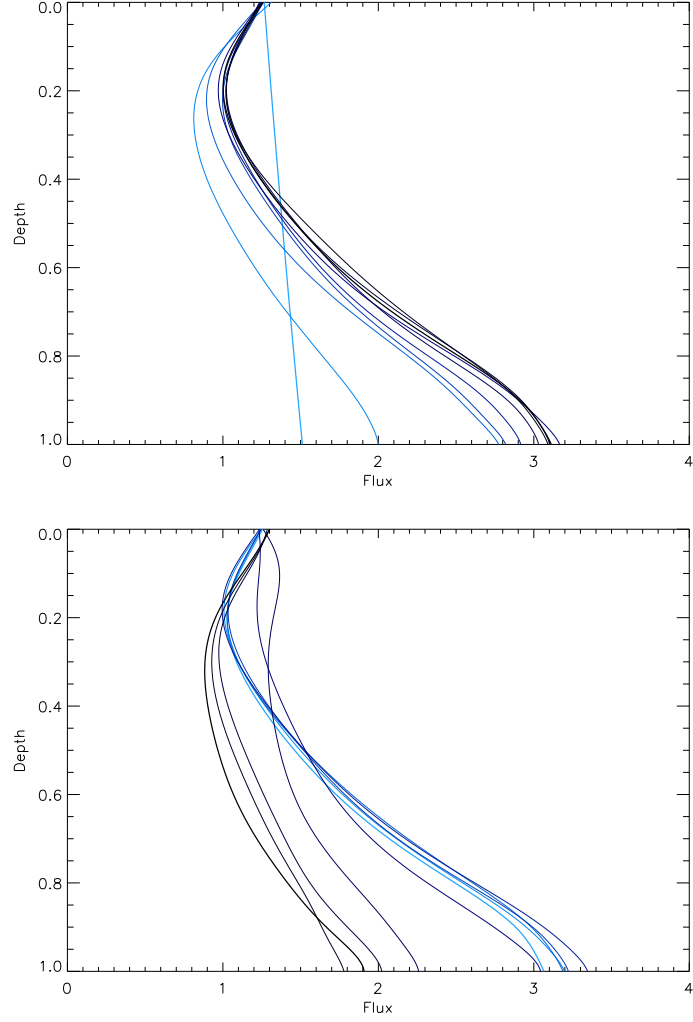


Figure 5.19: Plots of horizontal flux (i.e. $\langle B_x \rangle$) against z for various different times. Graphs are shown at equally spaced time intervals, with the lighter colours indicating earlier times, and darker colours later times. The top picture shows the initial calculation (in the narrow box) while the bottom picture shows the second calculation (with the wider box).

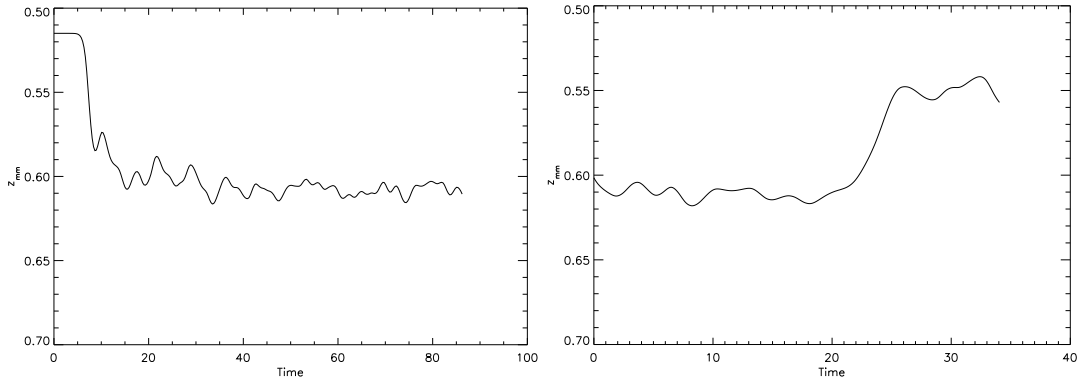


Figure 5.20: The centre of magnetic mass \bar{z}_{mm} against t , for the original run (left-hand picture) and for the run in the widened box (right-hand picture).

the latter case a cellular structure is found. This is somewhat unfortunate, since it is the wider box that corresponds to the more realistic situation (the narrower box being rather strongly constrained by the periodicity in the y -direction); however, the pattern found within the wider box is a rather uninteresting cellular pattern, looking a lot like field-free granulation. To model real penumbrae it seems clear that we will require a somewhat stronger field, in order to induce a transition to a new form of ‘penumbral’ convection. (This is what happened in the narrower box, but we would like to find a way of obtaining that pattern, or one like it, within a wider, less constrained computational box.)

5.6.2 Preventing the flux from escaping

Therefore, we would like to find a way to prevent the flux from escaping through the upper boundary. In this section we look at one way of achieving this aim.

Since the simulations take a considerable amount of time to run, we decided, before proceeding any further, to speed things up by using a reduced numerical resolution (of $129 \times 64 \times 144$, instead of the previous resolution of $257 \times 128 \times 144$) for all of our runs from this point onwards. This decision allows more runs to be performed, at the expense of possibly losing some accuracy. In this case we felt it was more important to be able to perform multiple runs than to have the highest accuracy possible, because we were mainly interested in the general patterns produced rather than the detailed nature of the solutions. In particular, we wanted to be able to search through parameter space (as far as possible) for any evidence of a ‘sunspot-like’ regime, in which a difference between ‘umbral’ and ‘penumbral’ convection could be seen, as was discussed above.

A second factor of note is that, we encountered problems (in some of the runs presented below) in which numerical instabilities would develop in the vertical direction. In particular, high-wavenumber vertical oscillations would appear (especially in the variables ρ and u_z). Where this occurred, it was corrected fairly straightforwardly, by occasionally applying a ‘smoothing’ procedure, in which high-wavenumber modes were filtered out (similarly to the dealiasing already being applied in the horizontal directions).

Imposing a perfectly conducting upper boundary

With these points in mind, we then looked for a way to run a simulation in which the flux was somehow constrained not to escape from the simulation domain. The simplest way to achieve this is to change the upper boundary condition to a perfectly conducting boundary (instead of matching to a potential field), which would prevent flux from either entering or leaving the box. We chose to run a simulation with this modified boundary condition (and with all other parameters being unchanged). This new run was started from scratch (i.e. from the motionless equilibrium state plus a small temperature perturbation) and the larger box size was used.

The final state for this new run is shown in Figures 5.21 (showing the temperature) and 5.22 (showing y -averaged magnetic field). In this result, we do *not* see a roll-like pattern in the penumbra; instead, a cellular pattern, similar to the previous result, Figure 5.17 (lower panel), is seen. So, although our change to the boundary conditions did prevent flux from escaping, it appears that this in itself was not sufficient; if we want to see more anisotropic patterns, then we will also need to arrange for there to be more flux within the domain at the beginning of the calculation.

The simplest way to do this would be to increase the parameter Q . However, this would increase the field strength within both the umbral and penumbral regions of the box, whereas we only want to increase the field strength in the penumbra. The problem with increasing the field strength in the umbra is that the field is already rather strong here, and we are in danger of shutting down the convection completely if Q is much increased (whereas ideally we want a simulation showing convection within both the umbra and the penumbra). To see this, note that the ‘effective Q ’ (i.e. $Q|\mathbf{B}|^2$, where $|\mathbf{B}|$ refers to the dimensionless field strength) is approximately 500 in the umbral region (although this does vary depending on exact position), compared to a value of about 50 to 100 in the penumbral region. Since the Rayleigh number is only around 25000, we

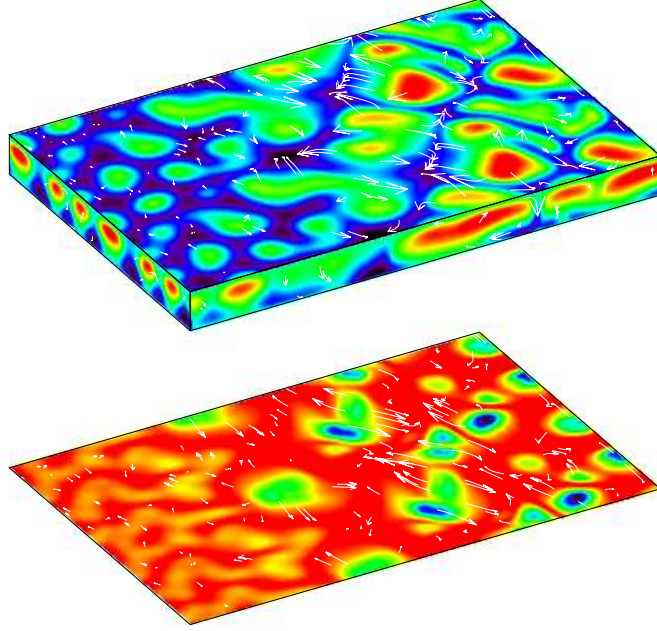


Figure 5.21: Temperature plot for the solution with the ‘flux sheet’ initial condition, in the wider box ($\lambda_y = 4$), and with the upper boundary condition changed to be perfectly conducting. (As usual, two copies of the domain are shown in the y -direction.) Note that the result shows a cellular pattern at the right-hand end (similar to Figure 5.17) rather than a roll pattern (as was seen in Figure 5.16).

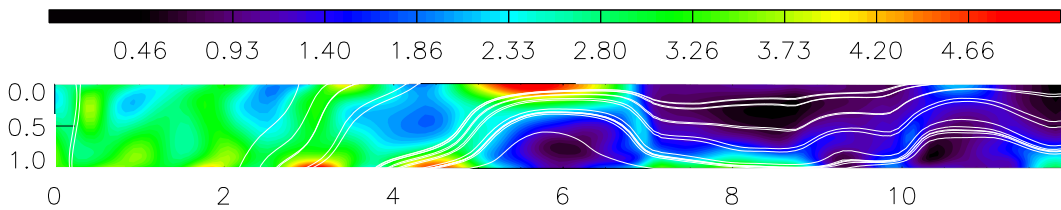


Figure 5.22: Plot of the magnetic field for the solution shown in Figure 5.21 (after averaging in the y -direction). The colours indicate $|\mathbf{B}|$ and the white lines are field lines. Notice how flux accumulates near the top boundary (in certain places); flux is still rising upwards as before, but this time it cannot actually move up and out of the box, because of the new boundary condition.

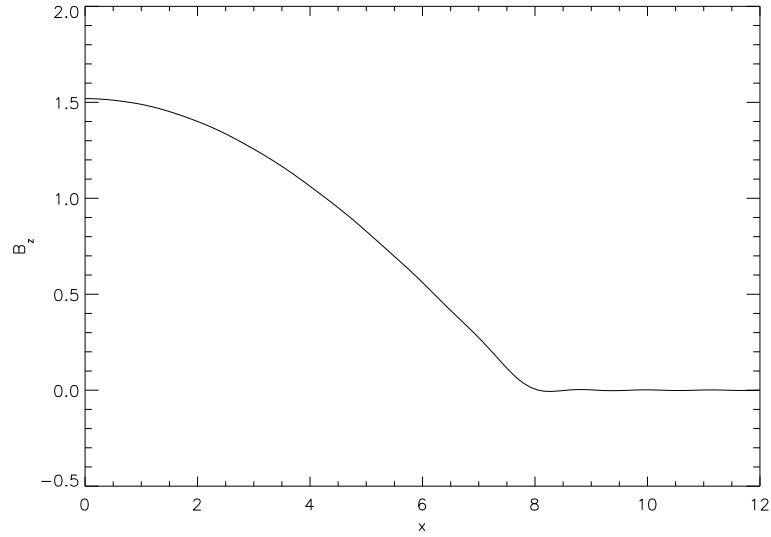


Figure 5.23: B_z as a function of x at the lower boundary, for the initial condition given by equations (5.35)–(5.36). (Compare Figure 5.14 above.)

would expect the convection to be shut down for an effective Q of about 1000 or above (very roughly); therefore, we cannot increase Q all that much without shutting down the convection in the umbra.

A new initial condition

To mitigate these issues, we looked at reworking the initial condition slightly, in order to reduce the contrast in field strength between umbra and penumbra in the model. This can be done by increasing x_0 in (5.26) above. In addition, we took this opportunity to slightly modify the initial condition (in order to reduce the contrast in field strength still further), by replacing (5.26) with the following new profile for B_z at the base of the layer:

$$B_z = \begin{cases} (\lambda'_x/x_0) \cos(\pi x/2x_0) & (0 \leq x \leq x_0) \\ 0 & (x_0 \leq x \leq \lambda'_x) \end{cases}. \quad (5.35)$$

This leads to the following replacement for equation (5.34) for A_k :

$$A_k = \frac{4\lambda'_x}{\pi^2} \left(\frac{2}{k} - \frac{1}{k - \lambda'_x/x_0} - \frac{1}{k + \lambda'_x/x_0} \right) \cos \left(\frac{\pi k x_0}{2\lambda'_x} \right). \quad (5.36)$$

The new choice (5.35) for B_z is similar to the previous one, but it gives a slightly less sharp transition between the umbral and penumbral fields (e.g. the previous ‘overshoot’ in B_z , where B_z became negative for a short distance, has been eliminated), and reduces slightly the difference in field strength between the two.

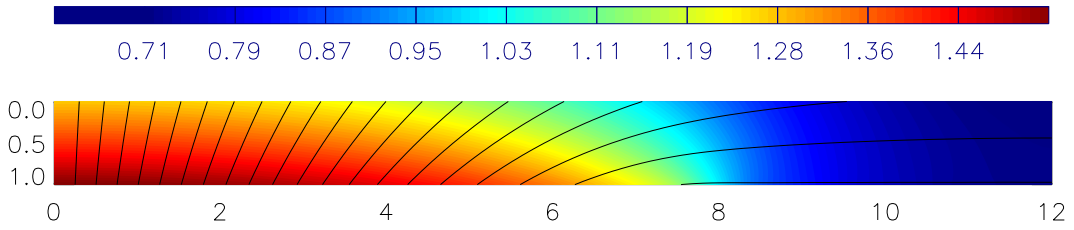


Figure 5.24: *Illustration of the initial condition given by (5.35)–(5.36). The colours indicate $|\mathbf{B}|$ and the curves are field lines. (Compare Figure 5.15 above.)*

The new B_z is plotted in Figure 5.23 (for the particular choice $x_0 = 7.9$), which may be compared to Figure 5.14 (solid curve) above. Note that the Gibbs phenomenon does not affect this new initial condition (because equation 5.35 does not contain a discontinuity) and so the number of modes N is irrelevant (provided it is sufficiently large).

The new field configuration itself is shown in Figure 5.24. Notice that the umbral region is somewhat wider than before, but in return we have a smaller contrast in the field strength between umbra and penumbra, meaning that we can have higher field strengths in the penumbra without necessarily needing excessive field strengths within the umbral region.

Using this new field configuration, runs for three different values of Q were performed, namely 500, 1000 and 1500. The results are shown in Figures 5.25–5.28. (Note that the Q values for these results are not directly comparable with the Q values for the previous results, because the new initial condition for \mathbf{B} is scaled differently.)

In the case $Q = 500$ (Figure 5.25), we see a hexagonal pattern of convection in the umbra (as usual). In the penumbra, a time-dependent, constantly changing pattern is seen. Sometimes the pattern looks fairly roll-like (lower panel of the Figure), and sometimes it looks more cellular (upper panel), albeit with some of the cells being stretched out along the x -direction. Therefore, our new initial condition has met with some success, because the penumbral pattern is now showing signs of anisotropy (and we have not lost the hexagonal umbral convection). However, in this regime the field is not quite strong enough to give a sustained pattern of rolls (or roll-like structures) in the penumbra.

When we increase Q to 1000, a different pattern is seen. The umbral convection is shut down for a field this strong. The penumbral convection initially takes the form of

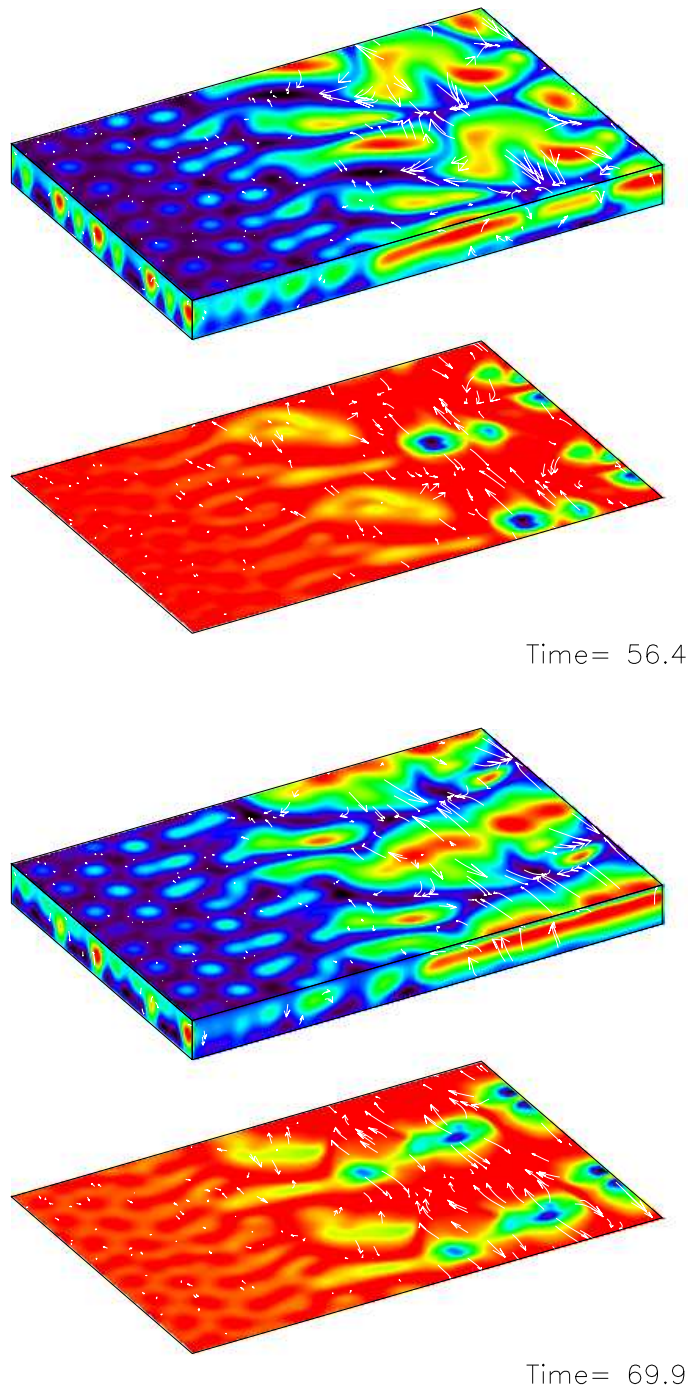


Figure 5.25: Results for the initial condition depicted in Figure 5.24, and with the perfectly conducting upper boundary condition. The result for $Q = 500$ is shown. The solution is time-dependent; snapshots at two different times are shown.

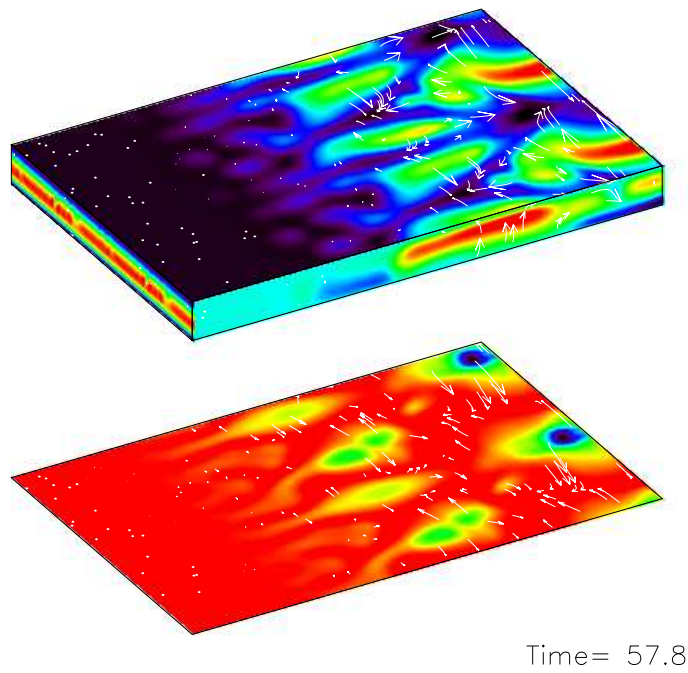
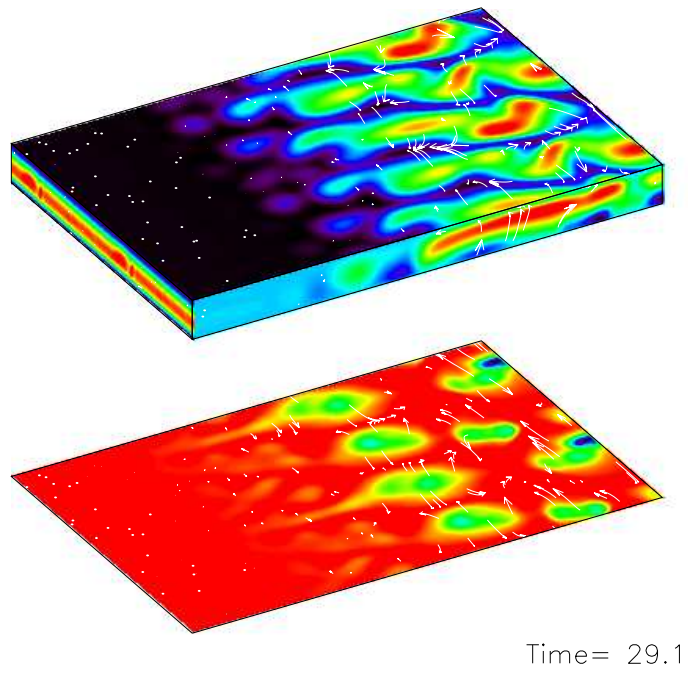


Figure 5.26: As Figure 5.25, but for $Q = 1000$. The upper panel shows the initial (transient) state and the lower panel shows the new state that appears after a certain time.

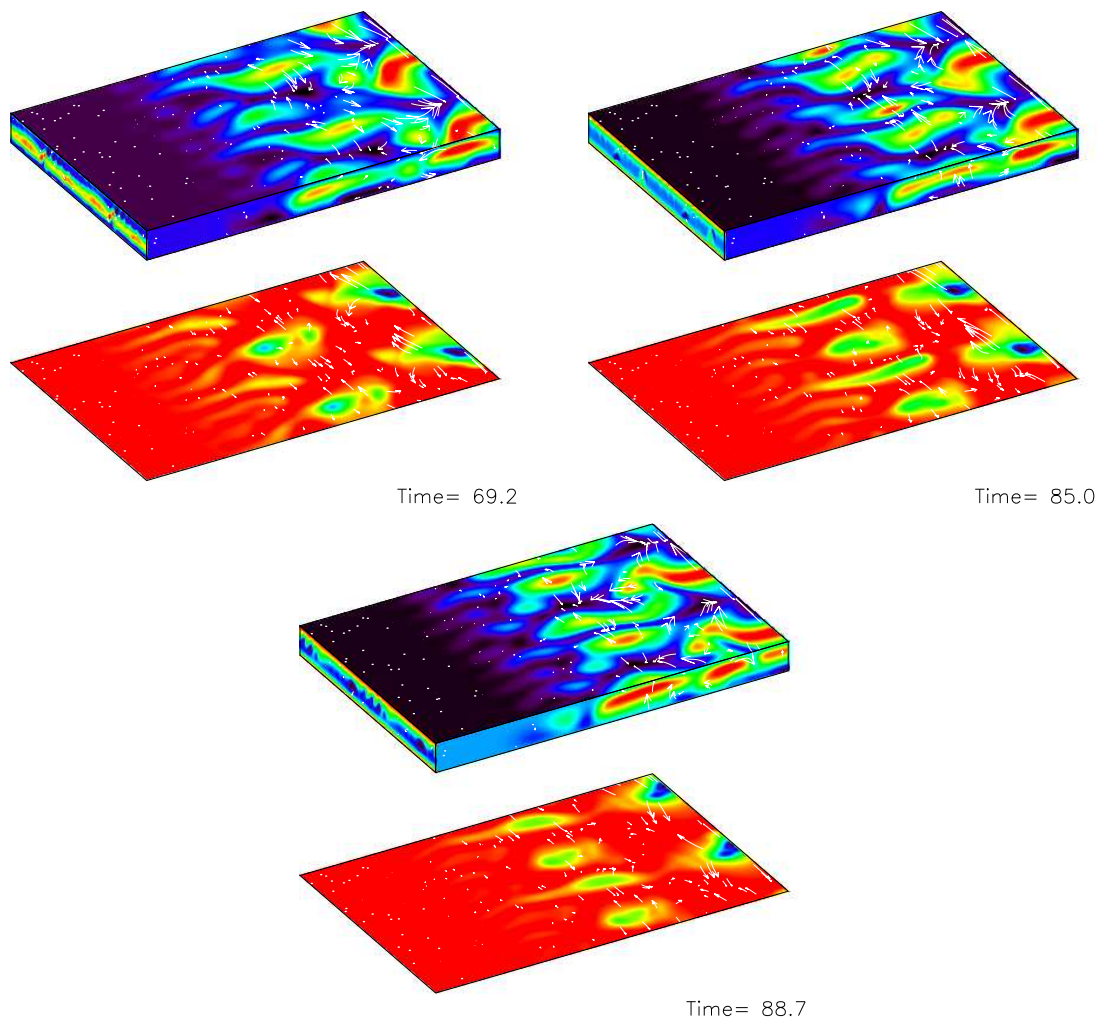


Figure 5.27: The continuation of the simulation for $Q = 1000$ beyond Figure 5.26. Snapshots are shown at three different times. The basic pattern (from Figure 5.26, lower panel) persists, although it is time-dependent in the details. In particular the rolls near the centre appear to wax and wane in amplitude in a chaotic manner.

rolls (Figure 5.26, top panel), but this is only a transient state and after a short time, the rolls break up into two separate groups (bottom panel, same Figure). We now see four smaller rolls near the centre, and two larger roll-like structures at the right-hand end of the domain (actually, there are really only two of the smaller rolls and one of the larger rolls, because the computational domain is only half of what is shown in the Figure).

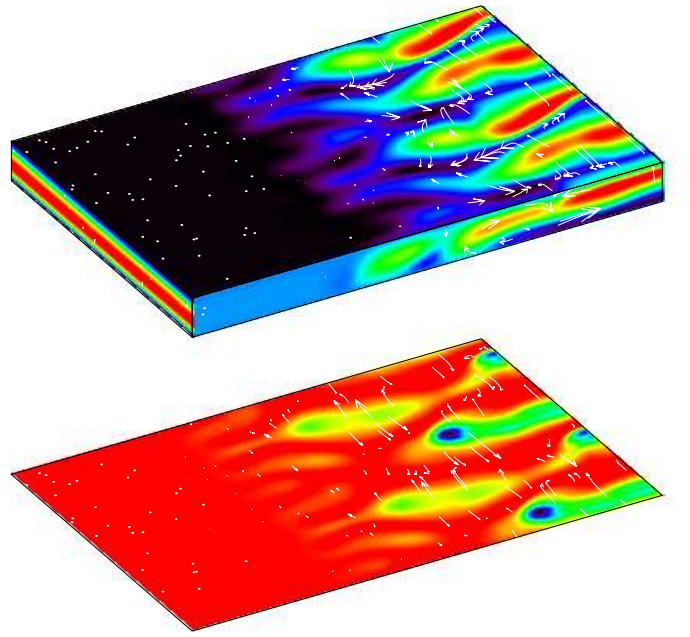
This basic pattern persists throughout the rest of the calculation. The fine details are however time-dependent: the smaller rolls (near the middle of the box) wax and wane in amplitude, while the larger rolls (at the right-hand end) undulate from side to side somewhat. These changes appear to be chaotic – there is no well-defined period of oscillation, for example; rather the rolls wax and wane in a seemingly random fashion. This time-dependence is illustrated in Figure 5.27, which shows snapshots of this final phase of the solution at three different times.

When $Q = 1500$ (Figure 5.28) we see much the same behaviour as for $Q = 1000$. The only major difference is that the evolution is somewhat slower this time, and it takes longer for the initial roll pattern (upper panel of Figure 5.28) to change into the final pattern (lower panel).

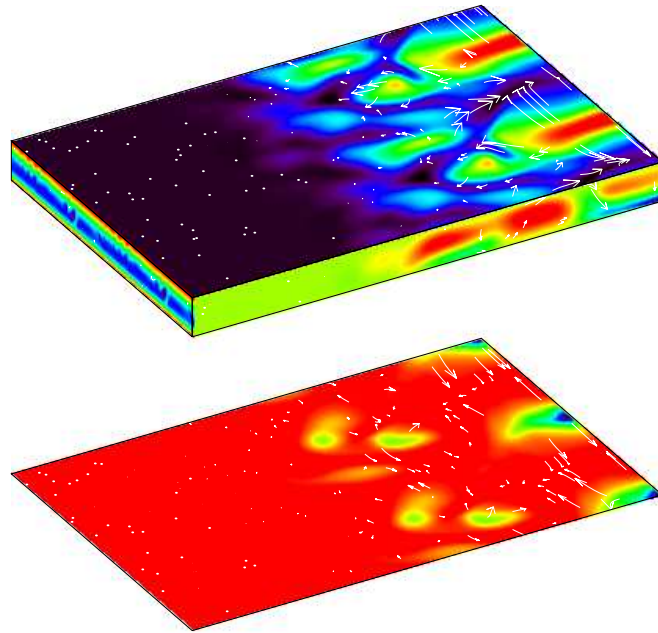
Unfortunately, we did not have time to run the calculation for $Q = 1500$ for quite as long as we would have liked; the calculation was stopped soon after the roll pattern seen in the lower panel of Figure 5.28 first appeared. We cannot say whether this pattern will continue to evolve in a similar way to the previous case (see Figure 5.27) or in some new way; although one might perhaps expect a less vigorous time-dependence in this new case, because of the higher field strength. (Indeed one might expect a transition from the chaotic time-dependence found for $Q = 1000$ to something more regular as Q is increased, or alternatively as R is decreased; however, we did not have time to investigate this idea.)

Note that for $Q = 1000$ and above, the field within the umbra has become too strong for convection to take place there. This is not what we intended (since a model of sunspot convection should ideally contain convection within both the umbra and penumbra), but it is still worth investigating these higher values of Q , since they show what happens to the ‘penumbral’ convection when the field strength is increased.

It seems therefore that our model has two different regimes. The first of these, found at lower values of Q , has convection within the umbra, and a time-dependent pattern within the penumbra. The penumbral pattern does show some anisotropy, but



Time= 26.0



Time= 84.4

Figure 5.28: As Figure 5.25, but for $Q = 1500$. The upper panel shows the initial (transient) state and the lower panel shows the final state.

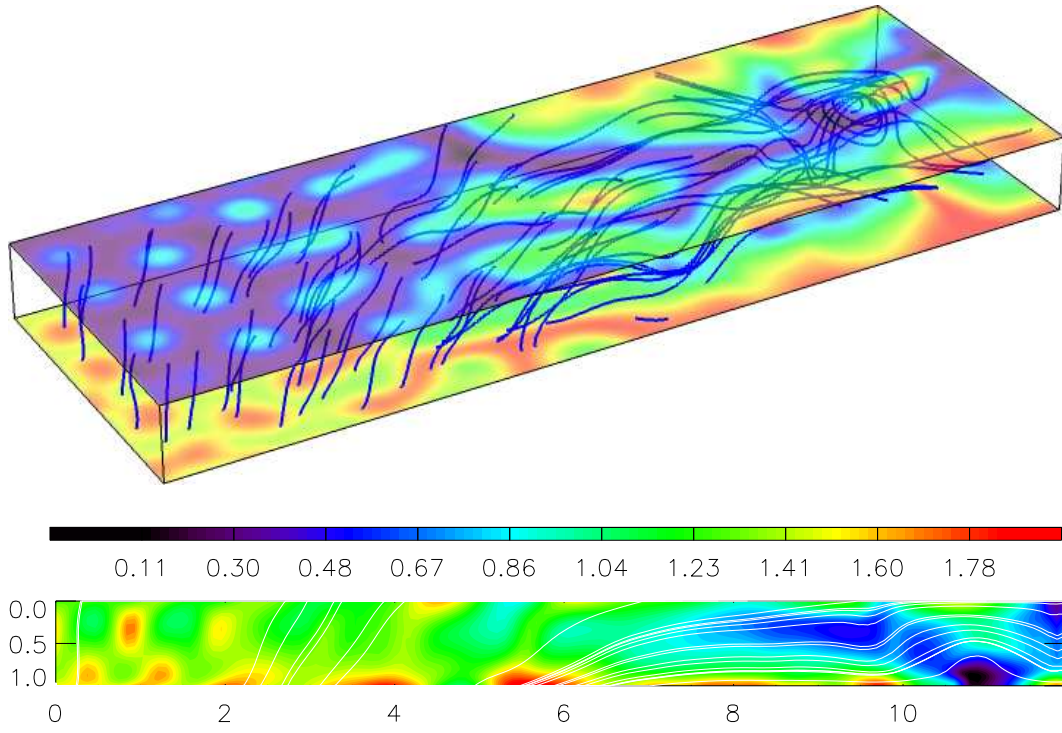


Figure 5.29: Magnetic field configuration for the $Q = 500$ run. (This is a snapshot corresponding to the lower panel of Figure 5.25.) The top panel shows the field lines in 3D, with temperature plots shown at the top and bottom surfaces for reference (with red-yellow-green indicating warm fluid and blue-black indicating cool fluid). The bottom panel shows the y -averaged magnetic field as a 2D plot. (The colour bar indicates the scale for $|\mathbf{B}|$ on the 2D plot.)

no coherent roll-like structures are found. In the second regime, found for higher Q , the penumbral field is strong enough to produce more recognizable field-aligned rolls in the penumbra, but convection is no longer present within the umbra.

The magnetic field structures for the cases $Q = 500$ and $Q = 1000$ are shown in Figures 5.29 and 5.30 (as both three-dimensional plots and y -averaged two-dimensional plots). The field is slightly more tangled and disordered in the former case than in the latter (as might be expected). Notice that in both cases, a significant amount of flux has risen to the upper boundary (visible as a bright green region near the top of the y -averaged plots); this indicates that flux is still rising upwards in the simulations, although this time it is prevented from leaving the box, and instead comes to rest against the upper boundary.

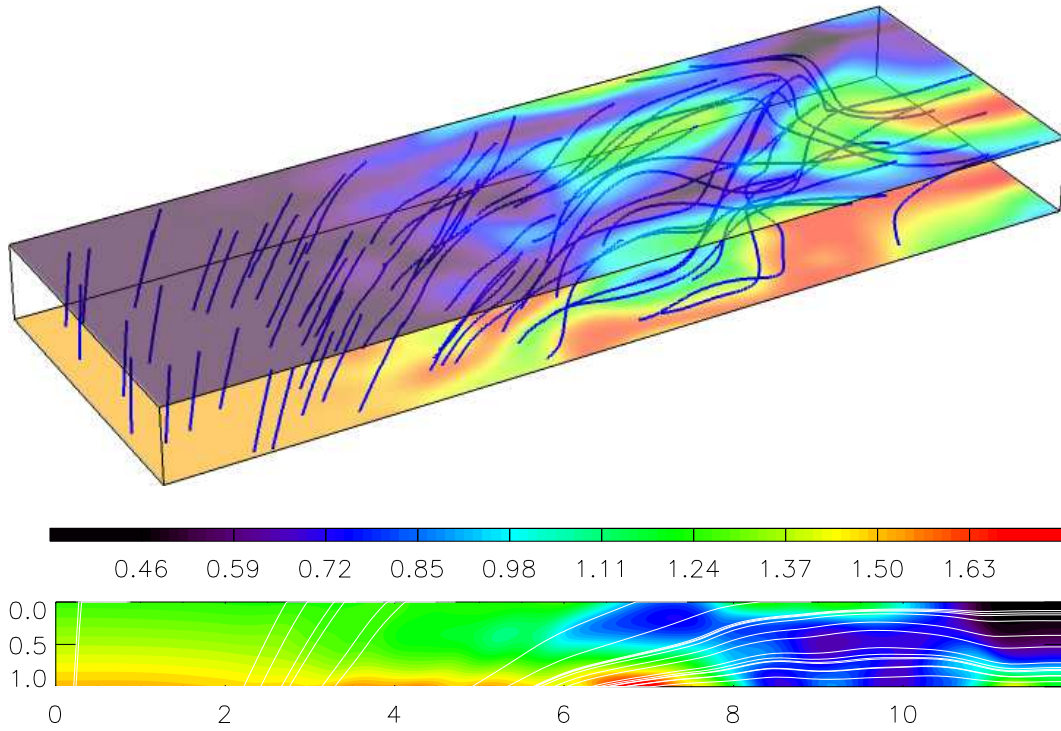


Figure 5.30: As Figure 5.29, but for the case $Q = 1000$. This Figure is a snapshot, taken at a time corresponding to the lower panel of Figure 5.26.

5.7 Conclusions

In this chapter we have looked at three-dimensional magnetoconvection simulations, with both uniform and non-uniform fields.

5.7.1 Runs with uniform fields

Initially we looked at runs where the magnetic field was uniformly inclined at an angle ϕ to the vertical. The aim was to look for changes in the convective pattern with varying ϕ . Although this has been briefly considered before by Hurlburt et al. (2000), they presented only a limited number of results, and used fairly small computational boxes. Here we have made a more thorough investigation, and built up a picture of the transitions that occur as both R and ϕ are varied (although, for reasons of time, we have considered only a single value of Q).

The results clearly show a transition from hexagonal convection for small ϕ , to convection in elongated cells (stretched out along the direction of tilt) for moderate ϕ , to rolls (aligned with the tilt) for larger ϕ . This was more or less in line with what we expected from the weakly nonlinear results of Chapter 3. One difference is that while

hysteresis (bistability) between rolls and hexagons was predicted in the weakly nonlinear analysis, we do not find anything similar in the fully nonlinear model of this chapter.

This can be explained in more detail by comparing Figure 5.2 (on page 161) with the lower part of Figure 3.9(b) (on page 83). Although these diagrams are not identical, we do see a region of rolls in Figure 5.2 for large ϕ or small R (corresponding to region R_1 in Figure 3.9), and a region of cells above this (corresponding to region H in Figure 3.9). According to Figure 3.9, the region of bistability would be found just above this; however, in Figure 5.2, we find only turbulent cellular solutions in this area. This may indicate that we have already left the weakly nonlinear regime before we ever get to the region where bistability was predicted, and hence we do not see the bistability in the fully nonlinear model. (Note however that we only looked at a limited range of parameters – and in particular only one value of Q – so it is always possible that bistability might be present at some other set of parameter values, that we did not investigate.)

Another feature that we find in the nonlinear simulations of this chapter is that the cellular patterns are sometimes subject to a ‘wavy’ modulation (e.g. Figure 5.2, bottom row). This kind of pattern could not be detected by the weakly nonlinear models of Chapter 3, because it involves modulations to the basic wavenumber of the pattern. It appears that we are seeing an instability of the basic, ‘straight’ cellular patterns; but we do not at present have an explanation for why this occurs.

One additional feature worth mentioning is that the results often show a difference between the convective pattern at the top of the box and the pattern at the bottom. See for example Figure 5.3, bottom left, where individual cells are seen near the top, but rolls are seen at the bottom. This is an interesting result since it illustrates that the pattern of convection at the surface of a sunspot (which is what we observe) may not be the same as the pattern in the deeper, sub-surface layers – something that must be borne in mind when interpreting observations of penumbrae.

In summary, the main result of the runs with uniform fields is to illustrate how the convective pattern changes for different values of ϕ , and in particular how it varies from a hexagonal or cellular pattern for small ϕ , to a pattern of rolls when ϕ is large.

The value of Q used to obtain these results corresponded to a strong field regime, and hence strongly anisotropic patterns were found, with the patterns aligning themselves with the field. An interesting future project might be to investigate how the results vary when smaller values of Q are used. In such a case we would expect to see a transition to more isotropic patterns, eventually leading to the granulation-like patterns that are

found in field-free convection, and it would be interesting to investigate the details of the transition between these two extremes.

5.7.2 Runs with non-uniform fields

The runs using non-uniform fields have shown some interesting results. Here we have run a number of simulations in a sunspot-like field configuration, in which the field varies from vertical at one end of the box, representing the umbra, to horizontal at the other end, representing the penumbra. The aim was to see whether a transition similar to the umbra-penumbra transition in a real sunspot could be observed, and to look at the kind of convective patterns found within each of these two regions of the simulation.

Summary of the results

We started off by using an ‘arch’-shaped magnetic field (section 5.5). The results here did indeed show a sharp transition between hexagonal patterns in the ‘umbra’ of the model and roll-like patterns in the ‘penumbra’. The hexagonal patterns correspond to what is expected in a near-vertical field (see e.g. Rucklidge et al. 2000 or Weiss et al. 2002), while the roll-like patterns can be compared to those found by Hurlburt et al. (2000).

Unfortunately, the region of roll-like patterns turned out to be fairly small in our ‘arch’ model, whereas we would actually like a model with a relatively large penumbral region, since this is where the more complex and interesting phenomena are found in real sunspots.

Therefore, we switched to using an initial condition based on a ‘flux sheet’ configuration (section 5.6). The results again showed a sharp distinction between ‘umbral’ and ‘penumbral’ forms of convection. (In fact, the presence of this transition was a robust feature that was seen in all of our simulation runs.) However, this time the actual pattern seen within the penumbra was different. Instead of roll-like structures, we saw a more disordered, almost granulation-like pattern (Figure 5.17). This was because the model had a somewhat reduced field strength in the penumbra, compared to the ‘arch’ model; this was not helped by the fact that much of the flux actually escaped through the upper boundary of the simulation (see e.g. Figure 5.19).

So, although the ‘flux sheet’ model does have a larger penumbral region (as desired), it also has a much stronger difference in the field strength between the umbra and the

penumbra, and consequently only a relatively weak magnetic field within the penumbra. It seems likely that a stronger field would be needed before more ‘penumbra-like’ (more anisotropic) forms of convection would be seen.

Our next step, therefore, was to look into ways of strengthening the magnetic field within the penumbral region. Two changes were made. Firstly, the potential field boundary condition at the top of the layer was replaced by a perfectly conducting condition; this prevented magnetic flux from escaping through the top boundary. Secondly, the initial condition was slightly rearranged, in order to strengthen the field within the penumbral region of the domain a little.

These changes seemed to have the desired effect, at least to a certain extent. For example, with a relatively low Q (of 500), we found that the penumbral convection pattern was more anisotropic than what was previously seen; however, the pattern was also somewhat disordered and time-dependent, with roll-like features continually appearing and disappearing. When Q was increased (to 1000 and above) a roll-like pattern was found in the penumbra, but unfortunately this time there was no convection within the umbra (because of the high field strength there). The penumbral pattern in this case showed an interesting ‘split’ structure with a pair of smaller rolls near the centre of the domain, and one larger roll-like structure at the right-hand edge (where the field was slightly weaker). The penumbral pattern was also somewhat time-dependent with rolls waxing and waning irregularly as a function of time. It would have been interesting to look at these roll structures in greater detail, but unfortunately we did not have time to do so.

Strengths and limitations of the model, and relation to other work

Although the simulations presented above present a reasonable first step towards a ‘unified’ model of sunspot convection (i.e. one containing both umbra-like and penumbra-like regions), there still remains much more work to be done. Perhaps the best way to illustrate the shortcomings of our results is to enumerate the features that we would like an ‘ideal’ model to have:

- convection within both the umbral and penumbral regions (this is important if we wish to investigate the interaction between the two forms of convection);
- a relatively strong field within the penumbral region, so that patterns do become noticeably anisotropic there;

- a relatively large penumbral region (compared to the size of the umbral region), so that we can investigate the more interesting penumbral behaviour in greater detail.

The ‘arch’ model does not satisfy the third point on this list. The ‘flux sheet’ model is better in this respect, since it allows control over the relative width of the umbral and penumbral regions. However, this model has a strong contrast in the field strength between the umbra and penumbra, which means that either the first or the second point on the above list would not be satisfied. The closest that we came to satisfying both points was the runs shown in Figures 5.25 and 5.26, but even here, one must choose between either a fairly weak field strength within the penumbra (Figure 5.25), or a field strength too strong for convection in the umbra (Figure 5.26).

Unfortunately, we have not found a model satisfying all three of the above points, although we hope that future, more complex models might be able to address this. One approach might simply be to look at higher Rayleigh numbers, which would allow us to increase Q to higher values while still having convection within the umbra. (Our runs have all been at relatively low Rayleigh numbers.) We did not do this because we encountered numerical problems; it was found that higher resolutions were required when higher Rayleigh numbers were used.

An alternative method would be to find another initial condition which produces a smaller contrast in field strength between umbra and penumbra. This is quite difficult to do if one restricts oneself to potential field initial conditions (as we have done). However, if one is willing to start from a non-equilibrium initial condition, then more interesting results might be obtained. (The downside of this would be that since the initial condition is not an equilibrium, it would rearrange into some new configuration right from the beginning of the calculation; although one could argue that this was already happening to some extent in our own results, as the magnetic field would eventually be rearranged somewhat once the convection got started.)

There is only one other published model so far that attempts to include both umbral and penumbral regions of convection within the same simulation; this is the cylindrical model described by Hurlburt et al. (2000).³ Here, the authors start with a ‘flux

³Here we are excluding two-dimensional or axisymmetric models, such as those of Hurlburt et al. (2002) or Botha et al. (2006), since, although these do a good job of describing the overall structure of umbrae and penumbrae, they do not show the detailed three-dimensional convection patterns, which is what we are mainly interested in here.

tube’ structure (the cylindrical equivalent of our ‘flux sheet’ models) and run a non-axisymmetric simulation within a wedge (consisting of 1/12 of a full cylinder). The results are similar to our later ‘flux sheet’ runs, in that a pattern of rolls is found within the penumbra, in this case appearing as radial spokes. These are not dissimilar to the rolls found in our own simulations, although our rolls are more disordered and time-dependent than those found by Hurlburt et al. (2000). In the umbral region of the model of Hurlburt et al. (2000), it appears that no significant convective motion is seen; in this sense their model suffers from the same weakness as our model, in that the first point on the above list is not satisfied. (It should be pointed out that the results of Hurlburt et al. 2000 are preliminary. According to Botha et al. 2006, the authors intend to carry out more detailed simulations of this kind, as the next stage of their research project.)

Magnetic buoyancy, flux pumping and boundary conditions

Another very important point concerns the maintenance of the magnetic field configuration against magnetic buoyancy. In our simulations, we set up a carefully constructed magnetic field at the beginning of the simulation; but there is no guarantee that this field structure will be maintained (in any recognizable form) once convection starts up.

Indeed, we found in some of our later simulations that a significant amount of flux escaped through the upper boundary of the computational box. This was the motivation for introducing a perfectly conducting upper boundary condition. However, real sunspots do not have perfectly conducting ‘lids’ over them, so the use of this boundary condition is somewhat unrealistic. On the other hand, we do need *some* method of ensuring that our chosen field structure is maintained against magnetic buoyancy. A proper solution (in the long term) would be to construct a model in which the flux could be held down self-consistently; but in the short term, the use of the conducting upper boundary condition allowed us to investigate the convective patterns that can occur within a sunspot-like field structure, without becoming overly concerned about how to maintain that structure against magnetic buoyancy.

This does of course raise the question of how fields within real sunspots are maintained against buoyancy. It has been suggested (e.g. Weiss et al. 2004) that turbulent flux pumping plays a key role here. In our simulations, we did find that flux was pumped downwards in some cases (for example, our ‘arch’ runs), but not in others.⁴ Clearly the

⁴It should also be pointed out that placing a stable layer underneath the main convecting region, as

balance between magnetic buoyancy and flux pumping is something that requires further investigation.

In fact, this is an idea that is currently being looked at in a calculation by Brummell et al. (2006). This calculation uses a similar field configuration to our own; the differences are that the Rayleigh number is much higher, and the box is narrower in the x direction, but deeper in the z direction (with a convectively stable layer being added underneath the main convecting layer). The aim of the calculation is to determine how effective flux pumping is at holding down the magnetic field against buoyancy. The convection in this calculation takes the form of field-aligned rolls, somewhat similar in appearance to the rolls seen in Figure 5.16 above; and preliminary results indicate that the magnetic flux is indeed pumped downwards. These results should help to clarify under what circumstances flux pumping is and is not effective within this kind of field configuration.

Filamentary structure

As a final point, we have not yet mentioned the intricate structure that is found within real penumbrae, such as the bright and dark filaments or the interlocking-comb magnetic structure. Nothing like this was seen in our model; instead of different structures representing bright and dark filaments, we see only the one ‘roll’ pattern within our penumbral region.

One theory that is emerging is that the filamentary structure is caused by an instability that leads to fluting around the edge of a sunspot and to the beginnings of a filamentary structure (a recent description of this idea is given by Weiss et al. 2004). For example, the calculation of Tildesley (2003a,b) (see also Tildesley and Weiss 2004) found an instability within an idealized Boussinesq configuration which led to a filamentary structure. The calculation of Hurlburt et al. (2000) also found a kind of fluting instability (this time within cylindrical geometry).

Therefore, it might be fruitful to look for similar instabilities (along the y -direction) within our model (or one like it). One point to note here is that Tildesley’s instability had a fairly long wavelength in the y -direction, which may indicate that a fairly wide computational box would be needed before any such instability would be seen.

was done by Tobias et al. (2001) for example, might change these results. For example, in cases where the flux was pumped downwards, we found a tendency for flux to build up against the bottom boundary (see e.g. Figure 5.10). If the flux could instead have been pumped down ‘through’ the bottom boundary into a convectively stable region underneath, then the simulation might have had very different results.