UNIVERSITY OF SOUTHAMPTON

MONTE CARLO SIMULATIONS OF RANDOM ANISOTROPY MAGNETS

by

David Richardson Denholm

A thesis submitted for the degree of

Doctor of Philosophy

Department of Physics

March 1995

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SCIENCE

PHYSICS

Doctor of Philosophy

MONTE CARLO SIMULATIONS OF RANDOM ANISOTROPY MAGNETS

by David Richardson Denholm

We have used a parallel computer to perform Monte Carlo simulations on the Harris-Plischke-Zuckermann lattice model of Random Magnetic Anisotropy, using the standard Metropolis algorithm. We have studied systems with two spin components in both two and three spatial dimensions, varying temperature, external magnetic field, anisotropy strength and other thermodynamic parameters which may possibly be relevant.

In the absence of randomness, the system reduces to the XY (planar rotator) model, and we find that the usual Kosterlitz-Thouless behaviour in the two-dimensional system survives the addition of randomness at intermediate temperatures, but at low temperatures the system shows some disordered features which might be identified with a spin glass phase. In three dimensions, the randomness destroys the long range order present in the pure system, replacing it with a phase with infinite susceptibility which exhibits algebraic order with a non-universal exponent. If the system is quenched, however, a phase with exponentially decaying correlations with properties predicted by Chudnovsky is found. For low randomness this apparently non-equilibrium phase is destroyed at a thermal depinning temperature dependent on the degree of randomness.

Contents

List of Figures

List of Tables

Preface

The original work in this thesis is as follows:

- The established Chudnovsky-Imry-Ma theory has been extended slightly in sections 2.8.3, 2.8.4 and section 2.8.9. The last of these extensions has been previously published [1].
- Chapter 4 describes the details of the implementation of the simulation on a parallel computer.
- Chapter 5 reports on the results from simulations on the HPZ Hamiltonian with two-component spins in two spatial dimensions, varying temperature, anisotropy strength and external field strength. Some of these results have previously been published in [1–3].
- Chapter 6 reports on the same system in three spatial dimensions. These results are to be submitted for future publication.

The ambiguity over the FWA phase (section 2.8.7) and the theory for the energy due to local fluctuations in spin directions (section 2.8.5) were pointed out by professor Alan Bray at the Ph.D. *viva voce*.

Acknowledgements

I am indebted to my supervisors Brian Rainford and Tim Sluckin, for first of all getting me into this research and (finally) getting me through it. Thanks in particular to Tim for knowing what to say when confidence was flagging.

Thanks go to Peter de Groot, John Chalker, Doug Newman and Wayne Saslow for useful discussions on the physics of the problem, and to Brian Carpenter and the High Performance Computing Group (formerly the Parallel Computing Support Group) of the university Computing Services for assistance in getting started with parallel computing. I must also thank the anonymous hordes who create and maintain the plethora of free software which was used extensively in this work; in particular to the maintainers of LaTeX, ghostscript and gnuplot, without which the production of this thesis would have been even more painful — I have tried to do my bit in return.

Thanks to family and friends for showing such remarkable patience during this endeavour, and to all at the physics dept. for making my time here so enjoyable.

Finally, I thank the SERC and the Southampton University Research Council for providing financial support — I think Al also deserves a mention here!

Chapter 1

Introduction

In this thesis an investigation of random anisotropy magnets, using Monte Carlo simulations performed on a parallel computer, is presented and discussed. Despite many years of study, the physics of random systems remains a topic of tremendous interest. To a lesser extent, the work was also intended as an appraisal of the utility of parallel computers in physics simulation, though the physical results presented herein stand independently from the computing-science aspects of the work.

In the remainder of this chapter, the general phenomenon of magnetism is discussed, and the experimental properties of random magnetic systems are reviewed. There is room to do little more than scratch the surface, and the reader is referred to standard references such as [4–7], which were used as source materials for this introduction. This is also true of statistical mechanics, largely used by physicists in the study of many-body systems such as magnets.

In chapter two, the properties of the model Hamiltonian introduced by Harris, Plischke and Zuckermann (HPZ) [8] for random anisotropy systems are explored. This is followed, in chapter 3, by a discourse of the classic Monte Carlo algorithm of Metropolis *et al.*, together with technical discussion of the problems of the technique. In chapter four, an introduction to parallel computers is given, together with specific details concerning our implementation of the Monte Carlo method. Because this dissertation is intended for an audience with a wide range of backgrounds, an attempt has been made to keep these chapters more or less independent. Readers already familiar with the theoretical aspects of the HPZ model can move straight to chapters three and four, while those with a Monte Carlo background can use chapter two as an introduction to the model, or chapter four as a basis for commencing simulation work on parallel computers.

The results of our extensive simulations on the random anisotropy model with two spin components in two spatial dimensions are presented in chapter five, while in chapter six the results from our preliminary work on the system in three dimensions are presented and discussed.

Finally, in chapter seven the work is summarised and suggestions are made for possible directions for future work.

A detailed description of this implementation can be found in the appendix, in the hope that other workers can continue to use the current program without the need to learn about parallel programming techniques in detail. A 'simulation control language' has been defined, so that a rich variety of simulations can be be performed without the need to modify the program source.

1.1 Magnetism and Magnetic Materials

For thousands of years, man has been fascinated by magnetism; the ancient Greeks were intrigued by the way the lodestone would attract pieces of iron, and the compass has long been used as a navigation aid by detecting the direction of the Earth's magnetic field. An important step in the study of magnetism was the forging of links between electricity and magnetism by the work of Oersted, Ampère, Faraday and others. All currents (moving charges) produce magnetic fields, and all magnetic dipoles are due to currents, microscopic or macroscopic.

All materials respond to an imposed magnetic field. The magnetic *susceptibility* is defined as the ratio of the resulting moment to the applied magnetic field. Some materials are *diamagnetic*, exhibiting a negative susceptibility, in opposition to the applied field, while other materials are *paramagnetic*, forming moments parallel to the field. A few materials, such as iron, retain their magnetic moment even after removal of the field. These are called *ferromagnets*. In fact *all* materials show diamagnetism, but this is masked by the much larger paramagnetic and ferromagnetic properties where they are present. There are also more esoteric forms of magnetic ordering, which will be discussed.

It is now well known that atoms are composed of negatively-charged electrons orbiting positively-charged nuclei, and the differences between the elements can be related to the differences between the electronic structure within the atoms.

For a complete understanding, Quantum Mechanics is required — a topic beyond the scope of this introduction. Suffice to say that magnetic materials are composed of atoms with a net dipole moment, and non-magnetic materials are composed of atoms with no net magnetic moment. (There are also effects due to the intrinsic magnetic moment of the atomic nucleus, but these moments are three orders of magnitude smaller than the electronic moments.) It is the interaction between the atomic moments (also called *spins*) that produces the interesting physics.

1.1.1 Paramagnetism

In a paramagnet, the microscopic magnetic moments of the atoms show a negligible interaction with each other. Consequently, in zero field, they are arranged randomly, and there is no net magnetic moment.

When an external field is applied to a single magnetic dipole, a torque is exerted upon that dipole, trying to align it with the field. This results in a net magnetic moment. Pierre Curie discovered that, at 'ordinary' fields and temperatures, many materials produce a magnetic moment directly proportional to the field, which increases as the temperature is reduced: Curie's law states that the magnetic susceptibility is given by

$$\chi = \frac{C}{T} \tag{1.1}$$

where C is a constant depending on the material under study. For large fields or low temperatures, the magnetisation approaches saturation once all the moments are aligned with the field. Thermal agitation tends to knock the dipoles away from alignment, so that a larger field is required to cause a given magnetisation as the temperature is increased (section 1.2.1).

In solids, the microscopic dipoles *can* interact with each other, but Curie's law is obeyed by dilute alloys for example, in which the magnetic atoms are spatially separated by non-magnetic atoms.

1.1.2 Ferromagnetism

The label 'ferromagnetism' derives from the Latin *ferrum*, meaning iron. The magnetic properties exhibited by iron, and similar elements, stem from the way

(a) Ferromagnetic ordering

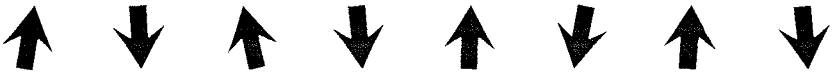

(b) Antiferromagnetic ordering

(c) Ferrimagnetic ordering

Figure 1.1: Schematic spin configurations for ferromagnetic ordering (a), antiferromagnetic ordering (b), and ferrimagnetic ordering (c).

the microscopic magnetic dipoles tend to align themselves with their neighbours. In this way, samples of the materials show long-range magnetic order, with spins aligned over long length-scales, as illustrated in figure 1.1(a). When a small field is applied, the dipoles are rotated towards the field direction as before, but in the case of ferromagnets, the rotation of the dipoles reinforces the force on the neighbouring dipoles, and a large net magnetisation results. Thus the net magnetic moment is very nonlinear in the external field.

Again, the full details of this *exchange interaction* which aligns the dipoles require quantum mechanics, and is beyond the the scope of this thesis. In general, the exchange interaction arises from electrostatic repulsion between negatively-charged electrons — quantum mechanics prevents two spins with aligned dipole moments from approaching closely, while it places no restriction on electrons in different quantum mechanical states. Thus aligned spins are on average further apart than misaligned spins, and therefore have a lower average coulomb energy.

It is found that above a certain temperature T_c known as the *Curie point*, long range ordering ceases, and ferromagnetic substances revert to paramagnetic behaviour, approximately obeying a modified form of Curie's law (equation 1.1) known as the Curie-Weiss law:

$$\chi = \frac{C}{T - T_c} \tag{1.2}$$

where T_c is the 'Weiss constant', and has dimensions of temperature.

Below T_c, the system shows a *spontaneous* magnetisation, in the absence of an external magnetic field. This spontaneous magnetisation is large at low temperatures, and decreases with increasing temperature, as indicated schematically in figure 1.2.

1.1.3 Phase Transitions

The first quantitative analysis was made by Weiss, who modelled the interaction of the neighbouring dipoles as a 'mean field' at each site, and was able to find a self-consistent solution. This technique is a very important one, being much used in statistical mechanics, and will be dealt with in more detail in section 2.3. While it unfortunately gets much of the detail wrong, it does at least qualitatively predict a transition from a disordered state to an ordered state as the temperature is lowered — a *phase transition.*

This is a specific example of a very general phenomenon. Many other systems show phase transitions, one common example being the transition from water to ice at 0°C. At first sight, these two phenomena appear quite unrelated, but in any system which undergoes a phase transition, one can define an *order parameter* which is a measure of order. In a disordered phase (such as water or paramagnetism) it is zero, and in an ordered phase (such as ice or a ferromagnet) it is non-zero — the spontaneous magnetisation is a suitable order parameter in a ferromagnet, for example. In practice, it is found that very general statements can be made about the way the order parameter behaves. The details of the system under investigation are rather unimportant, and general symmetry arguments can be invoked. One very important characteristic of a phase transition

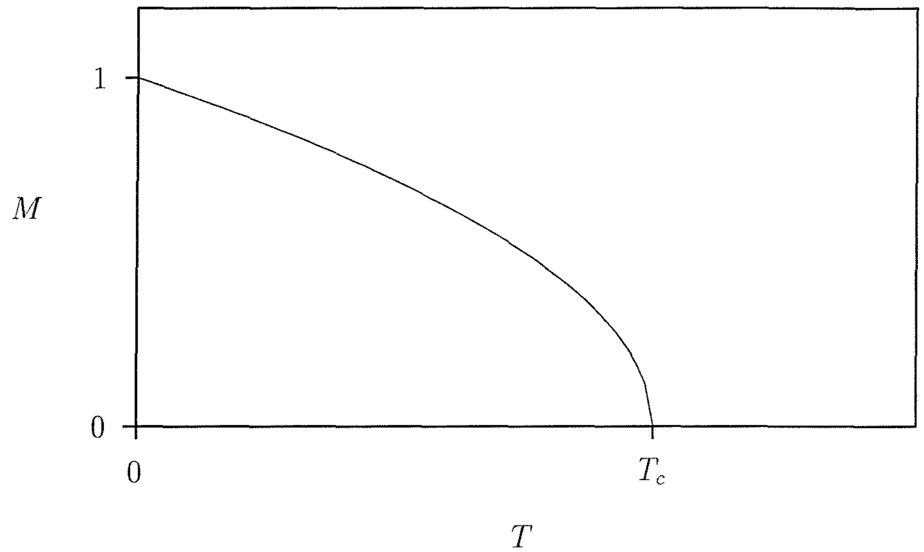

(a) Spontaneous Magnetisation

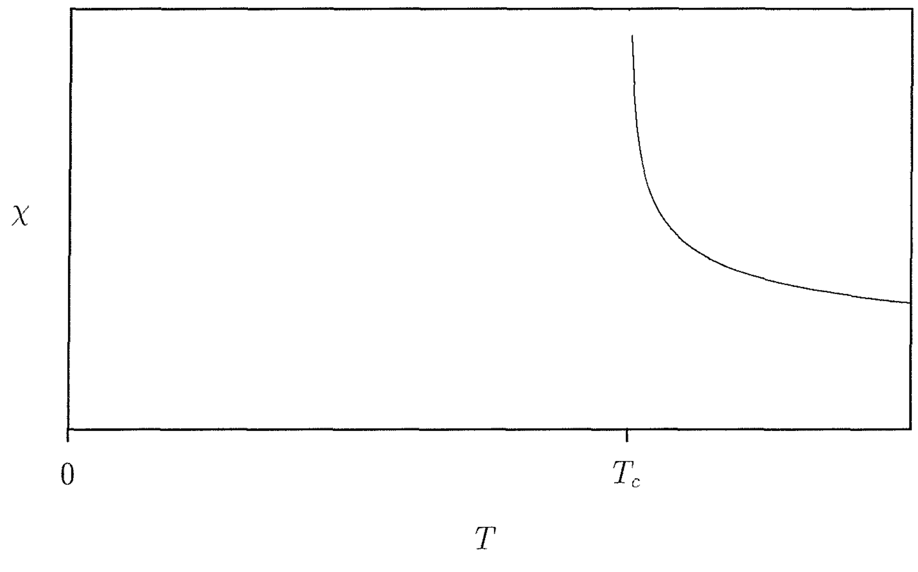

(b) Susceptibility

Figure 1.2: Spontaneous magnetisation (a) and zero-field susceptibility (b) of a ferromagnet. At low temperatures, the susceptibility depends on whether or not the applied field is in the same direction as the spontaneous magnetisation, and is not shown here.

is the value of certain *critical exponents* — it has already been pointed out that the magnetic susceptibility diverges at the Curie point. A number of other interesting quantities also show critical behaviour, which can be characterised quite generally by defining the following exponents.

$$C_v \sim |t|^{-\alpha} \tag{1.3}$$

$$M \sim |t|^{\beta} \tag{1.4}$$

$$\chi \sim |t|^{-\gamma} \tag{1.5}$$

$$M \sim H^{\frac{1}{\delta}} \tag{1.6}$$

$$\xi \sim |t|^{-\nu} \tag{1.7}$$

$$G(r) \sim r^{2-d-\eta} \tag{1.8}$$

where t is defined as $\frac{T}{T_c} - 1$ — it is a normalised temperature, relative to the critical temperature. H is the external magnetic field, M is the spontaneous magnetisation, and χ the magnetic susceptibility, in magnetic systems. C_v is the specific heat.

The correlation length ξ is a measure of the extent to which the spins are correlated. When a change is made to one spin, that change has an influence at a distance of the order of ξ. At low temperature in an ordered phase, the correlation length is very small, and the system remains predominantly ordered. (The system has long range order, but the correlation length is the typical size of fluctuations in the order, not the length-scale of the order itself.) At higher temperatures, these fluctuations grow, giving a reduced average magnetisation in a ferromagnet, say. At the critical temperature, the correlation length becomes infinite and fluctuations span the whole system, so that the magnetisation flips rapidly, averaging to zero. $G(r)$ in equation 1.8 describes the correlations at the critical temperature. Above the phase transition, the spins still prefer to align with their neighbours, through the exchange interaction, but now this order is limited to the (finite) correlation length. As the temperature is increased further, the correlation length decreases, giving smaller 'pockets' of order. This is discussed further in section 3.6.

The values of these critical exponents are a useful way of characterising phase transitions, but they are not independent — relationships can be derived relating them (section 2.4.1).

1.1.4 Dipolar Interaction and Domain Formation

Anyone who has tried to make two bar magnets lie parallel will find this ferromagnetic ordering counter-intuitive — macroscopic magnetic dipoles prefer to form *anti*-parallel alignment. This is because the exchange interaction which is the basis of ferromagnetic ordering is a very short range interaction, so that it couples only the nearest neighbour atoms together, but each interaction has a knock-on effect, and the order pervades throughout the sample. The dipolar interaction, on the other hand, though much weaker in strength at short distance, has a much longer range of interaction, decaying much more slowly with distance. For macroscopic magnets it is the dipolar interaction that causes antiparallel alignment, and this same mechanism accounts for the way samples of iron do not in general show the predicted spontaneous magnetisation at room temperature. In fact it is the combination of dipolar alignment and the reduction in the energy stored in the external magnetic field which gives rise to the domain structure commonly witnessed in iron (figure 1.3.) Real systems, particularly polycrystalline samples, forms many domains. These domains represent a compromise between the tendencies to align at short distances and to misalign at large distances; such compromise is a recurring theme in condensed matter physics.

1.1.5 Anisotropy

Another feature of magnetic systems is that of magnetic *anisotropy*. In our foregoing discussion, it was implied that the magnetic dipoles can rotate freely to point in any arbitrary direction. In a free atom, this is a reasonable assertion to make, but when the atom is part of a regular lattice in a solid, it is not in an isotropic environment. Thus we should not be surprised to find that the susceptibility depends on the direction of the magnetic field relative to the orientation of the crystal axes.

In fact, there is an electrostatic interaction between the electrons which contribute the magnetic moment and the electric field gradient which gives rise to the so called *crystal field* interaction. In the rare earth metals, for example, this anisotropy results in considerable differences in the magnetisation of the sample, depending on the relative directions of the field and the crystal axes. In particular, it is found that magnetic compounds exhibit an easy direction of magnetisation.

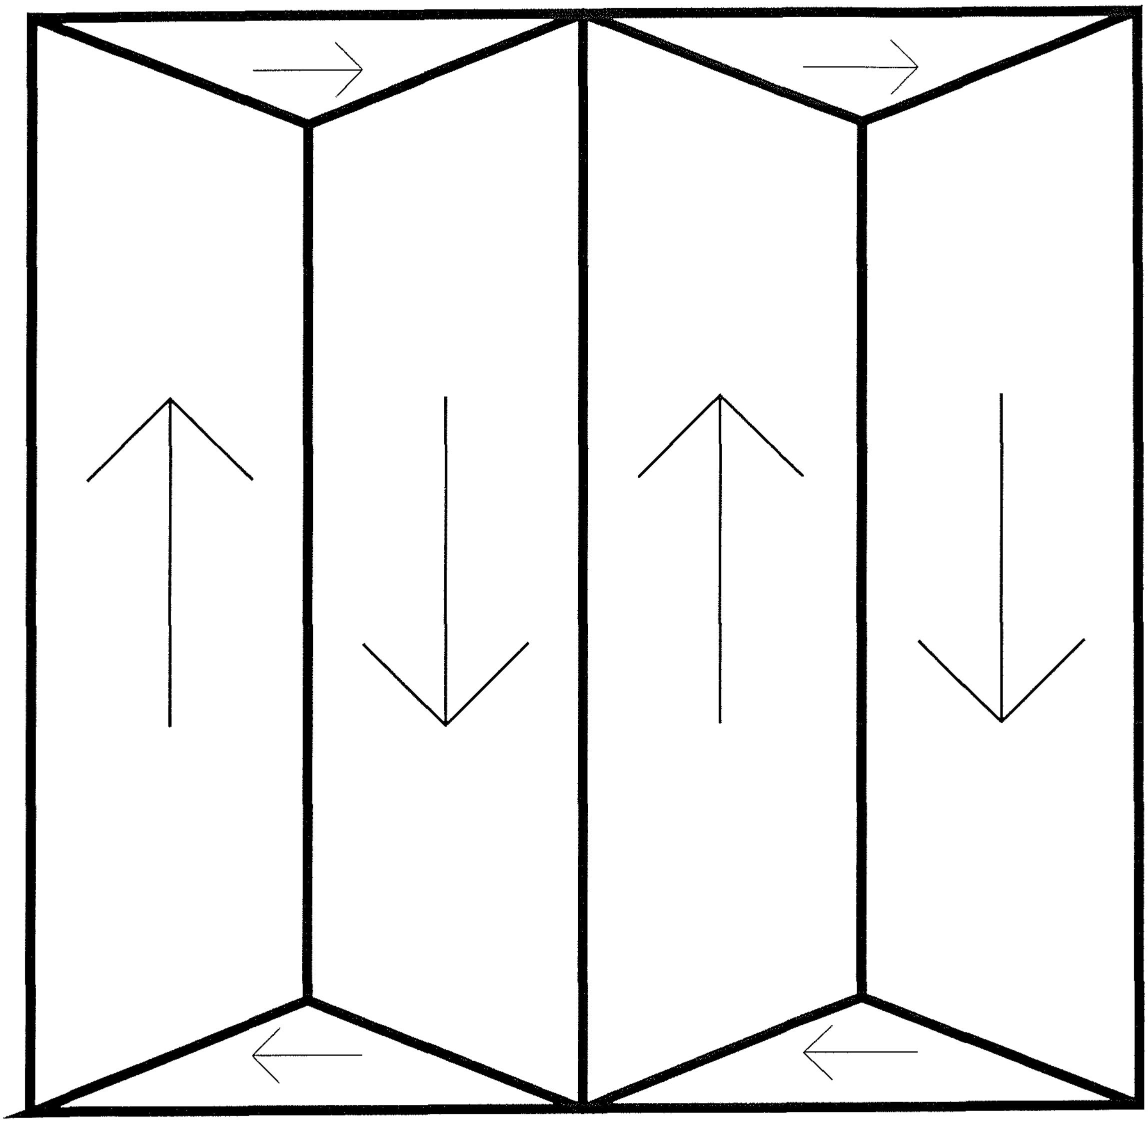

Figure 1.3: Domains in a ferromagnet. The energy stored in the external magnetic field is minimised by such a configuration.

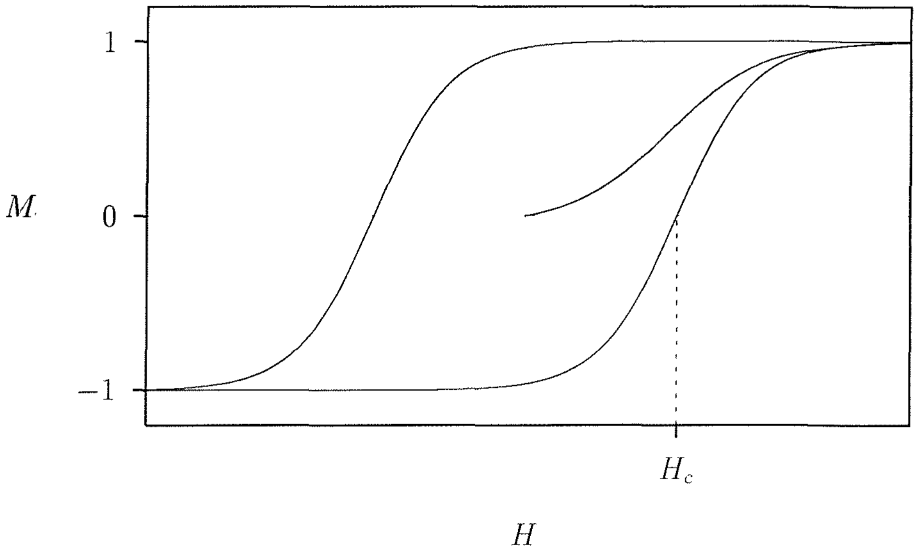

Figure 1.4: Hysteresis

This explains why, in a ferromagnet for example, the spins do not immediately rotate to align themselves with the external field direction, which would give a very large magnetic susceptibility.

1.1.6 Hysteresis

One hallmark of ferromagnetic materials is that of magnetic *hysteresis*, illustrated in figure 1.4. When a field is applied to an initially unmagnetised sample, the magnetisation follows an *initial magnetisation curve*. When the field is reduced. however, the magnetisation does not return down the same path, but maintains a large *remanent magnetisation* M_r even when the field is removed. Reversing the field does not immediately reverse the magnetisation; a field equal to the *coercive field* H_c must be applied before the sample reverses its direction of magnetisation.

This can be understood from the domain structure illustrated in figure 1.3: for very small fields, the domains parallel to the field grow at the expense of the domains of misaligned spins, in a reversible way (so that they will return along the same path when the field is removed). The anisotropy prevents the domains from simply rotating towards the field. As the field increases, the domain boundaries have to be dragged over crystal impurities, polycrystalline borders, pinning defects, etc., so that work is done and the change is irreversible. When the field is large, the domains rotate to saturation. This last régime is reversible, so that when the field is removed, domains form once more, but in a different configuration, in which the domains which were parallel to the field are favoured, giving the system a net magnetisation.

1.1.7 Antiferromagnets and Ferrimagnets

There are other magnetic systems which are closely related to ferromagnets on a microscopic scale, but which have very different macroscopic properties. *Antiferromagnets* such as Chromium show no spontaneous magnetisation, yet at microscopic level there is ordering similar to ferromagnets; but with the difference that neighbouring spins lie antiparallel, as illustrated in figure 1.1(b). Hence the moments cancel out, giving no net magnetisation.

At high temperatures, the system is again paramagnetic, obeying a Curie-Weiss law (equation 1.2), but with a negative Weiss constant. The susceptibility does not diverge at the *Néel temperature*, but there is a weak cusp. As with the ferromagnet, the low temperature susceptibility depends on the relative directions of the ordered atomic dipoles and the external field: when the field is perpendicular to the dipoles, they all tilt towards the field ($\chi_{\perp}$), but when the field is in the direction of the magnetisation ($\chi_{\parallel}$), there is little change in magnetisation unless the field is strong enough to overcome the exchange interaction.

In order to detect the magnetic order explicitly, one can use techniques such as *neutron scattering*, where a beam of neutrons is scattered from a target constructed from the material under study. Neutrons are uncharged, but have an intrinsic magnetic dipole moment which interacts with the microscopic magnetic structure in the sample, so that they can experience coherent diffraction from ordered arrays of magnetic moments. In practice it is necessary to distinguish the peaks which are due to interactions with the atomic nuclei from peaks due to the magnetic ordering in the neutron diffraction patterns.

Ferrimagnets are systems which show a spontaneous magnetisation like ferromagnets, but which have antiparallel alignment of the atomic dipoles at a microscopic level. This apparent inconsistency can be resolved by noting that these materials are composed of mixtures of atoms with unequal magnetic moments, arranged alternately in the lattice. All the atoms of one type lie parallel on one sublattice, which is antiparallel with another sublattice containing aligned atoms of the other type, as illustrated on figure 1.1(c).

1.1.8 Random Systems

In practice, materials with mixtures of ferromagnetic and antiferromagnetic interactions can be constructed. In such compounds there is competition between the different couplings, which complicates the physics. One such system is the class of alloys of the form $\mathrm{Eu_xSr_{1-x}S}$. With large x, this system has ferromagnetic nearest-neighbour coupling and antiferromagnetic next-nearest-neighbours. The Sr atoms are not magnetic, and serve simply to dilute the lattice, giving a random mix of ferromagnetic and antiferromagnetic interactions. Another system that has been widely studied consists of dilute solutions of magnetic transition metal impurities in noble metal (non-magnetic) hosts. Here, the exchange interaction between the magnetic impurities is mediated by the conduction electrons of the host by the RKKY interaction, after Ruderman, Kittel, Kasuya and Yosida [6]. The interaction oscillates in sign as it decays with distance, so that some pairs of magnetic dipoles interact ferromagnetically and others experience an antiferromagnetic coupling.

The crucial element of such *spin glasses* appears to be that the low temperature state is a frozen disordered one, rather than a uniform or periodic state described above. The competition between the different couplings and an element of randomness appear to produce this type of state, which will be discussed in more detail in section 2.6.

Experiments on samples of spin glasses reveal the following properties:

- a marked cusp in the magnetic susceptibility as the system is cooled through a critical temperature, but no spontaneous magnetisation at low temperatures, and no peaks in neutron scattering experiments which would indicate periodic order like antiferromagnetism. Nevertheless, the suppressed susceptibility does imply some sort of spontaneous magnetisation.

- remanent magnetisation at low temperature — the magnetisation of a sample which is cooled from above the freezing temperature in a non-zero magnetic field is higher than the magnetisation obtained when the same field is applied to the same system cooled in zero field. Equivalently, the remanent magnetisation on removal of the fields in the different circumstances is different.

- field dependent freezing temperature — the larger the field, the lower the temperature at which the above history-dependent magnetisation sets in; the Almeida-Thouless line separates glassy behaviour from *ergodic* behaviour (see below).

- when the susceptibility is measured using an oscillating field ('ac susceptibility'), the temperature at which the cusp appears is found to depend on the frequency of the of the field; raising the frequency increases the freezing temperature, which implies that the freezing is some sort of dynamic effect, and the timescale of the experiment influences the observed behaviour.

Random Anisotropy

Another way of introducing competing interactions is to work on amorphous samples of systems with large crystal fields. Without the long-range periodic structure of a crystal, the anisotropy discussed in section 1.1.5 varies from site to site through the system. Thus there is competition between the exchange interaction, which tries to align the neighbouring dipoles, with the anisotropy, which gives each dipole a *local* easy magnetic axis.

Amorphous alloys of rare earths in metals, such as DyCu and $TbFe_2$ show some of the properties of spin glasses, above: there is a cusp in the susceptibility [9] below which samples show history-dependent magnetisation [10]. The freezing temperature is field-dependent [11], but there is no indication of time-dependence in the freezing temperature [11].

It is these systems which have been studied in this work, but before discussing their properties, it will prove advantageous to review some of the background of how such systems are tackled theoretically.

1.2 Statistical Mechanics

Statistical mechanics is a tool used extensively in solid state physics, where properties of the system as a whole are inferred from consideration of the microscopic states of the system. A complementary theory is that of thermodynamics, where the macroscopic parameters of the system (such as magnetisation and susceptibility) are derived from general considerations such as conservation of energy. Links between these two theories have been established.

The fundamental building block for statistical mechanics is the *Boltzmann probability distribution*, which says that the probability of the system being found in a state i with energy E_i is proportional to $\exp\left(-\frac{E_i}{k_B T}\right)$ at temperature T; k_B is *Boltzmann's constant.* For brevity, it is customary to define an inverse temperature $\beta = \frac{1}{k_B T}$. Notice that statistical mechanics is based on *probability*, whereas mechanics traditionally involves following each particle deterministically. (Molecular dynamics is a numerical simulation technique where the forces *are* computed for each particle, and the equations of motion are integrated.)

If we imagine a collection, or *ensemble*, of identical copies of our system, and freeze the motion at some instant, we would find the systems in microscopic configurations reflecting this probability distribution. We define the quantities of physical interest as averages over all these configurations. This is the *ergodic hypothesis.*

As the name of the theory implies, we rely on the large number of particles involved to cancel out the fluctuations. For the most part, standard statistical theory says that fluctuations are typically proportional to the square root of the number of particles, and since macroscopic systems are typically 10^{23} particles, the relative fluctuations give errors smaller than 10^{-10}.

1.2.1 Example: paramagnetism

As a simple example, we can consider the paramagnetism due to a single magnetic dipole in a field, as discussed in section 1.1.1. The energy of a dipole in a magnetic field is given by

$$E = -\mathrm{m}.\mathbf{H} \tag{1.9}$$

and we can use Boltzmann's probability distribution to calculate an *average* value for the magnetisation m and susceptibility χ:

$$\begin{aligned} \langle \mathrm{m} \rangle &= \frac{\int \mathrm{m} \exp(\beta \mathrm{m}.\mathbf{H})}{Z} & (1.10) \\ \chi &= \frac{\partial \langle \mathrm{m} \rangle}{\partial H}. & (1.11) \end{aligned}$$

Here the angled brackets $\langle\rangle$ denote a thermal average, and Z is the *partition function*

$$Z = \int \exp(\beta \mathrm{m}.\mathbf{H}). \tag{1.12}$$

Formally, Z is merely the normalisation of the Boltzmann probabilities, but we can in fact calculate all the quantities of interest from the partition function alone. For example, equation 1.10 can be written

$$\langle \mathrm{m} \rangle = \frac{\partial}{\partial(\beta \mathbf{H})} \ln Z \tag{1.13}$$

from which it can be shown fairly easily that the zero field susceptibility is proportional to β, which is to say inversely proportional to temperature i.e. Curie's law (equation 1.1).

When we take a large number of *independent* dipoles, the result is unchanged, and the net magnetisation $\mathbf{M}$ is the sum of the microscopic contributions $\{\mathrm{m}_i\}$. As soon as interactions between the dipoles are considered, however, the complexity increases enormously; the energy of each microscopic state must now contain the mutual interaction of 10^{23} particles.

1.2.2 Use of Models

Because these problems rapidly become intractable, physicists must simplify the interactions to have any hope of finding a solution. One hopes that in constructing models for systems, the salient features can be captured without considering all the possible interactions. Typically, dipolar interactions (section 1.1.4) are ignored, and the exchange interaction is taken to be constant, acting only on nearest neighbours in a perfect lattice. The Heisenberg model for interacting magnetic systems takes the form

$$\mathcal{H} = -\sum_{\langle ij \rangle} J \mathbf{S}_i . \mathbf{S}_j - \mathbf{H} . \sum_i \mathbf{S}_i \tag{1.14}$$

with uniform exchange interaction strength J between spins (dipole moments) $\mathbf{S}_i$ and $\mathbf{S}_j$ at sites i and j, and uniform magnetic field $\mathbf{H}$. $\langle ij \rangle$ means sum over nearest neighbours only. J is positive if sites i and j have a ferromagnetic interaction [so that aligned spins have lower energy], and negative for an antiferromagnetic interaction, where spins prefer to lie anti-parallel. ($\mathcal{H}$ is called the *Hamiltonian*, and is just the operator equivalent to the total energy.)

Even this simplified system cannot be solved exactly for classical spins. An important landmark in the development of statistical mechanics was the exact solution of the *Ising model* in two dimensions. This system, inspired by quantum mechanics, allows each dipole moment to have only two orientations, fully parallel to or fully antiparallel to the external field. Thus each **S** in equation 1.14 is replaced by an integer taking the values ±1.

Despite its apparent triviality, the Ising model in two dimensions has two important qualities:

- the two dimensional Ising system *does* undergo a phase transition, from a paramagnetic high-temperature régime to a low-temperature ferromagnetic phase, and
- because it can be solved exactly, it is a useful testbed for techniques which are applied to more complicated models. The errors introduced by such techniques can be monitored.

Several statistical techniques will be illustrated in chapter 2 using the Ising model.

1.2.3 Phase Transitions and Broken Symmetry

It is useful to consider a phase transition from a statistical mechanics view. The Heisenberg model for a ferromagnet (Hamiltonian 1.14) undergoes a paramagnetic to ferromagnetic phase transition, developing a spontaneous magnetisation at low temperature. The Ising system in two dimensions behaves similarly. However, the energy of the system is unchanged if all Heisenberg spins are rotated through the same angle, or if all the Ising spins are flipped. Application of the above ergodic hypothesis would therefore produce a zero magnetisation at all temperatures, since any contribution from one configuration is exactly cancelled by a contribution from the configuration with all spins flipped through 180° degrees.

One can only conclude that the symmetry of the Hamiltonian has been ignored, and that the ergodic hypothesis has been violated. (The phase transition occurs in the model system, so we cannot blame some hitherto unconsidered interaction in the real system.) Our physically relevant measurements must now be limited to that part of 'configuration space' which is available to the system; this can be done formally by imposing a small field, and taking the limit of the result as the magnitude of the field is reduced to zero.

It should be noticed that the symmetry can only truly be broken for a system of infinite size, since a finite system will change its direction of magnetisation if a fluctuation in the local direction of the magnetisation reaches the order of the size of the system. Near the phase transition, these fluctuations can be very large, and only an infinite system is stable against such fluctuations since there is a non-zero probability of a fluctuation reaching any finite size. The phase transition is defined as the point at which the fluctuations become infinite (correlation length diverges in equation 1.7). In practice, real systems, with order of 10^{23} particles, do not flip on timescales of experiments, but the effect is important for the sorts of system sizes which are amenable to computer simulation. This will be discussed further in section 3.6.

Random Systems

For random systems, whether random exchange or random anisotropy, it is necessary to use a Hamiltonian which reflects the random variables. For spin glasses, with random exchange, the essential physics can be captured by replacing the term J in Hamiltonian 1.14 by J_{ij} whose sign and possibly magnitude various from site to site, giving the Edwards-Anderson Hamiltonian [12]:

$$\mathcal{H} = -\sum_{\langle ij\rangle} J_{ij}\mathbf{S}_i.\mathbf{S}_j - \mathbf{H}.\sum_i \mathbf{S}_i. \tag{1.15}$$

For the random anisotropy system, one postulated model Hamiltonian is the Harris, Plischke and Zuckermann (HPZ) random anisotropy model [8]

$$\mathcal{H} = -\sum_{\langle ij\rangle} J\,\mathbf{S}_i.\mathbf{S}_j - \sum_i D\,(\mathbf{S}_i.\hat{\mathbf{n}}_i)^2 - \mathbf{H}.\sum_i \mathbf{S}_i \tag{1.16}$$

where the $\hat{\mathbf{n}}_i$ is a unit vector in the easy direction at site i and D is the strength of the attraction to the local easy direction. The anisotropy energy is minimised when the spin lies in either direction along the local easy axis, and is a maximum when the spin is perpendicular. This is a simplification of a more general Hamiltonian where both D and J can vary from site to site.

In each case, it is necessary to average over the random variables J_{ij} or $\hat{\mathbf{n}}_i$ in order to arrive at physically meaningful results. But it is important to distinguish between the thermal (Boltzmann) average over the spin variables and the average over the (sample to sample) variations in the randomness. We must conceptually take an ensemble of systems with different random variables, and for each, perform the usual thermodynamic average. But as we have already seen, a single system with a particular set of random variables is in general too hard for us to study, yet now we have to find properties of systems with a particular *distribution* of random parameters. This will be discussed further in chapter 2.

Chapter 2

Theory

Before examining the properties of the HPZ model (equation 1.16), it is useful to study the Ising and XY models, in order to review some techniques used in solving statistical mechanics problems. The Ising system is exactly soluble in one and two dimensions, and is invaluable as a guide to the accuracy of these methods. The random anisotropy system in the large anisotropy limit might be expected to show some Ising behaviour. The XY (or planar rotator) model is the simplest continuous spin system, and is the small anisotropy limit of the random system studied in this work.

Both the Ising and XY models are special cases of the general Heisenberg model introduced in section 1.2.2:

$$\mathcal{H} = -J \sum_{\langle ij \rangle} \mathbf{S}_i . \mathbf{S}_j - \mathbf{H} . \sum_i \mathbf{S}_i \tag{2.1}$$

where J is the exchange interaction between nearest-neighbour spins $\mathbf{S}_i$ and $\mathbf{S}_j$ at sites i and j, and $\mathbf{H}$ is the applied magnetic field. This Hamiltonian forms the starting point for the discussion of both the Ising and XY models.

2.1 Ising Model

The Ising model is a very simple system in which each spin can take one of only two values, 'up' and 'down', and spins interact only with nearest neighbours. It is the classical analogue of the Quantum Mechanical spin-$\frac{1}{2}$ problem. In the limit of large anisotropy, where the spins are constrained to lie along the anisotropy axes, we might expect to find similar behaviour.

Hamiltonian 1.14 can be simplified to

$$\mathcal{H} = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j - H \sum_i \sigma_i \tag{2.2}$$

where σ_i is now an integer, taking values ± 1 and the exchange sum is taken over nearest neighbours only. The ground state in zero or uniform magnetic field clearly has all spins aligned (assuming $J > 0$), as this is the configuration with lowest energy. For $J < 0$, the spins are alternately up and down, in an antiferromagnet configuration.

2.1.1 Domain Formation

We can reformulate the one-dimensional model in zero field in terms of the bonds rather than the spins — each bond has energy 0 if the adjoining spins lie parallel, or $2J$ at a 'kink', or discontinuity, in the adjacent spins:

$$\cdots \uparrow\uparrow\uparrow\downarrow\downarrow\downarrow \cdots$$

The bonds are independent, and kinks appear at each bond with a probability $\exp(-2\beta J)$ where $\beta = \frac{1}{k_B T}$. Thus, in an infinite system, we expect kinks to appear at any non-zero temperature, destroying long range order.

In two [or more] dimensions, these kinks become lines [surfaces], surrounding regions of aligned spins, as illustrated in figure 2.1. Unlike the one-dimensional case, however, the energy of these domain walls is now proportional to the perimeter [surface area] of the regions they enclose. In order to form a domain of reversed spins of linear dimension$\sim \xi$, a domain wall of length ξ^{d-1} is required, as shown in the figure. Thus the energy of the fluctuation $\sim 2J\xi^{d-1}$. In one dimension, the energy is independent of ξ, as expected. But with $d > 1$, we would not expect to find large domains of reversed spins at low temperatures, implying that the system possesses a net magnetisation. At higher temperatures, of course, the increase in entropy associated with the domains would make them favourable, and so no net magnetisation is expected at sufficiently high temperatures.

The two-dimensional Ising model was indeed shown to be ferromagnetic below a critical temperature $T_c = 2.269J$ by Onsager [13]. Monte Carlo simulations find $T_c = 2.3$ ($d = 2$) and $T_c = 4.43$ ($d = 3$), in energy units of J.

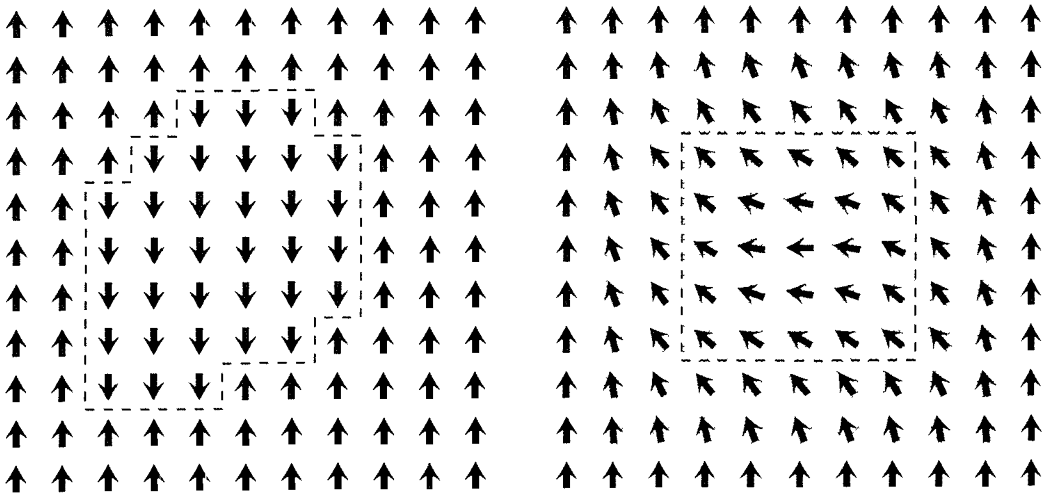

Figure 2.1: Comparison of domain formation in the Ising (left) and XY (right) models in two dimensions. In the Ising system, domains have a sharp domain wall, and the energy is proportional to the perimeter of the wall, while in the XY model, the dotted line is a nominal domain wall, enclosing the spins rotated by more than 45°, but in practice the exchange energy is distributed over all the shaded spins.

2.2 The XY Model

The XY (or planar rotator) model is an Heisenberg system (equation 2.1) with two-component spins constrained to lie in the XY plane. Its properties are presented here since we may expect to find similar properties in the limit of small anisotropy in our random anisotropy system.

We can recast equation 1.14 for the XY system in the form

$$\mathcal{H} = -J \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j) - H \sum_i \cos\theta_i \tag{2.3}$$

where θ_i is the angle between the spin at site i and the magnetic field (or some arbitrary direction, in zero field). Once again, the exchange sum is over nearest neighbours only.

2.2.1 Domain Formation

Consider the formation of a domain of size ξ, with spins rotated from the bulk direction by an angle θ (figure 2.1). The surface area is again of order ξ^{d-1}, but in the case of continuous spins, the discontinuity is not abrupt as that of the Ising system, but is spread over a number of spins, of order ξ, as shown in the figure. Thus, the deviation between adjacent spins $\sim \theta/\xi$, and the energy per bond is $1 - \cos(\theta/\xi) \approx \frac{1}{2}(\theta/\xi)^2$ for small angles. Thus the energy of the domain varies like $\frac{1}{2}\theta^2\xi^{d-2}$. In three or more dimensions, domains are inhibited at low temperatures, and spontaneous magnetisation can occur.

In one or two dimensions, large domains are not inhibited at any non-zero temperature — there may be no long-range order. However, in the presence of an applied field, full alignment is possible at $T = 0$K, implying an infinite zero-field susceptibility. In the real systems discussed in section 1.1.2, the dipolar interaction and crystal field anisotropy prevented free rotation of domains, but these interactions have been neglected here.

2.2.2 Spin Wave Analysis

A more formal approach is to calculate the amplitude of spin waves in the system. We again suppose that angles between adjacent spins are small, so that we can use $\cos\theta \approx 1 - \frac{1}{2}\theta^2$ in equation 2.3:

$$\begin{aligned} \mathcal{H} &\approx \frac{J}{2}\sum_{\langle ij\rangle}(\theta_i - \theta_j)^2 && (2.4) \\ &= \frac{J}{2}\sum_{\mathbf{R}}\sum_{a}\Big(\theta(\mathbf{R}) - \theta(\mathbf{R}+\mathbf{a})\Big)^2, && (2.5) \end{aligned}$$

where $\mathbf{a}$ is summed over the nearest-neighbour vectors. Following [14], we take Fourier transforms...

$$\begin{aligned} \theta(\mathbf{R}) &= \frac{1}{\sqrt{N}}\sum_{k}\theta_{\mathbf{k}}e^{i\mathbf{k}.\mathbf{R}}, && (2.6) \\ \text{where} \quad \theta_{\mathbf{k}} &= \frac{1}{\sqrt{N}}\sum_{\mathbf{R}'}\theta(\mathbf{R}')e^{-i\mathbf{k}.\mathbf{R}'}, && (2.7) \\ \text{so} \quad \mathcal{H} &= \frac{1}{2}\sum_{k}J_{\mathbf{k}}|\theta_{\mathbf{k}}|^2, && (2.8) \\ \text{where} \quad J_{\mathbf{k}} &= 2J\sum_{a}(1 - \cos\mathbf{k}.\mathbf{a}). && (2.9) \end{aligned}$$

Thus, in $\mathbf{k}$-space, in this small-deviation limit. the variables are decoupled. The energy is now expressed in terms of *spin waves* of wave vector $\mathbf{k}$ and amplitude $\theta_{\mathbf{k}}$. (A cross-section through the centre of the right hand side of figure 2.1 can be used to illustrate the nature of a spin wave.) The equipartition of energy gives

$\frac{1}{2}J_{\mathbf{k}}\langle|\theta_{\mathbf{k}}|^2\rangle = \frac{1}{2}T$, so the intensity of each possible wave is proportional to the temperature. The wavelength of the spin waves is closely related to the size of the domains discussed in section 2.2.1: at low temperatures, only very gradual changes in the magnetic direction are permitted, and large domains result, while increased temperature allows smaller domains to be formed.

From this spin wave analysis, various useful functions can be calculated [14]. One important result is the spin correlation function (section 3.5)

$$\begin{aligned} g(\mathbf{R}) &= \overline{\langle\cos\left(\theta(\mathbf{R}+\mathbf{R}') - \theta(\mathbf{R}')\right)\rangle} && (2.10)\\ &\sim R^{-\lambda T} \qquad \text{(in two dimensions)} && (2.11)\\ &\sim e^{-\lambda T(1+B/R)} \qquad \text{(in three dimensions)}. && (2.12) \end{aligned}$$

The correlation function in two dimensions (equation 2.11) decays to zero extremely slowly — there is no true long range order, but the system is "almost" magnetic. Thus considerable care is required when performing numerical simulations on finite systems.

2.2.3 Berezinsky-Kosterlitz-Thouless theory

Berezinsky [15] and Kosterlitz and Thouless [16] realised that an important excitation in the two-dimensional XY model is the *vortex*, a topological defect, illustrated on figure 2.2. As we go round a contour enclosing vortices, the spin angles change by $\pm 2n\pi$; n is the algebraic sum of enclosed vortices, using sign to distinguish the two possible polarities shown in the figure.

Consider a circle centred at an isolated vortex, with radius R large in comparison to the lattice spacing a. The change in angle of 2π is distributed over $2\pi R$ spins, so that the total energy in the bonds around the circle $\sim \frac{J}{R}$ (c.f. section 2.2.1), and the total energy within the circle due to the vortex is of order

$$J\int_a^R \frac{1}{r}dr = J\ln\frac{R}{a}. \tag{2.13}$$

Thus the energy of an *isolated* vortex diverges as we consider $R \to \infty$. At higher temperatures the extra entropy carried by a vortex can make isolated vortices favourable: inside the circle, there are πR^2 possible positions for the vortex centre, and the entropy $\sim \ln R$. A more careful analysis shows that $F = E - TS$ changes sign at a temperature $T_c = \pi J$.

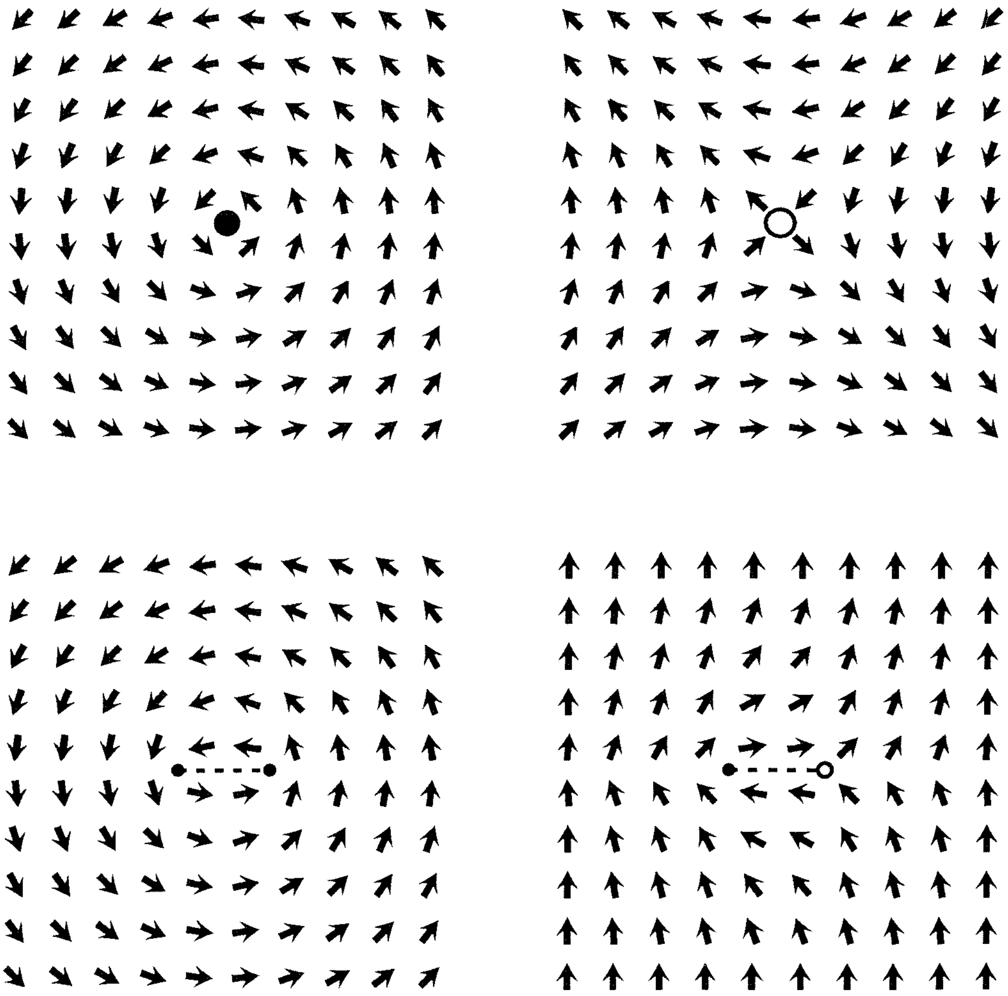

Figure 2.2: The upper two figures depict vortices of opposite polarity: running the figures together would cancel their effect at large distance, as shown in figure 2.3. It should be noted that the anti-vortex depicted on the right is *not* simply the spins going 'the opposite way' — flipping all spins in the left figure by 180° reverses the direction of rotation of the vortex, but does not change the polarity of the vortex. The lower two figures depict pairs of half-vortices — the left pair combine to form a full vortex equivalent to the one above it, while the right pair cancel, with no effect at long distance.

The impact of a vortex-antivortex *pair* with separation R is small at large distances (figure 2.3), so that the energy is finite, again of order $\ln R$, and such pairs can be created at low temperatures. As the separation R between vortices in a pair increases, at higher temperatures, other vortices can 'screen' the effect of one vortex on its partner, thus lowering the temperature at which the pairs dissociate. Kosterlitz and Thouless reformulated the problem as a gas of charged particles, modelling the interactions between pairs as a dielectric constant. Using a mean field theory on this system (section 2.3), the phase transition from 'topological order' was found to be at $T_c = 1.35J$. Monte Carlo results give $T_c = 0.89J$ [17], but agree qualitatively with the nature of the phase diagram.

Also shown on figure 2.2 are half-vortices: around a circle enclosing a half vortex, the spins go through an angle $\pm\pi$. Two half-vortices of the same polarity are equivalent to a full vortex, while half-vortices of opposite polarities cancel, as expected. Half-vortices tend to appear as thermal excitations; a full vortex can extend itself over several sites, forming a string of frustrated spins between the half-vortices, as shown — frustrated because there is no orientation available to satisfy the exchange interaction with all neighbours.

Kosterlitz [18] went on to study the critical properties of the system using renormalisation (section 2.4). The main result of this was the calculation of the critical exponents $\eta = \frac{1}{4}$ and $\delta = 15$. ($\eta = \lambda T$ in equation 2.11).

Vortices also have a part to play in the three-dimensional XY model, but as vortex *strings*. Because these are one-dimensional structures, rather than point defects in the two-dimensional model, they are not expected to play a significant role, although there is at present no theory which takes these vortex strings into account.

Finally, it should be noted that the XY model displays an anomalous 'bump' in the specific heat *above* the phase transition temperature.

2.3 Mean Field Theory

Mean field theory can be viewed in (at least) two ways: it is the *exact* solution for a system with infinite-range interactions (i.e. each particle interacts with every other with equal strength), or it is the *approximate* solution for finite-range interactions, ignoring fluctuations.

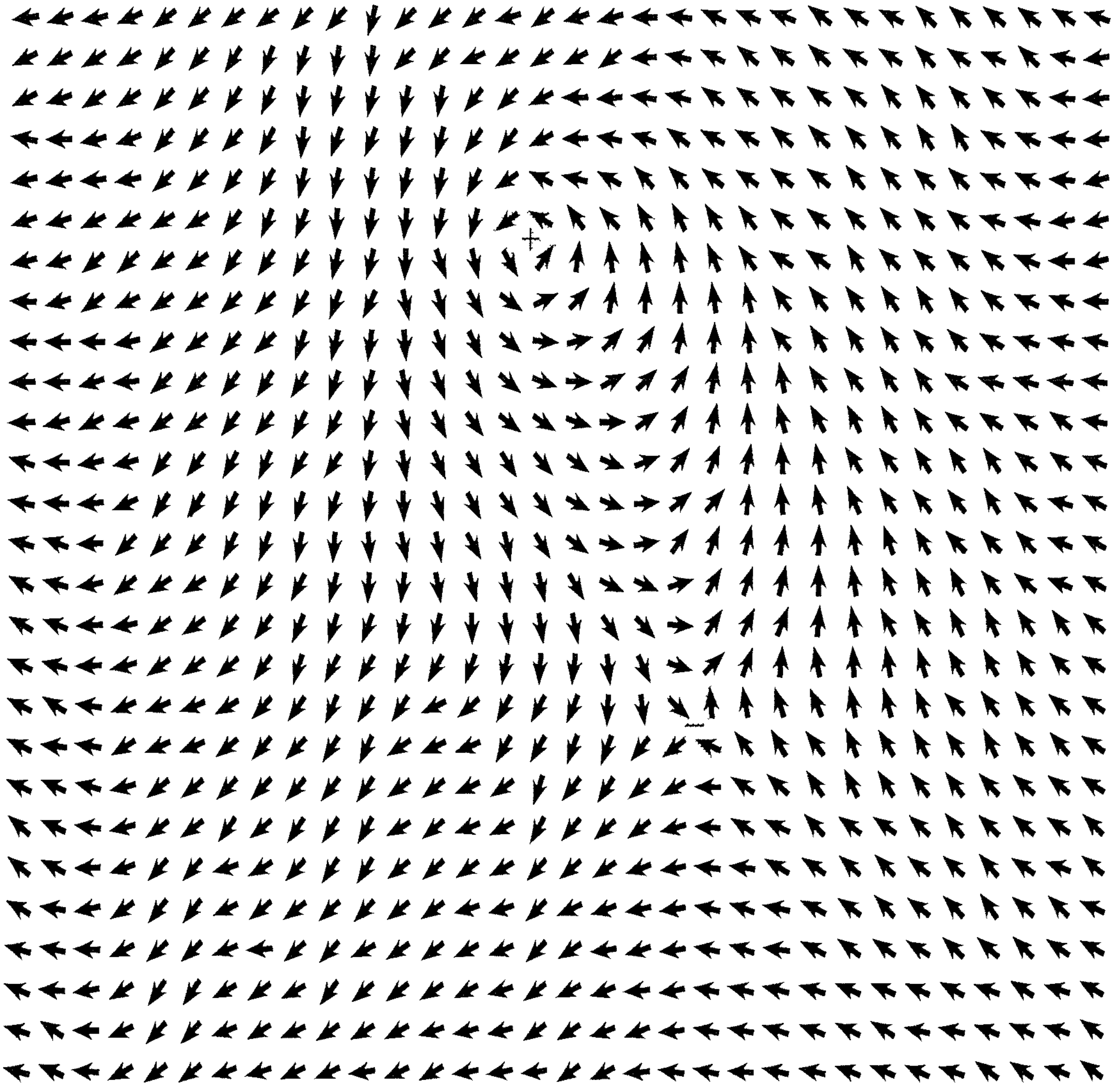

Figure 2.3: A pair of vortices (highlighted), one of each polarity. Notice how, far from the vortices, all spins point towards the left, so that the vortex pair does not have any long-range impact.

2.3.1 Ising model

In the Ising model, for example, equation 2.2 can be rewritten in the form

$$\begin{aligned}\mathcal{H} &= -\sum_i \left(H + J\sum_j S_{i+j}\right) S_i && (2.14)\\ &\equiv -\sum_i H_i' S_i. && (2.15)\end{aligned}$$

In this form, we can see that particle i moves in a local effective magnetic field H_i'. If we now take the (spatial) *mean* effective field, we find each spin moving independently in a (uniform) field

$$\overline{H'} = H + zJ\overline{S}, \tag{2.16}$$

where z is the number of neighbours (in the approximation view). This simple system has a partition function

$$Z = \exp\left(\beta\overline{H'}\right) + \exp\left(-\beta\overline{H'}\right) \tag{2.17}$$

and a (thermal) mean magnetisation

$$\begin{aligned}\langle S\rangle &= \tanh\left(\beta\overline{H'}\right) && (2.18)\\ &= \tanh\left(\beta\left(H + zJ\overline{S}\right)\right) && (2.19)\end{aligned}$$

at inverse-temperature β. This can be solved graphically, plotting the lines

$$\begin{aligned}y &= \tanh x \qquad \text{and} && (2.20)\\ y &= \frac{xT - H}{zJ}, && (2.21)\end{aligned}$$

as illustrated in figure 2.4. In zero field, the curves always meet at $x = 0$, where the gradient of equation 2.20 is 1; if this is greater than T/zJ, the curves will cross again, indicating a spontaneous magnetisation and hence ferromagnetism.

In summary, therefore, the mean field theory predicts a (second-order) phase transition at $T_c = zJ$. This is clearly in disagreement with the known behaviour for the one dimensional model (section 2.1), and significantly overestimates the critical temperature for two and three dimensions. The 1-d catastrophe occurs because mean field theory neglects fluctuations such as our kinks. (Or, in the infinite-range view of mean field theory, an isolated kink has infinite energy, as all the spins to one side interact unfavourably with those on the other side.) In two and three dimensions, it at least predicts the correct qualitative behaviour, though with too high a critical temperature.

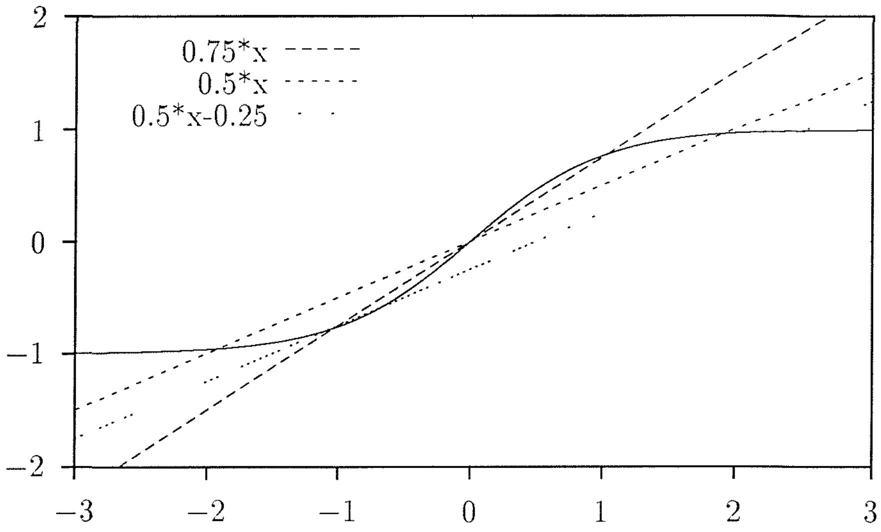

Figure 2.4: Form of graphical solution for equations 2.20 and 2.21. Changing the temperature changes the gradient of the straight line, while changing H moves the line in the y direction. The lines must cross once, and may cross up to three times.

2.3.2 XY Model

A mean field theory for the XY model follows section 2.3.1 fairly closely, but because this is a continuous system, equation 2.18 becomes

$$\begin{aligned} \langle\cos\theta\rangle &= \frac{\int_0^{2\pi}\cos\theta\exp\left(\beta H'\cos\theta\right)d\theta}{\int_0^{2\pi}\exp\left(\beta H'\cos\theta\right)d\theta} && (2.22)\\ &\equiv \frac{I_1(\beta H')}{I_0(\beta H')}, && (2.23)\\ \text{with}\qquad I_n(x) &= \int\cos^n\theta\exp\left(x\cos\theta\right)d\theta && (2.24)\end{aligned}$$

Again, a solution can be obtained by plotting self-consistent

$$\begin{aligned} y &= \frac{xT-H}{zJ} && (2.25)\\ \text{and}\qquad y &= \frac{I_1(x)}{I_0(x)}. && (2.26)\end{aligned}$$

for various temperatures.

Equation 2.26 is similar in form to equation 2.20, so that a graphical solution resembles figure 2.4. As before, we use the gradient of equation 2.26 at $x = 0$ to determine whether the lines cross more than once in zero field:

$$\left.\frac{d}{dx}\frac{I_1(x)}{I_0(x)}\right|_{x=0} = \frac{I_2(0)}{I_0(0)} = \frac{1}{2} > \frac{T}{Jz}. \qquad (2.27)$$

Thus mean field theory predicts ferromagnetism in the XY model for $T < dJ$ (since $z = 2d$). This is qualitatively correct for three dimensions, but, as before, in lower spatial dimensions, fluctuations cannot be ignored. In this case, as $\langle \cos\theta \rangle$ decreases, $\langle |\sin\theta| \rangle$ increases, and the *direction* of the local field changes, so that there is no long range order.

The evidence therefore tends to suggest that mean field theory is at best misleading, and at worst just plain wrong ! However, it has the big advantage of being simple to formulate, and does give an indication of how to proceed. This is especially true in very complex systems, when mean field theory is the only tractable theory. It is less useful in low spatial dimensions, where fluctuations are more important. In general, there is a *lower critical dimension* : in fewer spatial dimensions, the gain in entropy produced by introducing a domain wall exceeds the gain in energy at any (non-zero) temperature.

2.4 Space Renormalisation

Renormalisation is the procedure of changing the scale of the system, and comparing the properties of the new system with those of the original system. The procedure can be carried out discretely, by 'decimating', or summing over some of the sites using the Kadanoff construction [19] , or by making a continuous transformation. If the transformed lattice has the same structure as the original, then the partition function for the reduced lattice should have the same form as the original, but with an interaction strength K' different from the original K. Clearly, we can iterate this process. At infinite temperatures, all the sites are independent, so that changing the scale in this way cannot make any physical difference. Likewise, at zero temperature, in a ferromagnetic system, all sites are ordered, so that again there is no difference in macroscopic properties when we renormalise. These are called fixed points, and as we iterate our renormalisation procedure, we may head towards one or other of these fixed points.

At high temperature, for example, above any phase transitions, the correlation length is finite. The correlation length in the decimated lattice must be smaller, since the system is correlated over the same physical distance, which now corresponds to a smaller number of lattice spacings. smaller number of new-lattice-spacings). This is equivalent to a move to a higher temperature or a weaker interaction. Thus as we iterate, the system tends towards the infinite-temperature / zero-interaction fixed point.

Near a critical point, the system correlation length diverges, and since there is then has no natural length scale, we might expect to find a *critical fixed point* , where the system is again unchanged by the renormalisation, requiring $K' = K$.

For example, in the 1-d Ising model in zero field, we might decimate the lattice by explicitly performing the sum over the even spin sites in the partition function.

We first rewrite Hamiltonian 2.2 in a dimensionless form

$$\mathcal{H}' = \beta\mathcal{H} = K\sum_i s_i s_{i+1} \tag{2.28}$$

with a reduced interaction $K = \beta J$. Explicitly summing the even spins,

$$\begin{aligned}
Z(K,N) &= \sum_{\{s_i\}=\pm 1} \exp(\mathcal{H}') && (2.29)\\
&= \sum_{s_1=\pm1}\sum_{s_3=\pm1}\cdots\left(\sum_{s_2=\pm1}\sum_{s_4=\pm1}\cdots\exp(\mathcal{H}')\right) && (2.30)\\
&= \sum_{s_1,s_3,s_5,\cdots}\left[\left(e^{K(s1+s3)}+e^{-K(s1+s3)}\right)\left(e^{K(s3+s5)}+\cdots\right)\cdots\right] && (2.31)\\
&= \sum\left(\Phi\exp K's_1s_3\right)\left(\Phi\exp K's_3s_5\right)\cdots && (2.32)
\end{aligned}$$

where equation 2.32 follows by choosing K' and Φ such that each term agrees with terms in equation 2.31 for $s_i = \pm 1$ and $s_{i+2} = \pm 1$:

$$\begin{aligned}
\Phi\exp K' &= 2\cosh 2K \text{ when } s_i = s_{i+2} = \pm 1 && (2.33)\\
\Phi\exp -K' &= 2 \qquad\quad \text{when } s_i = -s_{i+2} = \pm 1. && (2.34)
\end{aligned}$$

Thus, with $\Phi = 2\sqrt{\cosh 2K}$ and $K' = \frac{1}{2}\ln\cosh 2K$,

$$Z(K,N) = \Phi^{\frac{N}{2}} Z\left(K', \frac{N}{2}\right) \tag{2.35}$$

and the partition function for the decimated system has been restored to its original form, with a reduced interaction strength — because it is $\ln Z$ that is of physical relevance, $\ln\Phi$, an additive constant, is irrelevant.

As we iterate, $K^{(n+1)} < K^{(n)}$, so that the *renormalisation flow* is always towards the $T = \infty$ fixed point for this system. The only solution for $K' = K$ is at $K = \infty$, the $T = 0$ fixed point. Hence, there is no critical fixed point, in agreement with our earlier arguments presented in section 2.1.1.

2.4.1 Scaling Hypothesis

The preceeding renormalisation algorithm can be expressed as a generalised scaling hypothesis, which allows relationships between the scaling exponents in section 1.1.3 to be defined.

We renormalise the lattice in blocks of L^d spins in d dimensions. (L was 2 in the example above.) In the vicinity of the critical point $t = h = 0$, we can derive relationships for renormalised parameters t' and h'. The simplest relationship consistent with the symmetry requirements as $h \rightarrow -h$ and the properties of a critical fixed point ($t' = h' = 0$ when $t = h = 0$), is

$$h' = hL^x \tag{2.36}$$

$$t' = tL^y. \tag{2.37}$$

The *scaling hypothesis* supposes that the free energy of the system is dominated by a singular term near the transition. For the renormalised system, that term must scale like

$$\Phi(t', h') = \Phi(L^y t, L^x h) = L^d \Phi(t, h) \tag{2.38}$$

because we have renormalised in blocks of L^d spins. As we have seen above, the correlation length of the renormalised system is reduced by a factor L:

$$\xi(t', h') = L^{-1}\xi(t, h). \tag{2.39}$$

By definition (equation 1.7), $\xi(t, 0) \sim |t|^{-\nu}$, giving $\nu y = 1$ in equation 2.39. Choosing $L = |t|^{-\nu}$, we find

$$\Phi(t, 0) \sim |t|^{d\nu} \tag{2.40}$$

But simple thermodynamics relates the behaviour of the free energy to that of the specific heat in zero field: from equation 1.3, $\Phi(t, 0) \sim |t|^{2-\alpha}$, giving *Josephson's law*

$$\nu d = 2 - \alpha. \tag{2.41}$$

Similar arguments give other equalities:

$$\alpha + 2\beta + \gamma = 2 \quad \text{Rushbrooke's law} \tag{2.42}$$

$$\alpha + \beta(\delta + 1) = 2 \quad \text{Griffiths' law} \tag{2.43}$$

$$(2 - \eta)\nu = \gamma \quad \text{Fisher's law.} \tag{2.44}$$

2.5 Replica Method

In random systems, not only are we required to perform the usual statistical mechanical thermal averaging, but we must also average over the randomness within and between samples. In experiments on samples of real systems, the globally-measured properties are an average over many possible local configurations of the random variables, so that fluctuations between different samples should be small.

We can write down the partition function Z of our system, but it is $\ln Z$ which is of physical relevance, and therefore which must be averaged over. The replica method makes use of the identity

$$\ln Z = \lim_{n \to 0} \frac{Z^n - 1}{n} \tag{2.45}$$

to allow us to deal with the randomness by averaging Z^n, which is somewhat easier in general (though in some cases it is perfectly possible to obtain the average of $\ln Z$ directly).

$$Z^n\{J\} = \sum_{s_i} \exp\left(-\beta \sum_{\alpha}^{n} H\left(S^\alpha, \{H\}\right)\right). \tag{2.46}$$

The α indices are called replica indices — conceptually, we have made n 'replicas' of our system. Provided that, after we have performed our average over the randomness and made whatever approximations that are justified, the resulting Hamiltonian for the n *coupled* replicas is well behaved in n, we can take the limit $n \to 0$ and can find our sample–average of $\ln Z$.

However, while the Hamiltonian is symmetric in the replica indices α, the lowest energy may be a state in which this symmetry is broken, in an analogous way to broken symmetry in ferromagnets. In this scenario, called broken replica symmetry, it is not clear that taking the limit $n \to 0$ is valid [20].

2.6 Spin Glasses

The random system which has been (and remains) most studied to date is the spin glass [7]; as described in section 1.1.8, these systems have competing ferromagnetic and antiferromagnetic interactions — the exchange interaction J_{ij} between spins i and j in Hamiltonian 1.14 is random, certainly in sign, and possibly also in magnitude. The origin of this random exchange in dilute magnetic

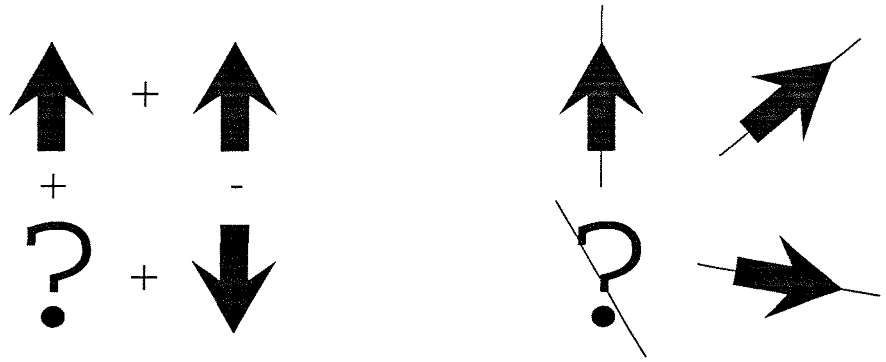

Figure 2.5: Frustration in (left) the Ising spin glass and (right) the Ising limit of the random anisotropy model (section 2.7). The '?' indicates that neither orientation satisfies all local bonds. The right picture illustrates a natural site for a half-vortex (section 2.2).

systems is thought to be due to the oscillatory RKKY interaction [6]; not only the magnitude, but also the sign, of the exchange interaction changes with distance, and in dilute or amorphous systems, the separation between spins is random.

The crucial factor appears to be *frustration* — the local variations in the exchange contrive to provide no possible configuration where all the ferromagnetic and antiferromagnetic interactions are satisfied at once. This is illustrated in figure 2.5. Instead, there are many *inequivalent* ground states [as opposed to the *degenerate* ground states in the Ising model (section 2.1) for example], with large energy barriers between these (relatively) low-energy states. The system becomes non-*ergodic* — it is trapped, or frozen, by these local minima, and is unable to sample the whole of phase-space. Because the energy barriers are not infinite, a non-zero temperature is required to assist the system in escaping from some of the energy minima. The volume of phase space available to the system depends on both the temperature and on the time-scale of the measurements being made — glassiness is a dynamic phenomenon, since the probability of the system crossing an energy barrier depends on the temperature. However, all systems can cross any finite barrier given sufficient time. (In an infinite system, the energy barriers can become infinite.)

While a non-random system is characterised by a global order parameter (magnetisation), the degree of freezing in a spin glass is measured using a local order parameter $m_i = \langle S_i \rangle$ at site i; the signature of the frozen state is that the spins do not deviate significantly with time at a local level, even though there may be no long-range order.

Because the freezing is a dynamic effect, a dynamic spin glass order parameter

$$q(t) = \overline{\langle \mathbf{S}_i(0).\mathbf{S}_i(t) \rangle} \tag{2.47}$$

is defined, with the Edwards-Anderson order parameter defined by

$$q_{\mathrm{EA}} = \lim_{t \to \infty} q(t) \tag{2.48}$$

A crucial question is whether random anisotropy systems (the subject of this study) exhibits glassiness.

2.7 Random Anisotropy Model

The rational for the RAM model was described in section 1.1.8 — the general Hamiltonian introduced by Harris, Plischke and Zuckerman (HPZ) to take into account the crystal field interaction in an amorphous compound is (equation 1.16)

$$\mathcal{H} = -\sum_{\langle ij \rangle} J_{ij} \mathbf{S}_i.\mathbf{S}_j - \mathbf{H}.\sum_i \mathbf{S}_i - \sum_i D_i \left(\mathbf{S}_i.\hat{\mathbf{n}}_i \right)^2 \tag{2.49}$$

where D_i is the strength of the anisotropy at site i in random direction $\hat{\mathbf{n}}_i$, uncorrelated with neighbouring sites. It is common to make the approximation $D_i = D$ in addition to the usual approximation of $J_{ij} = J$ for nearest neighbours (section 1.2.2).

For the two-component spins studied in this work, we can rewrite Hamiltonian 1.16 in the form

$$\mathcal{H} = -J \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j) - H \sum_i \cos\theta_i - D \sum_i \cos^2(\theta_i - \phi_i). \tag{2.50}$$

Generalising slightly:

$$\mathcal{H} = -J \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j) - H \sum_i \cos\theta_i - D \sum_i \cos\left[p(\theta_i - \phi_i) \right] \tag{2.51}$$

where p is the number of equivalent easy directions, equally spaced about ϕ_i. $p = 1$ corresponds to the random field model, where each site is subject to a local random field. $p = 2$ is the usual RAM model, with random uniaxial anisotropy, and larger values of p can be used for larger numbers of easy directions. In the infinite anisotropy limit, these correspond to the 'clock model'.

The interesting physics comes from the competing interactions: the exchange term tends to align the neighbouring spins, while the anisotropy term tries to direct each spin along the local easy axis. The spins can choose to align themselves with each other, with a large energy cost with respect to the anisotropy axes, or they can line up along the local axes, which carries an energy cost due to the misalignment with each other. Almost inevitably, the system must compromise, and it is this compromise that leads to interesting behaviour.

When D is very small, where the exchange dominates, the spins might be expected to be very nearly aligned in the direction of the external field (formally introducing a small symmetry-breaking field if necessary). In this situation,

$$\overline{D\cos^2(\theta_i - \phi_i)} \approx \frac{D}{2} + \text{more interesting part (see below)}, \tag{2.52}$$

and in order to compare energies with the energy of the pure system ($D = 0$) as a function of small D, it is convenient to redefine Hamiltonian 2.50 in the form

$$\mathcal{H} = -J\sum_{\langle ij\rangle}\cos(\theta_i - \theta_j) - H\sum_i \cos\theta_i - D\sum_i\left(\cos^2(\theta_i - \phi_i) - \frac{1}{2}\right). \tag{2.53}$$

in order to emphasise the more interesting part.

Conversely, when D is very large, the anisotropy dominates, and the spins might be expected to be very nearly aligned with the local axes, in which case

$$\overline{D\cos^2(\theta_i - \phi_i)} \approx D - \text{more interesting part (see below)}. \tag{2.54}$$

In order to compare energies as a function of large D, and to prevent $E \to \infty$ as $D \to \infty$, it is convenient to redefine the Hamiltonian in the form

$$\mathcal{H} = -J\sum_{\langle ij\rangle}\cos(\theta_i - \theta_j) - H\sum_i \cos\theta_i - D\sum_i\left(\cos^2(\theta_i - \phi_i) - 1\right). \tag{2.55}$$

In the infinite D (Ising) Limit, the spins are *constrained* to lie along the easy direction(s) — for $p = 1$, there is a unique ground state, while for $p > 1$ the exchange term can still influence which of the p equivalent directions the spin adopts. For $p = 2$ we can rewrite equation 2.55:

$$\mathbf{S}_i = \sigma_i \mathbf{n}_i = \pm \mathbf{n}_i \tag{2.56}$$

$$\mathcal{H} = -\sum_{\langle ij \rangle} J'_{ij} \sigma_i \sigma_j - \sum_i H'_i \sigma_i \tag{2.57}$$

where

$$J'_{ij} = J_{ij} \mathbf{n}_i . \mathbf{n}_j \tag{2.58}$$

and

$$H'_i = \mathbf{H}_i . \mathbf{n}_i \tag{2.59}$$

Equation 2.57 resembles an Ising spin glass (section 2.6), and there can be *frustration*, as shown on figure 2.5 — the configuration of the local anisotropy axes shown forms a perfect 'pinning centre' for a half vortex (figure 2.2).

2.8 Chudnovsky-Imry-Ma

The main analytic work on the HPZ model is the application of an Imry-Ma [21] argument by Chudnovsky and co-workers [22–25], reviewed in [26]. The following applies to zero temperature, unless otherwise stated.

As before (section 2.2.1), we consider the formation of a domain of size ξ lattice sites, containing $N \sim \xi^d$ spins misaligned with the surrounding spins (figure 2.1). We consider separately the single-spin energy contributions for the anisotropy energy E_{anis}, the exchange energy E_{exch}, and the energy due to the external field E_{field}. This is essentially a dimensional argument, so we are not interested in numerical factors.

2.8.1 Anisotropy

In a very large domain, we expect the anisotropy axes to be distributed isotropically, but in smaller domains, statistical fluctuations may be significant. For a domain of size N, in which all spins are aligned at an angle θ, for a given set of anisotropy axes $\{\phi_i\}$, we are interested in E_{anis} as a function of θ. This is illustrated in figure 2.6 where $E_{\text{anis}}(\theta)$ has been compared for three different domain sizes for particular choices of $\{\phi_i\}$.

We neglect signs, and use the small-D form of the Hamiltonian (equation 2.53) since we are calculating the energy compared with the fully-aligned state.

$$E_{\text{anis}}(\theta) = \frac{D}{N} \sum_i \left(\cos^2(\theta - \phi_i) - \frac{1}{2} \right) \tag{2.60}$$

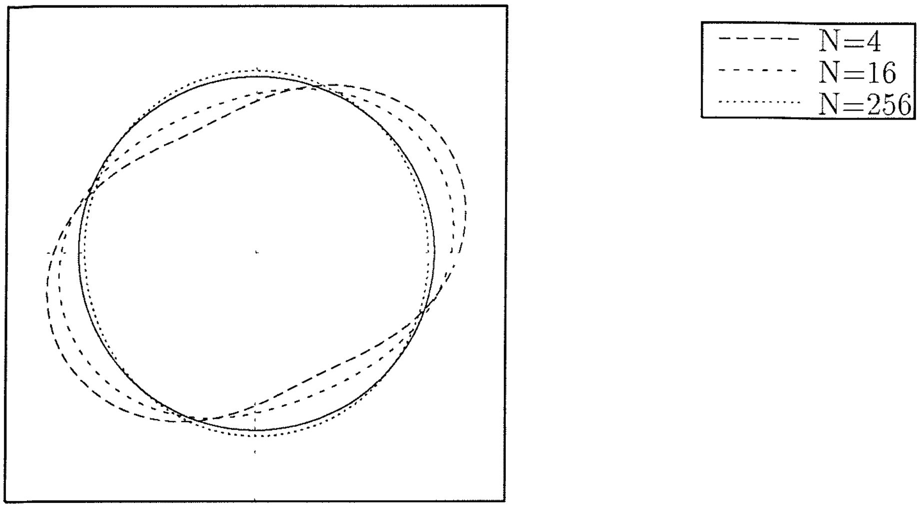

Figure 2.6: Polar plot comparing $E_{\text{anis}}(\theta)$ for three different domain sizes, clearly showing a preferred domain easy-axis, with energy difference dependent on domain size. Solid line is a circle of radius 1, representing an infinite domain, for comparison.

$$
\begin{aligned}
&\sim && \frac{D}{N}\sum_i \cos 2(\theta-\phi_i) && (2.61)\\
&= && \frac{D}{N}\sum_i (\cos 2\theta \cos 2\phi_i + \sin 2\theta \sin 2\phi_i) && (2.62)\\
&= && \frac{D}{N}\left(\cos 2\theta \sum_i \cos 2\phi_i + \sin 2\theta \sum_i \sin 2\phi_i\right) && (2.63)\\
&\equiv && \frac{D}{2}\frac{A}{N}\cos(2\theta-2\Phi) && (2.64)\\
\text{where} \quad A\cos 2\Phi &= && \sum_i \cos 2\phi_i, && (2.65)\\
A\sin 2\Phi &= && \sum_i \sin 2\phi_i, && (2.66)\\
\text{and} \quad A &= && \sqrt{\left(\sum_i \cos 2\phi_i\right)^2 + \left(\sum_i \sin 2\phi_i\right)^2}. && (2.67)
\end{aligned}
$$

A is precisely the displacement after a random walk of N steps, each step of unit length in random direction — this has the well known result $A \sim \sqrt{N}$. Thus $E_{\text{anis}}(\theta)$ fluctuates with an amplitude $\sim D/\sqrt{N}$, and a domain of spins of size$\sim \xi^d$ stands to reduce energy by $\sim D\xi^{-\frac{d}{2}}$ by rotating towards local easy direction Φ, but at a cost of misalignment along the domain walls. As $N \to \infty$, $E_{\text{anis}} \to 0$, while as $N \to 1$, $E_{\text{anis}} \to\sim D$, as expected.

2.8.2 Exchange

The exchange energy at the surface of a domain in a continuous system was considered in section 2.2.1: the energy of a domain of size ξ rotated through an angle of one radian is of order $\frac{1}{2}\xi^{d-2}$, so that the energy per spin is given by

$$E_{\text{exch}} \sim \frac{1}{2}\xi^{-2}. \tag{2.68}$$

2.8.3 Zero Field

In zero field, combining the anisotropy and exchange terms, and minimising for θ and ξ, we can calculate the energy per spin.

$$\begin{aligned} E &= E_{\text{anis}} + E_{\text{exch}} && (2.69) \\ &\sim -D\xi^{-d/2}\cos 2(\theta - \Phi) + J\xi^{-2}. && (2.70) \end{aligned}$$

The minimisation with respect to θ is trivial — the spins lie along $\theta = \Phi$, the average easy direction for the domain.

In more than 4 dimensions, E_{exch} always wins, $\xi \to \infty$, and the system is ferromagnetic. When $d = 4$, ξ is either 0 or ∞, depending on the ratio of $\frac{D}{J}$. A more careful analysis would be required to determine various numerical factors which have been disregarded.

In general, we are interested in cases when $d < 4$, and E has the form shown in figure 2.7. Equation 2.70 has a minimum...

$$\begin{aligned} \frac{\partial E}{\partial \xi} &= \frac{d}{2}D\xi^{-\frac{d}{2}-1} - 2J\xi^{-3} = 0 && (2.71) \\ \xi^{-\frac{d}{2}-1+3} &= \frac{4}{d}2J/D && (2.72) \\ \xi &\sim (J/D)^{\frac{2}{4-d}} && (2.73) \\ &= (D/J)^{-\frac{2}{4-d}}, \qquad \text{with energy} && (2.74) \\ E_{\text{min}} &\sim D\left(\frac{D}{J}\right)^{\frac{d}{4-d}}. && (2.75) \end{aligned}$$

Thus small D implies large clusters containing many spins ($D \to 0$ implies an infinite cluster, i.e. ferromagnetism — this argument does not consider the topological complexities of the two-dimensional XY model described in section 2.2). This system has been dubbed the *correlated spin glass*, or *correlated speromagnet* [24] — each cluster lies along its collective easy axis Φ, and the neighbouring

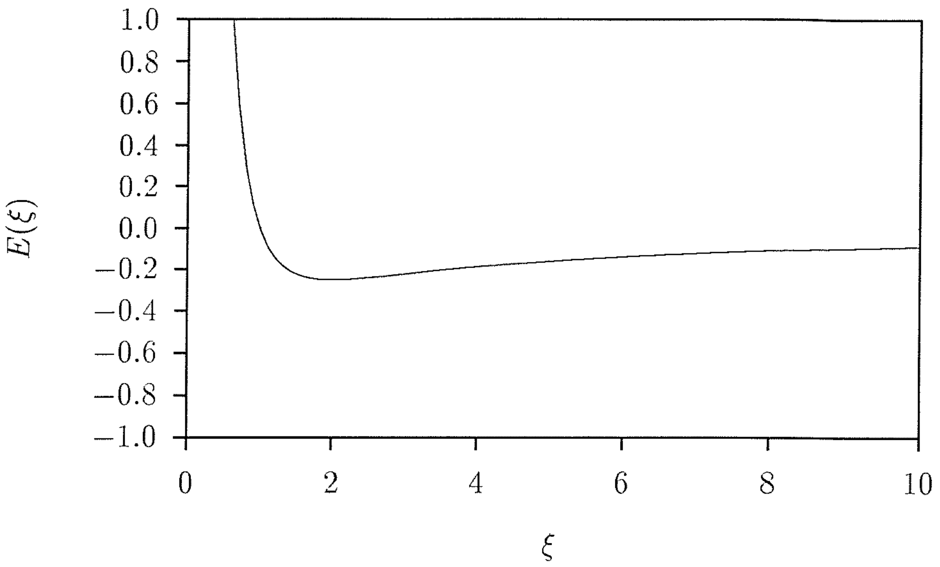

Figure 2.7: Form of equation 2.75, showing an energy minimum at the optimal domain size ξ.

clusters influence the orientation along this axis . This is analogous to the speromagnet or Ising spin glass (section 2.7), but for correlated domains of spins, rather the individual spins.

The spins in the CSG régime are arranged in a similar way to the illustration of domain formation in the two-dimensional (pure) XY model (figure 2.1); the direction of the local magnetisation varies smoothly through the sample, correlated in domains. The presence of random anisotropy, no matter how small, causes the breakup into domains, and therefore destroys any long-range order.

Large Anisotropy

As D increases, the clusters shrink, until they reach the order of the lattice spacing, in the *Ising limit.* In this limit, neighbouring spins influence only the orientation of a spin along its local axis (Hamiltonian 2.57). This is the (uncorrelated) speromagnet using Coey's terminology [27].

Using a simple extension to the ideas of Chudnovsky *et al.*, we can calculate the order of magnitude of the deviation of the spins from full alignment with the local anisotropy axes in this large anisotropy régime. Using the large D form of the HPZ Hamiltonian (2.55), each spin θ_i is pulled a small angle $\delta\theta_i$ from the local easy direction ϕ_i, releasing exchange energy proportional to $J\delta\theta_i$, but at a cost of an increase in the anisotropy energy, proportional to $D\cos^2\delta\theta_i$. Minimising the change in energy, we find

$$\begin{aligned} \delta\theta_i &\sim J/D, && (2.76)\\ \delta E_{\rm min} &\sim -J^2/D. && (2.77) \end{aligned}$$

2.8.4 Thermal Depinning

The energy minimum given by equation 2.75 is not very deep. A simple argument resulting from a collaboration with Prof. Wayne Saslow suggests that at temperatures of the order of $E_{\rm min}$, the thermal energy may allow the system to lift itself out of the low temperature Imry-Ma state, into some other régime. That is, the system is no longer 'pinned' by the configuration of the anisotropy axes above a temperature

$$T_{\rm depinning} \sim D\left(\frac{D}{J}\right)^{\frac{d}{4-d}}. \qquad (2.78)$$

2.8.5 Local Energy Fluctuatuons

While the spins within an Imry-Ma domain lie on average in the direction of the collective easy axis, there are fluctuations. The contribution to the energy can be calculated using an analysis similar to that used for spin-waves (section 2.2.2).

Writing the Hamiltonian in a continuous form,

$$\begin{aligned} \mathcal{H} &= \int \left[-D\cos^2(\theta-\phi) - J\left(\nabla\theta\right)^2\right] d^dx \\ \frac{\partial\mathcal{H}}{\partial\theta} &= J\nabla^2\theta + D\sin 2(\theta-\phi) \\ &= 0 \end{aligned}$$

Taking Fourier transforms,

$$\begin{aligned} -k^2\theta_k + D\left(\sin 2\phi\right)_k &= 0 \\ \theta_k &\sim \frac{D\left(\sin 2\phi\right)_k}{k^2}. \end{aligned}$$

Summing the contributions to the energy over the available modes,

$$\sum_{k=a}^{\xi} k^2\theta_k\theta_{-k} \sim \sum_{k=a}^{\xi} \frac{D^2}{k^2}\left(\sin 2\phi\right)_k\left(\sin 2\phi\right)_{-k}$$

$$
\begin{aligned}
&\sim \quad D^2 \log \frac{\xi}{a} \sim D^2 \log \frac{J}{D} && d = 2 \\
&\sim \quad D^2 && d = 3
\end{aligned}
$$

2.8.6 Small Field

If we now add a small external field $H \ll J$ as a perturbation, we can consider the extent to which the spins will move, and hence calculate the magnetic susceptibility. We neglect the exchange interaction, since neighbouring spins will usually be moved by similar amounts in the same direction, and we assume that the cluster size ξ is unchanged. In moving through a small angle $\delta\theta_i$ towards the field, the energy released is $E_{\text{field}} \sim H\delta\theta_i$. Again, we must minimise E to find the position of the compromise:

$$
\begin{aligned}
\delta E_i &= \delta E_{\text{anis}} + \delta E_{\text{field}} && (2.79) \\
&\sim +D\xi^{-\frac{d}{2}}\delta\theta_i^2 - H\delta\theta_i && (2.80) \\
\frac{\delta E}{\delta\theta_i} &= 0 && (2.81) \\
\delta\theta_i &\equiv \frac{H}{H_c} \qquad \text{where} && (2.82) \\
H_c &= D\left(\frac{D}{J}\right)^{\frac{d}{4-d}}, \qquad \text{with magnetic susceptibility} && (2.83) \\
\chi &= \frac{\delta m}{\delta H} = H_c^{-1} \sim D^{\frac{-d}{4-d}}. && (2.84)
\end{aligned}
$$

This argument therefore predicts a *finite* magnetic susceptibility at low temperatures.

2.8.7 Large Field

In this régime, the field can no longer be regarded as a small perturbative effect for fields larger than the H_c defined by equation 2.83. We can, however, reformulate the analysis so that it is the *anisotropy* that is the perturbation, causing small deviations from perfect alignment in the field direction. As before, the largest gain in energy is made by limiting the size of the domains, but at a cost in exchange energy — we suppose typical small deviations of θ from the field direction, correlated in domains of size ξ. This picture has been dubbed the *ferromagnet with wandering axes* in the literature [24]. The field now defines which direction along the collective easy axes the clusters in section 2.8.3 lie, and the spins change direction smoothly through the sample.

The exchange energy is unchanged (but we must now explicitly consider the change in spin direction):

$$E_{\text{exch}} = J\xi^{-2}\theta^2, \tag{2.85}$$

while the energy cost per spin for deviation from the field is given by

$$E_{\text{field}} = H(1 - \cos\theta) \approx H\theta^2. \tag{2.86}$$

The domains cannot move all the way to the local average easy axes, but lose an energy proportional to θ in tipping an amount θ towards them :

$$E_{\text{anis}} = -D\xi^{-d/2}\theta. \tag{2.87}$$

In the original work [24], it was asserted that, at the minimum energy, the three terms would be of similar magnitude, giving

$$\begin{aligned} \xi &\sim \left(\frac{J}{H}\right)^{\frac{1}{2}}, && (2.88)\\ \theta &\sim D\left(\frac{J}{H}\right)^{1-\frac{d}{4}}, && (2.89)\\ \delta m &\sim D^2\left(\frac{J}{H}\right)^{2-\frac{d}{2}}. && (2.90) \end{aligned}$$

However, formally minimising the total energy with respect to both θ and ξ gives a different result. Minimising the total energy with respect to θ,

$$\begin{aligned} E &= \left(H + J\xi^{-2}\right)\theta^2 - D\xi^{-\frac{d}{2}}\theta \\ \frac{\partial E}{\partial \theta} &= 2\left(H + J\xi^{-2}\right)\theta - D\xi^{-\frac{d}{2}} \\ &= 0 \\ \theta &= \frac{D\xi^{-\frac{d}{2}}}{2\left(H + J\xi^{-2}\right)} \end{aligned}$$

and substituting into the expression for E,

$$E = -\frac{D^2\xi^{2-d}}{4\left(J + H\xi^2\right)}.$$

which is always minimised by $\xi \to \infty$ for $d = 2$ and $d = 3$.

2.8.8 Very Large Field

As before (section 2.8.3), ξ cannot shrink below one lattice spacing in equation 2.88 — for very large fields $H > J$, we ignore the exchange term, and minimise the field and anisotropy terms. With no exchange to consider, these results are independent of dimensionality d.

$$\theta \sim \frac{D}{H} \tag{2.91}$$

$$\delta m \sim \left(\frac{D}{H}\right)^2, \tag{2.92}$$

with a field dependence different from equation 2.90, which may be detectable.

2.8.9 Thermal Effects at High Fields

A further extension to the ideas of Chudnovsky *et al.* is to consider also deviations from saturation in a large field due to a non-zero temperature, rather than the randomness, so that we minimise the free energy $F = E - TS$ rather than E. If thermal excitations allow a spin to deviate from saturation by angle up to $\delta\theta$, the volume of phase space available is proportional to $\delta\theta$, and the entropy $S \sim \log\delta\theta$. Neglecting the exchange and anisotropy terms in the very large field limit, we minimise

$$F = H\delta\theta^2 - T\log\delta\theta \tag{2.93}$$

with respect to $\delta\theta$, and find that

$$\delta m \sim \delta\theta^2 \sim \frac{T}{H}. \tag{2.94}$$

This effect will dominate equation 2.92 when $T > D^2/H$. Conversely, because of the differing field dependencies, equation 2.94 will always dominate for sufficiently large field $H > D^2/T$.

2.8.10 Summary

In two dimensions, a correlated spin glass ($D < J$) has four régimes:

Region	Lower field limit	Upper field limit	δm	ξ	description
a	-	D^2	-	D^{-1}	C.S.G.
b	D^2	J	D^2/H	$H^{-1/2}$	F.W.A.
c	J	D^2/T	D^2/H^2	1	very large field
d	D^2/T	-	T/H	1	entropy-induced

and there are two régimes with $\delta m \sim H^{-1}$, separated by a régime (c) with $\delta m \sim H^{-2}$.

In three dimensions, the table reads

Region	Lower field limit	Upper field limit	δm	ξ	description
a	-	D^4	-	D^{-2}	C.S.G.
b	D^4	J	$D^2/H^{\frac{1}{2}}$	$H^{-\frac{1}{2}}$	F.W.A.
c	J	D^2/T	D^2/H^2	1	very large field
d	D^2/T	-	T/H	1	entropy-induced

Further, as we change the relative sizes of the parameters D, J and T, we expect the crossovers to change:

- for a speromagnet, $(D > J)$, region b vanishes.
- at high temperature $(T > D^{\frac{4}{4-d}})$, region c vanishes. In two dimensions, there is nothing to separate the two H^{-1} régimes.
- at very high temperature $(T > J)$, region b vanishes,

2.9 Topological Defects (Vortices)

It is useful to consider explicitly the effect of the random anisotropy on the vortices which are significant in the (pure) XY system (section 2.2). For a sufficiently large system, there is a significant chance of natural 'vortex pinning centres', where a vortex is trapped in a potential well, as illustrated on figure 2.8. In order for the vortex centre to move one lattice spacing, at least one spin near the vortex centre must move 90° from a good alignment with the local axis to a perpendicular configuration, as shown. (The spins that are not adjacent to the vortex centre do not suffer such a violent change in angle, and so do not play a significant role.)

We might suppose that at a low temperature, with less thermal energy than the depth of the pinning centres, vortices can become trapped. As the temperature is increased, they may become free to move around the system and mutually annihilate as usual.

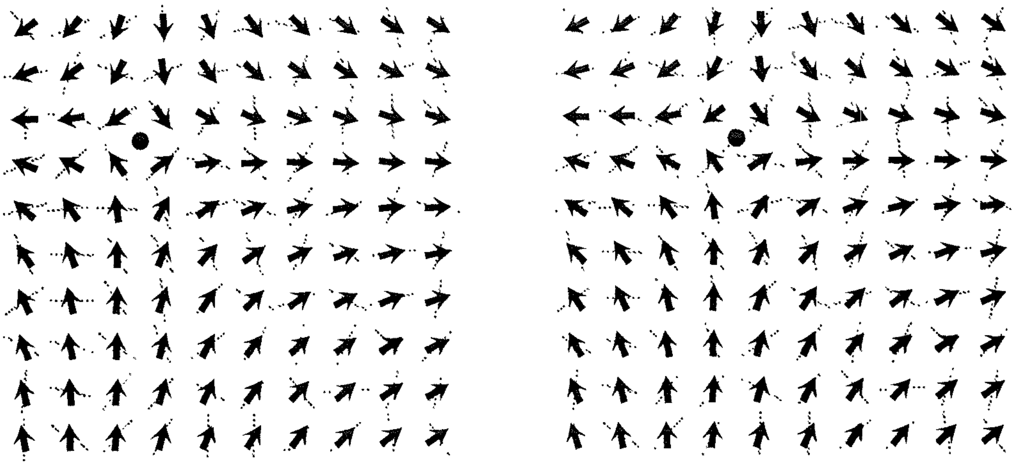

Figure 2.8: Illustration of the energy required for a vortex to escape from a pinning centre. The grey lines depict the local anisotropy axis at each site, and the vicinity of the vortex is highlighted. The only significant change between the two configurations is that the vortex centre has moved, and two spins in the centre of the highlighted region have turned through 90°. While on the left, all the spins around the vortex are almost aligned with the local anisotropy axes, the two central spins are badly misaligned with the anisotropy axes, bringing an energy cost proportional to the anisotropy strength D.

2.10 Other Theoretical Analyses

Early work [8, 28–30] used mean field theory (section 2.3) to study the HPZ Hamiltonian. The original work by HPZ [8] found the system to be ferromagnetic, but with reduced spontaneous magnetisation at zero temperature and a reduced transition temperature. Callen, Liu and Cullen [29], starting from a fully aligned state, calculate the extent to which each spin will be pulled towards the local anisotropy axis. This gives the remanent magnetisation after removal of a large magnetic field, but the authors assert that in the absence of any aligning field, the local magnetisation will rotate smoothly through the sample, in domains whose size depends on D, much like the Chudnovsky-Imry-Ma picture. Indeed, the calculations reproduce the results presented in section 2.8 where the limits $\xi \to 1$ have been reached.

Harris and Zobin [28] considered also the transverse components of the magnetisation which mean field theory neglects (section 2.3.2), and supported the view that the system is ferromagnetic for small anisotropy strengths, but predicted a spin glass state for large D; they confirmed these results by Monte Carlo simulation. Patterson, Gruzalski and Sellmyer [30] perform a numerical calculation of local mean-field equations, and report that the spin-glass phase found by [28] is metastable, and that the aligned state has lower energy. They show that to first order in D, the transition to the paramagnetic phase is unaffected by the anisotropy.

Derrida and Vannimenus [31] have studied the model in the infinite range limit (which is equivalent to mean field theory) and find that the low temperature state is ferromagnetic, and that the transition to the paramagnetic state is unaffected by D for all anisotropy strengths.

The replica trick (section 2.5) has been used by Chen and Lubensky [32] to show that a spin glass state is possible at low temperatures below some $T^* \propto D$. Several other authors have used the replica trick to transform the problem into a translationally invariant one, and hence apply renormalisation techniques (section 2.4).

Dotsenko and Feigelman [33] explicitly consider the role of vortices in a two-dimensional XY model with p-fold anisotropy of random magnitude; for $p \geq 2$, the system cannot be cooled below a temperature $T^* = 4\pi/p$, so that the system at lower temperatures show the same correlation properties as at $T = T^*$. That is, the system is renormalised back to T^*, and shows no spin-glass state at lower temperatures. The system is pure-XY-like for intermediate temperatures $T^* < T < T_{\mathrm{KT}}$. For the $p = 2$ case in which we are interested, $T^* > T_{\mathrm{KT}}$, so that no phase transition is expected for this case. These authors have also studied the three-dimensional HPZ model with two spin components in [34] and find that it does have a low-temperature spin-glass state.

Cardy and Ostlund [35] use renormalisation techniques to study the two-dimensional XY model with p-fold anisotropy. In a similar way to [33], they find that the Kosteritz-Thouless phase survives the addition of the anisotropy provided that $p \geq 3$, and that below a temperature $T^* = 4\pi/p^2$ the system is glassy. Le Doussal and Giamarchi [20] point out that the fixed point found in [35] at low temperatures is unstable against replica symmetry breaking (section 2.5), and predict that the spin correlations go like $T\log(r)$ above T^* and $T^*\log(r)$ for $T < T^*$, in a similar way to the prediction of [33].

Bray and Moore [36] study the $p = 2$ system in two dimensions in the limit of infinite anisotropy, finding an Ising-like transition at zero temperature to a non-ferromagnetic ground state. Evidence for glassy behaviour is presented.

Aharony and Pytte [37] calculate the magnetic equation of state to leading order in D, and find that the low-temperature state has no spontaneous magnetisation but has infinite magnetic susceptibility in zero field, and the spin correlations are found to decay as a power law with distance. The authors later consider the discontinuous change from a spontaneous magnetisation $M > 0$ for $D = 0$ to a nonmagnetic state for $D > 0$ [38]. They conclude that the susceptibility probably agrees with equation 2.84, but do not rule out the possibility of an infinite susceptibility. This latter analysis does not cover the two-dimensional XY model which does not have a spontaneous magnetisation even for $D = 0$.

2.11 Previous Computational Results

Early work by Chi and Alben [39] studying the HPZ model on a random lattice model found a magnetic ground state, despite the presence of the randomness. This result was later discounted as a finite-size effect by later molecular-dynamics simulations [40]; both studies used rather small systems with 10^3 spins.

Similarly, Harris and Sung [41] initially found that a ferromagnetic state had lowest energy, whatever the strength of the random anisotropy. Later work by Harris [42] found that spins could be flipped to give a lower energy, and that the energy barriers between the two states were approximately linear in anisotropy strength D. These results are based on small systems of 512 spins.

Jayaprakash and Kirkpatrick [43] performed extensive simulations in the Ising limit, for both two-component spins in two spatial dimensions and three-component spins in three spatial dimensions. They found a low temperature phase with short-range ferromagnetic order, and a peak in specific heat, but no critical behaviour, in the crossover to high temperature. They also found hysteretic behaviour in magnetic fields, and remanent magnetisation which is a signature of magnetic behaviour, together with a coercive field which decreases with increasing temperature.

Chakrabarti [44] studied this Ising limit further, measuring the spin glass order parameter, and studying the effects of finite size, concluding that the low temperature phase is glassy. This work has been criticised [45], since the ground state did not appear to have been correctly reached, so that the results have been cast into doubt due to the simple single-spin update algorithm used.

A great deal of work has been done by Fisch [45–48], again in the Ising limit, and predominantly at zero temperature. A study of the ground state in the Ising limit in three dimensions is presented in [45]. Finite size scaling of the magnetisation for systems up to 16^3 shows power-law decay of the spin correlations for two-component spins, suggesting infinite magnetic susceptibility, in agreement with the predictions of [37]. Exponential decay is indicated for the spin correlations of three-component spins, giving a finite susceptibility. This work is extended in [46] to explore the ground states of systems with higher orders of anisotropy. The work was further extended to larger systems and non-zero temperatures in [47], though still in the infinite anisotropy strength limit. The ground state energies and magnetisations are calculated, and spin correlation functions are found to obey a $k^{-2.4}$ law, in agreement with neutron-scattering experiments. A transition to a paramagnetic phase is found around $T/J = 1.91$. This work is repeated for threefold anisotropy in [48].

Reed [49–51] has studied the XY model with finite random anisotropy, though our results in chapters five and six indicate that he was unfortunate in his choice of $D = 1$. A study of the system in six-fold anisotropy [49] finds an intermediate Kosterlitz-Thouless régime, and a low temperature régime in which the parameters are renormalised back to T^* as predicted by [20, 33]. A comparison of p-fold anisotropy for $p = 2, 3, 4, 5, 6$ [50] is made for systems up to 16^2. The $p = 2$ system is found to be non-quasiferromagnetic, below a transition temperature of $T/J = 1.1$. Finite-size scaling is used on the three-dimensional XY model with two- and three-fold anisotropy [51], and the system is found to be non-quasiferromagnetic below $T_c = 2.2$.

Finally, Dieny and Barbara [52] have studied the two dimensional XY model with weak two-fold anisotropy at zero temperature but in non-zero magnetic fields, to find first magnetisation curves and hysteresis loops. They conclude that the effects of vortices are important. Saslow and Koon [53] have performed similar calculations on three-component spins in three dimensions.

Chapter 3

Monte Carlo Simulation

As described in section 1.2.2, even a small number of interactions between the variables of a system can render it intractable to theory. Although it is often possible to make simplifications without losing the essential physics, numerical simulation is one way of calculating the properties of problems of arbitrary complexity. The problem, however, is that the large number of degrees of freedom available to typical systems results in an *enormous* phase space over which to sum.

Consider, for example, a simple 8 by 8 Ising system (section 2.1), with 64 spins and therefore a total of $2^{64} \approx 10^{19}$ possible configurations. Even if we could evaluate the energy of 10^{12} configurations per second, the calculation would take seven months — however, at the time of writing, not even the fastest computer in the world can achieve 'tera-flop' performance (10^{12} *fl*oating point *op*erations per second), so we couldn't even simply accumulate the results that fast!

The situation gets even worse when we wish to investigate a system with *continuous*, rather than discrete, variables. We must discretise the parameters, but of course the total number of configurations grows rapidly with the number of quanta we consider — if just sixteen positions are made available to each spin in an XY model (section 2.2), only a 4 by 4 system can be computed in the same seven months.

3.1 Illustration I - the Ising Model

Let us reconsider the one-dimensional Ising model: reviewing the exact solution in terms of the 'kinks' introduced in section 2.1, we can sort the configurations into order of increasing energy, and hence find an expression for the density of states.

Suppose there are $N+1$ spins in the chain, and therefore N bonds at which we can put a kink. (We ignore here the degeneracy in the direction of the first spin, but could have included it formally by allowing a kink before the first spin, giving $N+1$ possible positions for the kinks.) There is one configuration (the ground state) with no kinks, and N places where we can put the first kink, so that there are N configurations with energy $2J$. There are now $N-1$ bonds where we can activate a second kink, giving $N(N-1)$ — except that we have counted each configuration twice, as the kinks are indistinguishable. In general, there are $g(k)$ configurations with k kinks, where

$$\begin{aligned} g(k) &= \frac{N(N-1)(N-2)\ldots(N-k+1)}{k(k-1)(k-2)\ldots(2)(1)} & (3.1)\\ &= \frac{N!}{(N-k)!k!} & (3.2)\\ &\equiv \binom{N}{k}. & (3.3)\end{aligned}$$

The partition function can now be written

$$Z_N(\beta) = \sum_{k=0}^{N} g(k)e^{-2\beta kJ} \qquad (3.4)$$

where $\beta = \frac{1}{kT}$. Table 3.1 shows the individual contributions to the partition function for a system with $N = 32$ at a reduced temperature $\frac{T}{J}$ of 0.25. It is clear that the partition function which is accumulated in the last column is correct to 6 significant digits when terms with only 0, 1 or 2 kinks are considered. This amounts to only 529 of the 2^{32} (four billion) total configurations. At a temperature of $0.5J$, we need terms up to 7 kinks, or 0.1% of the total configurations.

In general, while $g(k)$ grows very rapidly with k, the exponential Boltzmann factor decays even faster and the product $g(k).e^{-\beta E(k)}$ shows a narrow maximum. For configurations near the ground state, the reluctance of a system to take on a state with higher energy is compensated by the large number of states available

k	$g(k)$	$\sum_{i=0}^{k} g(i)$	$g(k)exp(-2k\beta J)$	$\sum_{i=0}^{k} g(i)exp(-2n\beta J)$
0	1	1	1	1.000000
1	32	33	0.0107348	1.010735
2	496	529	5.58174e-05	1.010791
3	4960	5489	1.87247e-07	1.010791
4	35960	41449	4.55403e-10	1.010791
5	201376	242825	8.55517e-13	1.010791
6	906192	1149017	1.29147e-15	1.010791
7	3365856	4514873	1.60918e-18	1.010791
8	10518300	15033173	1.68694e-21	1.010791
9	28048800	43081973	1.50908e-24	1.010791
10	64512240	107594213	1.16435e-27	1.010791
11	129024480	236618693	7.81192e-31	1.010791
12	225792840	462411533	4.58606e-34	1.010791
13	347373600	809785133	2.36685e-37	1.010791
14	471435600	1281220733	1.07756e-40	1.010791
15	565722720	1846943453	4.33776e-44	1.010791
16	601080390	2448023843	1.5461e-47	1.010791
17	565722720	3013746563	4.88151e-51	1.010791
18	471435600	3485182163	1.36464e-54	1.010791
19	347373600	3832555763	3.37315e-58	1.010791
20	225792840	4058348603	7.35518e-62	1.010791
21	129024480	4187373083	1.40994e-65	1.010791
22	64512240	4251885323	2.3649e-69	1.010791
23	28048800	4279934123	3.44929e-73	1.010791
24	10518300	4290452423	4.33915e-77	1.010791
25	3365856	4293818279	4.658e-81	1.010791
26	906192	4294724471	4.20696e-85	1.010791
27	201376	4294925847	3.13617e-89	1.010791
28	35960	4294961807	1.87869e-93	1.010791
29	4960	4294966767	8.69284e-98	1.010791
30	496	4294967263	2.91612e-102	1.010791
31	32	4294967295	6.31129e-107	1.010791
32	1	4294967296	6.61626e-112	1.010791

Table 3.1: Contributions to Ising partition function by states with different numbers of 'kinks' k, at temperature $T = 0.25J$.

with that energy, but if the energy is too high, the Boltzmann factor wins. The log of the number of states is the entropy of the system, and this is why it is the free energy $F = E - TS$ that is usually minimised in statistical mechanics: as the temperature is increased, the entropy S has a more important effect, via the density of states $g(k)$.

Unfortunately, however, we do not usually know the form of g in advance, and thus we do not know from the energy of a configuration alone, whether it contributes significantly to the partition function sum. We must just count 1 each time we meet a configuration with a given energy. But as has been seen, most of the configurations make negligible contribution. We need some way of selecting the important configurations if we are to evaluate the partition function with a sensible amount of computer power and in a sensible time.

3.2 Review of numerical integration

The simplest numerical integration technique is to sample the function at regular intervals, i.e. explicitly treating integration as a summation. Graphically, we find the area by summing lots of thin rectangular strips.

$$\int_a^b f(x)\,dx \approx \sum_{i=0}^{n} f(a + i.\delta x)\delta x \tag{3.5}$$

where δx is $\frac{b-a}{n}$. The smaller we make δx, the more accurate is the approximation, at a cost of more time. (This is true up to a limit — when δx is very small, rounding errors can lead to a loss of accuracy.) There are enhancements to this technique, which amount to considering more points at a time, and interpolating the function between the points.

Another way of implementing the same algorithm is to use random numbers to choose the points at which to sample:

$$\int_a^b f(x)\,dx \approx \frac{a-b}{n} \sum_{i=0}^{n} f(x_i) \tag{3.6}$$

with x_i chosen uniformly on $[a, b]$. This has the advantage that a result is immediately available after only a few points have sampled, and sampling need be continued only until it is deemed that sufficient accuracy has been established (generally by monitoring the effect of additional measurements).

3.2.1 Importance Sampling

When the function is sharply peaked, however, it does not make sense to use a constant strip width — using narrow strips would waste too much time at the 'uninteresting' parts of the function, but wide strips would give too little accuracy at the 'important' parts of the function. Let us make the strip width w a function of x:

$$\int_a^b f(x)\,dx \approx \sum f(x_i)w(x_i) \tag{3.7}$$

where $x_i = a + \sum_{j=0}^{i} w(x_j)$ and $\sum_{j=0}^{n} = b - a$.

We now write our function as the product of two others, for generality:

$$\int f(x)\,dx \equiv \int g(x)h(x)\,dx \tag{3.8}$$

where $h(x)$ is a smooth 'envelope', following roughly the shape of f, and $g(x)$ contains the fine details. We want to make our strip width $w(x)$ large when h is small, and narrow when h is large: $w(x) = \frac{1}{h(x)}$, for example, might be suitable:

$$\begin{aligned} \int_a^b f(x)\,dx &\approx (b-a)\frac{\sum g(x_i)h(x_i)w(x_i)}{\sum w(x_i)} && (3.9)\\ &= (b-a)\frac{\sum g(x_i)}{\sum w(x_i)} \qquad n = \sum w(x_i). && (3.10) \end{aligned}$$

This can be interpreted another way: our rectangles have uniform width δx and height $g(x)$, but where w is small, we count more of them in *the vicinity* of each point; that is, the *density* of samples is large when w is small. Thus, our final result is

$$\int_a^b f(x)\,dx \approx \frac{(b-a)}{n}\sum_{i=0}^{n} g(x_i), \qquad \text{since } n = \sum_i w(x_i) \tag{3.11}$$

When we now move to the alternative algorithm of choosing the samples randomly, the form of equation 3.11 is unchanged if we draw n samples on $[a, b]$ with a probability distribution $\frac{1}{h(x)}$. It is sensible to choose a function $h(x)$ that allows easy random number generation, such as a gaussian, or we will waste as much time generating the numbers as we hoped to save in the first place!

3.2.2 Application to Statistical Mechanics

In statistical mechanics, we are interested in averaging physical quantities over the configurations $\{x_i\}$:

$$\overline{A} = \int_{\{x_i\}} A(\{x_i\}) e^{-\beta E(\{x_i\})} \Big/ \int_{\{x_i\}} e^{-\beta E(\{x_i\})} \tag{3.12}$$

If we choose $h(\{x_i\}) = e^{-\beta E(\{x_i\})}$ in equation 3.11, then

$$\begin{aligned} \overline{A} &\approx \sum A(\{x_i\}) / \sum 1 && (3.13) \\ &= \frac{1}{n} \sum A(\{x_i\}) && (3.14) \end{aligned}$$

Thus, provided we can draw configurations at random with a Boltzmann probability distribution $e^{-\beta E(\{x_i\})}$, we can find thermal averages of as many quantities as we like, using a simple average over the generated configurations.

3.3 Metropolis Algorithm

In most of the following, we use *un*normalised probabilities such that the ratio of two numbers gives a relative probability, but the area under the distribution is not explicitly chosen to be 1.

One way of generating random *numbers* with an arbitrary distribution $p(x)$ is to (completely) 'enclose' $p(x)$ under another distribution $q(x)$ which *can* be generated (such as a gaussian, or a top-hat function which is constant everywhere that $p(x)$ is non-zero), in a similar way to the rewriting of a function as the product of two others in section 3.2.1. Then generate x_i from the $q(x)$ distribution, and accept x_i as a number drawn from $p(x)$ with probability $p(x_i)/q(x_i)$ — i.e. draw another random number z uniformly on $[0, 1]$, and accept x_i if $z < p(x_i)/q(x_i)$. Repeat until a number is accepted.

Conceptually, this can be thought of as throwing darts at a graph of the distribution. If the dart lands under the curve, we accept the x-coordinate as a random number, otherwise we rethrow. [The piece of paper represents the distribution $q(x)$. If we miss the piece of paper, then we have failed to draw on $q(x)$, which is a smaller distribution than the whole room which encloses it...] Clearly, the enclosing function should be chosen to pass as closely as possible over the real distribution, since any darts landing above the distribution are 'wasted'.

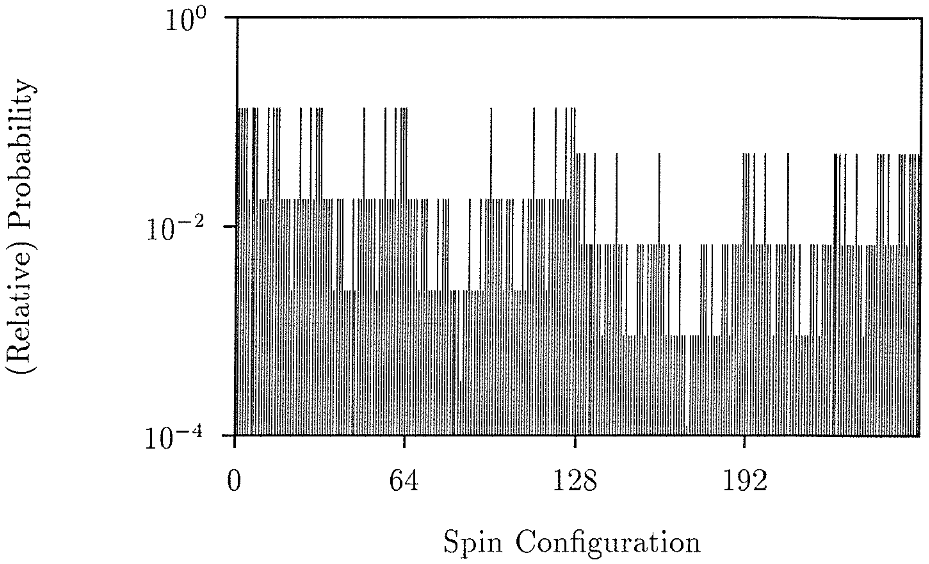

Figure 3.1: (Unnormalised) probability distribution for all possible configurations (unsorted) for 8-particle Ising system at $T = J$, shown on log scale to compress the range.

When we apply this technique to real statistical mechanics systems, however, we cannot draw the configurations with anything other than a uniform distribution, as we cannot sort the configurations into any sort of order of increasing probability (in general — if we could, we would not need to run a simulation in the first place!), and we cannot draw random configurations from anything other than a (reasonably) smooth distribution. (If we were to generate configurations by beginning with random values for the phase variables, then imposing constraints as we find the energy to be growing too large, a strong bias would be introduced into the phase space sampled.)

Consider, for example, the one dimensional Ising system with N spins. Without any prior knowledge, we might use the binary representation of a random integer drawn from $[0, 2^N - 1]$ to set the N spins, but for $N = 8$, for example, this gives a very bumpy probability distribution, even at high temperature, as shown in figure 3.1 for temperature $T = J$. We can only draw states i for the system at random, uniformly from $[0, 2^N - 1]$, and accept each one with a probability $e^{-2J\beta K_i}$ where K_i is the number of kinks in state i. Clearly this is very inefficient; even at this high temperature, most darts thrown at the graph would not lie under the required distribution.

Looking at this from a different point of view, all we are really doing is simple sampling, but instead of counting $e^{-\beta E}$ for each one, we count 1 with probability $e^{-\beta E}$, and we have gained nothing.

3.3.1 Markov Chains

Metropolis *et al.* [54] hit upon the idea of generating the configurations not from scratch, but by building new configurations by slight modifications to previous ones, since the energy cannot change *too* much if only small changes are made. In the Ising model, for example, successive configurations might be different by the flip of a single spin. Provided we don't do something drastic with our choice of transformation, this will automatically take into account the density of states term in the partition function sum, since if there are many configurations with a given energy, the system will repeatedly try to get into such a state. For example, from a state with one kink, there are $N-1$ nearby states with two kinks, but only one adjacent state with no kinks, so the system will make many more attempts to get to a state with higher energy, of which a proportion determined by the Boltzmann factor will be accepted. If a suitable operator M can be defined such that the configurations generated by its application have the Boltzmann distribution, then M can be applied over and over again, to generate a chain of correctly distributed configurations — M is a *Markov* operator, generating a Markov chain of states.

Borrowing notation from quantum mechanics, we might sort all the possible configurations into order $E_0 \leq E_1 \leq E_2 \leq \cdots$, and write a vector containing the probabilities that the system is in a given configuration. A conceptually easier way of looking at this is to suppose we have a very large number of different configurations (either the time progression of one system, or the instantaneous states of a large ensemble of systems), and the vector counts the *number* of systems in a given state. We can track the effect of M on the system by operating it on the probability vector:

$$\begin{pmatrix} p_1' \\ p_2' \\ \vdots \end{pmatrix} = M \begin{pmatrix} p_1 \\ p_2 \\ \vdots \end{pmatrix}.$$

M is now a matrix, with M_{ij} the transition probability from configuration j to configuration i.

We can now impose our constraints that configurations are drawn with Boltzmann distribution — p_i, the probability that the system is in configuration i with energy E_i, must be (proportional to) $e^{-\beta E_i}$, and the column of Boltzmann factors must be an eigenvector of the operator M, so that the result of operating M on a correct Boltzmann distribution preserves the probabilities — this requires that the number of systems leaving state i is balanced by the total number entering state i, so that the number of systems in state i is unchanged:

$$\sum_j M_{ji} p_i = \sum_j M_{ij} p_j . \tag{3.15}$$

Metropolis *et al.* go further, applying the *principle of detailed balance*, requiring that the individual terms in the sum balance; the number of systems making the transition from configuration i to j is exactly equal to the number making the inverse transition. Of the p_j configurations in state j, a fraction M_{ij} are moving to state i, and a fraction M_{ji} of the p_i in state i are moving to j. Constraining these to be equal:

$$\begin{aligned} M_{ij} p_j &= M_{ji} p_i && (3.16) \\ \frac{M_{ij}}{M_{ji}} &= e^{-\beta \Delta E_{ij}} && (3.17) \end{aligned}$$

where ΔE_{ij} is the energy difference between configurations i and j. One common way of achieving this is

$$M_{ij} \sim \begin{cases} \frac{1}{\tau} & E_i < E_j \\ \frac{1}{\tau} e^{-\beta \Delta E_{ij}} & E_i > E_j \end{cases} , \tag{3.18}$$

where the τ term is normally left out, such that a trial step to a state with lower energy is *always* accepted, but can be used to reduce this probability, to allow finer control over the system dynamics.

It is also necessary to show that the system tends towards the Boltzmann distribution no matter how it is started. Suppose that, at any point along the chain of configurations, configuration i has a higher probability than it should. At the next iteration, there will be more probability flowing out of i than is flowing into it, since the outward flow is proportional to the current level. Thus, even if we start our chain with all parameters in a certain state (all spins up, say), so that the probability is 1 for that configuration and 0 for all others, or in a random configuration (all probabilities equal), we will tend towards equilibrium as we iterate our Markov process — of course, it could still take an infinite number of iterations to get there! In fact, it is found empirically that the natural timescale τ for relaxation of fluctuations grows as some power z of the system correlation length — near a phase transition, this results in a relaxation time

$$\tau \sim L^z \tag{3.19}$$

for finite systems (section 3.6), and simulations exhibit *critical slowing down.*

In practice, we usually discard a large number of configurations, to allow the system to equilibrate, ensuring that we are not too far from the correct Boltzmann distribution of state population probabilities. Also, if we are running several simulations for different values of some thermodynamic parameter, such as temperature or magnetic field, we might start each subsequent simulation with the final configuration of the previous run, since that might be expected to be closer to the correct distribution than we could expect to achieve starting from either a random or an *a priori* configuration. Care must be taken, however, to avoid any hysteretic behaviour at the critical temperature, which is really just a symptom of insufficient equilibration.

Because, by definition, subsequent configurations are very similar, differing in only one or two of the phase space coordinates, it is customary to include not every configuration in our thermal averages, but to miss out a number of states between measurements, particularly if a measurement is relatively time-consuming; typically, we perform one, or several, Monte Carlo updates for every coordinate in phase space, such as each spin or particle position — we might include one configuration in the average every 5 MCS (*M*onte *C*arlo steps per *s*pin).

3.4 Illustration II

For definiteness, let us consider the simple system of a particle moving in two dimensions with an energy proportional to the displacement r from the origin. We represent the position $\mathbf{r}$ of the particle using cartesian coordinates x and y.

We might define our Markov operation with the following algorithm:

i: Choose trial step Δx and Δy to a 'nearby' position $\mathbf{r}'$

ii: Calculate the energy change ΔE to move from $\mathbf{r}$ to $\mathbf{r}'$

iii: Calculate the probability of change $p = e^{-\beta\Delta E}$

iv: Draw a random number z on $[0,1)$

v: If $z < p$ then accept move : $\mathbf{r} \to \mathbf{r}'$

(If the energy cost is negative, then our 'probability' p is greater than one — we could explicitly limit it to one, but the implementation of an acceptance with a given probability will always succeed for $p > 1$.)

It is important to emphasise that if we fail to accept a new state, we must include the unchanged state in our averages — if we failed to do this in a simple two-state system, for example, then the averages would *always* show equal probabilities for the two states, regardless of the energy levels — it would be the time spent in the lower energy state before changing to the excited state that contains the Boltzmann distribution, as the algorithm presented would *always* immediately step down (unless $\tau < 1$ in equation 3.18).

There are two parts making up M_{ij} — the probability that a system in state i will attempt to move to state j, and the probability that such an attempt is successful (equation 3.18). Generally, all the nearby states are made equally likely to be chosen, but care must be taken to avoid bias.

For this system, polar coordinates would have seemed a more natural choice — indeed, only one parameter would then be required to describe the system state. However, the density of states is proportional to r in polar coordinates, making it more difficult to sample space without biasing towards positions closer to the origin. It is possible to compensate for this by including an extra factor $g(\mathbf{r}')/g(\mathbf{r})$ in the Boltzmann update probability, for a trial step from $\mathbf{r}$ to $\mathbf{r}'$. In this way, attempts to move towards smaller values of r are discouraged, to compensate for there not *automatically* being more attempts to move to larger values of r as the density of states requires. With cartesian coordinates, however, the problem is simplified as the density of states $g(\mathbf{r})$ is uniform.

Similar care is required in simulations of magnetic systems with three-component spins, represented using spherical polar coordinates.

The choice of a trial step depends, of course, on the details of the system being simulated. In an Ising system, for example, a trial step might be the flip of a single spin, or it might be the *exchange* of two spin states — the latter conserves the magnetisation; it can be generalised in a continuous system to moving two coordinates, one in the opposite direction from the other, in order to meet a constraint of keeping some physical quantity fixed. In a Heisenberg (continuous) spin system (Hamiltonian 1.14), a trial step might be a small change in the direction of one spin, while for moving particles, as in this example system, we make a small displacement of each particle each step.

In a system with many state variables, we can either make trial changes to each coordinate in some predetermined order (for example, we might proceed 'typewriter fashion' through a lattice, from left to right, top to bottom), or choose both which coordinate to change and its new value at random.

The choice of what makes a state 'nearby' is usually dynamic — if we try to jump to states with a very large energy cost, these attempts are usually rejected, and although we will calculate the correct equilibrium quantities, it may take a long time. Similarly, if we make only very small jumps, they will nearly always be accepted, but it will take a long time to move any appreciable distance in phase space, and again, it may take a long time to reach equilibrium. A typical heuristic approach is to choose a step size such that the number of accepted trials at roughly half — this may be done dynamically, using feedback to manipulate the step size as the simulation proceeds.

10^4 update trials were performed on the example system, with different (fixed) step sizes. The table in figure 3.2(a) summarises the results, showing the step size (the side of the square centred at $\mathbf{r}$ from which $\mathbf{r}'$ is chosen uniformly at random), the number of accepted moves, χ_r^2, the sum of the squares of the deviations from the expected probability distribution $r \exp(-r)$ after histogramming the values of r in bins of width 0.1, and χ_θ^2, a measure of how anisotropic the distribution of sampled points is.

Scatter graphs for the sampled phase space points are shown in figure 3.2 for the step sizes indicated in the table — it can be seen that for very small steps, there are many points, but they are not statistically independent, and they do not cover phase space isotropically. For very large steps, the points are certainly statistically independent, but there are not enough of them to give good statistics. For a 'reasonable' step size, the results appear to be correct.

step	accept	χ^2_r	χ^2_θ	Fig
0.1	9879	0.95	1.80	
0.3	9656	0.10	0.36	b
0.5	9423	0.04	0.11	
5.0	5099	0.03	0.04	c
10	2744	0.08	0.09	
15	1531	0.18	0.18	
25	600	0.56	0.42	d

(a) Table of step sizes

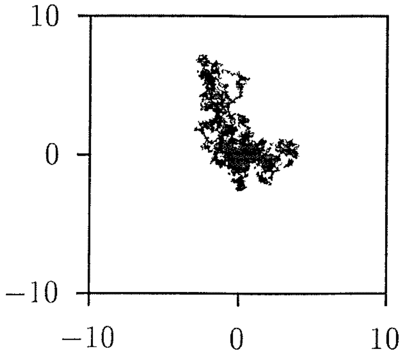

(b) step=0.3

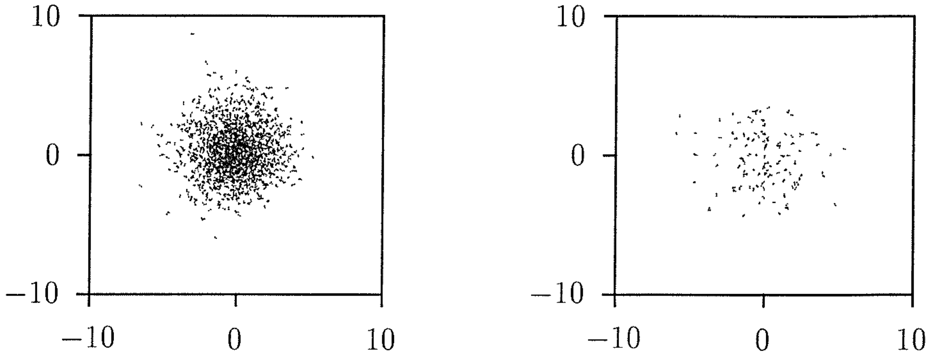

(c) step=5

(d) step=25

Figure 3.2: Results (a) and scatter graphs (b,c,d) for the simple system described in the text. For a small step size (b), the system does not progress far through phase space, and the results are anisotropic. For large step size (d), samples are isotropic, but there are too few events for good statistics. A reasonable step size (c) such that roughly half the trials are accepted gives good results.

3.5 What We Can Measure

Quantities such as the energy, magnetisation, pressure and density can be calculated as a simple average of the Boltzmann-weighted states generated by the Metropolis algorithm, by equation 3.11. It must be noted, however, that a simple (vector) average of the magnetisation of a spin system with continuous symmetry is not a useful commodity to record; the instantaneous magnetisation can undergo global rotations with no energy cost, and in a sufficiently long simulation, the magnetisation will average to zero. It is customary to record the average of the magnitude of the magnetisation, or the root mean square (RMS) magnetisation.

3.5.1 Correlation Functions

Both spatial and temporal correlation functions can be calculated. In the literature, spatial correlations are generally recorded using the form

$$g(\mathbf{r}) = \langle m(\mathbf{r}' + \mathbf{r}) m(\mathbf{r}') \rangle - \langle m(\mathbf{r}' + \mathbf{r}) \rangle \langle m(\mathbf{r}) \rangle \tag{3.20}$$

where m is the local order parameter and the average is over all $\mathbf{r}'$.

Here the second term removes any 'accidental' correlation — for example, in a magnetic system in a uniform external field, the spins will all be aligned with the external field, and hence with each other, but we might regard this correlation as accidental, and $g(\mathbf{r})$ above will reflect only *direct* correlations (where fluctuations in one spin are reflected in the other). In a ferromagnet in zero field, the correlation function decays rapidly at high temperature, but the correlation length (that distance over which $g(r)$ decays significantly) grows as the temperature is lowered (fluctuations increase, but the second term remains zero above the critical temperature). The correlation length diverges at the transition ($g(r) \sim r^{-(d-2+\eta)}$), then decreases as the second term in equation 3.20 increases below the critical temperature — here, equation 3.20 records the correlations in the response to fluctuations. (The XY model in two dimensions is critical at all temperatures below T_{KT}, so that $g(r) \sim r^{-(d-2+\eta(T))}$).

If it is the *nature* of the long-range order in zero external field that is under investigation (as in this work), then it is more useful to record the 'raw' correlation function

$$g(\mathbf{r}) = \langle m(\mathbf{r}' + \mathbf{r}) m(\mathbf{r}') \rangle . \tag{3.21}$$

Time-separated measurements

$$T(t) = \left\langle \overline{m(\tau)m(\tau+t)} \right\rangle - \left\langle \overline{m(\tau+t)m(\tau)} \right\rangle \tag{3.22}$$

are used in the study of spin glasses (section 2.6) or other dynamic systems to measure decay times. Here, the bar denotes a spatial average, while the angled brackets denote a thermal average, over different instantaneous configurations. If abundant storage is available, the combined measurement

$$g(\mathrm{r}, t) = \left\langle \overline{m(\mathrm{r}' + \mathrm{r}, \tau + t)m(\mathrm{r}', \tau)} \right\rangle - \left\langle \overline{m(\mathrm{r}' + \mathrm{r}, \tau + t)m(\mathrm{r}', \tau)} \right\rangle \tag{3.23}$$

might be made.

Because these calculations involve comparing each particle with every other particle, the calculation costs can become excessive, and such 'expensive' measurements are generally made less frequently than simple ones requiring nearest-neighbour comparisons only — one way of reducing the cost of the spatial correlation calculation is by taking advantage of *Fast Fourier Transforms* to map the expression into a product in Fourier space, requiring only a time proportional to $N \log N$, rather than the N^2 required for a direct calculation. If there is abundant on-line storage available, configurations can be stored for later analysis, on different hardware such as a vector processor, which is less suited for the Monte Carlo work in general.

3.5.2 Fluctuation-dissipation

The *fluctuation-dissipation theorem* can be used to calculate the susceptibilites of the system's state, such as magnetic susceptibility χ or specific heat C:

$$\begin{aligned}
Z &= \sum_i \exp(-\beta U_i) && (3.24)\\
\frac{\partial}{\partial\beta}\ln Z &= \frac{1}{Z}\frac{\partial Z}{\partial\beta} && (3.25)\\
&= -\frac{1}{Z}\sum_i U_i \exp(-\beta U_i) && (3.26)\\
&= -\overline{U} && (3.27)\\
\frac{\partial^2}{\partial\beta^2}\ln Z &= \frac{1}{Z}\frac{\partial^2 Z}{\partial\beta^2} - \frac{1}{Z^2}\left(\frac{\partial Z}{\partial\beta}\right)^2 = \left\langle U^2 \right\rangle - \left\langle U \right\rangle^2 && (3.28)\\
\frac{\partial}{\partial\beta} &= -kT^2\frac{\partial}{\partial T} && (3.29)
\end{aligned}$$

$$\frac{\partial}{\partial\beta}\frac{\partial}{\partial\beta}\ln Z = kT^2\frac{\partial}{\partial T}\langle U\rangle = kT^2 C \tag{3.30}$$

$$C = \frac{1}{kT^2}\left(\left\langle U^2\right\rangle - \langle U\rangle^2\right). \tag{3.31}$$

Similarly,

$$\chi = \frac{1}{kT}\left(\left\langle M^2\right\rangle - \langle M\rangle^2\right). \tag{3.32}$$

Self-averaging Quantities

Measurements on quantities such as the internal energy become very accurate in the thermodynamic limit — fluctuations are of the order $\sqrt{N}$, which become insignificant for large N. However, simulating larger systems necessarily requires more computing effort — the same total computer time can produce fewer statistically independent measurements. Quantities such as energy, magnetisation, etc. can be shown to be *strongly self-averaging*; the errors in the measurements are not increased if the system size is increased at a cost of number of independent measurements.

Measurements of specific heat and susceptibility by the fluctuation-dissipation theorem, however, are themselves based on fluctuations, and do not show self-averaging. The errors in the measurements are not driven down simply by increasing the system size, so that it is better to run longer simulations on smaller systems than to increase the system size.

It is necessary to take into account critical slowing down near a phase transition (equation 3.19), since a larger system requires that more steps are required to ensure statistical independence of the measurements.

3.6 Finite size effects

3.6.1 Boundary Conditions

If we are interested in the bulk properties of some system of interacting particles, we must be aware of surface effects: in a 'real' system, with typically 10^{23} particles, surface effects are (usually) small, whereas a three dimensional cubic lattice of side (optimistically) 100 units has 6×10^4 (6%) of the 10^6 sites on the surface — the problem is even worse for the smaller systems which can at present be handled in practice. One possibility is to accept this problem, with particles on the surface having fewer interactions that those in the bulk — this is termed *free boundary conditions.* One way of reducing the effect is to discount the particles at the surface, but computing effort is then wasted on what may be a large number of sites which contribute nothing to the results. This waste can be minimised by updating the surface sites less often, or even not at all (giving *fixed boundary conditions*).

A more common criterion to apply is *periodic boundary conditions*, whereby the system is replicated, and 'tiles' an infinite space — whenever we move out of the system, we also move out of one of the duplicates back *into* the original system, as indicated in figure 3.3. For a cubic lattice system, for example, this simply means that the opposite edges are connected to each other, or a one-dimensional Ising system is closed to form a ring. It should be noted, however, that in the latter case, the choice of periodic boundary conditions precludes the presence of an odd number of 'kinks' (section 2.1.1), and hence interferes with the correct statistical results being obtained in this case. Similarly, periodic boundary conditions in a two-dimensional XY model (in which the spins are formally on the surface of a torus) prohibit the formation of a single vortex. In cases such as this, the system is said to be *incommensurate* with the lattice.

Anti-periodic boundary conditions, where opposite edges interact with the 'opposite' sense from that of the bulk (such as an antiferromagnetic interaction across the boundaries in a ferromagnet), are usually reserved for studying special cases: for example, the dynamics of an isolated vortex may be studied in an XY system with antiperiodic boundary conditions.

3.6.2 Effect of Finite Size on Critical Behaviour

In the following, periodic boundary conditions are assumed. Provided that all the length scales within the system are much smaller than the system size, the behaviour of 'real' systems can reasonably be inferred from the simulation results. As the critical temperature is approached, however, and the correlation length grows to the system size, the fluctuations described in section 1.1.3 can cause critical behaviour above or below the critical temperature: at temperatures slightly below T_c, the fluctuations over a finite correlation length are nevertheless able to propagate through the entire system, so that the order parameter averages to

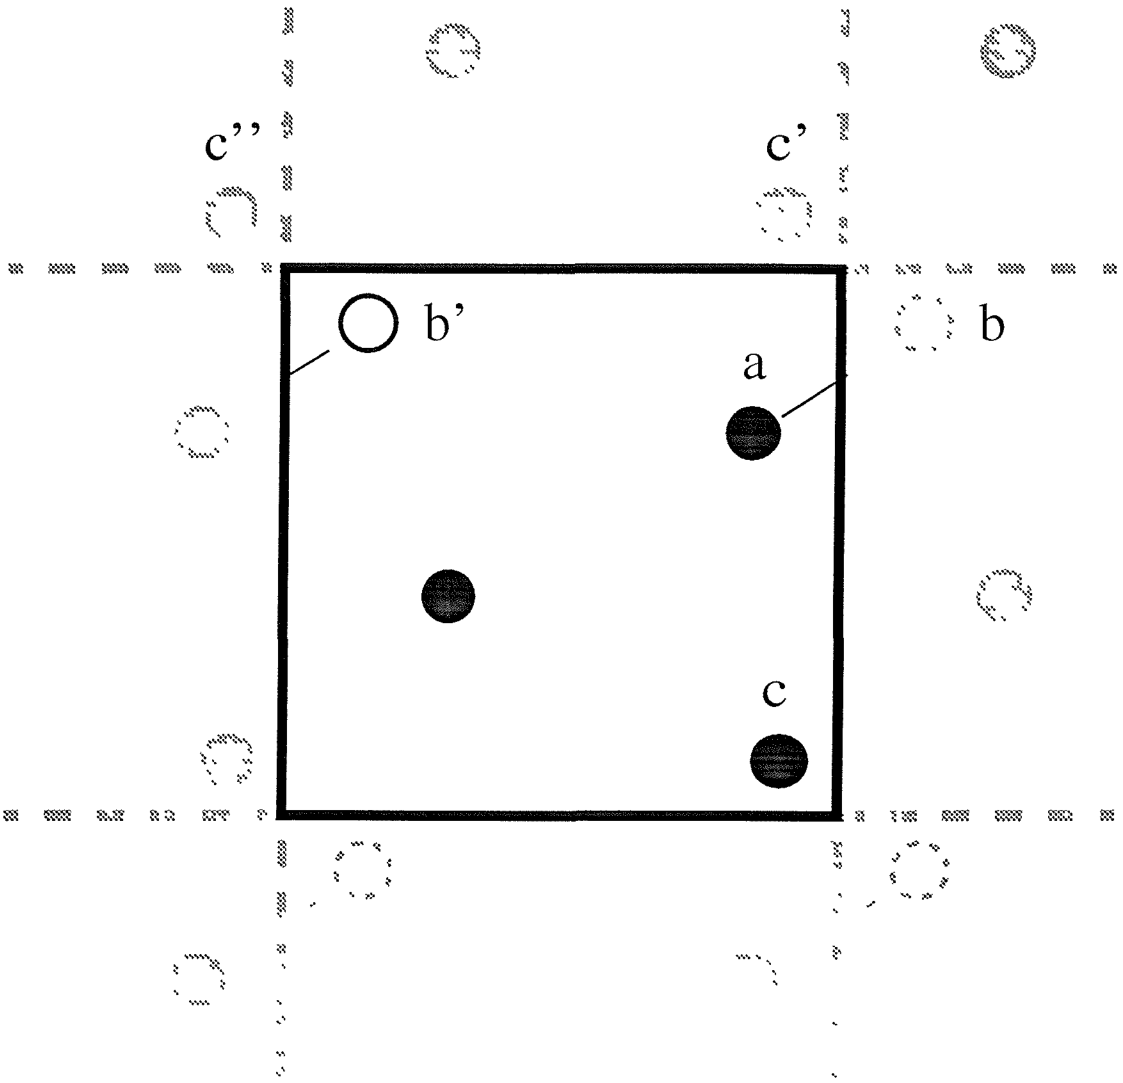

Figure 3.3: Periodic Boundary Conditions. Particle a makes a trial move to b, outside the system, and is mapped back to b'. Notice that in calculating the old energy (particle at position a), the image c' is used, and in calculating the energy at b, the image c'' is used instead of c.

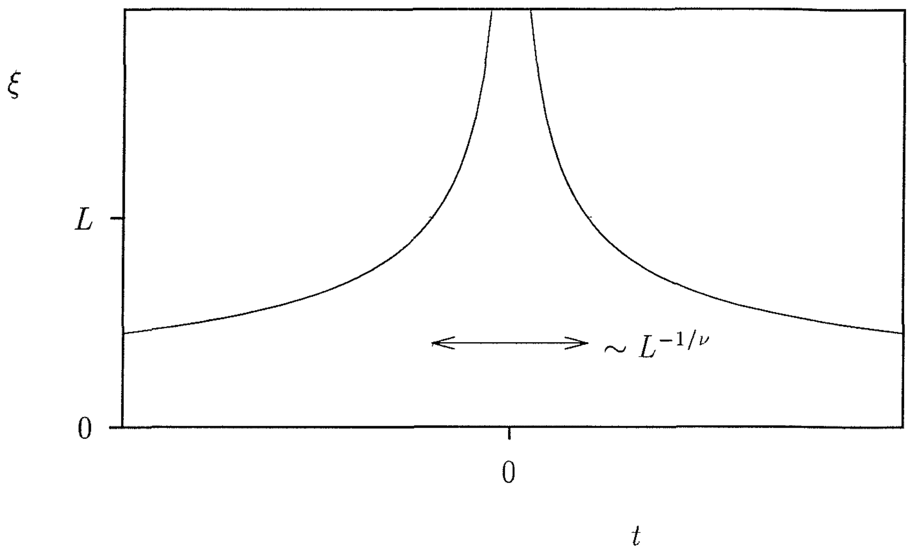

Figure 3.4: Illustration of the effect of finite size on critical behaviour

zero; above the critical temperature, the 'pockets' of order described previously fill the finite system, so that instantaneous measurements imply long range order (which changes direction with time, giving a zero for algebraic average of order parameter; the root-mean-square of the order parameter shows an anomalous non-zero value).

A finite size scaling theory [55] amounts to using the system size L as a variable in section 2.4.1, giving an energy scaling equation of the form (equation 2.38)

$$\Phi(t, L) = L^{-d}\Phi(tL^{1/\nu}). \tag{3.33}$$

Using $\xi \sim |t|^{-\nu}$, $tL^{1/\nu}$ can be reinterpreted as a function of L/ξ, so that the problem can be considered in light of the relative values of L and ξ. L is irrelevant if it is much larger than ξ, but near the critical temperature, it becomes important.

The problem is illustrated schematically in figure 3.4. Since ξ diverges as $|t|^{-\nu}$ near $t = 0$, the range of temperatures over which the correlation length exceeds the system size can be calculated, as shown on the figure. We might expect any critical behaviour to be 'smeared out' across this range of temperatures.

In practice, the *critical amplitudes* need not be the same above and below T_c, though scaling suggests that ν should be. This results in an additional shift of the centre of the range of critical temperature, again proportional to $L^{-\nu}$. This gives rise to specific heat curves of the form of figure 3.5 as a signature of finite size effects in Monte Carlo. By plotting the data on scaled x and y axes, the

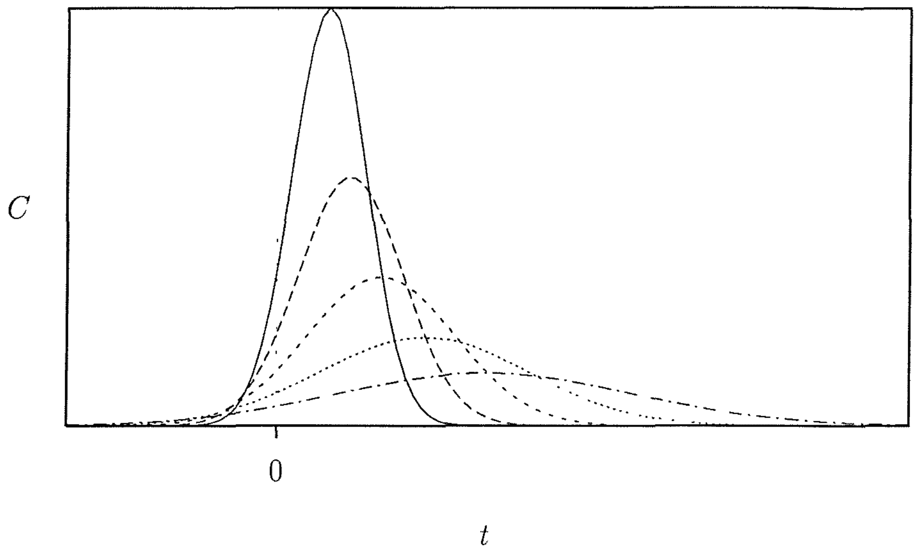

Figure 3.5: Typical effect of finite size on specific heat.

curves can be superposed, as will be shown in section 6.1.

3.7 Monte Carlo Techniques

Various techniques have evolved to improve the efficiency and accuracy of Monte Carlo simulations.

3.7.1 Multispin and Cluster Updates

Because the correlation length diverges near a critical temperature, excitations appear in clusters of increasing size. (For example, the excitations in the Ising system discussed in section 2.1.1 form misaligned domains with energy typically ξ^{d-1} in d dimensions — this energy can be equated to the thermal energy k_BT at temperature T to a first approximation.) It becomes increasingly difficult, therefore, to move around in phase space by making trial changes to only one spin at a time; while the flip of a spin may be valid thermally, its acceptance is unable to affect the organisation of the domains. For this reason, multi-spin-flips are frequently implemented, often in combination with single-spin-flip trials; as the name implies, attempts are made to move several spins together, so that the change is more likely to have a long-reaching effect before the move is undone at a later trial.

Taking this idea further, whole clusters can be updated in one group [56, 57]; the size of the domains is not fixed, but is chosen by the system itself, giving a direct measure of the correlation length and therefore of the magnetic susceptibility. The treatment by [56] is more intuitive, whereby the entire system is broken into domains, and one is chosen for a flip attempt, but [57] is more easily implemented; in the latter formulation, one site is chosen as a seed - it will definitely be updated. Starting from this site, the system explores outwards, offering to flip each adjacent spin. Accepting the flip at the neighbouring site is the default, since the update is chosen such that all the energy within the cluster is preserved; the only change is at the interface to adjoining domains. A spin can elect to reject the flip attempt, remaining outside the cluster, using the usual Metropolis acceptance criterion (equation 3.18); the spin will remain outside the cluster, in its old state, if that reduces the total system energy (given that the neighbours within the cluster *are* going to flip), or with a Boltzmann probability if the energy would be increased. This will be discussed further in section 4.4.5.

3.7.2 Histogram Method

Conventionally, it is necessary to run simulations at a large number of separate temperatures near a critical temperature in order to build up a sufficiently dense graph of discrete energies and specific heats in the vicinity of a phase transition. The histogram method [58] makes use of additional information collected from a simulation at *one* temperature, and allows all the important commodities (energy, specific heat, magnetisation, etc.) to be extrapolated continuously to any nearby temperature and/or magnetic field, etc. (though to collect data allowing extrapolation in more than one variable at once requires an enormous amount of storage).

In most simulations, only the average of the energy, and its standard deviation (to allow specific heat calculation by equation 3.31) are stored. But if the full distribution is recorded (in principle by also calculating higher moments, but in practice by binning the measurements, to form a histogram), the density of states in the vicinity of the sampled phase space can be determined (to within a constant factor), and once the density of states is known, everything else can be calculated (section 3.1).

The simulation generates a stream of measurements $\{E_i\}$ of the energy of the system. Assuming the system has been properly equilibrated, etc., these measurements are samples of the distribution $p(\beta, E) = g(E)\exp(-\beta E)$ (ignoring normalisation of p), and since we know the (inverse) temperature β, we can infer $g(E)$, and so the distribution of energies $p(\beta', E)$ at a nearby temperature β':

$$\begin{aligned} g(E) &= p(\beta, E)\exp(\beta E) && (3.34) \\ &= p(\beta', E)\exp(\beta' E) && (3.35) \\ p(\beta', E) &= p(\beta, E)\exp(\beta - \beta')E. && (3.36) \end{aligned}$$

Thus we can transform our sampled $p(\beta, E)$ to a new temperature by simply multiplying the measured (discrete) data by an exponential, without any need to first fit the measured data to a curve — we merely need to normalise when calculating any averages using $p(\beta', E)$.

$p(E)$ is likely to be approximately Gaussian, and multiplying a Gaussian by an exponential simply shifts that Gaussian to a new mean (again ignoring the normalisation) by boosting the values at one side of the original mean more than the other side, as illustrated on figure 3.6. (Of course, it is precisely the information about the *deviation* from Gaussian-ness that is useful here.) However, our measured $p(E)$ is discrete, and will necessarily be noisy at the tails (due to quantisation errors in the binning) — if we simulate for longer, we get more samples here, but end up with a few samples further out. This means that we can extrapolate (shift) the probability distribution only within the 'bell' of the distribution (the region with a large number of binned samples); if we try to extrapolate too far, we end up trying to infer system parameters from only a handful of events, since as is clear from the figure, the data at the tail is given an enormous boost by the exponential function; the situation depicted is slightly exaggerated, and is at the extreme of the range of reliable extrapolation.

In order to extrapolate other quantities, such as magnetisation, with temperature, it is sufficient to maintain a sum of the measurements of that quantity with each energy bin, and compute a weighted average with the shifted energy at each new temperature.

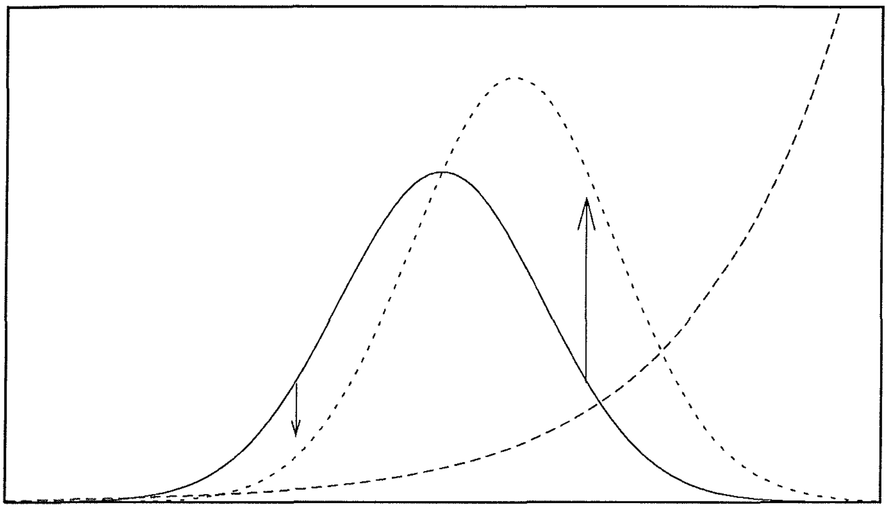

Figure 3.6: Illustration of the way in which multiplying a Gaussian by an exponential results in a shifted Gaussian. Real experimental data will be only approximately Gaussian, and the technique will use all the available information to track the change in temperature accurately. (Assuming a Gaussian would result in a shift proportional to temperature.) Notice in particular that a large shift (rapidly-growing exponential) will boost the tail of the distribution, where statistical accuracy is poor — this illustration depicts about the largest shift that can be made reliably.

Chapter 4

Implementation Details

4.1 Parallel Computers

Although Monte Carlo simulation dramatically reduces the computation required to perform measurements on statistical systems, enormous amounts of CPU time are still required on all but the biggest supercomputers. One approach that hardware manufacturers have taken to try to meet the demand is the production of parallel computers, whereby many processing units cooperate in a computation.

The University of Southampton has made a significant commitment to parallel computers, both in terms of hardware (having several transputer-based parallel computers, and one of the first of the second-generation Meiko Computing Surfaces in the world) and personnel (having set up a Parallel Computing Support Group, now called the High Performance Computing group, to assist users in writing code to make the best use of the available resources). In addition, the University has a close association with the Parallel Application Centre at Chilworth.

At present, a programmer must explicitly distribute the computation over the network of available processors, and must work hard to ensure that the system is *load-balanced* (that is, each processor has a fair share of the workload), though software tools will almost certainly take over this tedious task in the future. There are a number of ways in which the code can be broken up:

- Run the sequential program independently on each processor, each copy operating on a different set of parameters. This is actually the most efficient way of using parallel computers, and requires least programming effort, provided each processor has sufficient resources to run the whole program in a reasonable time. For Monte Carlo simulations, this corresponds to running the simulation for several different temperatures, say, on several workstations.

- Operate several programs sequentially, each applying small transformations to a stream of data. An analogy is a factory production line, where incomplete products pass down a conveyor belt, and each worker performs their simple task until a complete product emerges at the end. Once the first product has reached the end of the line, finished products emerge at the rate of the slowest step. In a Monte Carlo simulation, one processor might feed out trial movements, another might calculate the energy cost of each step, and so on. This *algorithmic parallelisation* was studied in [59], and is useful for an inhomogeneous network of computers, but it can be difficult to balance the loads effectively — it will be impossible in general the break up the work into pieces such that each processor is kept busy at all times; there will inevitably be delays where processors are waiting for results from another which has more to do, and the whole calculation can proceed only at the rate of the slowest step.

- Run the same program on each processor, each copy operating on a subset of the entire system, with programs (explicitly) exchanging information across the boundaries. This is ideal for calculations on regular lattices, where each program might control a subset of the lattice sites, and is sent updates from neighbouring processors when changes are made at the common edges. *Geometric parallelisation* was chosen for this work.

4.1.1 Communicating Sequential Processes

All these strategies are encompassed under the *communicating sequential processes* [60] model of a program — a program is viewed as a number of sequential units, operating concurrently with the only time constraint that any pair are synchronised when they require to exchange information — even a conventional program may be viewed in this way, either trivially as a collection of one sequential process, or as a number of processes, the beginning of each being synchronised with the end of the previous one (as each step operates on data received from the previous step). Any step which does *not* require to wait for data from the previous step can be run concurrently with that previous step, and in this way, a program can gradually be parallelised (though it is usually better to begin with a parallel algorithm from the outset). Even on sequential computers, cpu-intensive activities can often be run concurrently with file-accesses or floating-point calculations, (section 4.4), as is usually done under multitasking operating systems, and modern RISC processors such as the DEC Alpha can issue several non-interacting instructions on each clock cycle ('superscalar' architecture) — compilers for such machines will attempt to reorder the instructions to maximise this overlap where possible.

4.1.2 Transputers

Although the above techniques (particularly the first) can be applied to any network of workstations, for example, hardware manufacturers have designed processors optimised for parallel computation. Inmos' T414 transputer implements a 10 MIP (10 million instructions per second) CPU with four dedicated high-speed links, allowing large numbers of these processors to be coupled to form large computer networks. The T800 transputer adds a 1 M-flop floating point unit, and arrays of these processors make formidable computing machines.

The links transfer data by DMA (direct memory access), so that the CPU can continue to do useful work on other processes while data is transferred, intervening only to initialise and finalise the transfer — there is a small impact on processing speed in practice, since the CPU has to share access to the memory.

Although it can be inconvenient that processors can communicate only with those others which are physically coupled, this can be overcome by either adding a software "harness" to route messages through intervening processors, or with extra hardware to change the link connections dynamically. The next generation of transputers is to have on-chip message routing, so that messages are transparently routed through each transputer without assistance from the CPU. An alternative scheme, where any processor can communicate directly with any other along a bus structure, has the disadvantage that the communications bandwidth is limited, and can saturate, negating any advantage in adding further processors; the bandwidth available in a point-to-point communications network scales as processors are added.

In this work, the actual simulation code requires communications with only neighbouring processors (since the system being simulated has nearest-neighbour interactions only), so that the code suffers little delay due to communication overheads in practice. The transputer is able to proceed with other (independent) calculations while data is transferred, though in practice we were not able to take advantage of this except for simulations on very large systems (section 4.4).

4.2 Hardware Used

The group acquired a Meiko Computing Surface, containing a T414 'master' transputer with 3MB memory, a T800 transputer with hardware for full-colour high-resolution graphics (which proved very useful for monitoring the progress of simulations) and 32 'worker' T800 transputers, each with (a rather limiting) 256kB memory.

Like many other transputer arrays, this system has neither operating system nor filestore — instead, the system is 'hosted' by another machine, which is used for developing and debugging the code, 'booting' the code onto the system and providing access to the host's filesystem. The Meiko came with hardware suitable for hosting the system from a IBM PC compatible, which is, unfortunately, not the most suitable host:

- The PC, in general, does not have a large filespace, severely limiting the data that can be stored from a run — correlation data, in particular, places a significant strain on the resources.
- The host is required to be continuously available during a set of simulations, and because a PC cannot, in general, multitask, it is locked up for the duration of the simulation. Software such as MS-Windows, which does allow multitasking in a limited way, is not sufficiently stable to be expected to run for the duration of a simulation without crashing.

When the group purchased a VAX cluster, the former problem was solved by networking the PC to the VAX, allowing the VAX filespace to be accessed transparently by the PC, and hence by the Meiko. Means of "borrowing" the PC from the Meiko were devised, so that the PC could be used for useful work during the day, but it proved quite frustrating when the connection could not be re-established securely, and data was consequently lost.

The problem was finally resolved when the VAX joined the department's ethernet network — a PC dedicated as a LAN bridge was required both for security and to keep the volume of network traffic under control, and it was realised that the Meiko could be *physically* hosted from this PC, and a simple software fix would allow communications between the Meiko and VAX along ethernet, allowing the VAX to be the *logical* host. Because the VAX is a multitasking system, it can serve the Meiko with virtually no impact on its main functions, but there are several additional benefits, such as access to the excellent queueing facilities available with VMS, and now that the VAXcluster has joined the Internet, the Meiko is effectively accessible from anywhere in the world (subject to the usual security checks, of course !).

One slight complication is that the VAX and transputer use different representations for floating point numbers; the communication protocol was *defined* to use IEEE format numbers, so the VAX had to translate from its internal representation to IEEE; fortunately, this can be done using simple bit manipulation.

4.3 Algorithm

We have used a standard Metropolis single-spin update algorithm — because we are using a parallel computer, several spins are in practice updated simultaneously, but such spins have no direct interactions, so that the results are exactly equivalent to a sequential computer updating the first spin visited by the first processor, then the first spin visited by the second processor, and so on through all the processors and all the spins.

There has been some criticism [45] that the single-spin Monte Carlo method does not adequately sample the phase space of the system, and hence does not reach true thermodynamic equilibrium. However, because it is believed that it is the inability of experimental spin glass systems to explore the whole of the available phase space (section 2.6) that leads to interesting behaviour, we take the view that if our simulations do not reach equilibrium, this may reflect analogous behaviour in real systems.

We rewrite Hamiltonian 1.16 in a simplified form, absorbing spin magnitudes into J, and representing spin direction at site i using a single parameter θ_i. The anisotropy $\hat{\mathrm{n}}_i$ at site i is represented instead as angle ϕ_i, the field direction is θ_{field}

and primed quantities are in units of J — primes will be dropped for all of the following discussion:

$$\begin{aligned} \mathcal{H} &= -J\sum_{\langle i,j\rangle}\mathbf{S}_i.\mathbf{S}_j - D\sum_i(\mathbf{S}_i.\hat{\mathrm{n}}_i)^2 - \mathbf{H}.\sum_i\mathbf{S}_i & (4.1)\\ \mathcal{H}' &= -\sum_{\langle i,j\rangle}\cos(\theta_i-\theta_j) - D'\sum_i\cos(\theta_i-\phi_i)^2 - H'\sum_i\cos(\theta_i-\theta_{\text{field}}). & (4.2)\end{aligned}$$

The temperature T is expressed in units of J/k_B, so that the Metropolis update criteria is a simple test based on the ratio of ΔE to T.

Each update attempts to rotate a single spin through either 180° (a spin 'flip' — useful for the large anisotropy, Ising limit (section 2.7)) — or through some random angle, whose maximum magnitude is dynamically adjusted to keep the fraction of accepted trials at approximately half (section 3.4).

A *cycle* is the fundamental unit of the simulation, and consists of a number of sweeps through the lattice, attempting a spin rotation at each site, followed by a number of sweeps attempting a spin flip at each site. Measurements of system state are made after every cycle during the simulation (section 3.3.1), after delays of typically several thousand cycles for equilibration (section 3.3.1). Cycles in this work were usually composed of 5 sweeps, mostly rotation sweeps for small anisotropy, and typically 3 rotation and 2 flip sweeps for larger values of D.

In order to compare results with previous simulations performed in the large D limit (section 2.11), we can simulate systems with infinite D using a cycle composed purely of flip sweeps, initially aligning each spin along the local anisotropy axis. The parameter D is irrelevant in this case, since it is the update algorithm which constrains the spins to lie along the anisotropy axes, but it is useful to set $D = 0$ since the energy is still measured using equation 2.50. As described in section 2.7, for simulations at large D, a Hamiltonian

$$-J\sum_{\langle ij\rangle}S_i.S_j - D\sum_i\left((S_i.\hat{n}_i)^2-1\right) \tag{4.3}$$

is more appropriate, since $\langle U\rangle$ remains finite as $D\to\infty$ — it is necessary simply to add an energy D per spin to measurements made on equation 2.50 to allow comparison with results at infinite D as described above.

One other correction is necessary: for a finite anisotropy strength, at non-zero temperature, each spin is free to oscillate about its local energy minimum, contributing one degree of freedom, and, by equipartition of energy, an energy $\frac{1}{2}k_BT$ per spin. It is useful to add 0.5 to the specific heat measurements made at infinite anisotropy to allow comparison with the low temperature results.

4.4 Implementation details

Because no one transputer has sufficient memory or computational power to run the complete simulation for sufficiently large systems, it was necessary to distribute the program over all the transputers, and choose a geometric parallelisation algorithm (section 4.1), in which each transputer has full control over a subset of lattice points, and holds a copy of the adjacent spins at all times; the transputers send messages to each other to update neighbours' copies of the shared spins. Periodic boundary conditions were implemented trivially by configuring the transputer links to form a ring, as shown in figure 4.1, so that nothing was required of the software.

In order to ensure that no attempt was made to update adjacent spins on different transputers, each transputer was required to control at least two rows. The sites were visited in 'typewriter' fashion, sequentially along rows from top to bottom, in lock-step on all transputers; in this way, transputers did not have to send out individual requests for information, but simply transmitted the updated edges to the neighbouring transputer, at the same time receiving the opposite neighbour's updates, after completing each edge — because processes on a transputer are synchronised by such a data exchange, one transputer could not process a shared edge while its neighbour was still processing it. (With finer synchronisation, the program could have been written to allow one transputer to begin on a row before its neighbour had finished, but no time would be saved overall, since the delay would propagate around the ring, and that transputer would have to wait for longer at the beginning of the next sweep. In addition, the extra synchronisation requires extra communication, which all introduces more delays.)

The smallest (square) system for which all transputers are used has 64^2 lattice sites, though this was reduced to 62^2 after one of the transputers developed a hardware fault. For larger systems, when each transputer controls more than two rows, the intermediate rows can be processed while the data exchange is proceeding, taking advantage of the background communications available on the transputer (section 4.1.2).

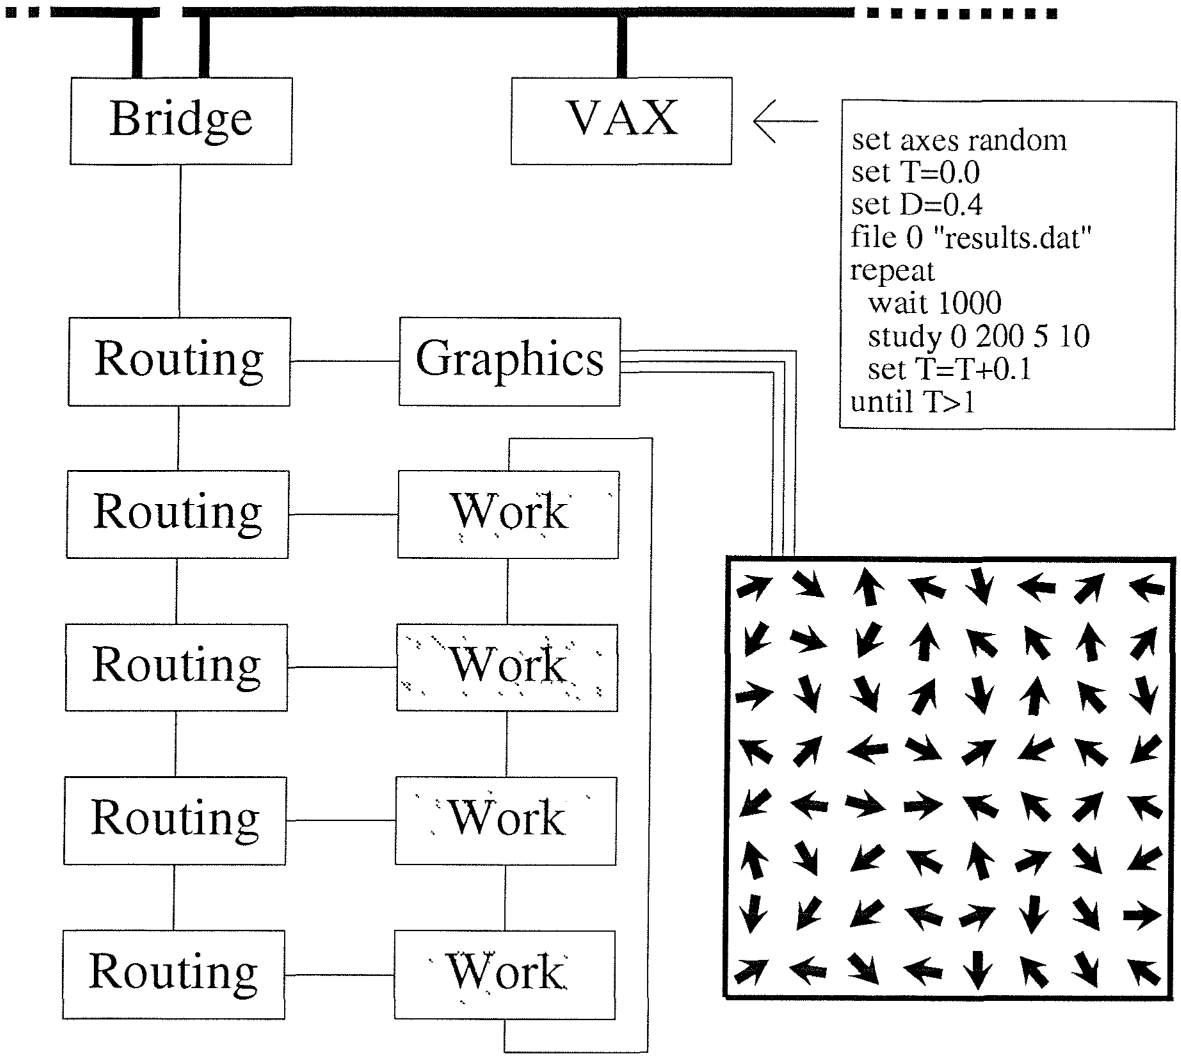

```
set axes random
set T=0.0
set D=0.4
file 0 "results.dat"
repeat
  wait 1000
  study 0 200 5 10
  set T=T+0.1
until T>1
```

Figure 4.1: The hardware layout

For smaller systems, a subset of the transputers was used — an alternative scheme, used in the code for the three-dimensional system, is to decompose the lattice into two sublattices, such that interacting neighbours are on different sublattices — for example, the squares of a two-dimensional system can be coloured black and white, like a chessboard. Now no two lattice sites of the same colour interact directly with each other, so trial moves could be attempted at all the white squares in the system in parallel, and each transputer can operate on a single row of the system — an update sweep now requires two passes, once operating on only the white squares, then, after exchanging data with the neighbouring transputers, all the black squares, before exchanging information again. This technique is also used on vector processors, where full speed operation requires that all the information required for a large number of update attempts be available at one time.

4.4.1 Driving the Simulations

At first, the program was written entirely in Occam (the language invented alongside the transputer, to allow easy implementation of the parallelisation and communication primitives available on the transputer), with the simulation parameters such as temperature, equilibration time, and the number of Monte Carlo sweeps between measurements coded as constants, and the changes to the various parameters during a set of runs hard-coded. For conventional programs, it is straightforward to make changes to such a simulation, to make the system sweep the magnetic field, for example, to measure hysteresis, or to sweep the temperature, to compare quenched and annealed results. For parallel programs, however, there is an added complication: the system requires extra processes to run in the background on each transputer, to route messages to the graphics transputer for display, and to and from the 'master' processor for averaging and filing. (Only the latter transputer can communicate with the host, and therefore access the filesystem.)

The code was therefore modified to make it 'event' driven - the simulation was broken down into a number of low-level primitive operations, and the workers acted on a stream of messages such as "increase the magnetic field", "run for x sweeps for equilibration", "measure the energy", so that only the co-ordinating process on the master transputer had to be recompiled to change the simulation — this is much closer to a conventional program, but rather than writing subroutines to implement the steps of the simulation, messages were sent to the workers, telling them what to do next.

Although the main loop for the simulations could easily be built from low-level primitives, it was realised that overhead in communicating a sequence of "wait n sweeps" and "measure" instructions made the system inefficient — too much time was wasted waiting for the next command. Thus one "primitive" invoked a simulation for a large number of cycles, with the workers sending out measurements at predefine intervals. (One can make analogies with CISC versus RISC CPU design issues.) Essentially, the program operates in two modes: slave mode, where parameters are set, and simulate mode (the 'study' command). Unusual simulations can still be built up from the other primitives, of course (section A.7).

When a compatible C compiler became available, it became possible to make the controlling program read in a description of the required simulation from a text file, and send the appropriate messages to the workers — a "simulation control language" was defined. (This would have been possible in Occam, but this sort of application is not one of its strengths; Occam's I/O is rather primitive in comparison to that of C.) In this way, the description of the simulation can be prepended to each results file, to minimise the risk of mixing up results when a large number of files are generated. The details of the control language are presented in appendix A.3. In addition, because the transputers were no longer required for recompiling the program for changes to the simulation, it became possible to prepare simulation descriptions in advance, reducing the turnaround time for subsequent simulations.

4.4.2 Quantisation of the Angles

At first, each θ_i was stored as floating point variable, thus stored as accurately as possible, but requiring expensive trigonometric function calls in the inner loop. Later, the angles were quantised into (usually) 1024 integral parts — the only constraint (given sufficient memory) is that it is a power of two, to allow the periodicity to be implemented using logical operators. Extensive tests were performed to find régimes where this change made any significant difference, but, even at very low temperature or very close to saturation, no quantitative change in behaviour could be detected.

By storing the angles as integers, all the computationally-expensive trigonometric and exponential functions can be precalculated and stored in tables, such that the update loop runs significantly faster — if we consider a single-spin trial step $\delta\theta_i$ from θ_i to θ'_i at site i, the change in system energy U is simply the change in the partial energy u_i, containing only terms containing θ_i:

$$\begin{aligned}
\exp(-\beta\,\delta U) &= \exp(-\beta\,\delta u_i) && (4.4)\\
&= \exp(-\beta u_i)/\exp(-\beta u'_i) && (4.5)\\
u_i &= -\sum_j \cos(\theta_i-\theta_j) - D\cos{(\theta_i-\phi_i)}^2 - H\cos(\theta_i-\theta_{\mathbf{H}}) && (4.6)\\
\exp(-\beta u_i) &= \prod_j \exp(\beta\cos(\theta_i-\theta_j)) \times \exp(\beta D\cos{(\theta_i-\phi_i)}^2)\\
&\qquad \times \exp\left(\beta H\cos\left(\theta_i-\theta_{\mathbf{H}}\right)\right) && (4.7)
\end{aligned}$$

and each of the terms in equation 4.7 is tabulated for all possible values of $(\theta_i-\theta_j)$. In fact, the tables run for two and a half cycles — the second cycle allows negative angles to be handled trivially by adding 2π before the lookup — $\theta_i-\theta_j$ can return values from -2π to 2π, so that $\theta_i-\theta_j+2\pi$ will fall within our 2-cycle lookup table. (In practice, a pointer is set halfway through the table, so that offsets from -2π to 2π can be fetched directly.) The extra half cycle allows the reciprocals of the exponentials to be calculated by adding an extra π before the lookup, for the exchange and field terms, since multiplication by the reciprocal is significantly faster than division in equation 4.5. The anisotropy term required a separate table for the inverses.

Thus, inside the main update loop, only multiplications are required, with trigonometry required only when the tables are updated after a change in any of the parameters T, D or $\mathbf{H}$ (outside the main loop). A tradeoff of speed against memory has been performed: in order to increase the speed in the innermost loop, we have had to set aside a great deal of memory for these lookup tables, limiting the maximum system size we can simulate.

One problem with this technique was that, at low temperatures and large field or anisotropy strength, the tables or intermediate products could overflow, even though the final energy change might be quite small — for example, suppose $T = 0.1$ and $D = 10$: a very small rotation from the local anisotropy axis might have $u_i = -10$ and $u'_i = -9.99$, so that the Boltzmann factor is $e^{0.1}$, but the intermediate term, e^{100} overflows the floating point unit.

This problem was finally overcome by storing the tabulated values in two parts: an integer power of e, and the exponential of the fractional part. For example, $e^{99.9} = e^{100-0.1} = e^{-0.1}e^{100}$, and would be tabulated as $(e^{-0.1}, 100)$ (or $(e^{0.9}, 99)$), and the intermediate product accumulates both an integer sum of the power, and a real product of the exponential part, which is guaranteed not to overflow. Thus the calculation returns the Metropolis acceptance probability M in two parts M_i and M_r, such that

$$M = M_r \times \exp(M_i).$$

While overflows are to be avoided at all costs, underflows are quietly truncated to zero, so it is necessary to modify slightly the Metropolis acceptance condition equation 3.18: for random number r between 0 and 1, we accept the update provided

$$\begin{aligned} r &\leq M_r \exp(M_i) \qquad M_i < 0 \qquad (4.8) \\ r \exp(-M_i) &\leq M_r \qquad\qquad M_i \geq 0 \qquad (4.9) \end{aligned}$$

Yet another table is used to lookup the exponentials of integers.

4.4.3 Further Optimisations

In order to take further advantage of the hardware concurrency, the innermost loop for the update sweep was coded in assembly — because the floating point unit and the integer unit can operate independently (synchronising only to exchange data), careful coding allows the integer unit to do useful work while a floating point operation proceeds, rather than idling while waiting for the result from the floating point unit. As described previously, the calculation was mostly table-driven, so while the floating point unit was performing a calculation, the integer unit could be calculating the offsets into the tables for the next interaction, preparing the next random number, implementing the periodic boundary conditions, and so on. Another advantage of coding in assembly is the greater control over which variables are stored in the CPU's registers for maximum efficiency. (In fact, the transputer doesn't use traditional CPU registers, but a stack — hand-assembly allows control over which variables are stored on the stack, though the instruction set doesn't have the full range of stack manipulation instructions that might have been expected.)

It should be noted that modern RISC processors are becoming increasingly complex, and superscalar processors are able to issue multiple instructions per clock cycle, as mentioned above. Without enormous care, it is possible that recoding a routine in assembly will result in code running more *slowly*, because compilers are much better able to reorder the instructions to achieve maximum throughput — even if hand-assembled code uses fewer instructions, it may still run slower than compiler output, since compiler may be able to make better use of the double-issue of instructions.

With all the above optimisations implemented, and the further refinement of careful ordering of the source code so that the code implementing the innermost loop was stored within the transputer's fast-access on-chip memory, each transputer averaged $25,000$ spin updates per second — since a single sine, cosine or exponential calculation was found to take around $20\mu s$ on the same hardware, a time of $40\mu s$ for the entire update calculation is very encouraging. Thus the complete system could update each spin in a 62^2 lattice one million times in under two hours. Nevertheless, we have had to run the Meiko box more or less continuously in order to make an investigation over an extensive range of the system parameters and simulation timescales. This is particularly true of the three-dimensional system since, in addition to the increase in the number of spins in the system, the larger communications overhead reduces the efficiency of the program somewhat.

4.4.4 Spin Correlations

The evaluation of spin-spin correlations (section 3.5) required careful thought: while the magnetic interactions in the model are local, so that each transputer requires little information other than its own subset of the spins for an update attempt, the spin-spin correlations require global information. One technique was devised, where an accumulator grid was broken down into rows in the same way as the lattice, and the pieces of the accumulator circulated around the transputers in one direction while the spin configurations were circulated in the opposite direction, but the code for this was very convoluted and hard to maintain. The controlling T414 transputer had sufficient memory to store the entire lattice, but with no floating point unit, the correlations would take too long.

When the VAX took over as the Meiko host, it was realised that, rather than running a dumb slave process on the VAX (allowing access to host files) and the controlling code on the root T414 transputer, the controlling code could be trivially recompiled to run on the VAX itself, with only the routines for communicating with the rest of the transputers rewritten. This done, the VAX essentially became another computational element in the network, and its enormous memory made it ideal for performing the spin-spin correlations.

In practice, we did not use the fast-fourier transform technique for calculating correlations (section 3.5.1) — because the system size was not always a power of two, use of the FFT algorithm would require padding the data up to a power of two, which would destroy the symmetry of the data. With a reasonably optimised algorithm, the VAX was able to measure the correlations of the two-dimensional system well within the required time, and in three dimensions, we correlated only a subset of the sites with all other sites for each configuration — it seemed better to perform some correlations on a large number of different configurations than to measure all the correlations on fewer configurations.

Correlations were recorded, averaged and stored in two or three dimensions — while most of the later analysis was performed on a one dimensional (circular/spherical) average, the raw data was kept on tape for reference.

4.4.5 Multispin and Cluster Updates

Because the spins are distributed over the parallel processors, with each processor 'owning' a subset of them, implementing an algorithm which updates pairs of spins at a time (section 3.7.1) would be rather difficult; while pairs within one processor's full control could be handled like any conventional multispin encoding, an update across a processor boundary carries a large communications overhead — the entire system would stall while the adjacent processors negotiate. An additional problem to be overcome is the guarding against simultaneous update of adjacent spins – because each processor controls at least two rows of spins in the two-dimensional simulations, single-spins are automatically locked against adjacent updates; for two-spin flips, messages would have to be sent to explicitly synchronise access to neighbouring spins. These problems are not particularly hard to solve, but all the extra communication breaks up the flow of execution. (The processors operate on their slice of the system in typewriter-fashion, as described earlier, but they are not lock-stepped down to the individual spin; they merely synchronise before starting on the top row and bottom rows, to exchange data with the adjacent processors. On average, each lattice site takes the same time to process, but allowing the processors to operate at their own speed allows the system to compensate for small fluctuations; when messages are forwarded along the communications harness, for example, a small amount of cpu time is borrowed from the simulation.)

The obvious algorithm would be for every other processor to lend its spins to the neighbouring process, which will update pairwise then return the spins. Care must be taken to ensure that all possible pairs get a chance at such update.

An alternative algorithm which could be investigated for the two-dimensional system would be to have every second processor doing vertical pairs of (its own) spins (to avoid attempts to update adjacent spins). Then after every such sweep, the entire lattice is rotated one step, so that next time, the same processors will be operating on different lattice sites (which are not equivalent, because of the anisotropy). (There is no need to shuffle the spins back after the sweep — all processors are equivalent.) This means that all possible two-spin pairs are processed, without the need for inter-processor negotiation. More careful thought will be required to ensure that there is no possibility for biasing the exploration of phase space.

Cluster Update

Because the RAM model is predicted to breakup into domains at low temperature, the cluster-update algorithm (section 3.7.1) would appear to be the obvious weapon to deploy in simulating this system. Implementing such an algorithm, however, is decidedly tricky with a geometric distribution on a parallel machine; growing a cluster requires a great deal of negotiation between the processors, and it is not clear whether this task can be coded efficiently in parallel, or whether it simpler just to send all the current configurations to one node (or perhaps the host VAX) for cluster update.

Flanigan and Tamayo have recently published a paper describing an implementation of a cluster update on parallel computers [61]; the details have not yet been fully examined, but it appears that it would be reasonably easy to incorporate this into the existing code.

In order to evaluate the cluster update technique, a simple (sequential) program was developed for experimentation. One further difficulty which emerged was that the cluster-update algorithm as outlined in section 3.7.1 relies on a symmetry in the Hamiltonian: once the cluster has been identified, it is transformed as a whole, under the assumption that only the energy at the surface of the cluster changes. For the random anisotropy model under investigation, however, the internal energy of the cluster can be preserved only if a 180° flip is performed on each spin in the cluster, since any other rotation or reflection would change the energy between each spin and the anisotropy at that site. However, we found that such a cluster update rule generated such a severe change that all spins elected to follow the site in the transformation, and the entire lattice simply flipped time and again. Even if this can be overcome (using bigger system sizes for example), it is evident that a cluster update is not sufficient alone to explore the whole of phase space, except in the infinite anisotropy limit (section 2.7. But while the cluster algorithm encourages larger system sizes, the auxiliary (traditional) spin updates required to move the spins relative to their local anisotropy axes will still suffer from the slowing down which the cluster algorithm tries to avoid.

4.5 What We Measure

After every cycle during a simulation, the instantaneous values of the following quantities are accumulated:

- Energy U and U^2
- Magnetisation **M**. M^2, and M^4
- 'Liquid Crystal order parameter' $Q_{\mathrm{lc}} = \overline{\cos 2(\theta_i - \phi_i)}$
- Number of vortices and half-vortices

Several sets of data are collected, and from each, averages are calculated for the energy, magnetisation and Q_{lc}. The specific heat is calculated both by numerical differentiation of the energy (fitting a quadratic to triplets of energy measurements and differentiating), and using the fluctuation-dissipation theorem (equation 3.31). Magnetic susceptibility is calculated using equation 3.32, since in general we do not have data at finite fields for numerical differentiation. Errors in the results are estimated by averaging over the independent data sets, and comparison can be with data extrapolated using the histogram method (section 3.7.2).

If histogramming has been enabled, the measured quantities are added into the accumulators closest to the measured energy — it is sufficient to record only an average of the other quantities at each energy in order to extrapolate the results to other temperatures.

4.5.1 Liquid Crystal Order Parameter

Q_{lc} gives a measure of the extent to which the spins are aligned with their local anisotropy axes ϕ_i. $Q_{\mathrm{lc}} = 1$ for full alignment (section 2.7) and 0 for uncorrelated. In addition, this gives a direct measure of the component of the system energy due to the anisotropy: for two-component spins, equation 2.60 gives

$$\begin{aligned} E_{\mathrm{anis}}(\theta) &= \frac{D}{N}\sum_i \left(\cos^2(\theta-\phi_i) - \frac{1}{2}\right) && (4.10) \\ &= \overline{D\cos^2(\theta-\phi) - \frac{1}{2}} && (4.11) \\ &\sim DQ_{\mathrm{lc}}. && (4.12) \end{aligned}$$

Thus, the predictions made for the energy at zero temperature can be applied here (since the two terms in equation 2.70 are equal at the minimum).

Similarly, for large anisotropy strength, equation 2.76 suggests that

$$\begin{aligned} Q_{\mathrm{lc}} &= \cos\left(1-\frac{J}{D}\right) && (4.13) \\ &\sim \left(\frac{J}{D}\right)^2 . && (4.14) \end{aligned}$$

4.5.2 Vortices

The code to count the vortices and half-vortices was added to the $2D$ program towards the end of the work, when most results had been gathered, so it was not deemed worthwhile to spend much effort in optimising it. A vortex is detected by explicitly adding up the angles round each plaquette (group of four adjacent lattice sites), and if it is $\pm 2\pi$ (or rather, $\pm q$ where q is the number of quantisation angles), then a vortex is counted. We do not bother to count positive and negative vortices separately, since the use of periodic boundary conditions precludes the presence of a net vorticity in the system.

The half-vortices are counted in a similar way, with no attempt being made to distinguish half-vortices which are half of a full vortex (spread by thermal excitation, perhaps) and those which are simple single-spin excitations (figure 2.9), resulting in a pair of opposite half-vortices as shown in figure 2.2.

We have not yet attempted to count vortices in three dimensions — while it would be no problem to identify the presence or absence of a vortex at each volume element in the system, it would be much harder to identify the number of vortex *strings* in the system.

4.5.3 Addition Measurements

Additional measurements can made on the system less frequently:

- at any time, the current spin configuration can be copied to a 'reference' array, such that a spin-glass order parameter (section 2.47) can be calculated by comparing all the later configurations with the reference spins; the simulation control language allows the enabling and disabling of this measurement, which is not particularly time-consuming, but does generate a large volume of data.
- if spin-spin correlations have been enabled, then the correlations of a (random) subset of the spins with every other spin are calculated. The simulation description specifies how many spins are correlated each time, and how how often — it was decided that it was better to correlate a subset of the spins on a large number of configurations that to exhaustively study a few configurations.

4.6 Real Time Graphics

In order to aid qualitative understanding of the system, it was decided to make use of the graphics hardware available to show the spin configurations continuously as the system evolved in time. It was hoped that a video of these graphics could be submitted as part of this thesis. to demonstrate the insight thus afforded to the system, but hardware problems have precluded this. It is hoped that the colour pictures in the following chapters will illustrate the effectiveness of the graphics.

The spins were depicted in two complementary ways:

- the palette on the graphics system was configured to provide a continuous sequence of colours from black gradually increasing in intensity to red, then through purple to blue, finally diminishing in intensity back to black. Black was used to represent spins at 0°, through to red at 120° and blue at 240°. In this way, the presence of domains of spins pointing in approximately the same directions could easily be seen.
- green arrows were superimposed over these colours, to depict the actual spins at each lattice site — these arrows could be switched off to emphasise the colours, and hence the domain structure.

Because periodic boundary conditions were chosen for the simulations, the system was drawn repeatedly to fill the screen. Crude control was provided to allow the view to be panned across the system, and allow the view to be magnified to concentrate on features such as vortices, or reduced to see the entire system on screen at once. These graphics proved invaluable in understanding qualitatively the behaviour of the system, as the colour pictures in the following chapters should convey.

A possible enhancement would be to rewrite the graphics routines to display the system state in an X-window on the host — this might be the first step in perhaps allowing some degree of interaction with the system, as well as improving remote access to the system, not only from elsewhere in the building (the Meiko has a rather noisy fan, and would ideally be situated in a dedicated machine room — a long video cable is an alternative solution), but also (in principle) from the Internet (section 4.2). One complication would be that, while the graphics hardware is dedicated to the simulation code (since there is no software running on the Meiko to allow multiple users to use the system concurrently), there would need to be some other means of communicating with the program, in order to tell it onto which screen to open the window, and so on.

Chapter 5

Two Dimensional XY Model with Random Anisotropy

Presented here are the results from extensive simulations on two-dimensional systems of up to 128^2 sites. Because there remains dispute in the literature over the nature of the low-temperature régime (section 2.10), we are particularly interested in measurements which test the predictions of the Chudnovsky-Imry-Ma (CIM) theory (section 2.8). We are also interested in whether the Kosterlitz-Thouless phase survives the addition of the random anisotropy.

If the low-temperature régime is indeed glassy, we might expect results which depend on the history of the system, as well as on the timescales of the experiments. For this reason, we compare the results from 'cooled' and 'quenched' simulations: cooled systems start at high temperature, and simulations are performed at successively lower temperatures, with each simulation beginning with the final configuration of the previous one; quenched systems begin at low temperature with a random spin configuration, as if quenched from high temperature, and each successive simulation is at a higher temperature. We compare these results to separate qualitatively régimes of reversible and irreversible behaviour; for quantitative analysis, very much longer experiments would be necessary.

We have tended to regard the configurations generated by the cooled systems as true equilibrium, with the final configuration the ground state, but great care should be taken in practice. The Imry-Ma-Chudnovsky argument implies that the system breaks up into clearly defined domains, but this is not the case: for small anisotropy strength D, the spins change direction smoothly throughout the sample, correlated over what we choose to call the domain size. But because the energy gain in forming domains is due to statistical fluctuations in the details of the local random configuration within the domain, there is not one unique way in which the system can segregate. The energy barriers in this case may be the energy required to change the domain 'centres'. If we cool the system through a glass transition, we may freeze the domain structure, and the system may be unable to change the layout of the domains, in order to find the true ground state.

In section 5.1 we consider the internal energy and specific heat of the system. CIM is a zero-temperature theory, so that the specific heat is not considered. In the two-dimensional (pure) XY model, there is a peak in the specific heat which occurs *above* the Kosterlitz-Thouless transition temperature; this is therefore not a good quantity to study for a quantitative measure of the effect of anisotropy on the KT transition. However, there remains dispute in the literature on exactly where the Kosterlitz-Thouless transition lies in the pure system. For this reason, we nevertheless use the specific heat as a qualitative indicator of the effect of the anisotropy.

In section 5.2 we present measurements of what we have called the liquid crystal order parameter, which gives a measure of correlation between each spin and its local anisotropy axis, as described in section 4.5.1. The spin glass order parameter is also a measure of how much of the internal energy is due to the anisotropy, which allows an explicit check on the individual terms in equation 2.70.

In section 5.3 we consider the form of the spin-spin correlations. The pure XY system is expected to show exponential decay $\exp(-r/\xi)$ with distance r above the Kosterlitz-Thouless transition, and algebraic decay $r^{-\eta}$ below T_{KT}. We perform a spherical average of the correlation data and fit it to a generalised form $r^{-\eta}\exp(-kr)$. Care must be taken in fitting the data: because of the periodic boundary conditions, spins half-way across the lattice are influenced by the images of the spin at the origin (figure 3.6.1), so that we must limit the range over which the fit is carried out.

In section 5.4 we consider the spontaneous magnetisation of our system. It must be emphasised that while CIM predicts destruction of any long range order, and that the pure system supports only quasiferromagnetism rather than true long-range order, our finite samples do show a net magnetisation. We investigate how the magnetisation is suppressed by the addition of anisotropy, and how this scales with system size, in an attempt to extrapolate to the thermodynamic limit. We have been unable to make a quantitative analysis of the magnetic susceptibility of the system, because of the lack of self-averaging of data calculated using fluctuation-dissipation (section 3.5.2).

We also study the magnetisation in the presence of a large external magnetic field, to verify the régimes summarised in section 2.8.10, and in particular to pinpoint the crossovers between the different behaviours.

Having investigated the 'static' properties (on the timescales of our simulations), we turn to dynamic properties in section 5.5. Vortices are an important consideration in the pure system, so it is natural to investigate their properties in the presence of the anisotropy. We also study hysteresis curves, since the M-H characteristic is an important criterion for the existence of magnetism, and the amorphous systems represented by this model show particularly interesting hysteresis curves. Finally, we study the magnetisation over very long simulations, to try to get a quantitative estimate of the relaxation time of the system.

We have concentrated on small anisotropy strengths, comparing results with the pure system in the $D \to 0$ limit, but we have also performed simulations with large anisotropy strengths, even going up to $D \to \infty$ in order to compare our results with previous simulations in this limit. Other work to date has concentrated on the zero temperature limit, which is also where the CIM predictions are valid.

5.1 Energy

The energy per spin for cooled systems is shown in figure 5.1(a) for various values of small anisotropy strength D. The results for $D = 0$ are in excellent quantitative agreement with [17]. The shape of the energy appears to change very little as a function of temperature, so that the energy shift due to the anisotropy is independent of temperature. Accordingly, the specific heat curves are superposed (figure 5.1(b)). Also shown on this figure are curves for specific heat at $D = 0.4$ for both smaller and larger systems. Again, there is no change at low temperature, though the peak at $T \approx 1$ does slightly sharpen and move to slightly lower temperature, as expected (section 3.6).

For $D > 1$, there are clear differences in the energy and specific heat, as illustrated in figure 5.2 (where we have transformed the energy to Hamiltonian 2.55

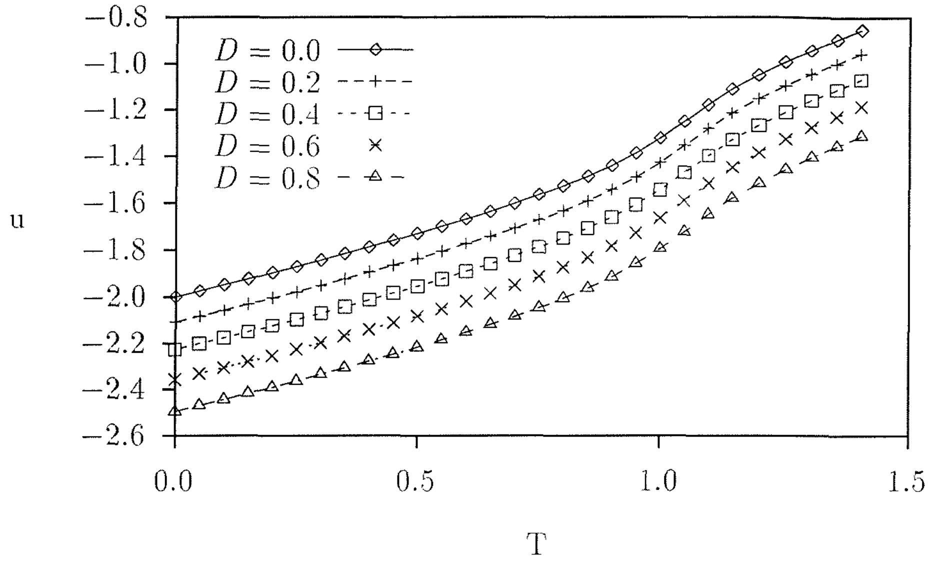

(a) Energy

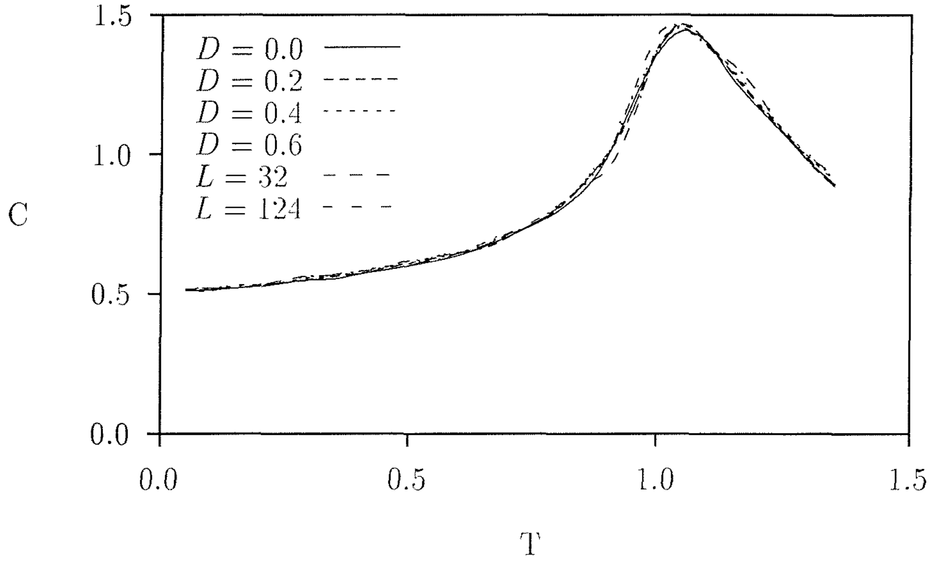

(b) Specific Heat

Figure 5.1: Single-spin energy (a) and specific heat (b) as a function of temperature for various anisotropy strengths D as indicated, for 62^2 systems. The specific heat curves in (b) are by numerical differentiation of those in (a), and appear to be independent of D for these small anisotropy strengths. Smooth curves are drawn through the discrete values to simplify the figure. Also shown on this figure are results for $D = 0.4$ for systems of 32^2 and 124^2 for comparison.

to ease comparison). The specific heat peak broadens significantly, and shifts to higher temperatures, reaching $T \approx 1.3$ as $D \to \infty$, in good agreement with [43].

This broadening seems reasonable for random systems, since different regions of the system might be expected to lose order at different temperatures, depending on the details of the local randomness. It is odd that the peak should move to higher temperatures as we increase the randomness, since we would expect a system with less order to require less thermal energy to disorder. It must be recalled, however, that the peak in the specific heat for the two-dimensional XY model appears *above* the transition temperature T_{KT}. The broadening could be a finite-size effect, but as we increase the anisotropy strength, the spins should be correlated over smaller distances, so that if the system is to show any sort of of finite size change, it would be as if the system was becoming larger, not smaller, in which case the peak would be expected to become sharper (section 3.6).

It in necessary to add 0.5 to the specific heat measurements made at $D = \infty$, since in this limit, the spins are constrained to lie along the anisotropy axes. At finite D, no matter how large, the spins are able to make arbitrarily-small oscillations about the axes — thus, there is one extra degree of freedom per spin, and so $\frac{1}{2}k_B T$ contribution to the specific heat per spin, by equipartition of energy. Because we have quantised the spins, for computational efficiency, it is possible for the small oscillations to be inhibited at very low temperature or very large anisotropy, though we have not seen this in practice. Because our system is discrete rather than continuous, the specific heat must eventually drop to zero at low temperatures, as required by the third law of thermodynamics. For the most part, our simulations are performed in the classical (continuous) limit of our discrete system.

Much more information can be conveyed in a three-dimensional plot — figure 5.3 shows specific heat as a function of both D and T for very large values of D. The contours under the surface of figure 5.3(a) suggest a slight shoulder which is barely perceptible — to make this clearer, the contours are shown in figure 5.3(b). Above $D \approx 5$, the main peak in specific heat has stabilised on the large-anisotropy limiting value of $T = 1.25$, though the height of that peak continues to decrease as D increases. A second, shallower, ridge appears to follow the line $D = 5T$ as indicated. This is merely noted in passing, as we have not performed detailed simulations in this part of the phase diagram.

The curves of specific heat from direct measurement by the fluctuation-dissip-

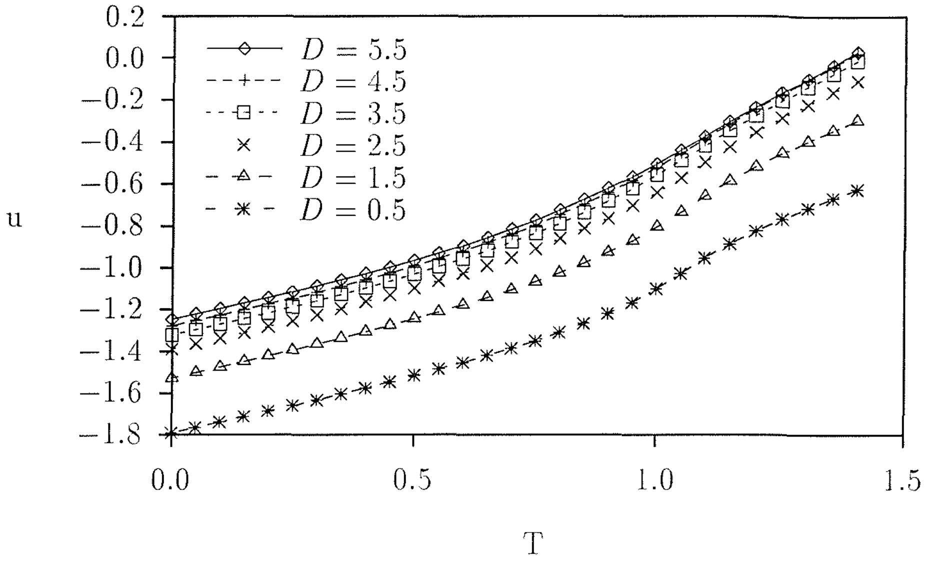

(a) Energy

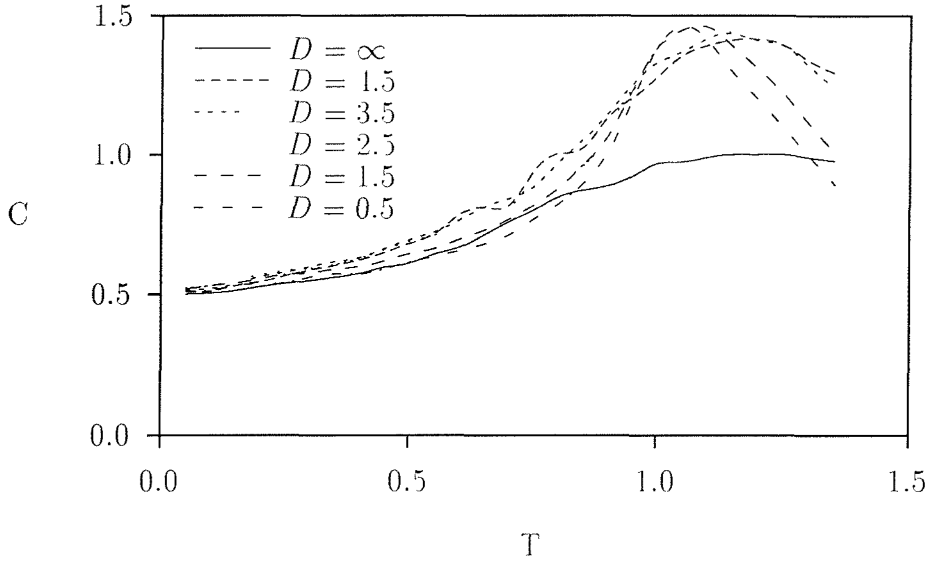

(b) Specific Heat

Figure 5.2: Energy (a) and specific heat (b) of cooled systems with larger anisotropy strengths D, as a function of temperature. The measurements in (a) have been transformed to Hamiltonian 2.55 to simplify comparison, and it is necessary to add 0.5 to the results for $D = \infty$, as described in the text. By $D \sim 5$, the specific heat curve has the same shape as for $D = \infty$, though is higher — the energy curve is not shown for $D = \infty$ in (a) to simplify the figure.

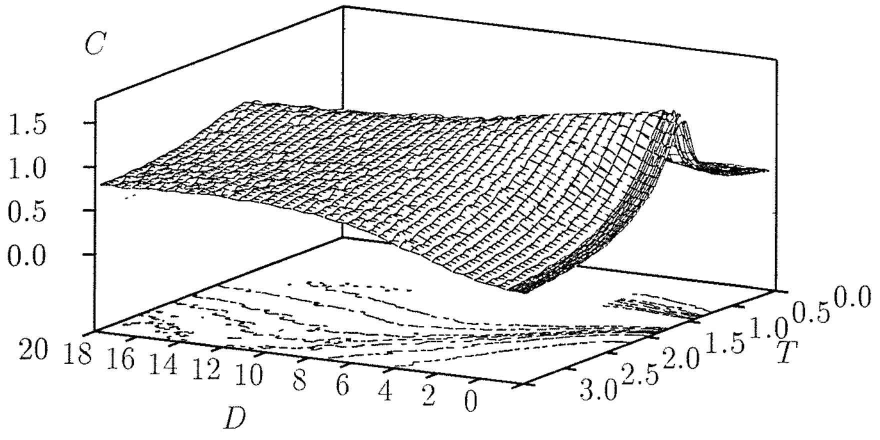

(a) Surface Plot

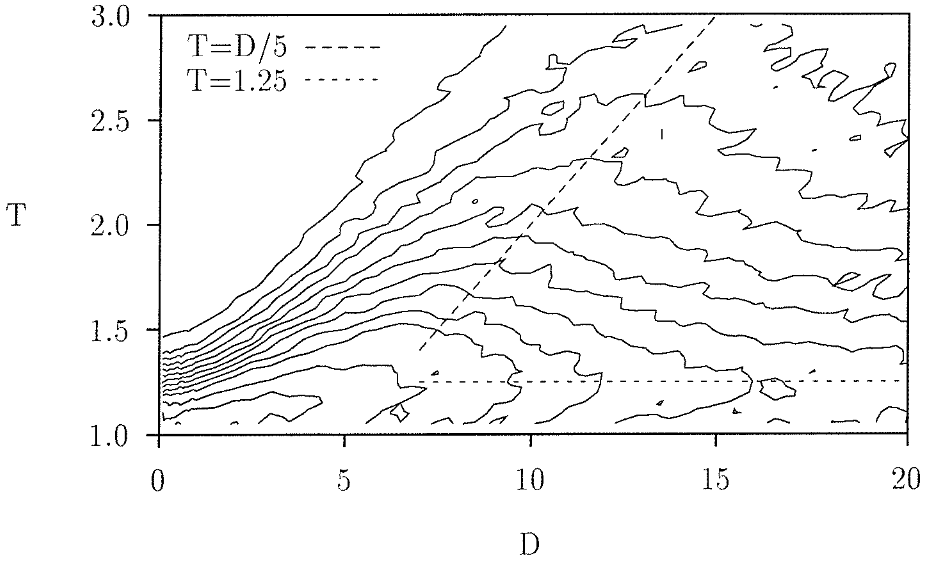

(b) Contour Plot

Figure 5.3: Surface plot (a) showing specific heat as a function of both temperature T and anisotropy strength D. The contours on the base on the base, reproduced on (b), imply a very gentle shoulder. as described in the text.

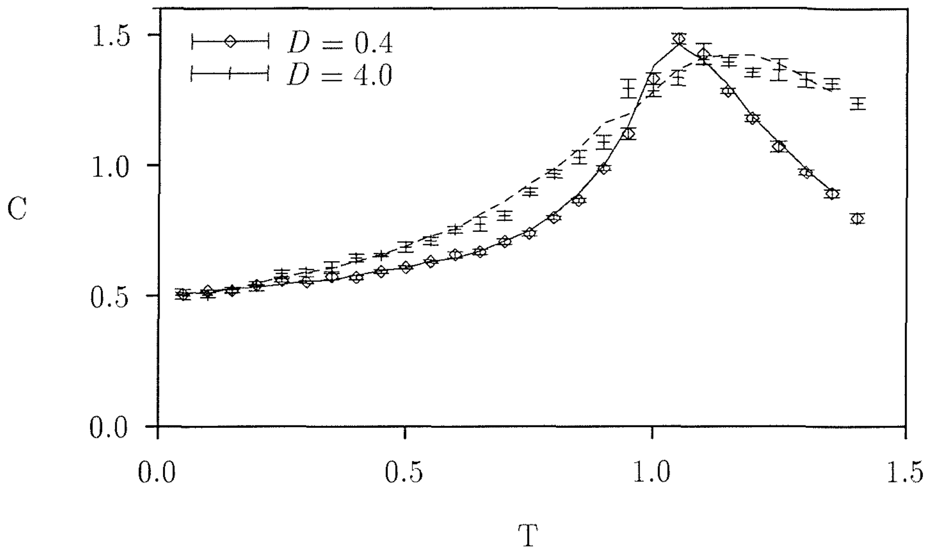

Figure 5.4: Comparing specific heat by direct measurement (points) and by numerical differentiation (lines). The points show some scatter, but lie on the lines to within experimental uncertainty

ation theorem (equation 3.31) are rather noisier than the results from numerical differentiation of the energy, as was also found by [17], but the curves do agree within experimental error (figure 5.4). This difference is due to the lack of self-averaging in the direct measurements (section 3.5.2).

In order to compare the change in energy as a function of D with the prediction made at zero temperature (equation 2.75), we must manipulate the measurements slightly, to transform to Hamiltonian 2.53. Figure 5.5 shows a log-log plot of $2+\frac{D}{2}-u(D)$ against D, with a gradient found to be around 1.75, to be compared with a prediction $u \sim D^2$ in two dimensions. While the data does not fit the form predicted by section 2.8.5, there may be a compromise between the two effects.

In the large anisotropy limit, in order to include measurements made at $D = \infty$, we transform the energy to Hamiltonian 2.55, and plot against $1/D$ in figure 5.6. For $D \to \infty$, the energy per spin at zero temperature is around -1.1, which corresponds to an energy per bond of $-0.55J$. We might expect $-\frac{2}{\pi}J$ if each spin were able to lie along the local anisotropy direction closer to alignment with the neighbouring spins, but because of the frustration, the alignment is worse, and the energy is increased. Fisch [47] similarly finds an energy of -0.5 per bond in simulations at $D = \infty$ for two-component spins in three dimensions. As D is reduced, energy seems to be reduced inversely proportionally with D, as

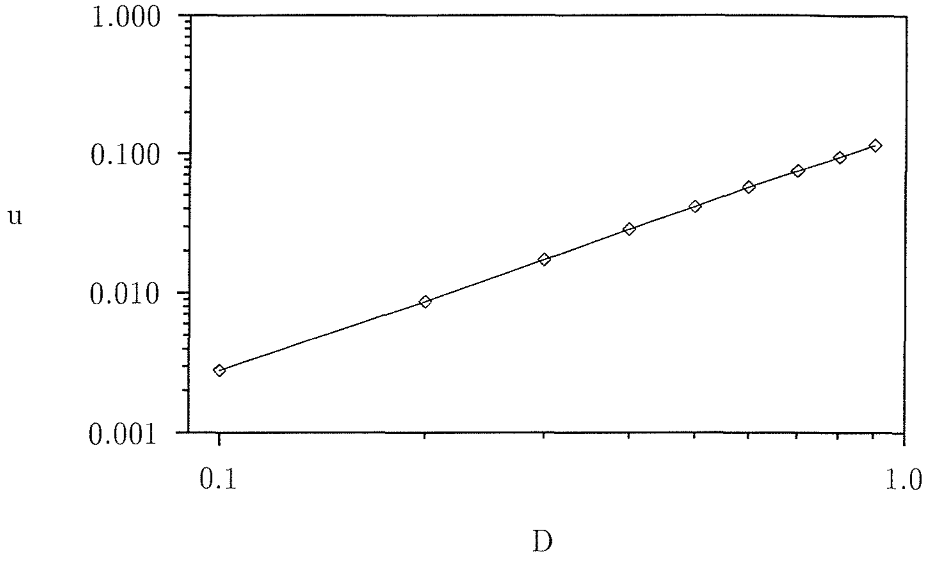

Figure 5.5: Modified single-spin energy as a function of anisotropy at zero temperature for small anisotropy strengths. We plot $2 - \frac{D}{2} - u(D)$ against D on a log scale as described in the text. A straight line fit has gradient 1.75, indicating $u \sim D^{\frac{7}{4}}$, compared with $u \sim D^2$ predicted by equation 2.75

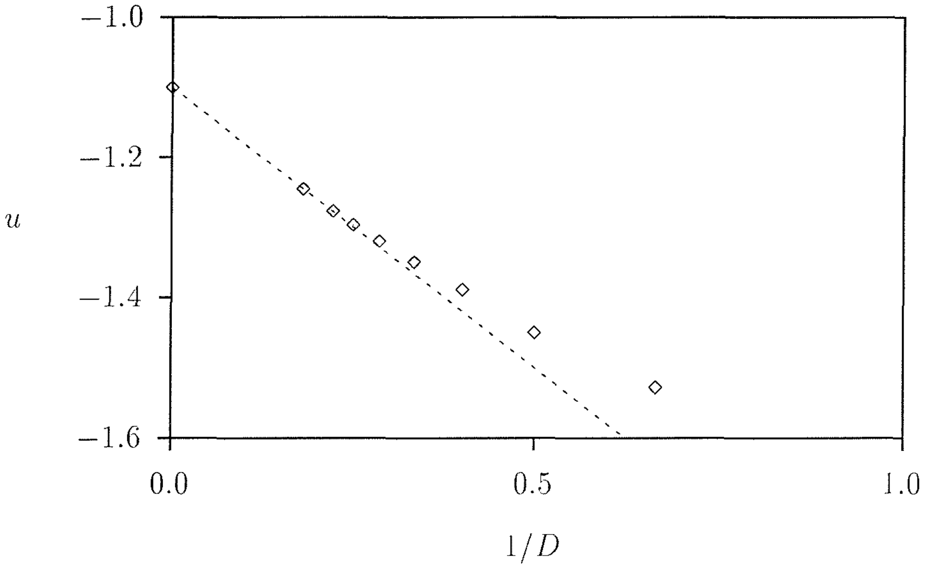

Figure 5.6: Modified single-spin energy as a function of anisotropy at zero temperature for large anisotropy strength. We plot $u(D) + D$ against $\frac{1}{D}$ as described in the text, and find approximately linear behaviour on the approach to infinite D.

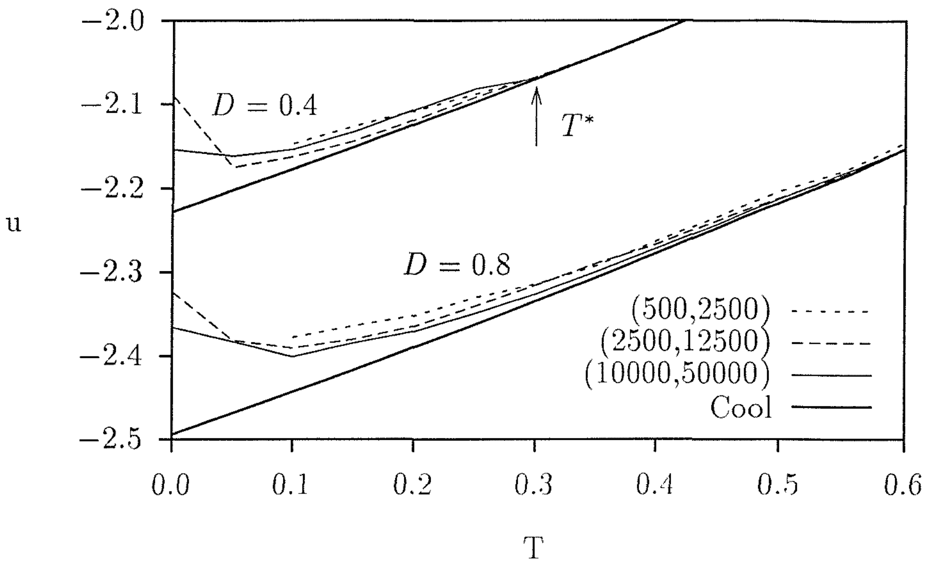

Figure 5.7: Comparing u(T) for cooled and quenched systems. The heavy line is for slowly-cooled systems. (x,y) represents x equilibration sweeps followed by y production sweeps.

predicted by equation 2.77.

Quenched Systems

Quenched systems appear to have higher energy at low temperatures than cooled systems — this indicates that the quenched system is unable to find the ground state, until it is warmed slightly, when it manages to *reduce* its energy and fall onto the cooled energy curve (figure 5.7).

At very low temperatures, the system makes large drops in energy for small increases in temperature — here, small gains in thermal energy allow it to cross the smaller of the energy barriers (section 2.6). Subsequently, the excess energy is lost very gradually — it seems that the system is caught in a deep energy minimum, and rather than crossing the final energy barrier into the region of the ground state configuration, the system remains in this energy well, and as thermal energy is gained, the system spends less time at the bottom of the energy minima, so that the average system energy becomes less dependent on the depth of the local minima. At sufficiently high temperatures, the system has sufficient thermal energy to be unaffected by the fine details of the random potential (section 2.8.4).

The extent to which this is a dynamic effect is unclear — for a finite system, the Metropolis algorithm has a non-zero probability of reaching *any* configuration, so if we run the simulations for long enough, the system *must* find the ground state. On the other hand, the fact that our system appears to have difficulty finding the ground state *may* indicate some sort of glassy behaviour, as discussed in section 4.3. Disregarding the very fast simulations, figure 5.7 does appear to show one temperature T^* for each D at which the rate of warming is irrelevant and the curves meet up — this might be taken to be a glass temperature (section 2.6, below which the system shows freezing. Of course, the exact details of the way the energy reaches the cooled value changes from run to run, and depends on the starting spin configuration and on the layout of the random axes; this will be discussed further in section 5.5.1.

If we only partially warm the system, to a temperature below T^*, then cool back to zero temperature, we find that the zero-temperature energy is between the original quench energy and that of the cooled system (figure 5.8). Once we warm beyond T^* and cool, the system energy follows the curve taken by a sample cooled from high temperature. Thus T^* seems to be an irreversibility temperature.

It is interesting that the measurements of the specific heat by the fluctuation–dissipation theorem do *not* differ significantly from the specific heat for the cooled system at all but the lowest temperatures (figure 5.9). This indicates that the system is exploring a region of configuration space similar to the region around the ground state, but as we have seen, the minimum is shallower.

5.2 Liquid Crystal Order Parameter

Inspired by work on liquid crystals, we have recorded what we have called the 'Liquid Crystal order parameter'

$$Q_{\mathrm{lc}} = \overline{\left\langle 2\left(\mathbf{S}_i.\mathbf{n}_i\right)^2 - 1\right\rangle}, \tag{5.1}$$

measuring the correlation between spin and anisotropy direction at each site — when there is a strong correlation, such as in the large anisotropy limit at low temperature, we would expect $Q_{\mathrm{lc}} \to 1$. For weak anisotropy or high temperature, $Q_{\mathrm{lc}} \to 0$.

Figure 5.11 shows this order parameter as a function of temperature for various values of anisotropy strength D, and as a function of D for various tempera-

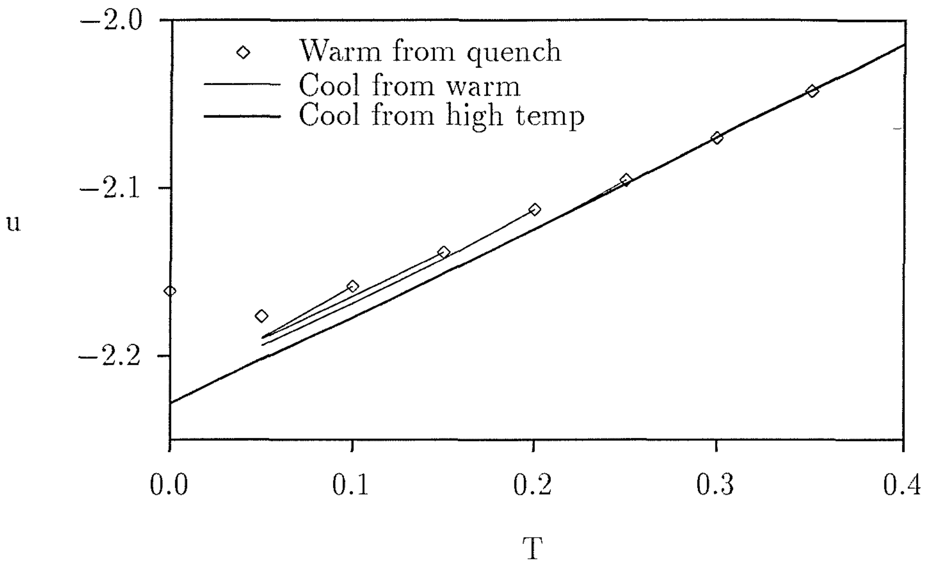

Figure 5.8: Investigation of reversibility for 62^2 and $D = 0.4$. The heavy line is for slowly-cooled systems. The system is started from a random configuration (quenched), then warmed (with 5000 equilibration and 50000 production sweeps). The diamonds show these points, which agree with figure 5.7. At each point, the final configuration is stored away, then the system is cooled once more. The lines in the figure show the energy as a function of temperature as the system is cooled from these various starting states. After each cooling run, the system is restored to the configuration before the cool, and warming proceeds. In an experiment on a real sample, the system would have to be quenched back to zero then warmed. It can be seen that above some T^*, the energy is a reversible function of temperature.

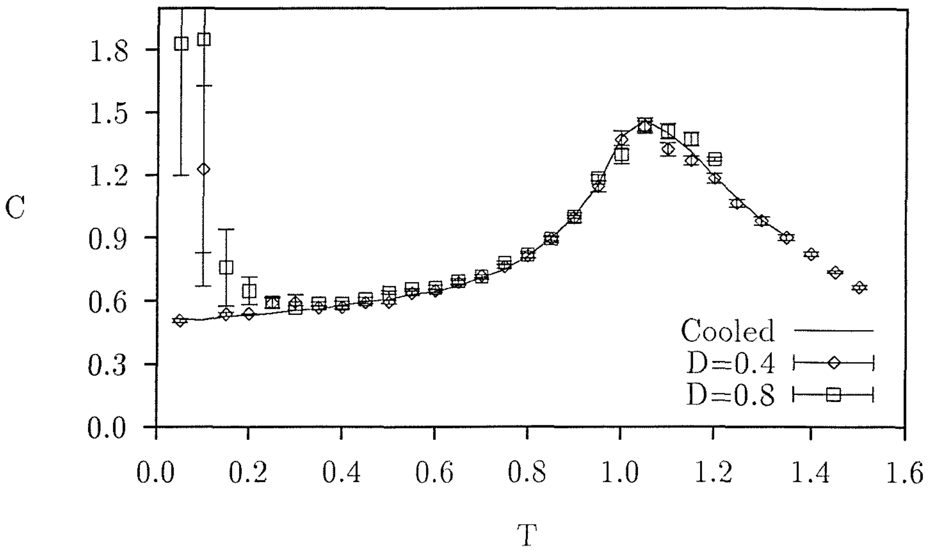

Figure 5.9: Comparing c(T) for cooled and quenched systems. The solid line is specific heat by numerical differentiation of energy for slowly-cooled systems. The points are measurements by fluctuation-dissipation, and do not appear to differ significantly from the cooled results, except at very low temperatures, where figure 5.7 shows large jumps in energy to be occurring.

tures T. Because these points appear to lie on smooth curves, considerable effort was made to find a universal form for the data. Fitting the data to the forms

$$Q_{\mathrm{lc}}(D) \sim D^{\lambda(T)} \tag{5.2}$$

$$Q_{\mathrm{lc}}(T) \sim \exp(-\mu(D)T) \tag{5.3}$$

gives $\lambda \approx 0.75$ at $T = 0$ (as shown on figure 5.11), rising approximately linearly with temperature until around $T = 1.2$ where λ levels off at around 1. This zero temperature result is consistent with the result of $U \sim D^{1.75}$ using equation 4.12.

Because we can also write equation 5.3 in the form $D^{\mu' T}$, or indeed as a product $D^{\mu_1 T}\exp(-\mu_2 T)$, we can combine the equations to arrive at

$$Q_{\mathrm{lc}} \sim D^{0.75+0.2T}\exp(-\mu_2(D)T). \tag{5.4}$$

Fitting this to the data shows that $\mu_2 \approx 0.7$, independent of D. Thus we can collapse $Q_{\mathrm{lc}}(D,T)$ onto $Q_{\mathrm{lc}}(1,T)$ in figure 5.11 by plotting $Q_{\mathrm{lc}}/D^{0.75+0.2T}$ against T, as shown in figure 5.12.

For $T > 1.2$, the fit is better using the simpler Q_{lc}/D. In equation 5.2, it may be possible to find a better form for λ which saturates at $\lambda = 1$ above

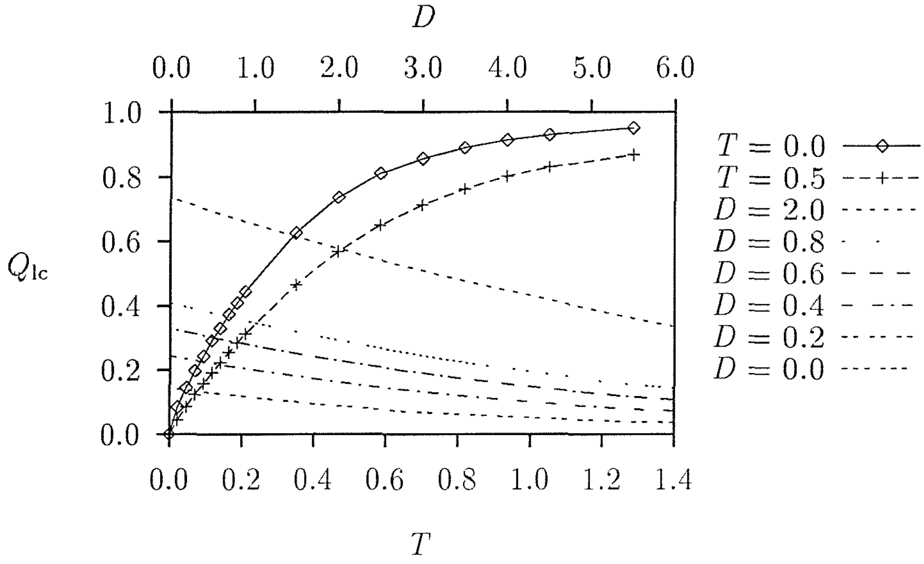

Figure 5.10: Liquid crystal order parameter as a function of temperature for various anisotropy strengths D (falling curves), and as a function of anisotropy strength for various temperatures T (rising curves). The falling points lie on curves of the form $Q_{\mathrm{lc}} \sim \exp(-\lambda(D)T)$ while the rising curves are fit by $Q \sim \mu(T)D$. The experimental points show anomalous behaviour at neither the T^* found in section 5.1, nor at T_{KT}.

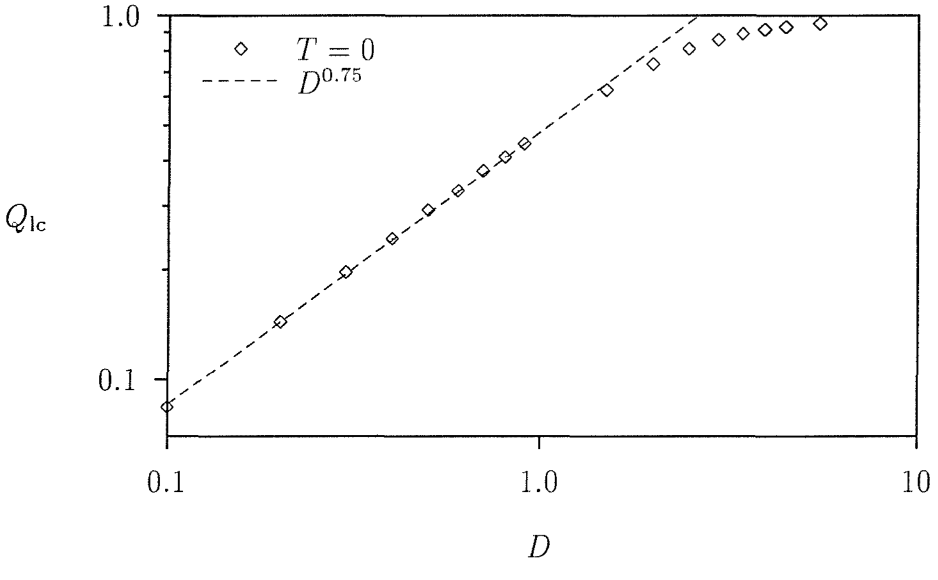

Figure 5.11: Liquid crystal order parameter as a function of D at zero temperature. For $D < 2$, Q_{lc} lies on the curve $D^{0.75}$, but it must fall away from this curve for larger D, since the order parameter cannot exceed 1.

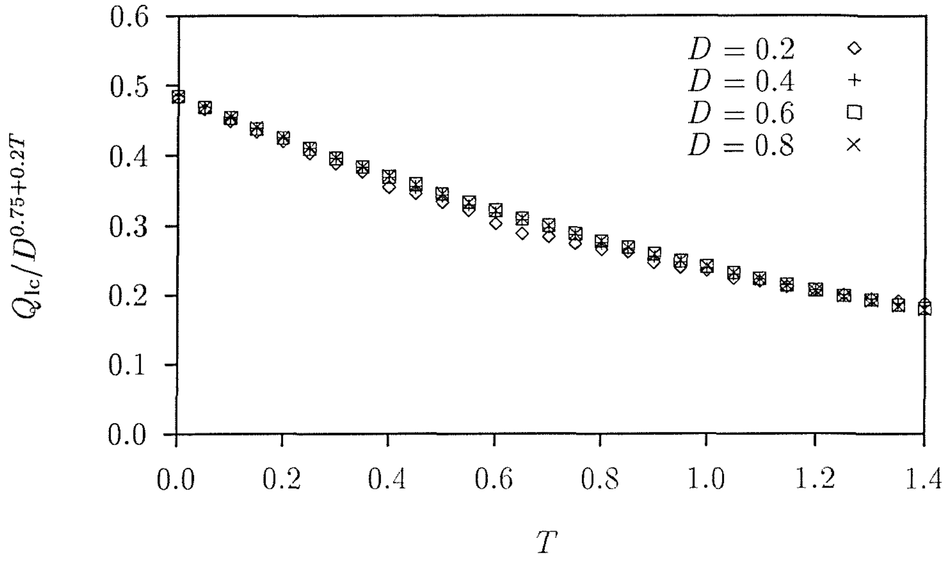

Figure 5.12: A plot of $Q_{\rm lc}/D^{0.75+0.2T}$ showing that the liquid crystal order parameter can be fitted to a universal curve, as described in the text.

$T = 1.2$. It had been hoped that finding a universal curve for $Q_{\rm lc}$ would allow some insight into the physics of the low-temperature state, but this has not yet been forthcoming.

For larger anisotropy strengths, $Q_{\rm lc}$ is predicted to approach 1 as D^{-2} (equation 4.14); figure 5.13 shows a log-log plot of $1 - Q_{\rm lc}$ against $\frac{1}{D}$, with the gradient of the straight line found to be 1.75.

Quenched Systems

For quenched systems at low temperatures, $Q_{\rm lc}$ is slightly higher than for the cooled system. Together with the results for energy of quenched systems (section 5.1), this suggests that the spins are 'pinned' by the anisotropy — this will be explored further in section 5.5.1.

5.3 Correlations

As described in section 4.5, we are able to measure spatial spin-spin correlations at regular intervals during simulations.

Because we are really interested in using the correlations to determine the

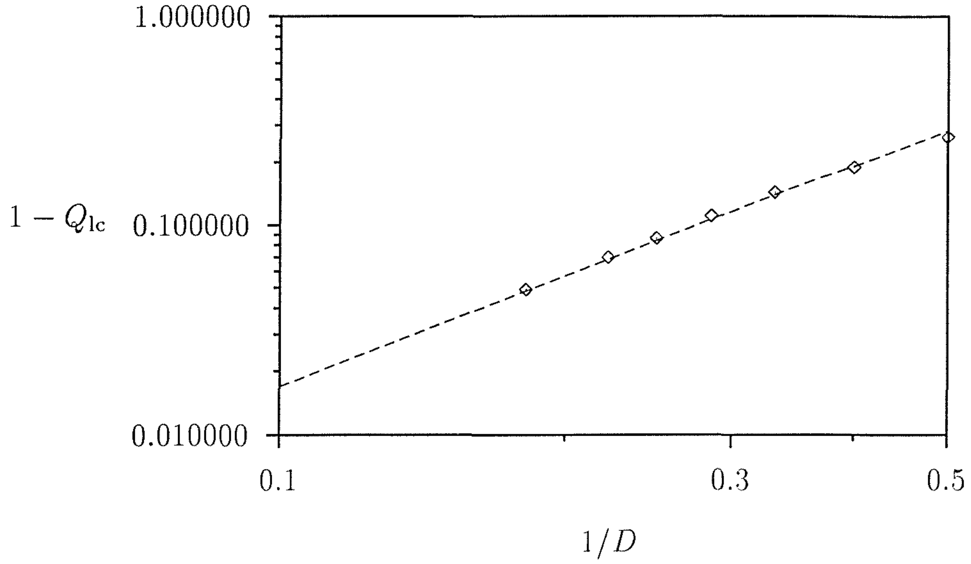

Figure 5.13: Behaviour of Q_{lc} at zero temperature in the large anisotropy limit. The straight line demonstrates that $Q_{\mathrm{lc}} \sim 1 - D^{-1.75}$.

nature of the long-range order, if any, we use the 'raw' correlation function (equation 3.21). For the two-dimensional (pure) XY model, the second term in equation 3.20 should anyway be zero at all non-zero temperatures, as there is no true long-range order in this model — we expect even less order in this random system.

As before (section 5.1), there appear to be three different régimes: a high temperature paramagnetic phase, an intermediate XY-like region with power-law decay, and a low-temperature régime. This is reflected in the spin-spin correlations — figure 5.14 shows the three different behaviours of $g(r)$ (circular average of $g(\mathbf{r})$) for a 62^2 system at $D = 0.4$ cooled from high temperature. Because of the periodic boundary conditions, the data is only valid out to $r = 31$ along the side of the square, but of course, along the diagonal, we can go out to $L/\sqrt{2}$ as shown.

The most striking feature at low temperatures is, of course, the *rise* in $g(r)$ for $r > 25$, but it should also be noted that for smaller r, the decay is no longer algebraic — figure 5.14(b) shows the same curves on a log-log scale, where algebraic decay would show up as a straight line. At very short distances, there do appear to be straight lines, but of course, most of the theoretical predictions for correlation functions are for the large r limit. Nevertheless, this may be evidence in support of the breakup into domains in section 2.8 — the correlations are XY-like

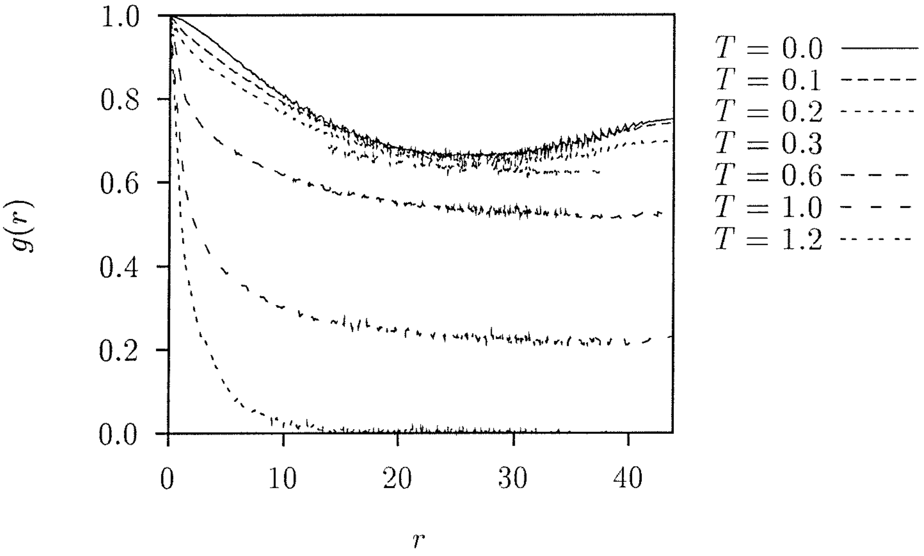

(a) Linear scale

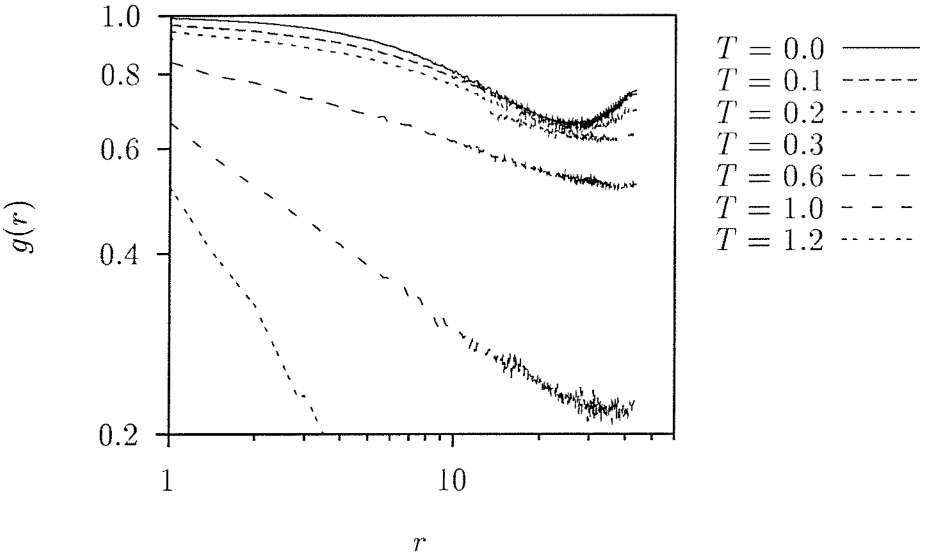

(b) Log scale

Figure 5.14: $g(r)$ vs r for cooled 62^2 system with $D = 0.4$, showing paramagnetic behavior (exponential decay) at high temperature, XY behaviour (power-law decay) at intermediate temperatures, and a distinct low-temperature behaviour. The same information is reproduced with a log scale in (b), showing straight lines (indicating power-law decay) at intermediate temperatures, but straight lines only at short distance at low temperature. The correlations would eventually approach zero for a sufficiently large system.

at short distances, and the correlations fall off more rapidly at larger distances, implying a domain-structure, perhaps. Our data supports the prediction that the spin correlations are independent of temperature at low temperature [20].

It is possible that the apparent XY behaviour is a feature of the limited system size, and that departures would show up at larger distance at all temperatures. However, there is no theoretical foundation for a temperature-dependent domain size in the Chudnovsky-Imry-Ma theory (section 2.8), and section 2.8.4 does predict that the breakup into domains will cease above some temperature T^*.

It appears that the spin correlations are XY-like provided that $g(r)$ has fallen by at least a third by $r = 25$ or so. At lower temperatures, where the pure system would be expected to be well correlated at this distance, the new behaviour takes over. This is understandable when the nature of the domains in section 2.8 is recalled: the domains are *not* due to some sort of correlations in the anisotropy axes; rather, they are due to the exchange interaction holding the spins together to a certain extent, *despite* the presence of the uncorrelated random axes. At higher temperatures, where the exchange interaction is not strong enough to keep the spins aligned to this degree in even the pure system, it is not surprising that the system cannot sustain the Imra-Ma-Chudnovsky domains with the addition of the random anisotropy.

We might suppose that the domains can be sustained provided $g(r)$ is greater than some constant over the domain size ξ:

$$\begin{aligned} g(\xi) &= \xi^{-\eta} && (5.5) \\ &= \text{const.} && (5.6) \\ \eta &\sim \frac{1}{\log \xi} && (5.7) \\ &\sim -\frac{1}{\log D} \qquad \text{using 2.74.} && (5.8) \end{aligned}$$

This holds only when ξ is greater than one lattice spacing, so that $D < 1$ and the right hand side is positive; it is only valid for small D, and indeed diverges as $D \to 1$. Using $\eta \propto T$ from equation 2.11,

$$T^* \sim -\frac{1}{\log D}. \qquad (5.9)$$

Before discounting the rise in $g(r)$ at large r as a finite size effect, or a feature of measuring $g(r)$ along the diagonal rather than along the edge, it is useful to

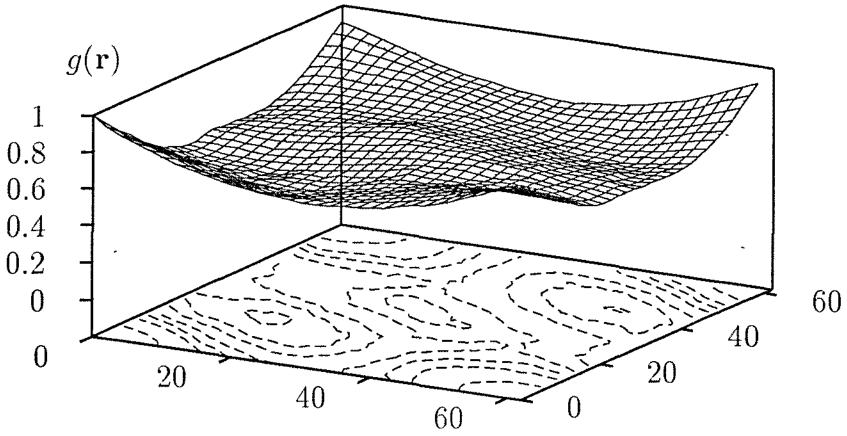

Figure 5.15: Three dimensional plot of $g(\mathbf{r})$ for 62^2 system with $D = 0.4$ cooled to $T = 0.1$. Contours on base are a guide to picking out features on the surface.

look back at $g(\mathbf{r})$ in two dimensions (i.e. without performing a circular average). Although this data is for one particular set of anisotropy axes, each $g(\mathbf{r})$ is a spatial average over all pairs separated by vector $\mathbf{r}$, so there has been some averaging over the details of the anisotropy. Figure 5.15 shows $g(\mathbf{r})$ for a 62^2 system with $D = 0.4$, cooled to low temperature. The system is clearly *not* circularly symmetric, and there does appear to be a domain structure, though adjacent domains are approximately aligned (reminiscent of the Ferromagnet with Wandering Axes discussed in section 2.8.7), so that the structure is not terribly clear. The results are similar to those from the studies of the one-dimensional XY model with random anisotropy in [62]. Figures 5.16 and 5.17 show the same system at higher temperatures, where the unevenness is reduced, suggesting again, perhaps, a changeover to XY-like behaviour.

We fit the correlations for a cooled 62^2 system to the form

$$g(r) \sim r^{-\eta} \exp -\frac{r}{\xi}, \tag{5.10}$$

and find results for η and ξ consistent with the qualitative description above, as shown in figures 5.18(a) and 5.18(b).

ξ^{-1} becomes non-zero above $T \approx 1$, so we conclude that the spin correlations change to exponential decay at high temperatures, as has been shown earlier. The apparently random values for η in this régime are simply features of the fitting

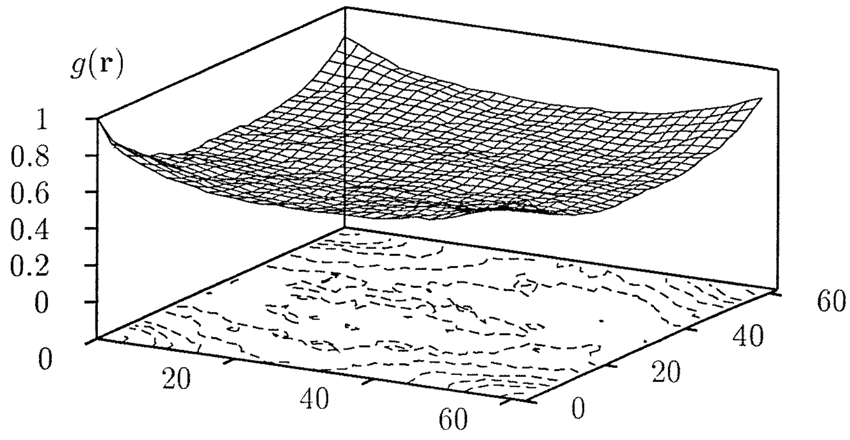

Figure 5.16: Three dimensional plot of system in figure 5.15 at $T = 0.3$, showing only a little unevenness.

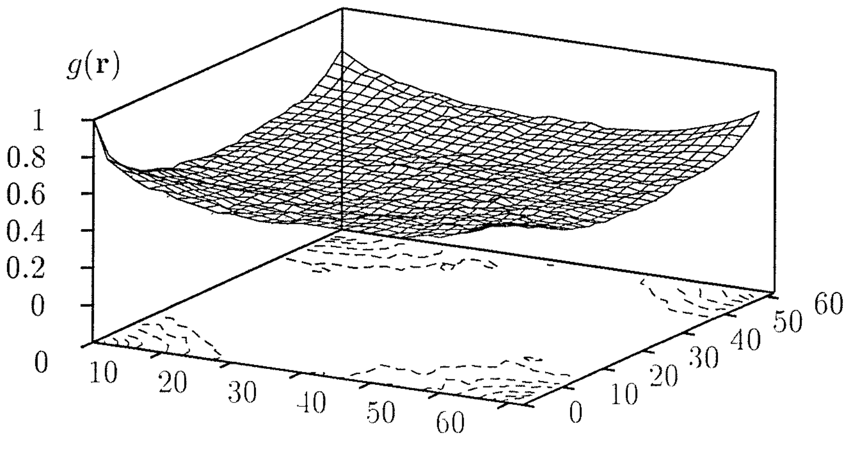

Figure 5.17: Three dimensional plot of system in figure 5.15 at $T = 0.6$, showing a much flatter surface.

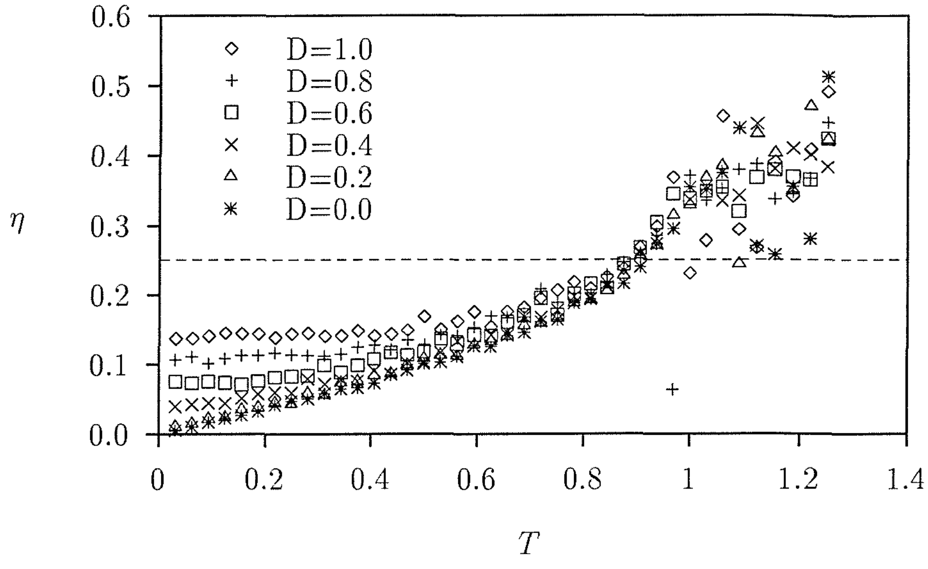

(a) η

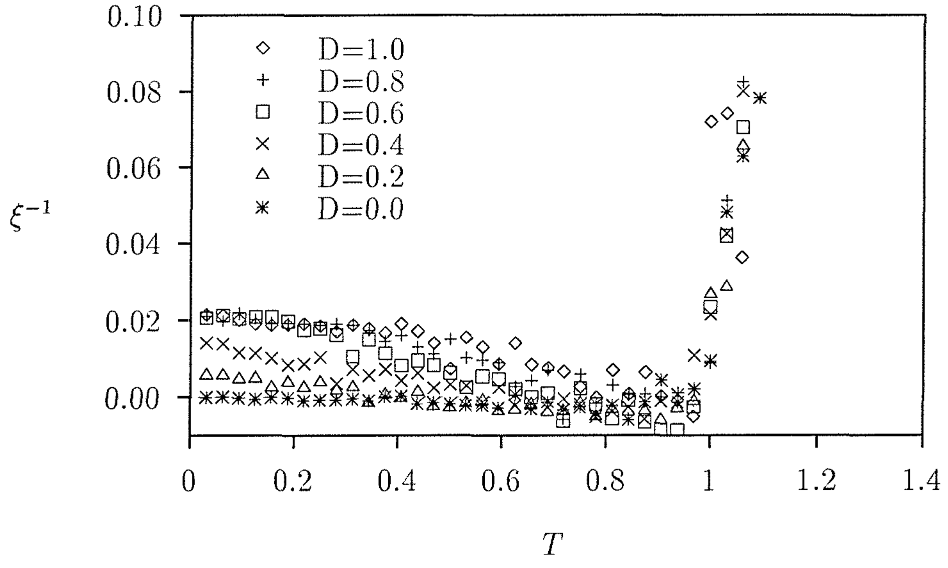

(b) ξ^{-1}

Figure 5.18: Fitted value of η and ξ^{-1} in equation 5.10 to $g(r)$ for 62^2 system. At high temperatures, all ξ diverge ($\xi^{-1} \to 0$) at the same temperature, indicating that the parametric to quasi-ferromagnet transition is unaffected by the anisotropy. (Random values for η are a feature of fitting process.) Decay is power-law at intermediate temperatures, with $\eta(D,T) \approx \eta(0,T)$ Poor fit at low temperature indicate that data no longer fits this form.

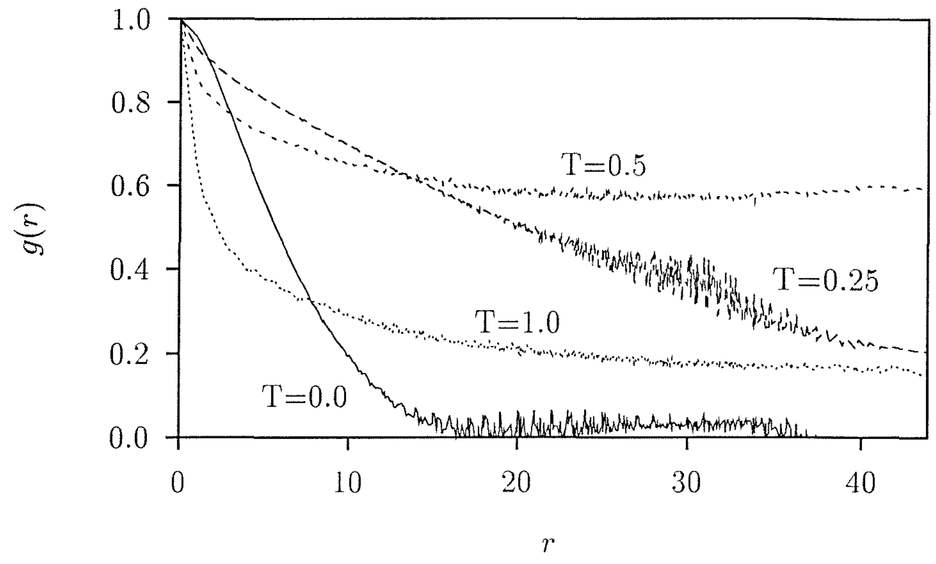

Figure 5.19: $g(r)$ for quenched 62^2 system with $D = 0.4$

process. For all the values of D shown, the fitted parameter ξ goes to infinity at a temperature just below J on figure 5.18(b), where η has a value around 0.25 (indicated by dotted line) as expected (section 2.2). Below this temperature, η for the pure system ($D = 0$) falls slowly with temperature to 0 at $T = 0$, again as expected. For non-zero D, η follows the $D = 0$ curve down to some non-zero temperature (which depends on D), below which it appears to level off, as predicted by [20, 33]. The fit at low temperatures should not be trusted quantitatively, since the data is no-longer a good fit to equation 5.10, as described earlier; this does, however, give us another measure of T^*, the temperature below which pure XY behaviour is lost.

Quenched Systems

The correlations for a system quenched then warmed again show three distinct régimes, but now the low-temperature states have rapidly-decaying correlations, as shown in figure 5.19, except perhaps for a slow decay at very short distance, as before. This is consistent with the conclusions drawn from the energy data in section 5.1: the system at low temperatures is pinned by the anisotropy, and cannot escape until the system is warmed somewhat. The fanning of the data at larger r again implies loss of circular symmetry, and perhaps a breakup into domains.

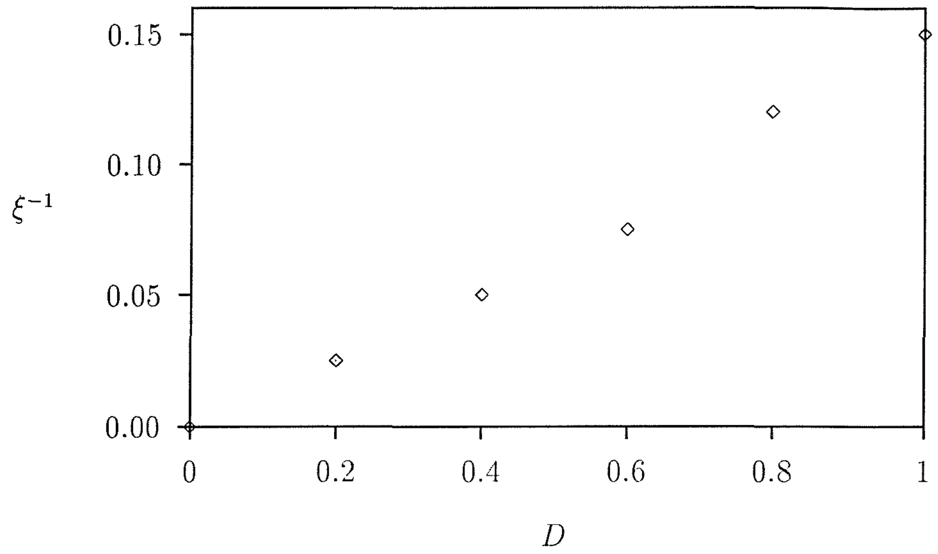

Figure 5.20: Domain size is inversely proportional to D, as predicted by equation 2.74

One possible picture is that, while the domains in the cooled system were almost aligned, here the neighbouring domains are misaligned, and the averaging implicit in the measurement has set most of the long distance correlations to zero — see section 5.3.1. The curve at $T = 0.25$ shows that even after the correlation length reaches the order of the lattice size, the system is not yet in the pure XY state — this is at least not *entirely* a finite-size effect.

In order to get some measure of the domain size, we have forced a fit to an exponential form, to find a correlation length as a function of D — any consistent way of fitting the data would be appropriate here. This measurement is made not at zero temperature, since our single-spin-update Monte Carlo algorithm is unable to cross even the smallest energy barriers to minimise the energy. Figure 5.20 shows that domain size is approximately inversely proportional to D, in good agreement with equation 2.74, despite the rather ad hoc measurement.

Dieny and Barbara [52] found that the spin correlations decayed from 1 to $\exp(-1)$ over a distance $\xi \approx 14\exp(-1.4D)$ at zero temperature — this does not diverge as $D \to 0$ as expected due to difficulty in eliminating all pairs of vortices in the pure system. Despite a slightly different definition of correlation length, our results are in good quantitative agreement, and the correlation length by this measure for $D = 0.4$ is consistent with the domain size suggested by figure 5.14.

5.3.1 Spin Configuration Snapshots

In this section, we present pictures of instantaneous spin configurations, as revealed by the real time graphics (section 4.6) — it has been hoped that a video showing the progression of the configurations through time could be made available, but that has not proved possible in practice. Figure 5.21 shows an instantaneous spin configuration for a 62^2 system with $D = 0.4$, slowly cooled to $T = 0.3$. The spins are predominantly aligned in the y direction, but this is almost certainly a feature of the finite size of the system.

The analogous plot for a quenched system is presented in figure 5.22.

Warming further, the vortices disappear, and we reach the configuration shown in figure 5.23

The role of vortices will be discussed further in section 5.5.1.

5.4 Magnetisation

It must be emphasised that, as has been discussed in chapter 2, the XY model in two dimensions has no true long range order, and can have no spontaneous magnetisation in zero field at non-zero temperature in the thermodynamic limit. Instead, it shows 'quasi' ferromagnetism — the spin-spin correlations show algebraic decay (equation 2.11), and the system is *almost* magnetic. Simulations of finite systems *do* show a spontaneous magnetisation at low temperature, however, because the spin correlations do not decay within the finite size of the lattice (section 5.3) — this measured magnetisation is expected to fall to zero in the thermodynamic limit.

We would expect the addition of randomness to further reduce the amount of order at low temperatures, so that there can be no net magnetisation in the random case either.

That said, we have been able to use the spurious spontaneous magnetisation of our finite systems to try to determine further the low temperature properties of the random anisotropy system. We record the root mean square magnetisation, so that we can follow the magnitude of the magnetisation — the pure system is degenerate in the direction of any net magnetisation, so that a vector average decays with time as the spins rotate, whereas a root-mean-square measurement

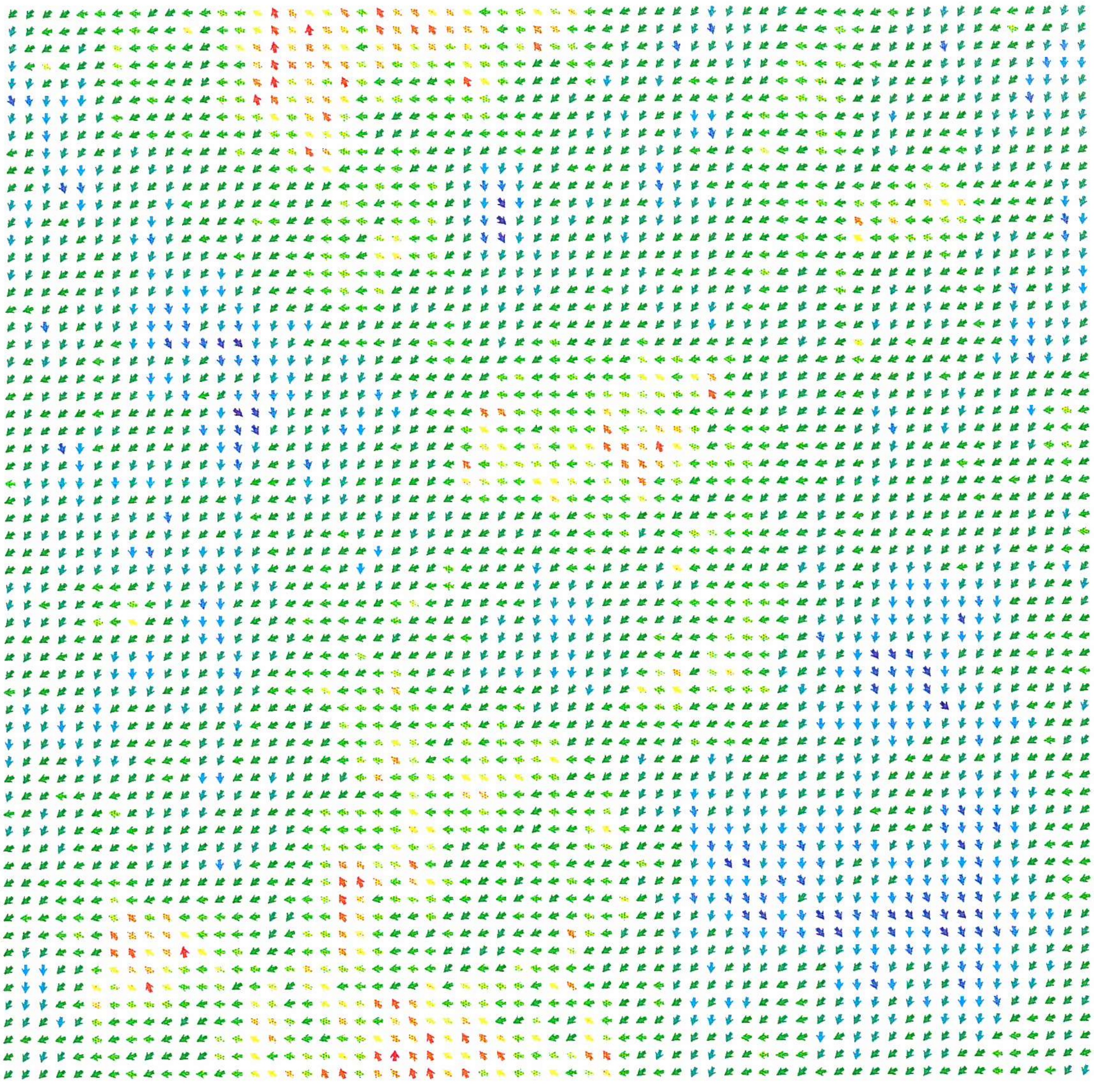

Figure 5.21: Instantaneous spin configuration for 62^2 $D = 0.4$ system cooled to $T = 0.3$

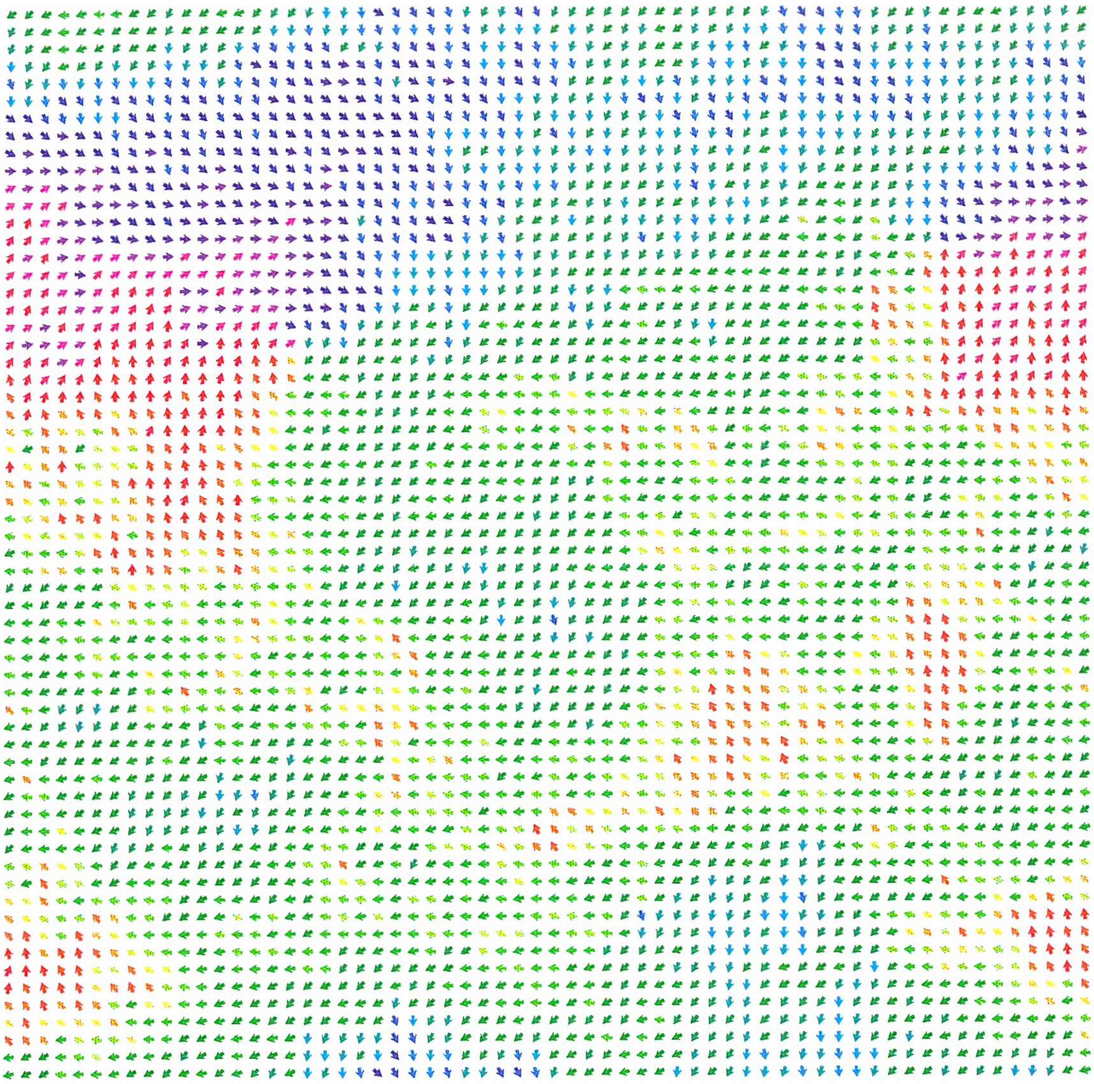

Figure 5.22: Instantaneous spin configuration for 62^2 $D = 0.4$ system quenched then warmed to $T = 0.3$

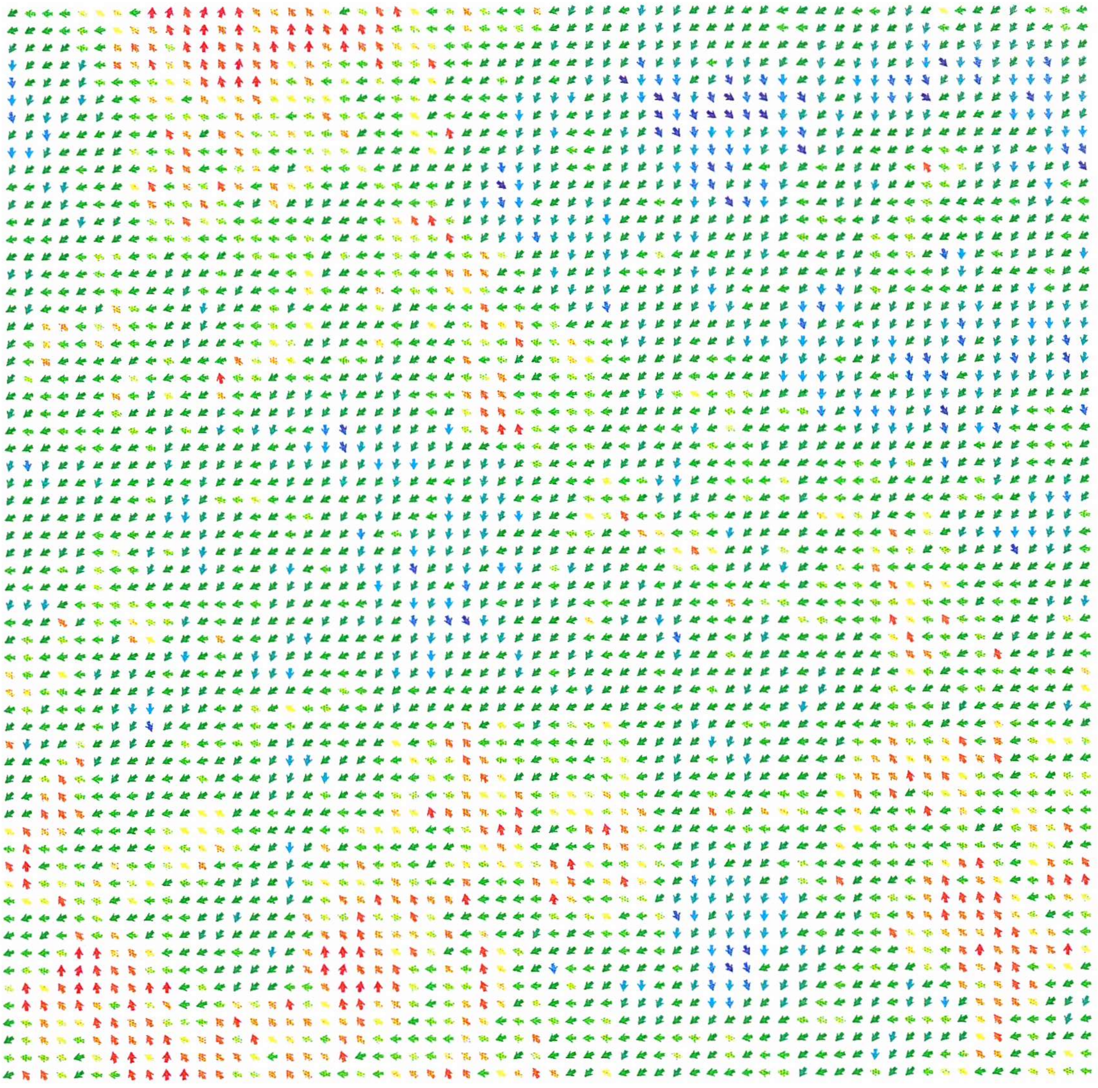

Figure 5.23: Instantaneous spin configuration for 62^2 $D = 0.4$ system quenched then warmed to $T = 0.5$

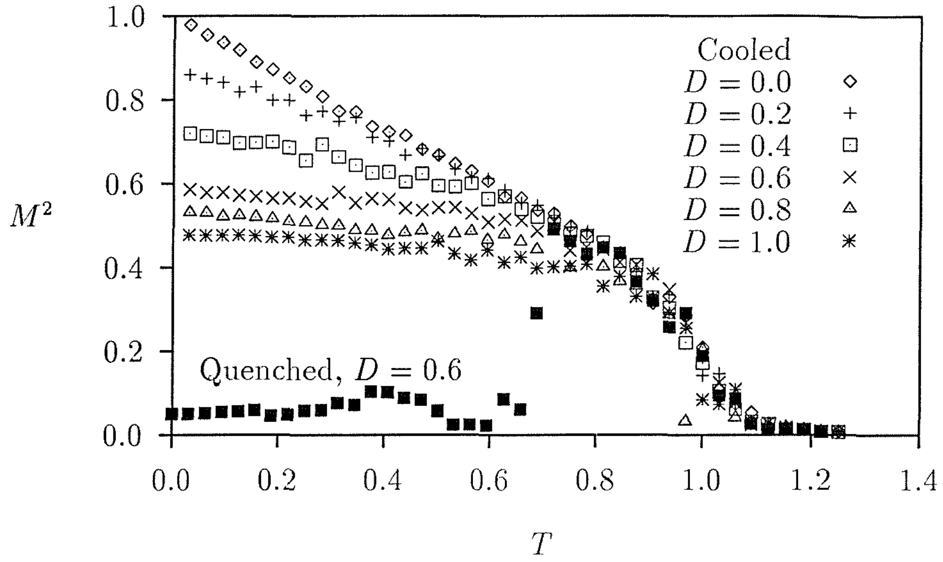

Figure 5.24: M^2 as a function of temperature for various anisotropy strengths D. As noted in the text, the magnetisation must go to zero in the thermodynamic limit, but it is useful nevertheless to investigate the behaviour of this non-thermodynamic variable. Also shown is the magnetisation from a system which we have quenched then warmed, showing non equilibrium behaviour.

should have a well-defined average.

5.4.1 Magnetisation of a Finite System

In figure 5.24, we compare M^2 as a function of temperature for various anisotropy strengths D for a 62^2 system. For $D = 0$, the magnetisation is small for $T > 1$, and below this temperature, it grows smoothly as the system is cooled, reaching $M^2 = 1$ at $T = 0$. These results are very similar quantitatively to those of the 60^2 system studied in [17]. As we introduce randomness, keeping $D < 1$, the magnetisation appears to follow the curve for $D = 0$ as we cool down to some $T^*(D)$, below which it levels off, so that the magnetisation is depressed at low temperatures.

This is just as would be expected from the behaviour of the spin-spin correlations in section 5.3 — at sufficiently high temperatures, the spin correlations are XY-like, and by equation 5.14 the magnetisation should behave similarly.

These results are consistent with [34], where the system at low temperatures is renormalised back to T^* and $D = 0$, and [20], where the spin-spin correlations are reported to be independent of temperature below T^*. For larger D, T^* has increased beyond $T = 1$, and there is no intermediate range of temperatures in which the system can be said to be XY-like.

Quenched Systems

Figure 5.24 also shows a typical result from a quench, where we find behaviour similar to that shown by the energy (section 5.1): at low temperature, the magnetisation is significantly suppressed, as is expected from the behaviour of the correlations in section 5.3, and only when we start to warm the system is it able to increase the magnetisation to match that of the cooled system — the system is apparently increasing its order as we increase the temperature, which may be an indication of energy barriers being overcome as the system gains thermal energy.

Magnetisation as a function of D

For small anisotropy strengths, we can compare the deviation from saturation with the predictions of [29]. Figure 5.25 shows the deviation from saturation for a 32^2 system at $T = 0$ in zero field, and the line indicates $D^{2.2}$, showing reasonable agreement with the prediction. The smaller system size tends to mimic the effect of cooling in a non-zero field — while at $T = 0$, the non-random system must be fully-aligned, the random system may need extra encouragement to align, and the smaller system size helps this. (Usually this is regarded as a problem, of course!)

5.4.2 Susceptibility

While we have recorded the magnetic susceptibility during our simulations, the lack of self-averaging which reduced the quality of the specific heat measurements (section 5.1) is evident also for the susceptibility measurements. Because of finite size effects, we have been unable to get the system into an equilibrium state with small spontaneous magnetisation, so that our results cannot be interpreted as zero-field susceptibilities. This has also prevented us from finding the susceptibility by explicitly differentiating the magnetisation with respect to applied field strength.

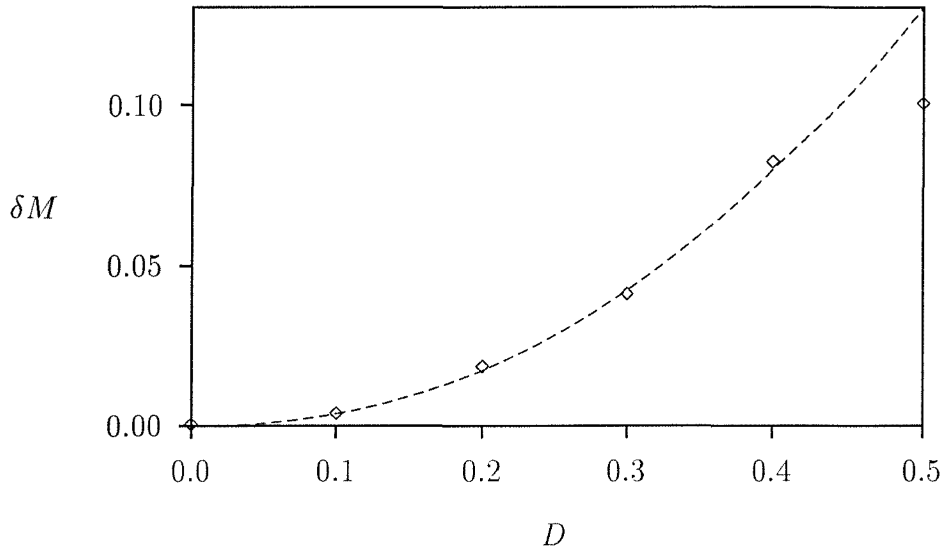

Figure 5.25: Deviation from saturation as a function of small D for a small (32^2) system at $T = 0$, in zero external field. Line shows $D^{2.2}$, compared with prediction of D^2 [29],

The power-law decay of correlations at intermediate temperatures indicate an infinite-susceptibility régime, and all we can say at present is that the susceptibility is strongly suppressed at low temperatures by the addition of the anisotropy. We cannot confirm the predictions in equation 2.84 directly.

5.4.3 Scaling Behaviour

The magnetisation and spin correlations are linked by

$$\begin{aligned}
\left\langle M^2 \right\rangle &= \left\langle \left(\sum_i \mathbf{S}_i\right) \cdot \left(\sum_j \mathbf{S}_j\right) \right\rangle &(5.11)\\
&= \sum_i \left\langle \sum_j \mathbf{S}_i . \mathbf{S}_j \right\rangle &(5.12)\\
&\approx L^d \left\langle \int_{r=0}^{L} g(r)\, d^d\mathbf{r} \right\rangle &(5.13)\\
\left\langle m^2 \right\rangle &\sim L^{-d} \left\langle \int_{r=0}^{L} g(r)\, d^d\mathbf{r} \right\rangle &(5.14)
\end{aligned}$$

so that the way the magnetisation falls to zero with L can reveal information about the nature of the spin correlations, and allows comparison to be made with the explicit measurements in section 5.3.

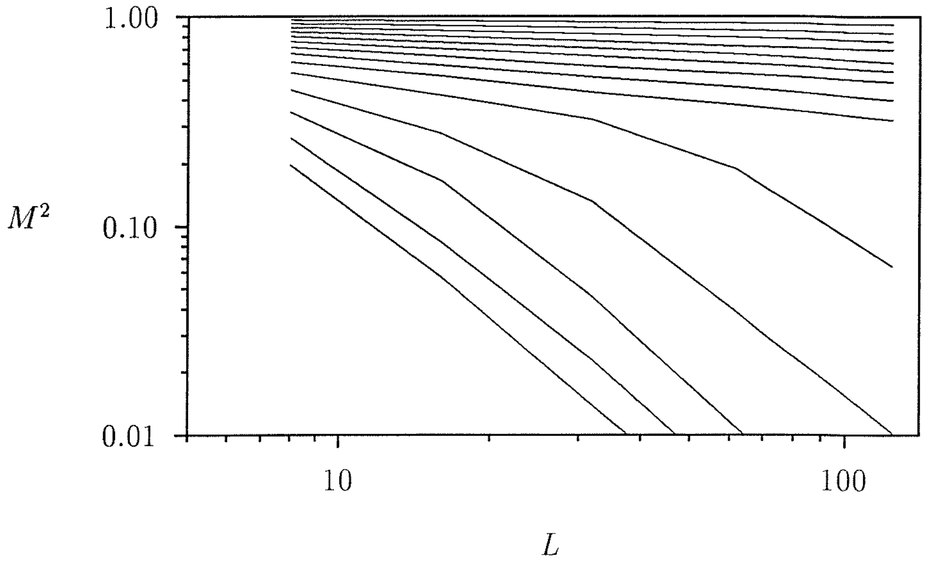

Figure 5.26: Scaling behaviour of the pure (nonrandom) system. Straight lines on a log-log plot indicate algebraic (power-law) decay.

When there is no long range order, and the system size is much larger than the correlation length (so that the integral in equation 5.14 is independent of L), the magnetisation falls as $L^{-\frac{d}{2}}$.

Using $g(r) \sim r^{-\eta}$ in equation 5.14, it can be verified that the magnetisation per spin scales as $L^{-\eta/2}$ for the non-random case, and we can investigate the scaling behaviour of the random magnet.

Figure 5.26 shows a log-log plot of zero-field magnetisation as a function of size for various temperatures for the pure system, clearly showing straight lines as expected, except for a region of transition from low-temperature quasiferromagnetism (algebraic decay) to paramagnetism (fluctuations) near $T = 1$, where the lines appear to be tending towards $m^2 \sim L^{-d}$ for large L. The fit of η to these curves is shown in figure 5.28, and shows η increasing from 0 at $T = 0$ to 0.25 around $T = 0.9$, in good agreement with [17]. The magnetisation scales as L^{-1} at high temperatures, as expected.

Figure 5.27 show the scaling behaviour for $D = 0.4$ — figure 5.27(b) concentrates on the low-temperature régime, which is not clear from figure 5.27(a). At low temperatures, the magnetisation is continuing to bend over at large L, so that we are unable to determine the behaviour quantitatively. If the Chudnovsky-Imry-Ma domain picture, supported by section 5.3, is to be believed, it might be

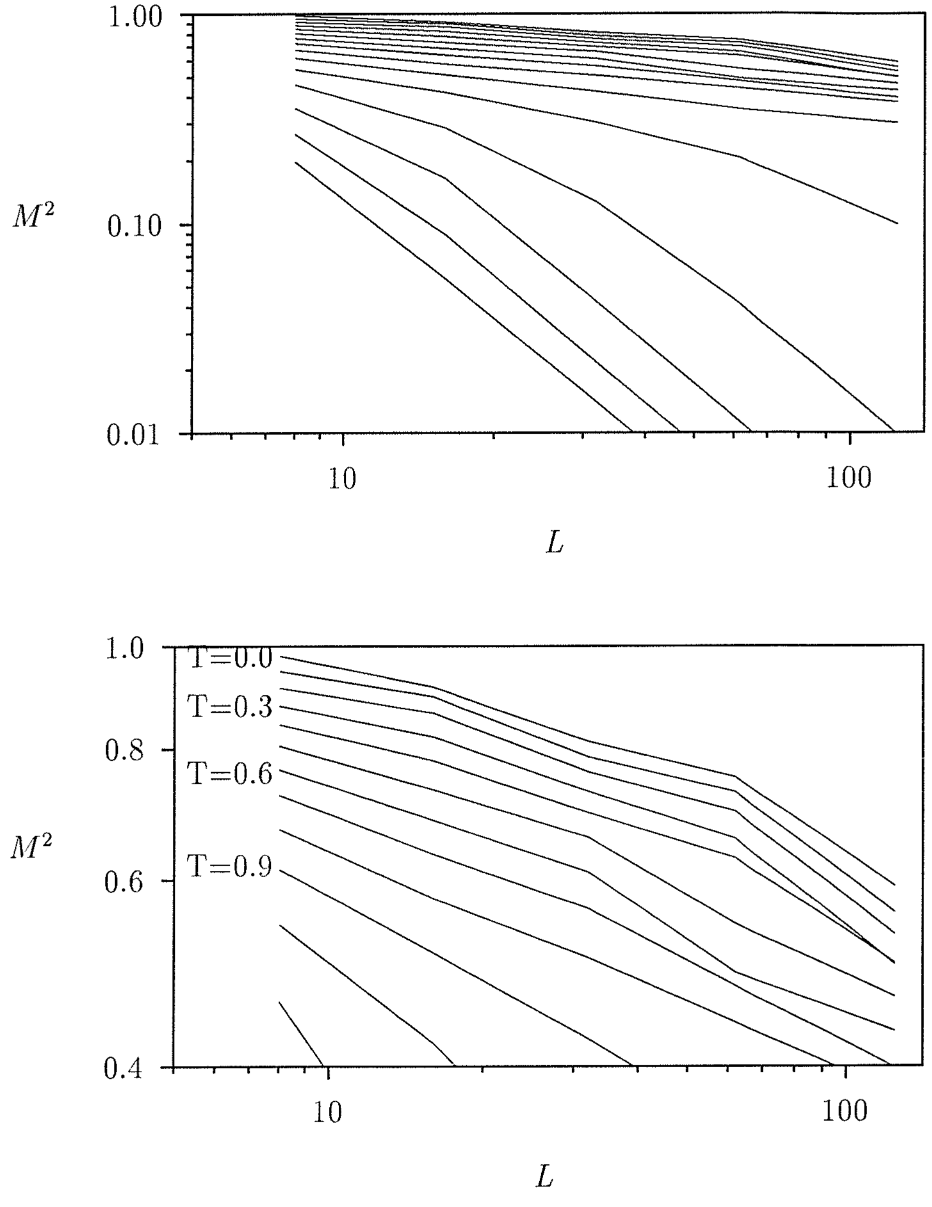

Figure 5.27: Scaling for $D = 0.4$ showing undetermined behaviour at low temperatures (clarified on lower figure), power law decay (straight lines) at intermediate temperature, and parametric behaviour (L^{-d}) at high temperature.

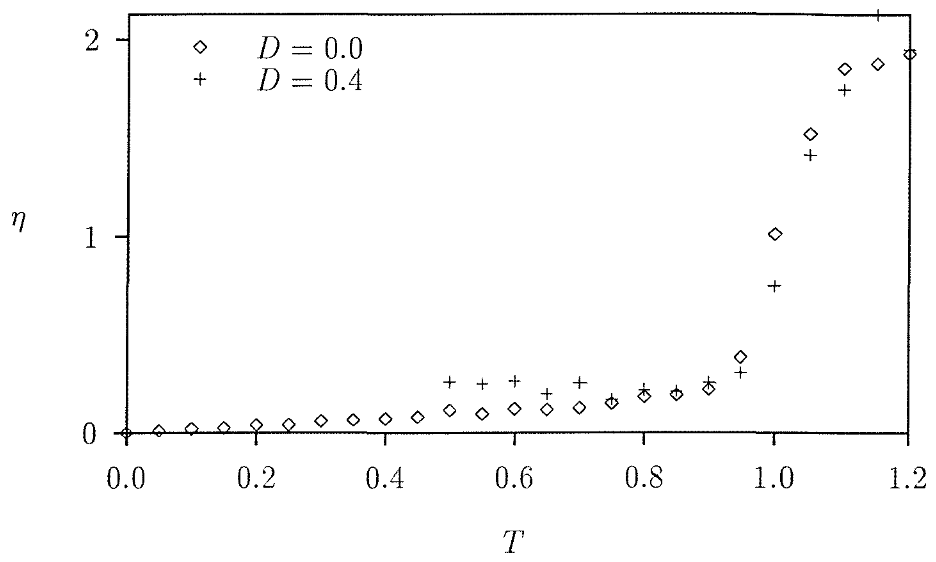

Figure 5.28: Comparing η from fit of $M \sim L^{-\eta/2}$ for $D = 0.0$ and $D = 0.4$. The horizontal lines show $\eta = 0.25$ which is the prediction of the pure system at T_{KT} (section 2.2.3), and $\eta = 2$ which is the expected behaviour for a paramagnet. No fit is shown for low temperature for $D = 0.4$ since the data clearly does not fit this form.

supposed that at these low temperatures, the system should scale as the number of domains, rather than number of spins — in this case, there are too few domains for the scaling behaviour to be determined, and simulations on much larger systems need to be performed. The shape of the low-temperature curves appears to be independent of temperature, so that the scaling curve can be separated into $f(T)g(L)$. At temperatures higher than $T \approx 0.4$, the scaling curves, though somewhat bumpy, appear to lie on approximately straight lines, implying re-emergence of pure-system $M \sim L^{-\eta/2}$ behaviour. The data does not conclusively show whether or not η is the same as the pure system — as shown in figure 5.28, at temperatures just below T_{KT} there is good agreement between the η from $D = 0.0$ and $D = 0.4$, while below $T \sim 0.7$, the values differ. Of course, while qualitatively, the data appears to lie on a straight line, determining the gradient of such noisy data is very error prone. At temperatures $T > T_{KT}$, the scaling behaviour is identical to the pure system.

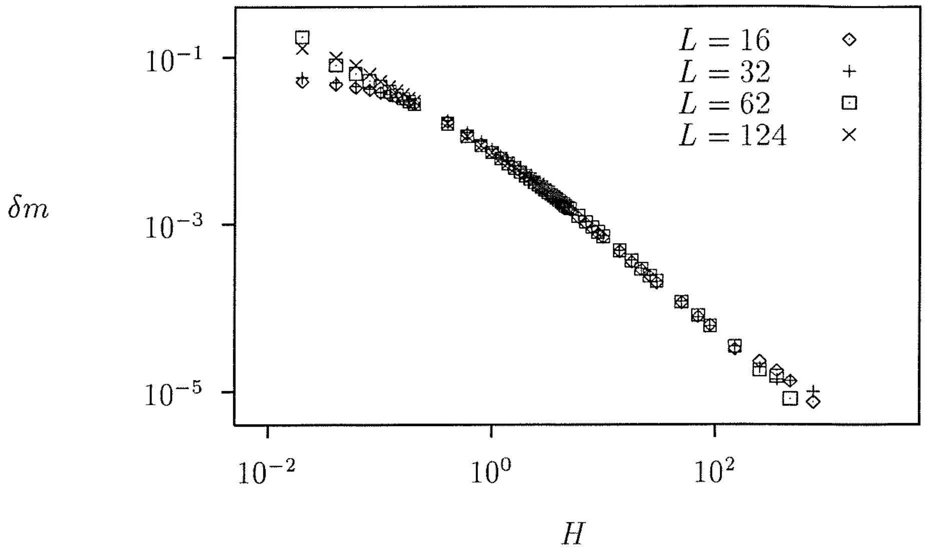

Figure 5.29: Deviation from saturation for various system sizes for $D = 0.5$ and $T = 0.01$, demonstrating system-size independence

5.4.4 Approach to Saturation

We have also made studies of the magnetisation as the system approaches saturation in large magnetic fields, to compare with the predictions made in section 2.8. All figures show deviation from saturation on a log-log plot, unless otherwise stated.

Figure 5.29 compares the approach to saturation for a number of system sizes with $D = 0.5$ and $T = 0.01$, demonstrating that, in the region of interest, the system size is unimportant. Thus all of the following is based on results from 16^2 systems, for computational ease.

The effect of thermal desaturation in an external field (section 2.8.9) can be demonstrated by plotting the deviation from saturation in the pure system, with $D = 0$ — as shown in figure 5.30, the magnetisation approaches 1 as $\frac{T}{H}$ for $H > 40$, as predicted by equation 2.94. In a similar way, a large anisotropy at low temperatures gives a desaturation $\delta m \sim D^2/H^2$ as shown in figure 5.31.

For $D < 5$, the approach is slower than H^{-2}, as can be seen from the figure: for $D = 0.5$, the approach is H^{-1}, which we interpret as FWA behaviour (section 2.8); for $D = 1$, the slope is around 1.4; the slope is 1.7 by $D = 2$ and has settled at 2 by $D = 5$.

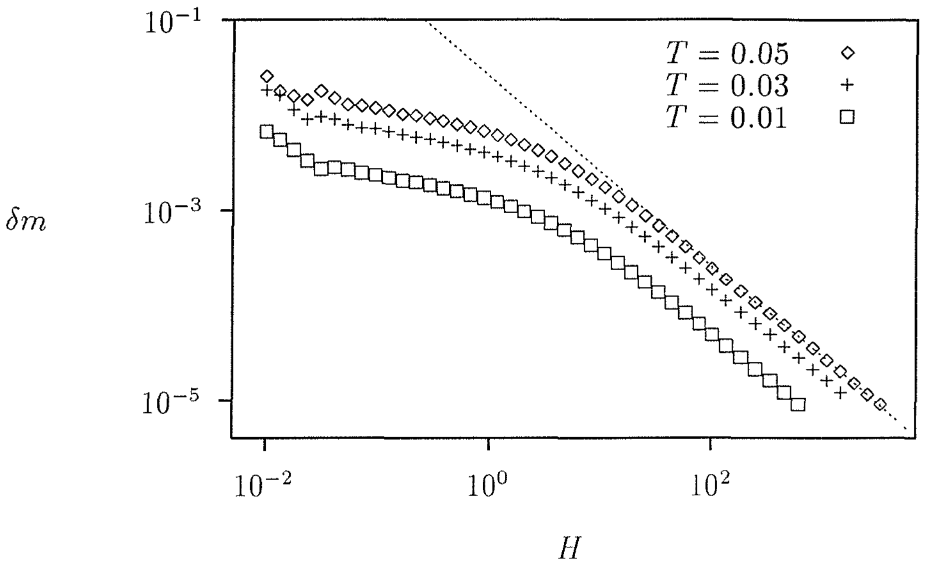

(a) As a function of H

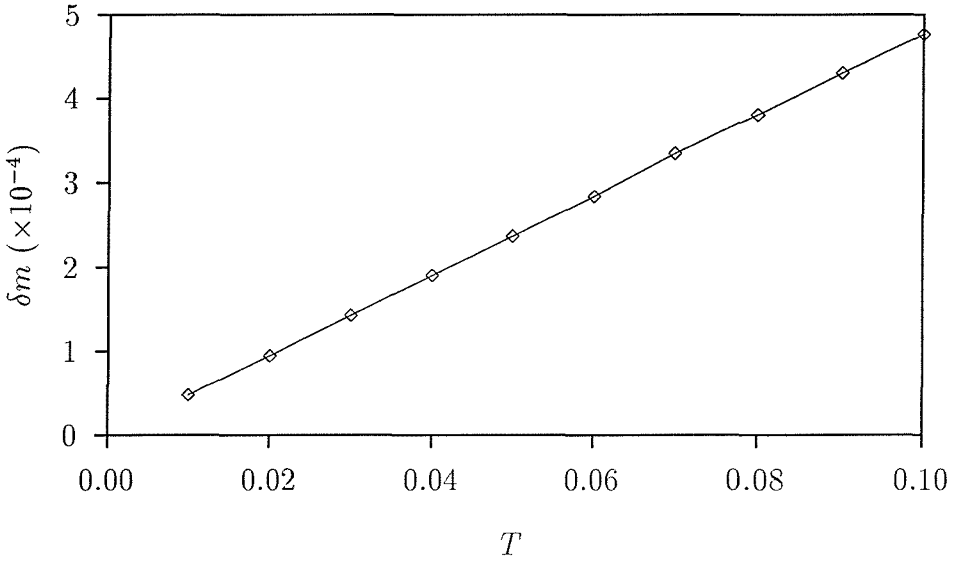

(b) As a function of T

Figure 5.30: The effect of temperature on high-field magnetisation of the pure system is to cause deviation from saturation (a) proportional to H^{-1} (dotted line) at sufficiently high fields, and (b) proportional to T (along $H = 100$).

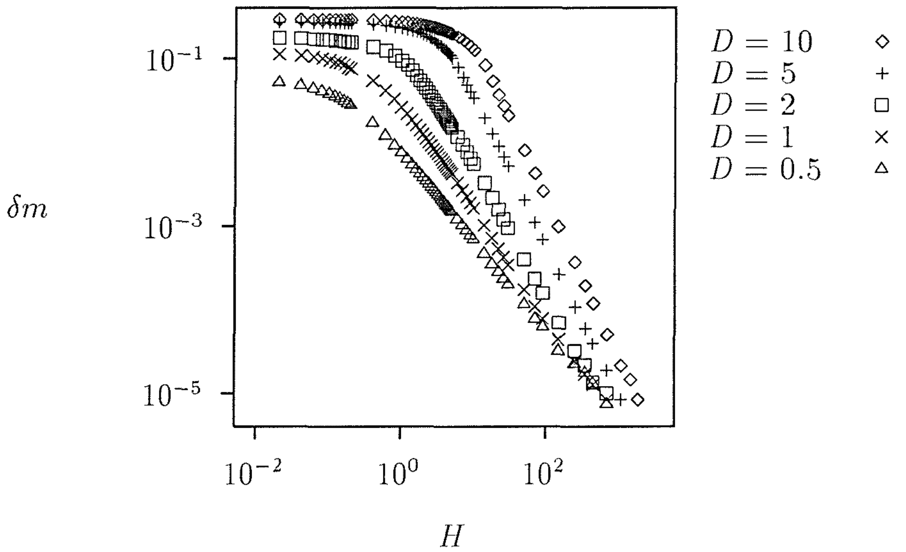

(a) As a function of H

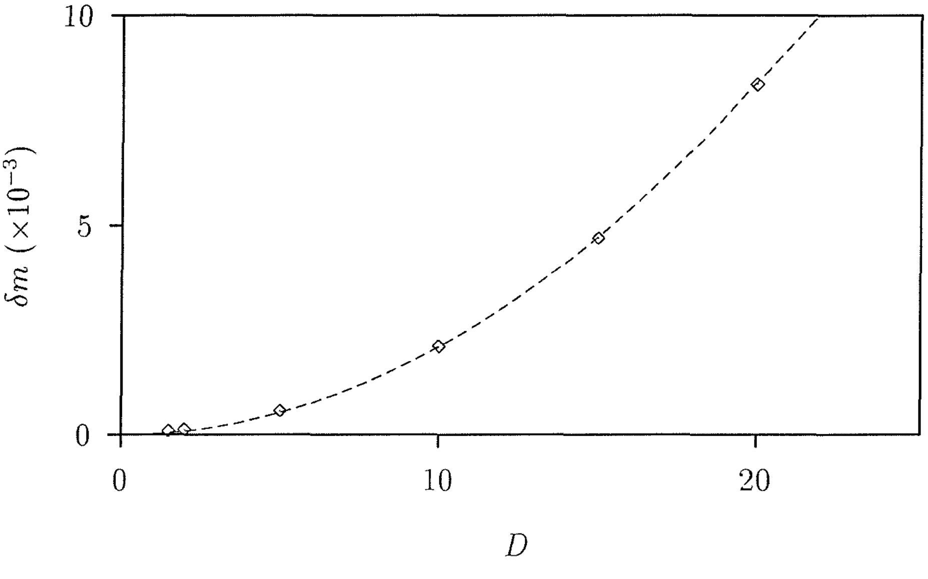

(b) As a function of D

Figure 5.31: The effect of (large) anisotropy on high-field magnetisation at low temperature and sufficiently high fields is to cause deviation from saturation (a) proportional to H^{-2} and (b) proportional to D^2 (along $H = 100$). The crossover to H^{-1} thermal demagnetisation can just be seen at very large fields and small D in (a). The D^2 behaviour is valid for $D > 5$, as described in the text.

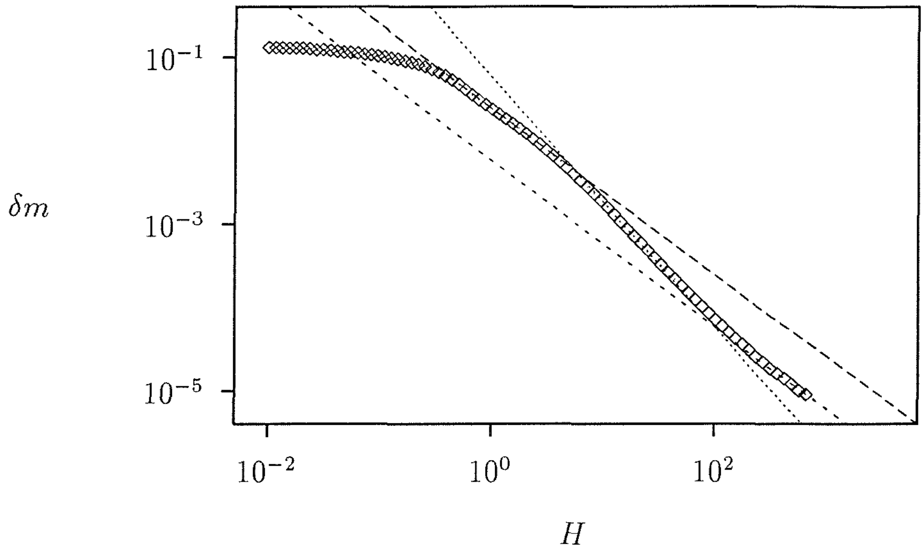

Figure 5.32: Demagnetisation in high field for $D = 1$ at $T = 0.01$, showing three régimes of δm vs. H behaviour

Figure 5.32 shows the demagnetisation for $D = 1$ and $T = 0.01$: there are three distinct high-field régimes as predicted in section 2.8, but the two H^{-1} regions are separated by a region with $\delta m \sim H^{-1.5}$, whereas the prediction was H^{-2} in equation 2.92. The low-field H^{-1} régime is interpreted to be the Ferromagnet with Wandering Axes (FWA) predicted in section 2.8.7, while the high field behaviour is thermal depinning as above. D is not sufficiently large for the intermediate régimeto show the expected H^{-2} behaviour, as described above.

It is almost possible to make out the changeover from H^{-2} to H^{-1} at high fields for small D in figure 5.31(a), but this crossover field can be picked out quantitatively by writing the deviation from saturation as the sum of the two contributions (equations 2.8.8 and 2.8.9):

$$\begin{aligned} \delta m &\sim \left(\frac{D}{H}\right)^2 + \frac{T}{H} && (5.15) \\ \delta m \frac{H^2}{T} &\sim \frac{D^2}{T} + H && (5.16) \end{aligned}$$

so that a graph of $\delta m\, H^2/T$ against H should give parallel lines of the form $H + H'(D,T)$, as shown on figure 5.32, and we might define the crossover to be where the two contributions are equal, i.e. where $H = H'$, by extrapolating back to $\delta m = 0$ at $-H'$, as shown on the figure. Because the lines are parallel, we can just use the y intercept at $H = 0$. The crossover field H' is plotted on a log-log scale against D in figure 5.34, showing the D^2 behaviour predicted in section 2.8.9

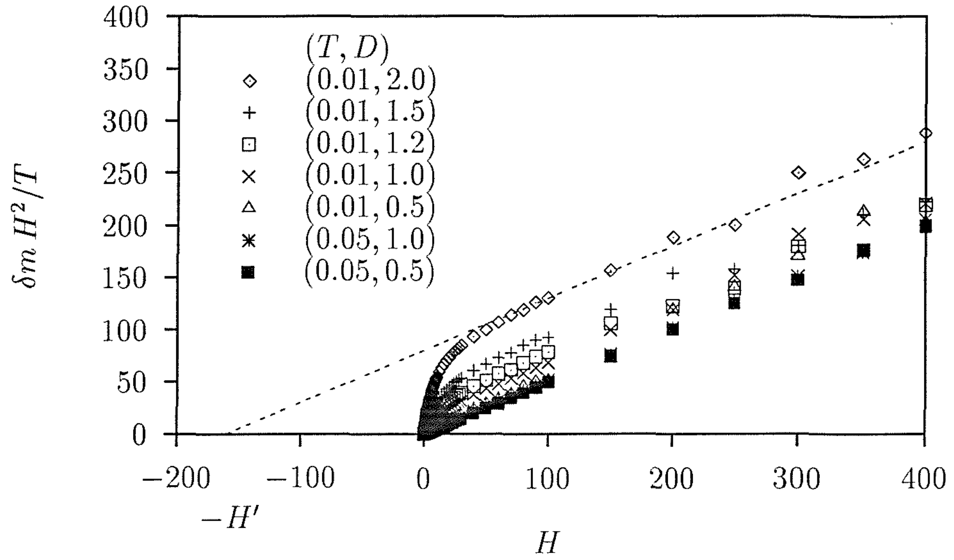

Figure 5.33: Graphical determination of the crossover between anisotropy- and thermal-induced demagnetisation - the crossover field is the x axis intercept H'

over a wide range of D. Figure 5.35 shows the inverse of the crossover field as a function of T for $D = 1$ — the intercept is rather small here, so that it is difficult to determine the intercept precisely, but the data supports a crossover field inversely proportional to T, as predicted.

5.5 Dynamic Effects

5.5.1 Role of Vortices

A great deal of insight into the nature of the low-temperature configuration has been afforded by the use of real-time graphics (section 4.6), allowing the progress of the system to be followed. The colour spin-plots in section 5.3.1 show how the choice of colours makes it easy to make out domains (spins with similar colour) and vortices (where abrupt colour-changes radiate out in all directions). In later simulations, code was added to give a quantitative measure the vortex density, as described in section 4.5.

For the cooled systems, the vortex density is plotted as a function of temperature for various weak anisotropy strengths D in figure 5.36. It is clear that there are no vortices present at low temperatures, and therefore that the pinning of vortices is *not* the mechanism behind the various changes in behaviour observed

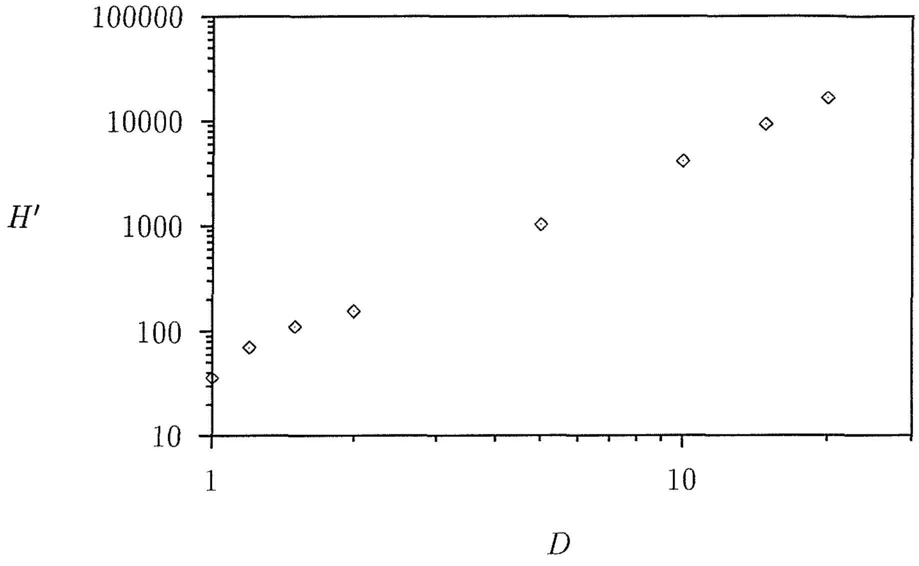

Figure 5.34: Crossover field H' as a function of D, using data at $T = 0.01$.

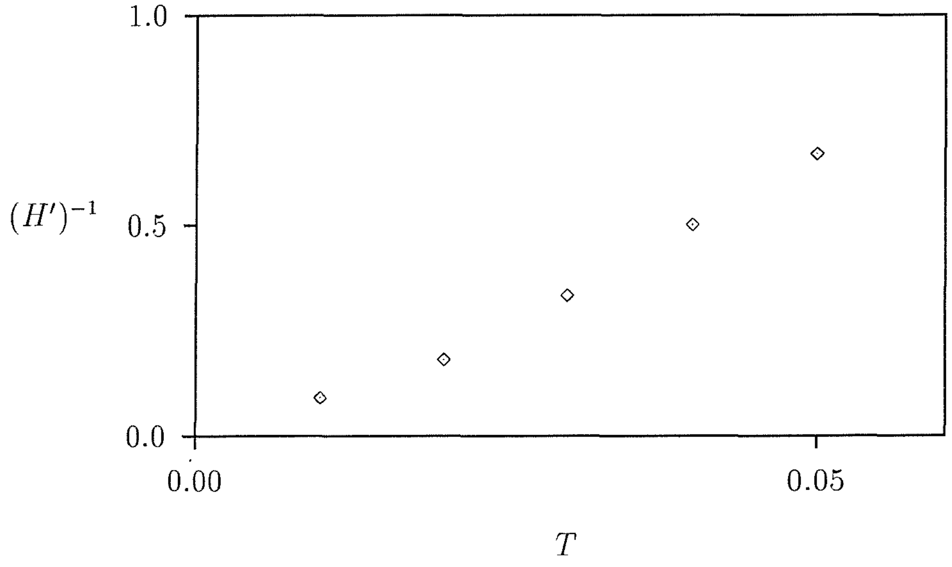

Figure 5.35: Crossover field H' as a function of T, using data at $D = 1$. The thermal desaturation above $T = 0.05$ overwhelms the effect due to the anisotropy, so that the intercept cannot reliably be distinguished from 0.

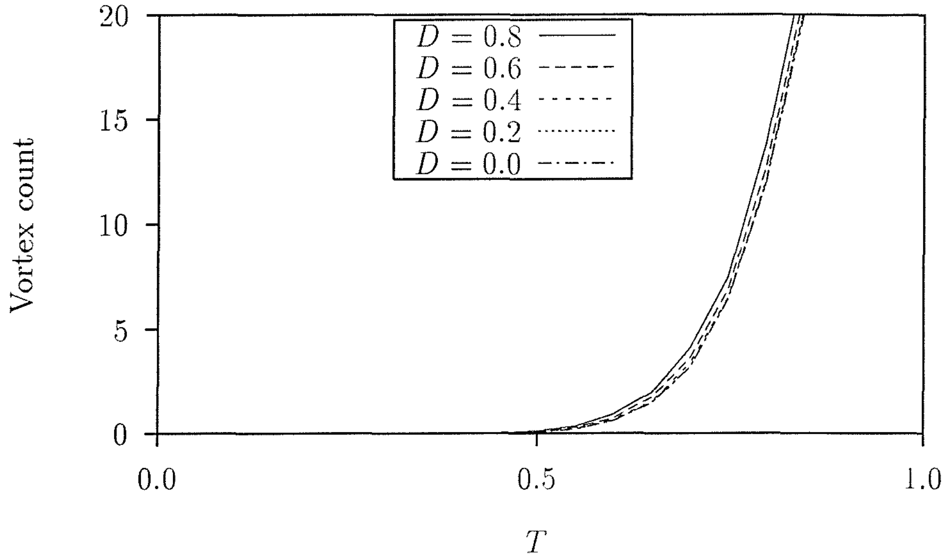

Figure 5.36: Vortex count versus temperature for various anisotropy strengths D for 62^2 system, showing that the vortices are abundant above the Kosterlitz-Thouless, but are absent below T_{KT}. This is further evidence that the transition from the paramagnetic to XY régimes is unaffected by the addition of weak anisotropy.

at low temperature.

On the other hand, figure 5.37 shows that if the system is started from a random (quenched) configuration, the vortices inevitably present in the initial configuration are unable to mutually annihilate at low temperatures; as was observed in the energy measurements (section 5.1), the system must be warmed to a degree dependent on the anisotropy strength D before the vortex density drops to zero. It is immediately apparent that vortices play a crucial part in the low temperature behaviour of quenched systems. It is also clear that the vortex density is not a well-defined variable; as before (section 5.1) it depends crucially on the starting configuration and the layout of the anisotropy axes for these finite systems; for infinite systems, as described in section 2.9, we would expect all possible local configurations of the random axes to appear, so that this sample to sample variation would be reduced.

It was noted in section 5.2 that the liquid crystal order parameter is higher than expected in quenched systems. This indicates that the spins are more closely correlated with the local anisotropy axes than for cooled systems. This supports the argument proposed in section 2.9 whereby the vortices are 'pinned' by particular configurations of the easy axes.

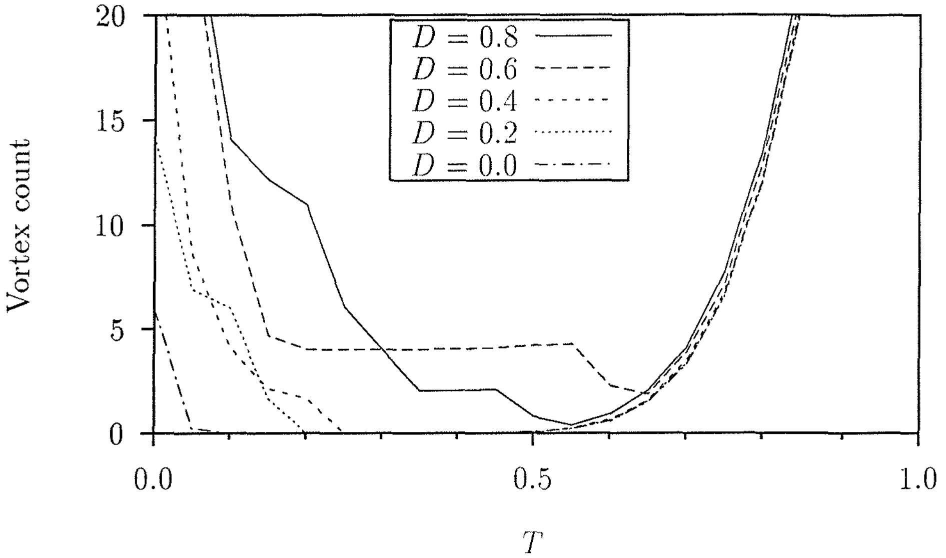

Figure 5.37: Number of vortices in quenched 62^2 system for various D

Because it is in the nature of vortices to annihilate in pairs at low temperature, it is reasonable to suppose that in a system started in a configuration containing vortices, the vortex count will decay at low temperatures; given that the vortex pinning centres discussed in section 2.9 have an energy of order D, the probability of a vortex escaping from a pinning centre might be $\sim \exp -\frac{D}{T}$ in unit time, giving a mean lifetime growing like $\exp \frac{D}{T}$.

Figure 5.38 shows the average time taken for all the vortices to annihilate from ten random starting configurations. While the number of vortices in the initial configurations were reasonably consistent (within 10%), there were very large fluctuations (factors of 10 and more) in the measured vortex lifetimes. It was necessary to cap each simulation at 100000 cycles, to avoid wasting cpu time on pathological cases; this makes it rather difficult to fit quantitatively. The trends, however, are clear; the lifetime decreases with increasing temperature, and increases with increasing anisotropy strength: and for temperatures $T > D$, the measured lifetime matches roughly the baseline for $D = 0$ (to take simulation dynamics into account).

5.5.2 Hysteresis Loops

Strictly speaking, hysteresis loops are a purely non-equilibrium phenomenon — if the system was always allowed to relax fully after each change in the external

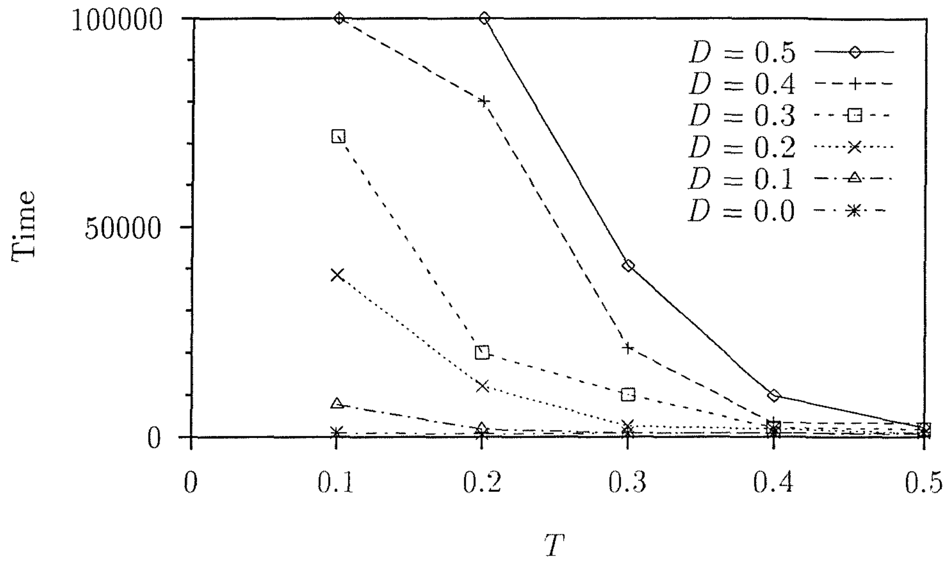

Figure 5.38: Vortex lifetime as a function of temperature and anisotropy strength. The simulations were truncated at 100,000 cycles, so that the times are suppressed slightly at the top end.

field, the effect would be gone. On the other hand, as described in section 1.1.8, the experimental systems which inspired this work do show hysteresis, because of the slow relaxation times.

In section 1.1.6, the hysteresis was accounted for in terms of domains, which were in turn ascribed to the dipolar interaction. In this system, we have neglected the dipolar terms from the Hamiltonian, yet we still find hysteresis — Figure 5.39 shows the sort of hysteresis loop produced by our simulations — because of the difficulties in starting a system from equilibrium in zero field with zero magnetisation, as noted in section 5.4, we have not attempted to find an initial magnetisation curve, as done in [52]. Instead, we start the spins from a fully aligned state in a large field, then slowly reduce and then reverse the field. In general, we perform the sweep on only one direction, since the system is invariant under a 180° rotation, so that following the reverse curve is identical to running a second sweep in the increasing direction — the curve is reversed and superposed to form a conventional hysteresis loop as required. Shown on figure 5.39 are various important quantities: the remanent magnetisation M_r is the magnetisation at zero field; the coercive field H_{co} is that field at which the magnetisation changes direction; and the *reversible field* H_{rev} is the field above which the hysteresis ends — it is the field above which the system exhibits reversibility.

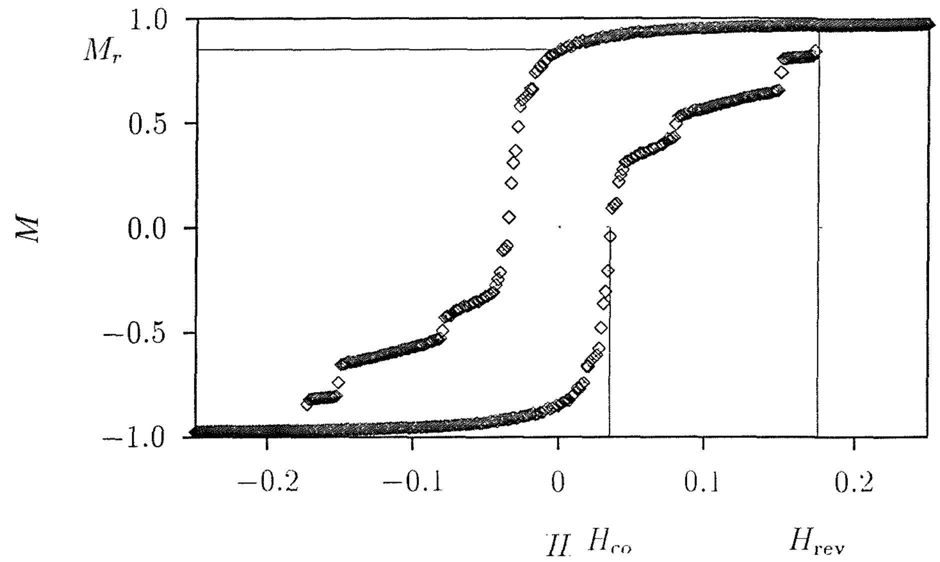

Figure 5.39: Typical hysteresis loop at low temperature for 62^2 at $D = 0.4$

The figure shows irreproducible jumps in the magnetisation at small reverse fields, bearing considerable resemblance to features seen on real systems (section 1.1.8). The long, bumpy tail at fields larger than H_{co} is not always present on the measured curves, but does appear sufficiently often to be significant — this tail is very similar to those seen by Dieny and Barbara [52] in simulations of this model at zero temperature, where the jumps in the tail are identified as due to the annihilation of vortices — we shall return to this below.

Because hysteresis curves are a dynamic effect, it is necessary to take care over the rate at which the magnetic field is swept, to ensure that the curves have converged — Saslow and Koon [53] show that earlier work by Chi and Alben (which suggested that H_{co} did not rise monotonically with D) was flawed by lack of convergence.

The value of the remanent magnetisation M_r appeared to change little with sweep rate (see section 5.5.3), so we ran the system from negative field to zero field once, storing this final configuration (shown in figure 5.40), then ran further simulations starting from this stored configuration for several different sweep rates. At each field value, 500 equilibration cycles and 2500 production cycles were performed, with each cycle consisting 4 rotation and 1 flip updates per spin. The field step was varied from 0.001 to 0.128 for each T and D.

Figure 5.41 shows the sequence of curves for a 62^2 system at $D = 0.4$ and

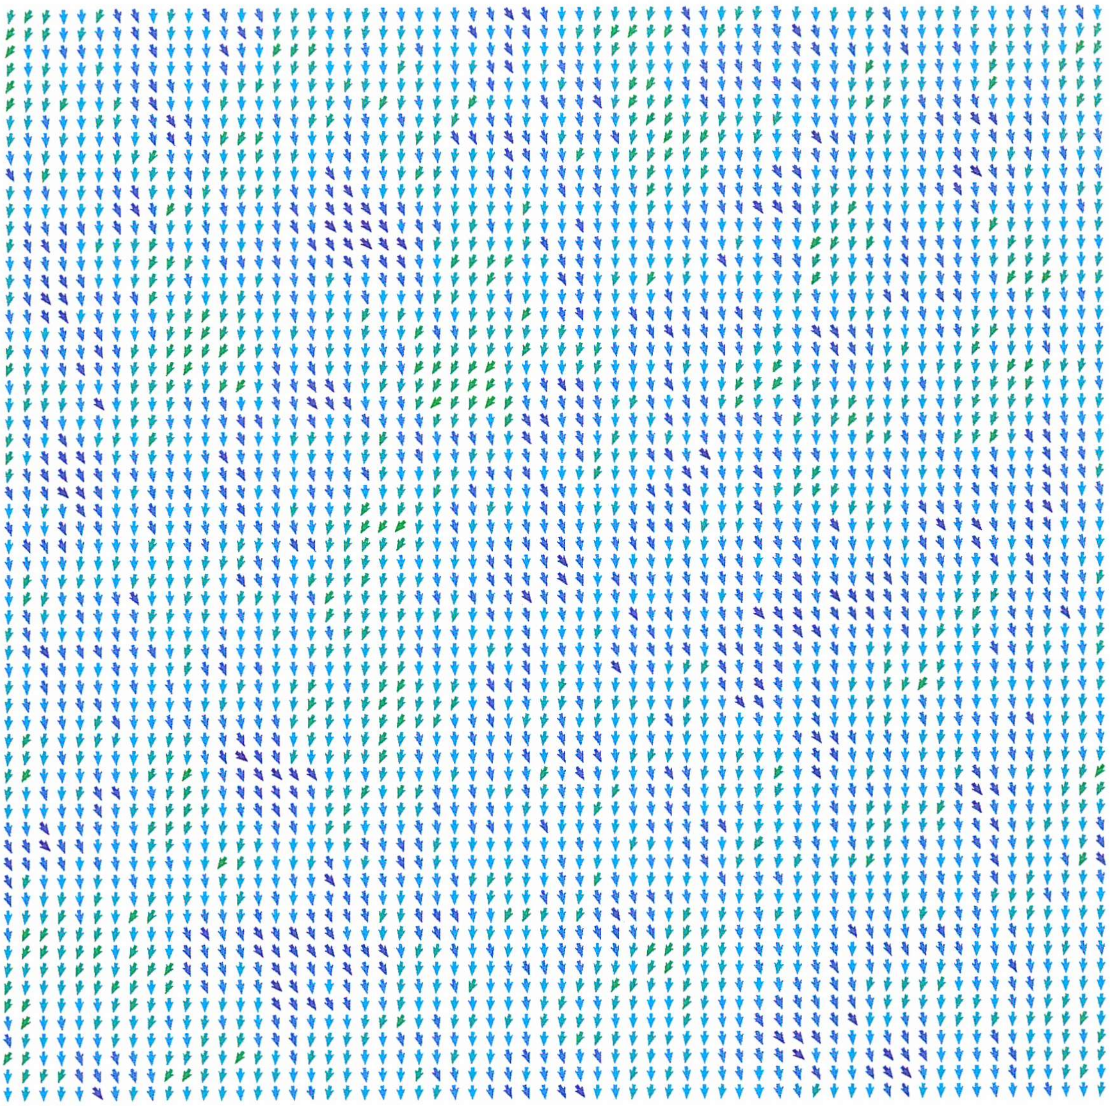

Figure 5.40: Instantaneous configuration of remanent magnetisation state for $D = 0.4$ and $T = 0.1$, showing small domains. Because of the small system size and long relaxation times, the system is still largely aligned in the direction of the initial field.

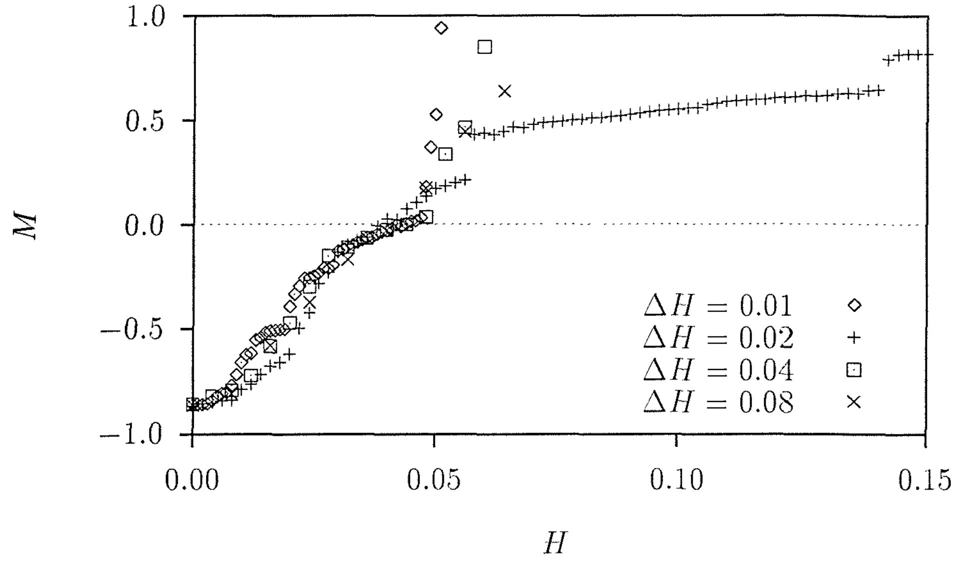

Figure 5.41: Test of convergence of hysteresis loops — hysteresis loops for various field sweep-rates as described in the text

$T = 0.1$ — although the details of the path from 0 to $H_{\rm co}$ changes, there is no doubt that the value of $H_{\rm co}$ is reasonably converged, and the configuration at $H_{\rm co}$ is very similar for different sweep rates: one such configuration is shown in figure 5.42. The energy of this configuration is approximately the same as for the cooled systems described in section 5.1, and lower than those of the quenched systems at this temperature, as before. There appear to be no vortices or half-vortices in this configuration. The differences in the path to $H_{\rm co}$ are taken to be characteristics of different runs over the system, with different random numbers, rather than features due to the differing step rate. The jumps may be due to local rearrangements of the breakup into domains (see section 5.3).

On the other hand, $H_{\rm rev}$ is clearly *not* repeatable, and varies widely from run to run. As was found by Barbara and Dieny at zero temperature [52], the way in which the domains collapse in the reversed field has a significant effect on the subsequent behaviour, in the same was as the starting configuration in a quenched system has a profound effect on the development of the system (section 5.1). The previous work at zero temperature found that the jumps in the tail were due to the successive elimination of the vortices formed by initial quench (for the initial magnetisation curve) or the domain reversals. While we find that topological defects play a part, the simulation shows a small number of *half*-vortices, increasing with field, but no significant effect from full vortices — this

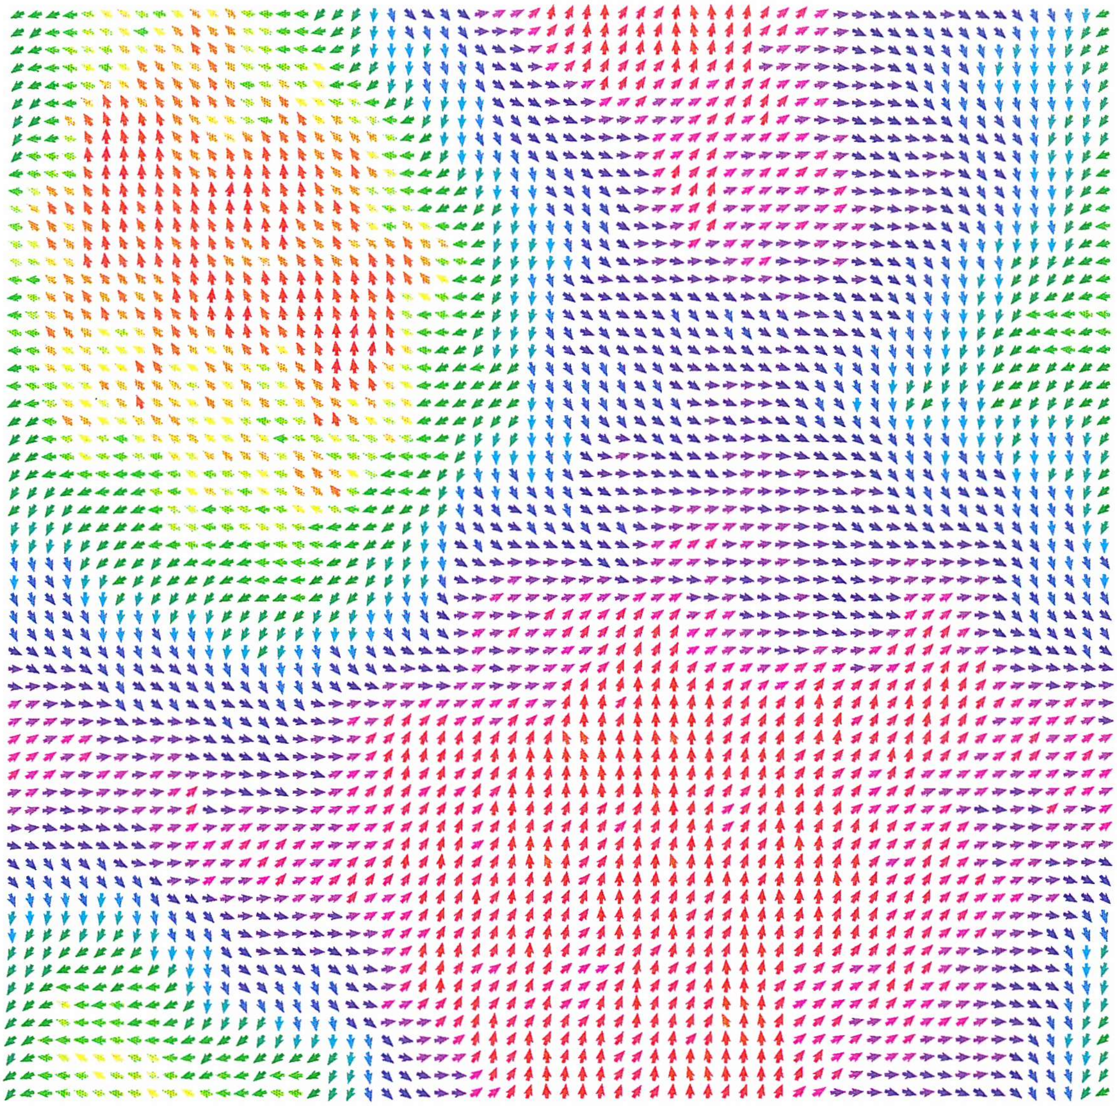

Figure 5.42: Instantaneous configuration at $H = H_{\mathrm{co}}$ for $D = 0.4$ and $T = 0.1$. The domains have grown in the reverse field, and are consistent in size with the domain size found in section 5.3

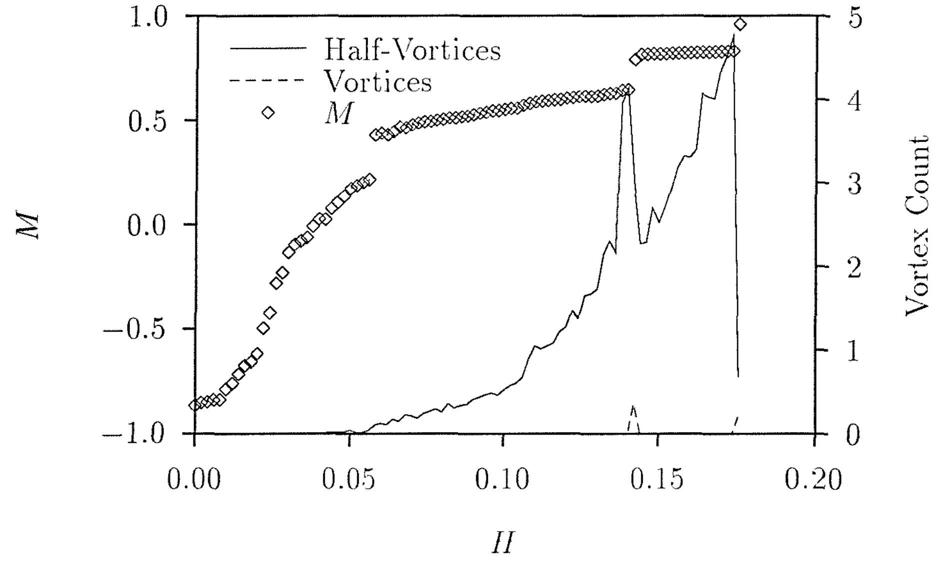

Figure 5.43: Role of vortices and half vortices in hysteresis curve. The half-vortices indicate an increasing stress in the system, which is resolved in small jumps in the magnetisation (with the brief appearance of full vortices).

may just be a feature of the way the code detects the vortices (section 4.5) or may be due to the effects of non-zero temperature on the vortices (so that the code does not recognise them as vortices. The presence of these half-vortices implies sharp domain walls, as depicted on figure 2.2, and an increasing strain in the system.

Coercive Field

While the reversible field H_{rev} above is probably of more interest, the coercive field H_{co} is somewhat easier to measure in practice: results are much more easily reproduced, as shown above; and it is easier to detect when the coercive field has been reached from within a simulation — a simple test using the primitives in section A.3 goes like

```
repeat
   ...
   getstate   # gets current state into system variables
until X > 0
```

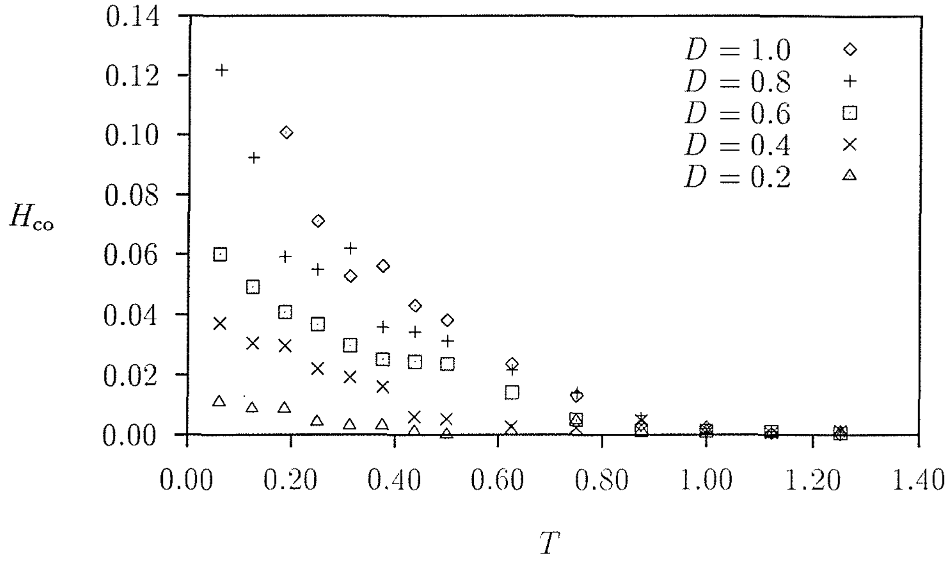

Figure 5.44: Coercive field H_{co} as a function of temperature T for different anisotropy strengths D. There appear to be drops in the coercive field strength at D-dependent temperatures.

whereas it is much harder to predict beforehand what magnetisation is required to be sure that the reversible state has been reached.

Figure 5.44 shows the coercive field H_{co} collected from simulations at various temperatures and anisotropy strengths. Qualitatively, the curves are as expected from previous results: the hysteresis loops get narrower as the temperature is increased, but gets wider as the anisotropy strength is increased. ([39] found the unexpected result that H_{co} did not increase monotonically with D for simulations of three-component spins in three dimensions, but this result was later refuted by [53].)

All hysteresis ceases above $T = T_{KT}$ for the anisotropy strengths studied as expected, since other evidence is that for small anisotropy, the system is XY-like at intermediate temperatures. For sufficiently small anisotropy strengths, anisotropy ceases at still-lower temperatures. On the curves for $D = 0.2$, $D = 0.4$ and $D = 0.6$, there appears to be a definite jump in $H_{\mathrm{co}}(T)$, which could be an indication of a change in behaviour similar to that found earlier.

Qualitatively, the curves can be extrapolated back to $T = 0$ to give approximate agreement with the zero-temperature simulations of [52].

5.5.3 Relaxation from Remanent State

We have run long simulations starting both from an aligned state (roughly equivalent to the remanent magnetisation state in section 5.5.2) and from a random starting configuration. During these simulations, we track both magnitude and direction of the net magnetisation (bearing in mind that this is a finite size effect — section 5.4), and we record the spin glass order parameter (section 2.6); both measurements afford a measure of the rate at which the system moves through phase space. For the first 1000 cycles (with each cycle comprising 10 rotation attempts per spin — section 4.3), the generated configurations are compared with the starting state; the 1000th configuration is then stored, so that subsequent spin glass ordering is compared with this intermediate result; essentially, we are allowing 1000 cycles for equilibration, but are recording this information rather than discarding it as usual.

Measurements of the direction of the net magnetisation, starting from an aligned state, are shown in figure 5.45, with the angle of magnetisation in the left figure and a probability distribution (histogram) on the right. At low temperatures, the system moves to and remains at the nearest local minimum; when the temperature is raised, it is able to find what appears to be a 'global' minimum direction (since the natural domain size found in section 5.3 is of the order of the system size. At higher temperatures still, it is able to flip between the two equivalent orientations along this minimum — it is possible in principle to use the time between flips as a measure of the system relaxation time, but we have been unable to fit this data quantitatively. Even at high temperatures, the system shows a marked preference for the easy direction. At low temperatures, the magnitude of the magnetisation is very slightly higher than the earlier results where the system is cooled from high temperature, which is as expected.

For a quenched (random) starting state, it is the magnitude of the magnetisation that is of interest — we are interested in the timescales required for the magnetisation of the quenched state to 'recover' (section 5.4.1). Figure 5.46 compares the magnetisation as a function of time for two temperatures for a 62^2 system with $D = 0.4$. With $T = 0.3$, the system rapidly finds one metastable state, where it remains for what is in fact a considerable amount of Monte Carlo time, before jumping to a region of phase space with magnetisation akin to the cooled system. Almost inevitably, the jump coincides with the annihilation of a pair of vortices, as illustrated on the bottom curve in the figure.

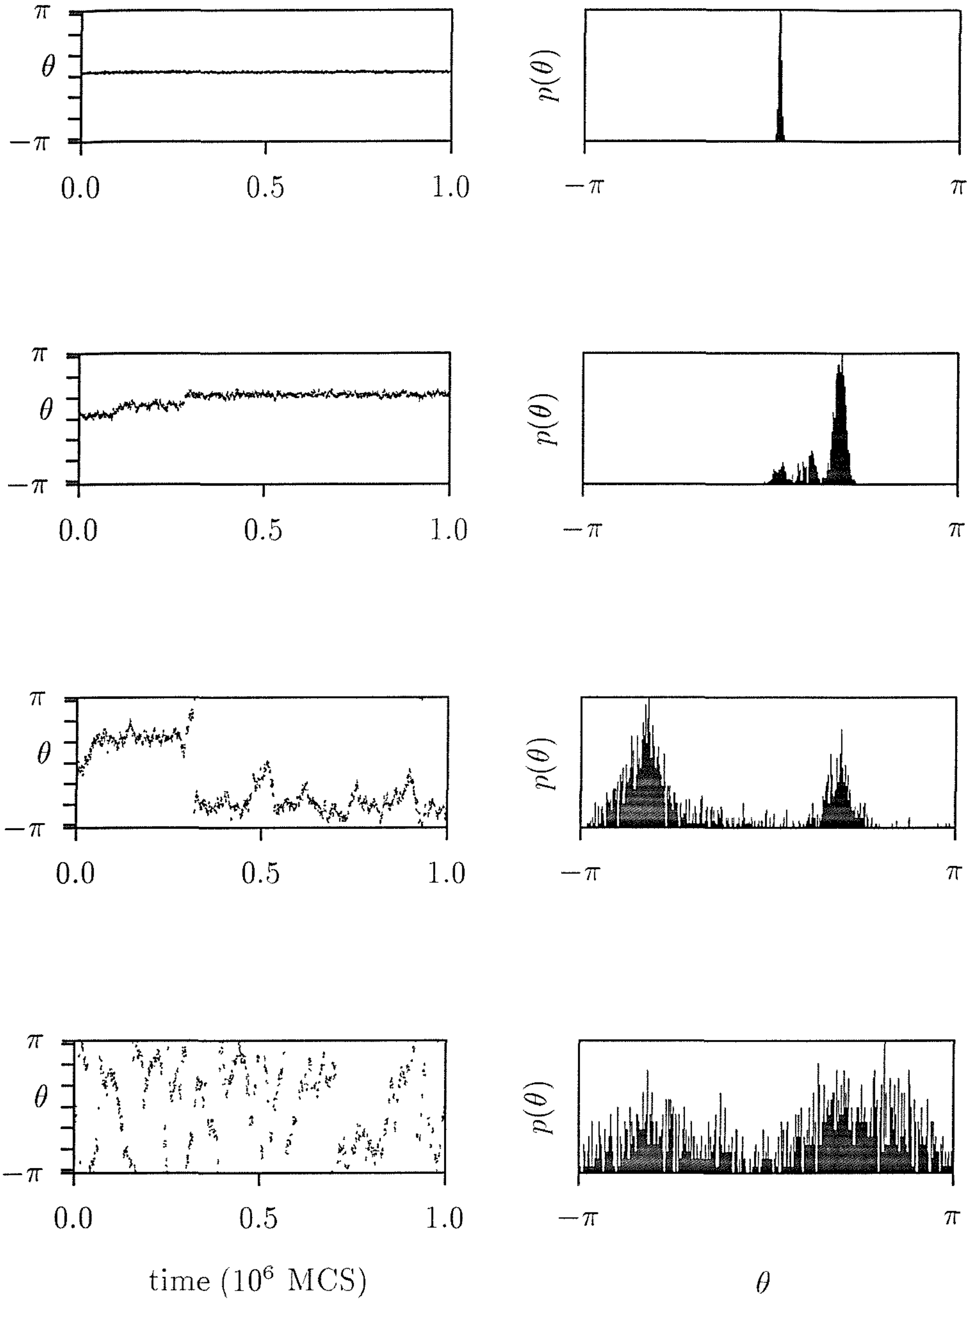

Figure 5.45: Angle of magnetisation with time (left) and a probability distribution (right) for 62^2 system with $D = 0.4$, for temperatures (top to bottom) 0.1, 0.4, 0.8, 1.0. At low temperature, the system finds nearest local minimum and stays there; at slightly higher temperatures, the system is able to find a deeper energy minimum. At higher temperatures still, the system flips between the two equivalent directions. Even at high temperatures, the probability distribution shows that there is still a tendency to favour these preferred directions.

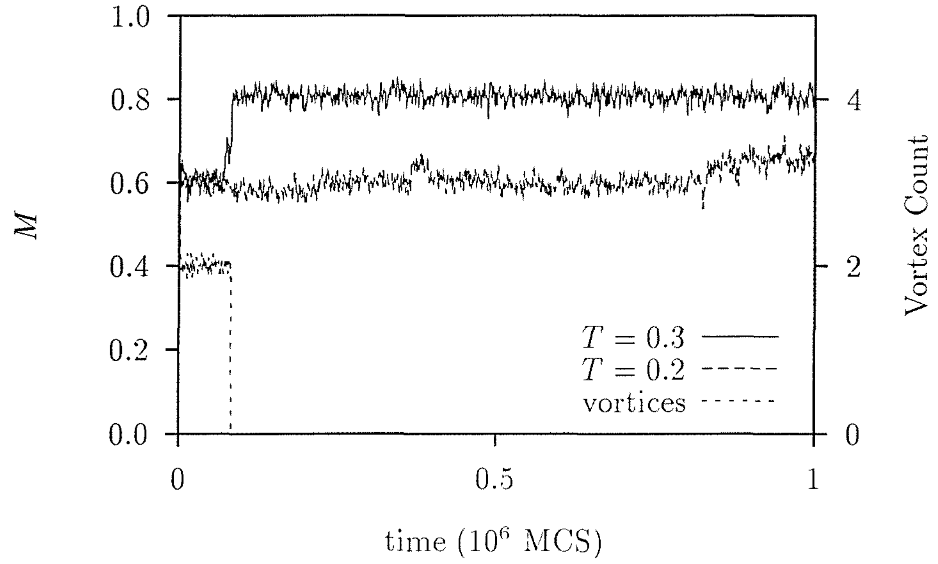

Figure 5.46: Magnetisation as a function of time from random starting state. The jump in the curve for $T = 0.3$ corresponds to the annihilation of a pair of vortices as shown in the lower curve.

At a lower temperature, the system finds what appears to be the same metastable state, but spends the whole of the substantial simulation without being able to reach the vicinity of the ground state (whose magnetisation is slightly higher than the final state of the $T = 0.3$ configuration). The magnetisation is certainly rising with time — even longer simulations would eventually allow the system to shake off the last pair of vortices (not shown).

Measurements of the spin glass order parameter (section 4.5) tell the same story: the system moves rapidly away from the initial random configuration, then appears to be trapped in a metastable state for long periods at low temperature. At higher temperatures, the system fluctuates more rapidly. Figure 5.47 shows the spin glass order parameter measurements as a function of time at different temperatures. The order parameter was measured with respect to the initial (random) configuration for the first 1000 sweeps, then the 1000th configuration is stored, and subsequent measurements are with respect to that configuration. If a later configuration is stored, the low temperature measurements reveal the system to be stuck in a metastable state, whereas the higher temperature curve again moves rapidly away, revealing that it is sampling a fairly large volume of phase space. We have been unable to make a quantitative analysis of the spin glass order parameter results.

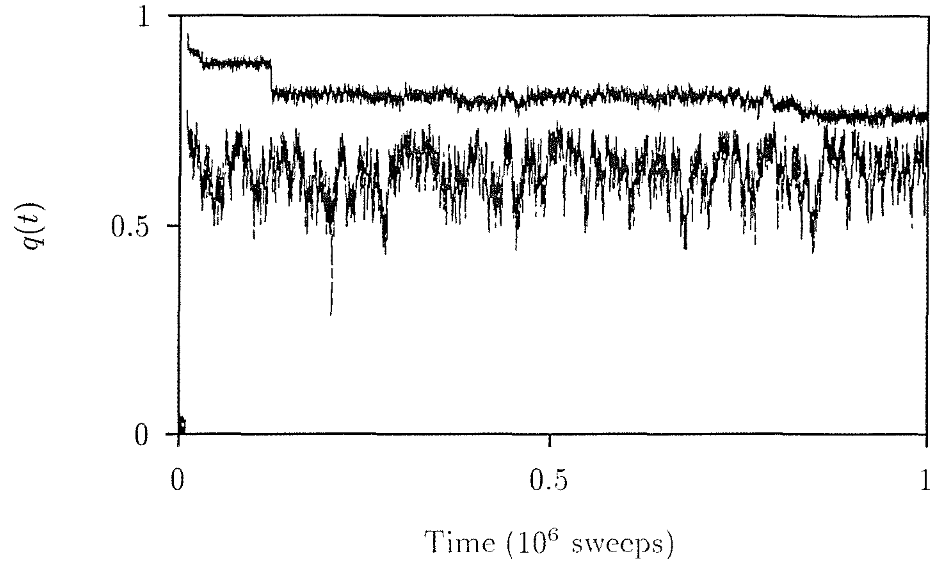

Figure 5.47: Spin Glass Order Parameter from a quenched start. The system rapidly moves away from the initial configuration (time up to 1000 sweeps) for both curves. For the upper curve (low temperature), the system appears stuck in a metastable state. The lower curve at higher temperature is sampling a larger volume of phase space; redefining the order parameter relative to a later configuration shows that the system is not stuck in a metastable state.

5.6 Summary

Our simulations show that, below a certain T^* which depends on D, the system shows irreversibility — the system parameters such as energy, magnetisation and liquid crystal order parameter have different values for cooled and quenched systems, and the spins show different correlation functions. Vortices appear to play an important part in this, as shown in section 5.5.1, since the graphs of the cooled and quenched quantities meet as the last vortices are annihilated from the quenched systems.

While we find this irreversibility for all system sizes, the quantitative details do depend on L, and for small systems the irreversible behaviour ceases at lower temperatures. This is as expected, since a Monte Carlo algorithm can explore a small system much more carefully, and a small system has much smaller energy barriers.

It is clear, however, that the differences in behaviour are not *entirely* due to the vortices. The scaling of the magnetisation and the spin correlations for the cooled systems indicate a marked change in behaviour at around the same temperature,

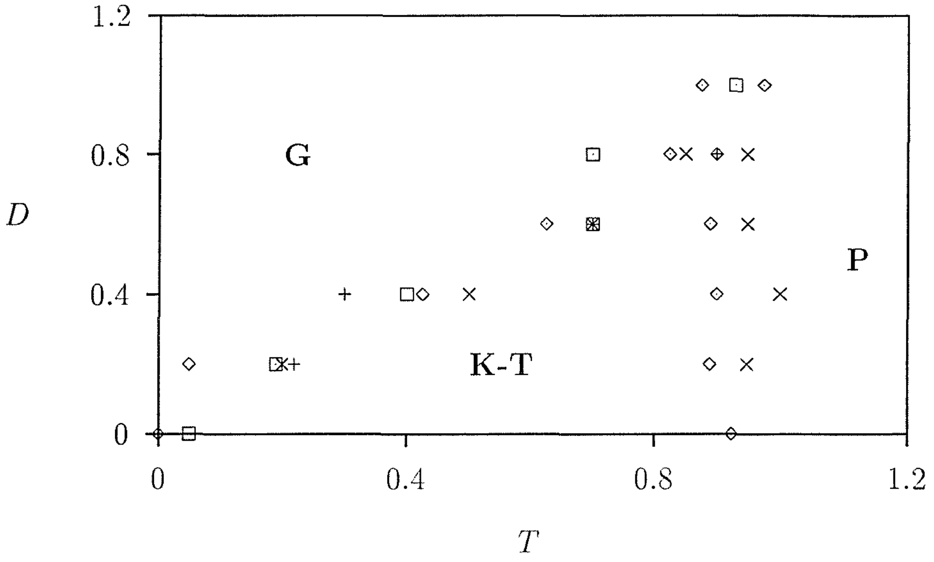

Figure 5.48: Tentative phase diagram for the two-dimensional XY model with small random anisotropy, collecting together measurements from cooled correlations ($\diamond$), quenched correlations ($+$), cooled and quenched energy and magnetisation ($\square$) and results from finite size scaling of magnetisation ($\times$). **P** represents the paramagnetic phase, **KT** is the Kosterlitz-Thouless régime, and **G** denotes the part of phase space which we call glassy.

with a change from Kosterlitz-Thouless behaviour to a distinct low-temperature régime. There are no vortices present in the cooled system to account for this change.

We collect together the various values of T^* from the previous sections to form a tentative phase diagram (figure 5.48):

- one measure of T^* is the temperature at which the values of the energy, magnetisation and liquid crystal order parameters coincide for cooled and quenched experiments. The cooled systems always have lower energy, implying that the quenched systems are trapped in local minima until they gain enough thermal energy to escape (c.f. section 2.8.4).
- It is possible to identify a region at intermediate temperature in which the spin-correlation function is algebraic, like the Kosterlitz-Thouless phase of the pure system. There is a fairly clearly marked change to a different (unidentified) régime at low temperature.
- The scaling behaviour of the magnetisation gives another indicator of T^*

as the system changes from power-law decay to a low temperature scaling form.

- At low temperatures in quenched systems, the spins decorrelate over much smaller distances than for systems cooled to the same temperature. The temperature at which the correlations become algebraic give an independent measure for T^*.

It is convenient to defer further discussion until the three-dimensional results have been presented, so that the two systems can be reviewed in parallel.

Chapter 6

Three Dimensional XY Model with Random Anisotropy

Presented here are initial results from simulations on the three-dimensional XY model with random anisotropy. Because the number of spins is very large for even rather modest system dimensions, we have not been able to perform such extensive simulations over as wide a range of the system parameters as we were able to cover for the system in two dimensions in chapter 5. In order to make even a 31^3 system run in a reasonable time, many sacrifices in terms of memory usage had to made (section 4.4.2), so that 31^3 was in fact the largest system we were able to study with the limited storage available on the main hardware — we were able to make brief use of a bigger parallel computer, with fewer compute nodes, but more memory per node, to simulate slightly larger systems; this time, the computer power, rather than the memory was the limiting factor.

Most of the following corresponds closely to the equivalent sections in the previous chapter. The most obvious difference between the results in two and three dimensions is that much larger values of D are required to make a significant effect on the system — because of the higher dimensionality, more paths exist between nearby spins, and the exchange interaction has a more significant effect. In the pure system, the exchange interaction in three dimensions is sufficient to produce long-range order at low temperatures, whereas in two dimensions, the exchange can only sustain quasi-ferromagnetism — the system *almost* orders. It should be easier to detect the destruction of true long-range order in three dimensions as opposed to the two dimensional quasi-ferromagnetism. Also, the effect of vortices in three dimensions is expected to be much less significant —

while vortices can still form, they take the form of vortex strings, as described in section 2.2.3.

As before, we contrast results from cooled and quenched systems in zero field, and we study energy (section 6.1), liquid crystal order parameter (section 6.2), correlations (section 6.3) and magnetisation (section 6.4), in order to investigate the predictions of the Chudnovsky-Imry-Ma theory in section 2.8.

6.1 Energy

Like the two dimensional system (section 5.1), the shape of the energy vs. temperature curve changes very little with D at small D, suffering only a shift dependent on D but independent of temperature, as shown in figure 6.1(a). Accordingly, the curves for the specific heat are superposed in figure 6.1(b).

For larger values of D, differences appear in the shape of $u(T)$, as shown in figure 6.2(a), and consequently in the specific heat curves, as shown in figure 6.2(b). The energy has been transformed to Hamiltonian 2.55, which is more appropriate for large anisotropy strengths (section 2.8.3), and a factor of $0.5k_BT$ has been added to the specific heat for $D = \infty$, to compensate for the loss of one degree of freedom per spin as before.

Here, the specific heat curves broaden and move to lower temperature, as expected and unlike the two-dimensional system. The curve for $D = \infty$ has a broad peak centred around $T = 1.95J$, in good agreement with the description in [47], though there was no graph shown in that work, so detailed comparison cannot be made. The energy of the ground state is in excellent agreement with the results presented in [47], despite the general criticism in that paper of the single-spin update algorithm used here.

6.1.1 Scaling

Because the peak in the (pure) XY model in three dimensions represents a genuine phase transition (whereas the peak in the specific heat for two dimensions occurs above the phase transition, as discussed in section 2.2.3), we can attempt to fit a scaling form to the data. $C(T, D)$ is taken to be of the form

$$C(T, D) = C_0 + D^{-\lambda} f\left((T - T_c)D^{-\mu}\right) \qquad (6.1)$$

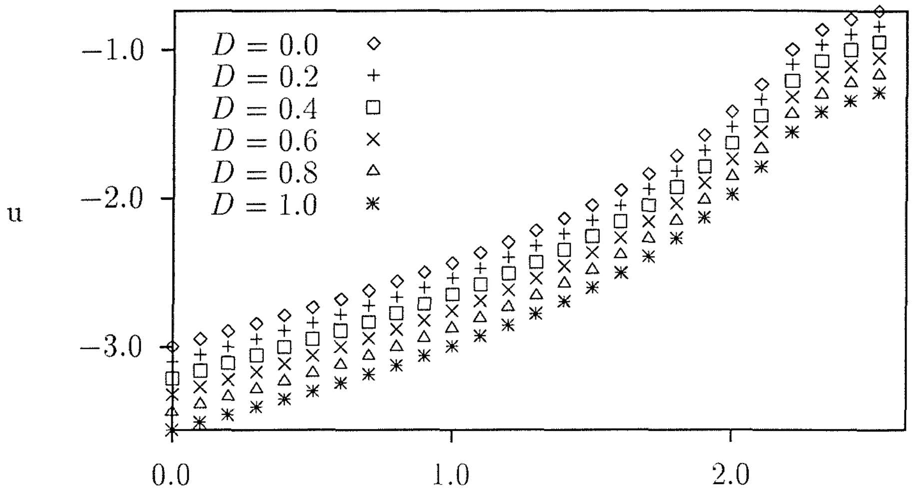

(a) Energy

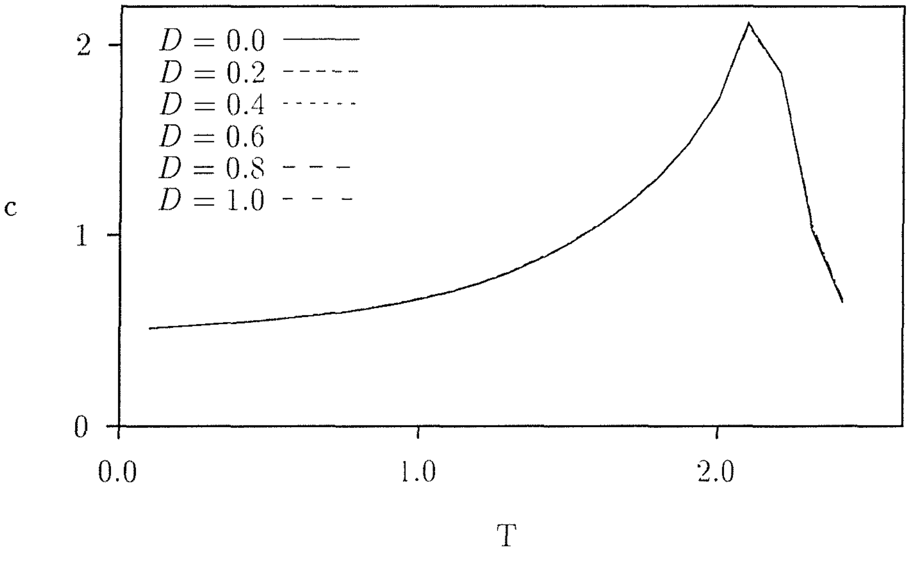

(b) Specific Heat

Figure 6.1: Energy (a) and Specific Heat (b) against temperature for cooled 31^3 systems with various small anisotropy strengths D. As before, the shape of the energy curve is unchanged with D, so that the specific heat curves are superposed.

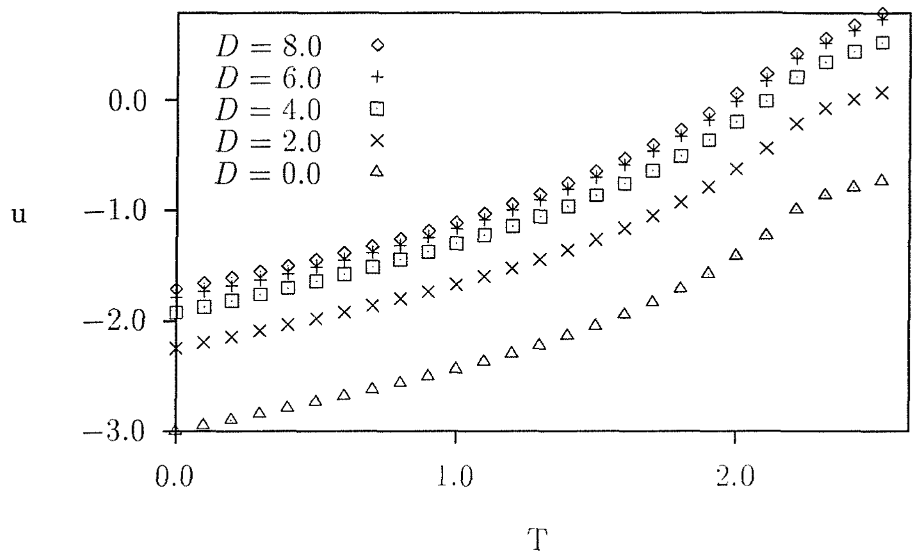

(a) Energy

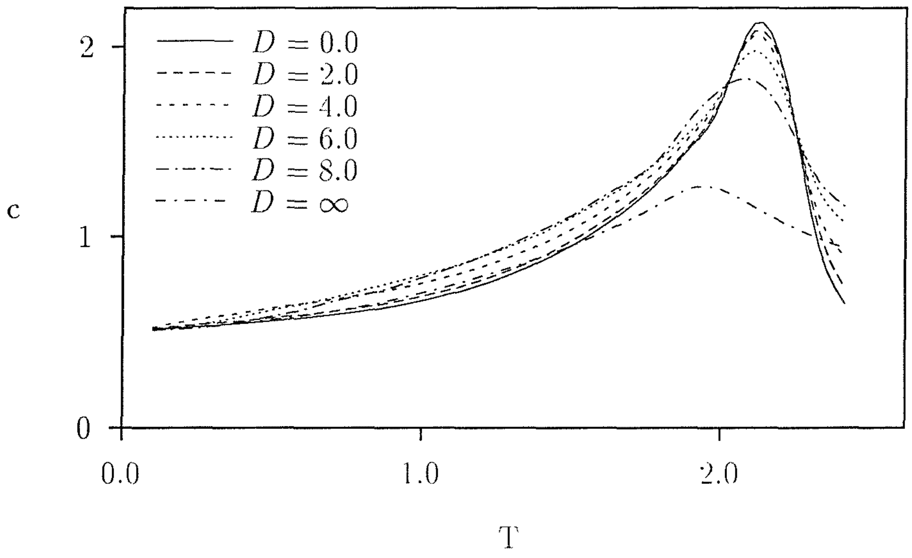

(b) Specific Heat

Figure 6.2: Energy (a) and Specific Heat (b) against temperature for cooled 31^3 systems with various larger anisotropy strengths D. The energy data is transformed to the large-D Hamiltonian (2.55), and the specific heat is plotted with smooth lines to simplify the comparison. A factor 0.5 is added to the $D = \infty$ curve as before

(section 2.4.1) so that a plot of $(C - C_0)D^{\lambda}$ against $(T - T_c)D^{-\mu}$ collapses all the data onto one curve, as shown on figure 6.1.1 with $\lambda = 0.25$, $\mu = 0.5$ and $T_c = 2.2$ (as found by [51]) — more data for larger system sizes is required to determine T_c more accurately, and without an accurate value for T_c, the exponents cannot be precisely determined. As noted for the two-dimensional case, it might be inferred that the change in the curve is due to some indirect effect on system size. But again, an increased anisotropy strength would result in a smaller correlation length, so that a given system size would better-represent a real system as the anisotropy is increased; this is the opposite from what is observed.

The specific heat data for $D < 3$ does not collapse onto the common curve, as can be seen from the figure (with points joined by straight lines for emphasis) — the system size L is also a scaling variable, and it may be that the rounding due to L outweighs that due to the anisotropy D for small D (since the curves for $D < 3$ in figure 6.2(b) are superposed without scaling).

Fitting this data to a Gaussian plus a linear term (for want of a better expression) we arrive at a curve

$$f(t) = c + a \exp\left(-\left(\frac{t+e}{b}\right)^2\right) - gt \tag{6.2}$$

with $t = T - T_c$ and $a = 1.2$, $b = 0.08$, $c = 0.65$, $e = 0.045$ and $g = 1.5$, as shown in figure 6.3(b).

It is possible to derive a relationship between the scaling exponents λ and μ in equation 6.1. As $D \to 0$, we expect to recover equation 1.3 ($C \sim |t|^{-\alpha}$).

Thus, for small argument x, $f(x) \sim x^{-\alpha}$, so that equation 6.1 becomes

$$C \sim D^{\lambda+\alpha\mu} t^{-\alpha}, \tag{6.3}$$

requiring $\alpha = -\frac{\lambda}{\mu}$.

Since α is very small for the three-dimensional XY system, we should expect to find $\lambda \ll \mu$, but if we apply this constraint above, we cannot acheive a satisfactory fit.

6.1.2 Comparison of Specific Heat by Differentiation and Fluctuation-Dissipation

The specific heats by numerical differentiation of the energy data and by the fluctuation-dissipation theorem (equation 3.31) are compared on figure 6.4. There

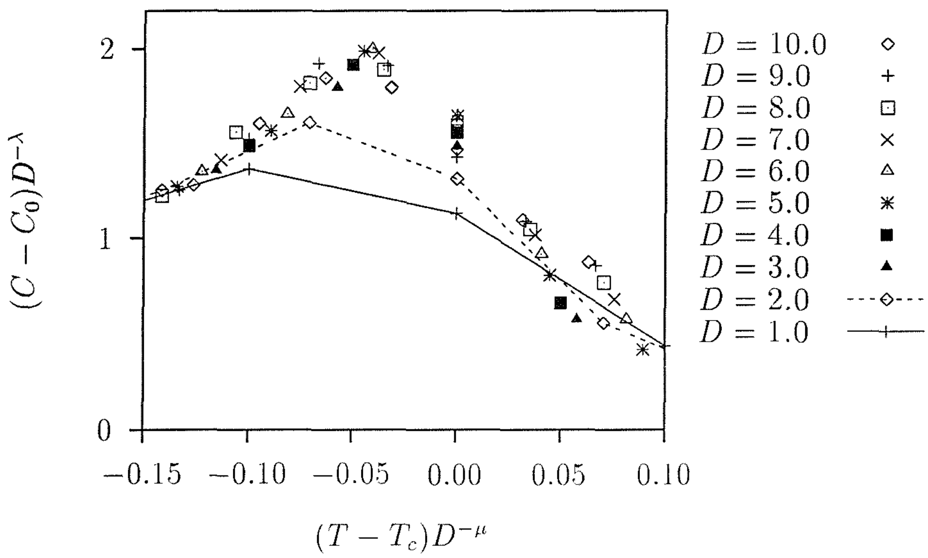

(a) Scaled Data

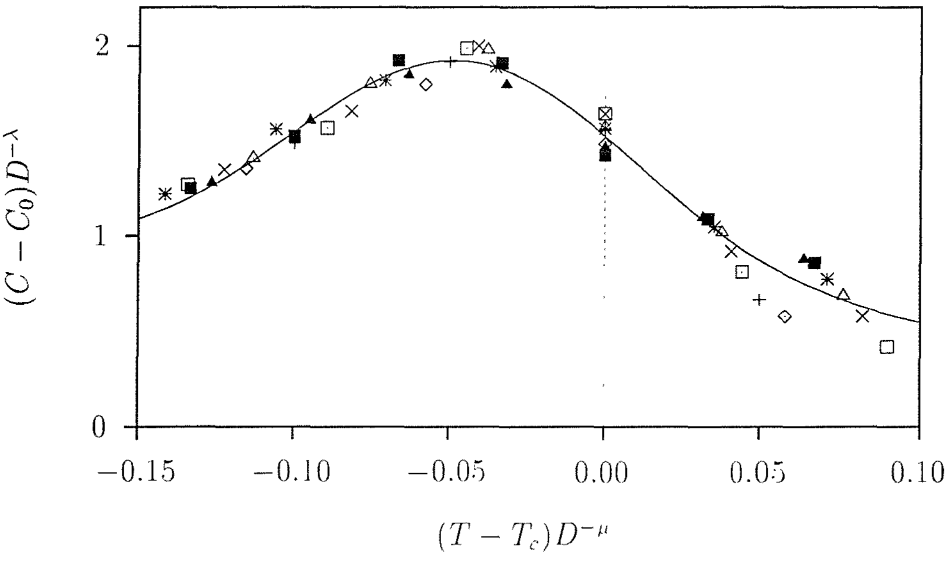

(b) Scaling Curve

Figure 6.3: Scaling curve for specific heat from figure 6.2(b). The data for low D does not fall on the scaled form in (a), as indicated with straight lines for emphasis. The data is fit to a scaling form (equation 6.2) as shown in (b).

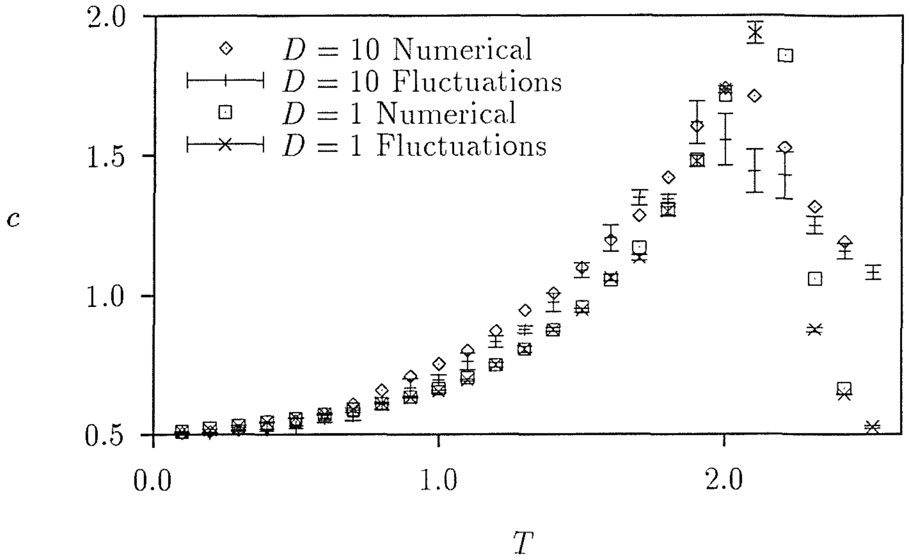

Figure 6.4: Comparing specific heat by fluctuation-dissipation and by numerical differentiation. Apart from near the phase transition, the curves are in good agreement at small D, but for larger D there is some disagreement for intermediate temperatures.

is good agreement (except near the phase transition) for $D = 1$, but there does appear to be some mismatch at intermediate temperatures for larger D. As before, there is no peak other than the (shifted) pure XY peak.

6.1.3 Dependence of Energy on Anisotropy Strength

In order to compare the energy as a function of D with the prediction made in equation 2.75, we transform the energy data to Hamiltonian 2.55, resulting in a curve of the form $\delta U \sim D^2$ as shown in figure 6.5, which is in poor agreement with the Chudnovsky-Imry-Ma theory, which predicts $U \sim D^4$ in three dimensions, but in excellent agreement with the local fluctation theory presented in section 2.8.5. It is possible that the small system dimension L is insufficient for the breakup into Imry-Ma domains to be fully realised, though a local energy term would always be expected to dominate a longer-range effect.

Figure 6.6 shows the behaviour of the energy as $D \to \infty$ at zero temperature — as before, the energy is transformed to the large D Hamiltonian (equation 2.55) to allow inclusion of the point for $D = \infty$. The approach to infinity is approximately linear in D^{-1}, as predicted by equation 2.77.

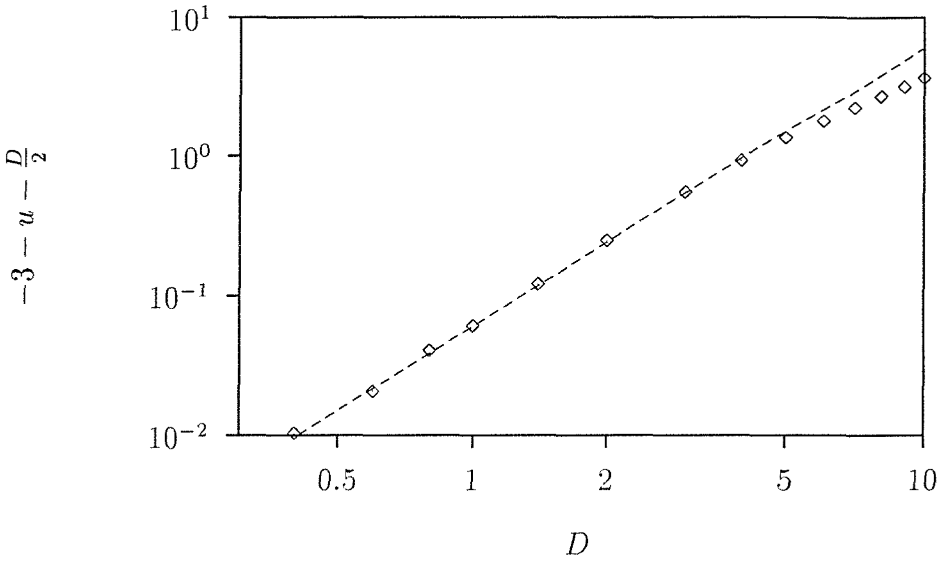

Figure 6.5: (Modified) Energy as a function of small D, with line showing D^2 fit for small D

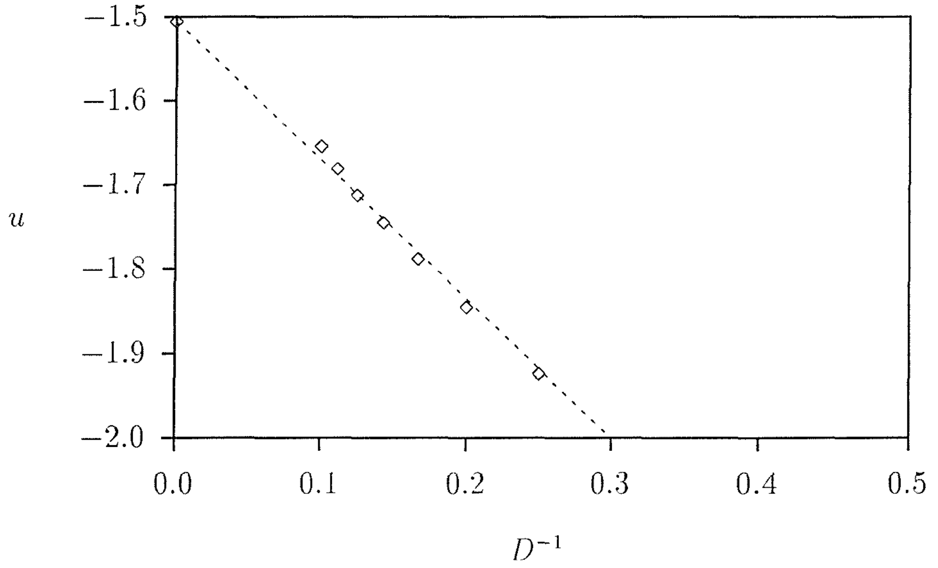

Figure 6.6: Energy is inversely proportional to D on approach to $D = \infty$ (using the large D Hamiltonian — equation 2.55).

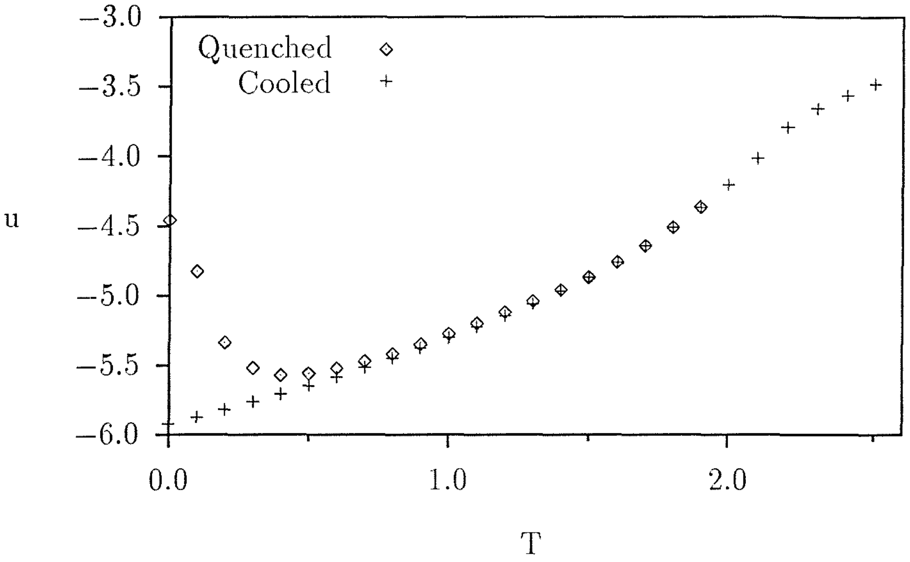

Figure 6.7: Comparing quenched and cooled energy for $D = 4$

Quenched Systems

As with the two dimensional system, we have also run simulations where we start from a random configuration at low temperature (as if quenched from high temperature), and slowly warm. As before, the energy *falls* as the temperature is increased — again, the interpretation is that the system is somehow frozen, and is able to reorganise only when sufficient thermal energy is available to surmount energy barriers. Figure 6.7 shows the sort of graph obtained from such simulations, which (apart from the larger D required to see the effect) are reminiscent of the two-dimensional simulations (figure 5.7).

As before, there appears to be a range of temperature over which the energy drops rapidly, followed by a gradual convergence to the energy from the cooled simulations, which is interpreted to be the result of the system exploring a deep potential well — as the temperature is increased, so that the system is sampling phase space away from the local energy minimum, the details of the region of phase space become irrelevant. The details do not appear to change significantly over a range of simulation timescales.

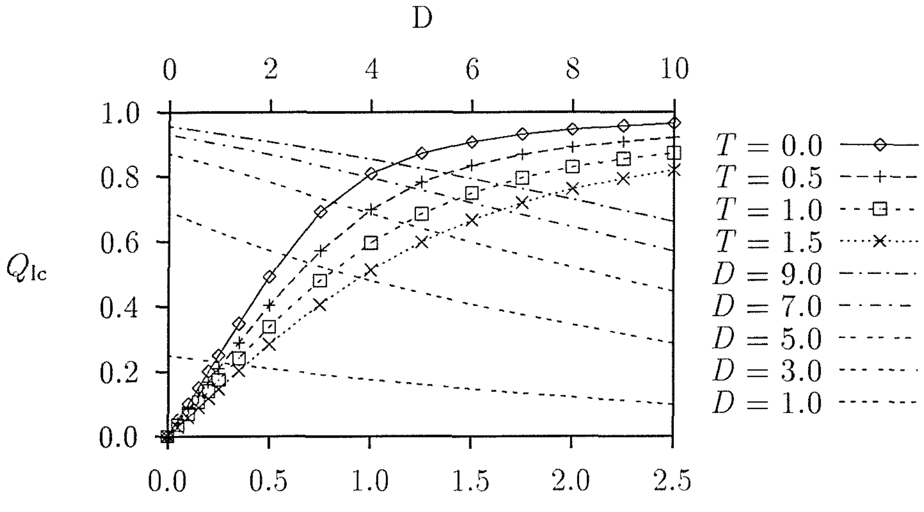

Figure 6.8: Liquid crystal order parameter with D and T.

6.2 Liquid Crystal Order Parameter

Figure 6.8 shows the liquid crystal order parameter measured for various D and T — the curves are similar qualitatively to those results for two dimensions (figure 5.10), but at higher anisotropy strengths, the curves appear to become convex rather then concave. (Data was not collected for Q_{lc} at such large anisotropy strengths in two dimensions.) We do not have a theory for this behaviour, and have been so far been unable to find a universal form for these three-dimensional data, but the approximately linear behaviour with D at zero temperature (uppermost rising curve on the figure) is compatible with the observation of $E \sim D^2$ in section 6.1, by equation 4.12. The linear behaviour ceases around $D = 3$, which is also the anisotropy strength above which the specific heat fulfilled the scaling relation for this system size (section 6.1). However, the quadratic nature of energy against anisotropy strength appears to persist until around $D = 4$.

$1 - Q_{\mathrm{lc}}$ is plotted against $1/D$ on a log scale, in order to verify the predicted D^{-2} behaviour for large D in equation 4.14; here, we find the straight line has a gradient of 1.9 for $D > 3$ in good agreement with theory.

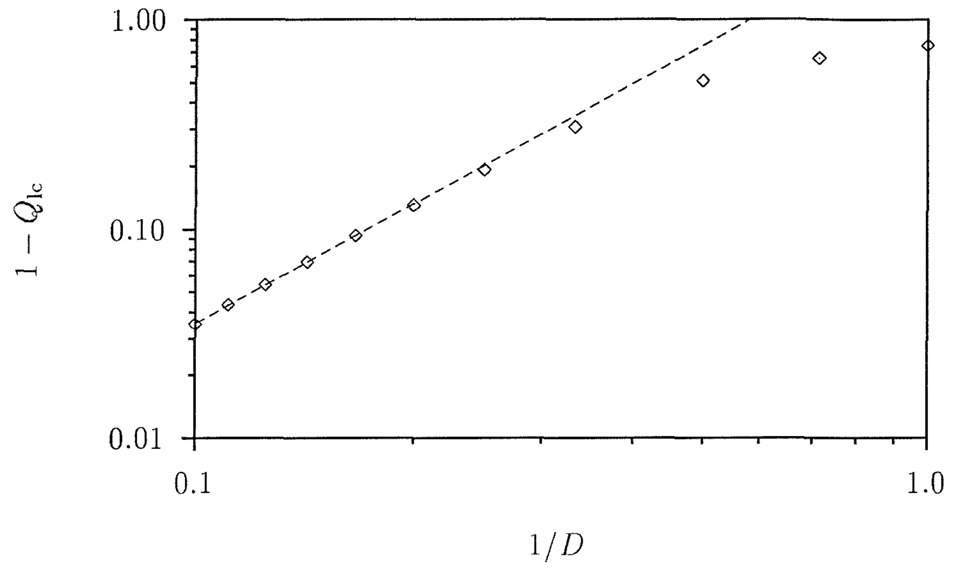

Figure 6.9: $Q_{\rm lc}$ on approach to the Ising limit. The dotted line illustrates $(D^{-1})^{1.9}$.

Quenched Systems

We find that at low temperatures, the liquid crystal order parameter is higher for the quenched system than the cooled system, as shown on figure 6.10, implying that the metastable state at low temperatures is in some way pinned by the anisotropy — the system is in a local minimum in which the spins are slightly more closely correlated to the anisotropy axes than in the cooled state, as before.

6.3 Correlations

The theory in section 2.2 predicts that the real-space spin-correlations in the pure three-dimensional XY model decay like $g(r) \sim a \exp(b/r)$; for even quite small r, this is indistinguishable from $g(r) \sim a + ab/r$ (since $\exp(x) \sim 1 + x$ for small x), as shown in figure 6.11(a). The fit to the data can be made even better at short distance by using a power smaller (more negative) than -1 in the second form.

Since we find that the behaviour changes with the addition of the random anisotropy, it was found to be convenient to use a generalised fit

$$g(r) \sim a + br^{-\lambda} \tag{6.4}$$

to the spherical average of the data.

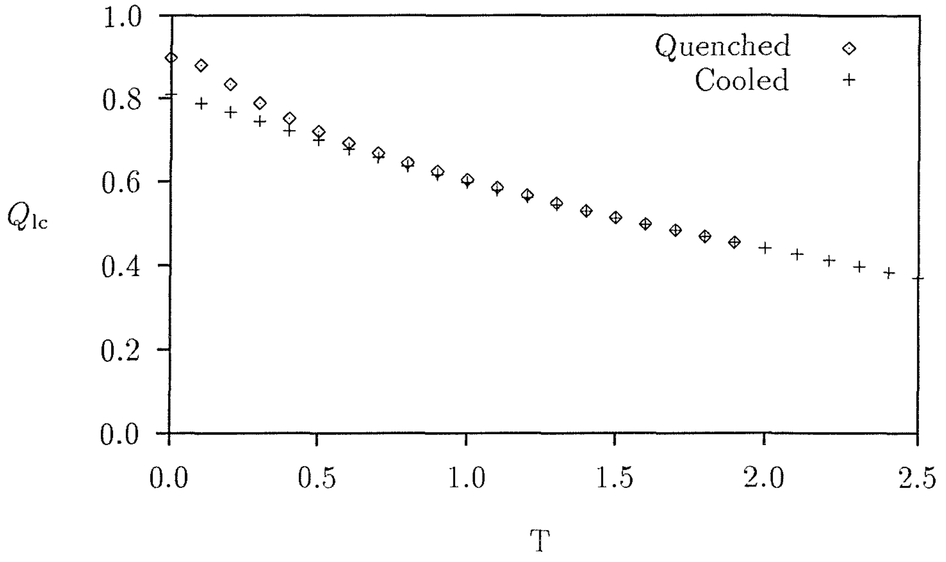

Figure 6.10: Comparing $Q_{\rm lc}$ for cooled and quenched systems for $D = 4$.

However, actually performing a fit with this form is very difficult, since the parameters are highly correlated and therefore the minimum of χ^2 in 'parameter space' is rather shallow. One way of fitting the data qualitatively is to plot on a log-log scale, subtracting different amounts until the points form a straight line — the appropriate value for the constant is fairly evident, since the curve on the log plot abruptly changes from concave to convex in the vicinity of the correct subtrahend.

To fit quantitatively, it is convenient to take Fourier transforms, so that the constant term is mapped to $\mathrm{k} = 0$, and

$$\begin{aligned} \int r^{-\lambda} \exp(-i\mathrm{k}.\mathrm{r})\, d^3\mathrm{r} &= k^{\lambda-3} \int (kr)^{-\lambda} \exp(-ikr\cos\theta)\, d(kr) d\theta d\phi &(6.5)\\ &\sim k^{\lambda-3}, &(6.6)\end{aligned}$$

since the integral on the right hand side of equation 6.5 is simply a number.

Thus, if the data fits this form, $3 - \lambda$ is simply the gradient of the straight line on the log-log graph of the fourier transform of the correlation data, as shown in figure 6.11(b); fitting this straight line is a much easier task, both because we have reduced it to a two-parameter fit, and because the χ^2 minimum is better-defined for a fit of this form. It is the small-k data which is important, since this corresponds to large distance in real space. The gradient on figure 6.11(b) is slightly larger (less negative) than -2, which corresponds to a slightly faster decay in real space, as noted above.

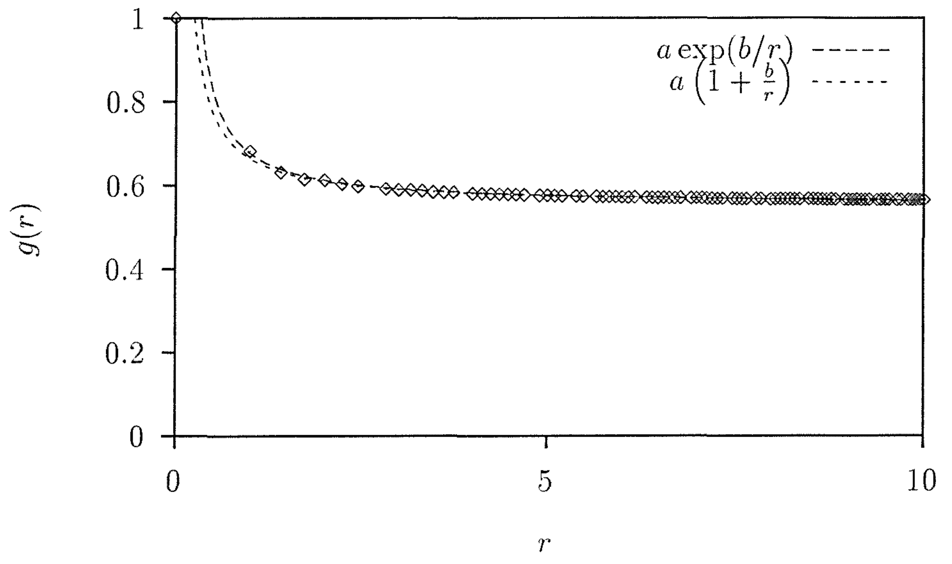

(a) Comparison of fit of $a \exp(b/r)$ and $a\left(1+\frac{b}{r}\right)$ to spin correlations in real space

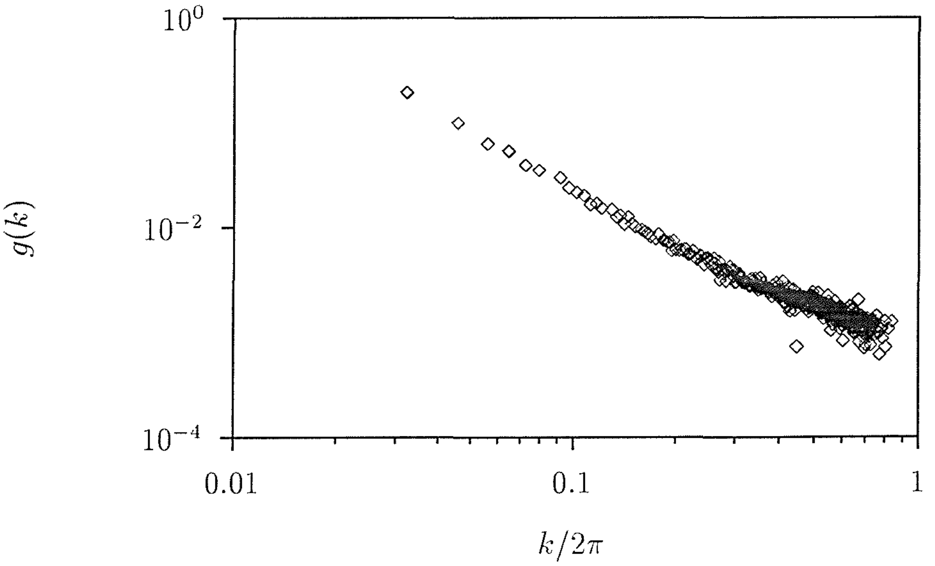

(b) Correlation data in k-space

Figure 6.11: Correlations of pure XY model in three dimensions. Fit is simplified by taking Fourier transforms, whereby the second form in (a) becomes a straight line with gradient approximately -2 on a log plot, as shown in (b).

6.3.1 Cooled systems

Figure 6.12 compares the correlation function at five temperatures for $D = 0$ (a) and $D = 4$ (b). As in the two-dimensional system, the correlation function for the random system changes very little below a certain D-dependent temperature. Unlike that system, however, the form of the correlation function does not change from intermediate temperatures, though the power does appear to change.

The dotted line on figure 6.12(b) shows a fit of the low temperature correlations to $a + br^{-0.7}$, though as was noted above, the minimum of χ^2 for a fit of this form is very broad, and the fit is quite good over a rather large range of parameters. The parameter a corresponds to a long range order — we find that this constant term falls with increasing system size, suggesting that in the $L \to \infty$ limit, we have simple power-law decay. This behaviour was also found in related work on the random anisotropy nematic spin model [63].

Figure 6.13 shows the fitted value for λ against temperature in equation 6.4 for various anisotropy strengths D. The fit is rather erratic, but it is fairly clear that the trend is for the power to level off at low temperature, and the power increases with increasing D. For $T > 2.2$, the fit produces the same parameters for all D — we identify this as the paramagnetic régime.

The fit gives the same results for all D for $T > 2.2$ — this is identified as the paramagnetic régime.

As noted in the introduction to this section, λ is expected to be close to 1 for $D = 0$. However, for this case, the constant a does *not* vanish with increasing system size, as we expect true long-range order for the non-random system. The data for $D < 1$ is extremely hard to fit. We attribute this to a finite size effect — the system is crossing over from a state with (approximately) power-law decay with a constant to a state with a different power-law decaying to zero (in the large L limit).

This will be discussed further in the section 6.4.2.

6.3.2 Quenched Systems

The spin correlations have a different form at low temperature in quenched systems, as shown in figure 6.14(a). As the system is warmed, the order increases as

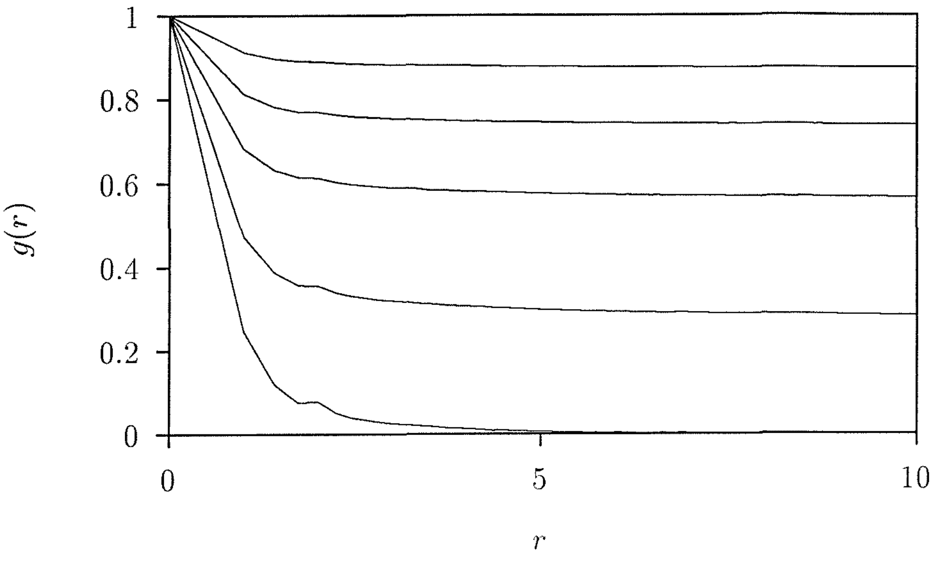

(a) $D = 0$

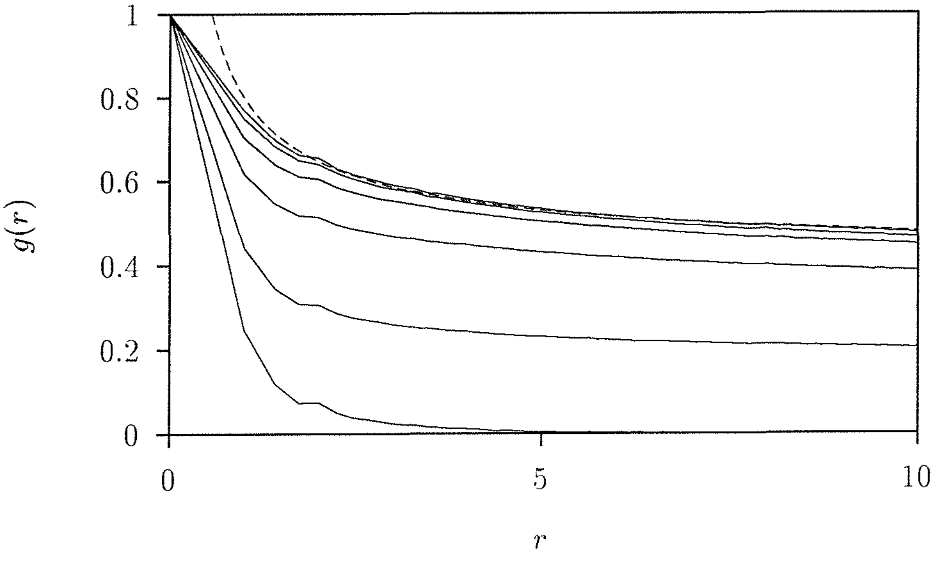

(b) $D = 4$

Figure 6.12: Spin correlations for (a) $D = 0$ and (b) $D = 4$. Temperature increases from 0 at top to 2.5 at bottom, in steps of 0.5. The low-temperature data shows 'bunching' analogous to the two-dimensional system. The dotted line on (b) shows the curve $0.4 + 0.4 \times r^{-0.7}$.

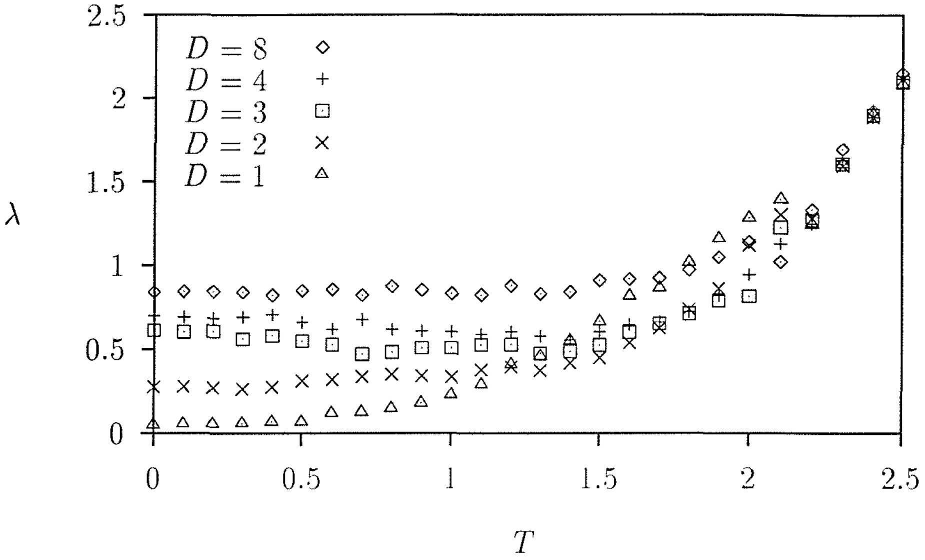

Figure 6.13: Fit of form $a + br^{-\lambda}$ to the spin correlations. Despite rather erratic values, the trend is clear: λ increases with increasing D.

the system overcomes energy barriers, until the system is able to adopt the same configuration as the cooled systems.

In order to obtain a measure of the domain sizes, to verify CIM predictions, we force-fit an exponential form to the low-temperature data, as before. This gives approximate domain sizes as shown in figure 6.14(b). The line shows a fit of the form $\xi^{-1} \sim D^{2.15}$, in good agreement with the CIM prediction of $\xi \sim D^{-2}$ (equation 2.74).

6.4 Magnetisation

Unlike the two-dimensional case, the (pure) XY model in three dimensions is expected to have true long-range order, and hence a spontaneous magnetisation at low temperature. However, only an infinite system can genuinely show a broken symmetry state; the magnetisation of a finite system can still rotate in space with time, so that care is still required during a simulation in order to ensure that the measured magnetisation $\mathbf{M}$ does not average to zero.

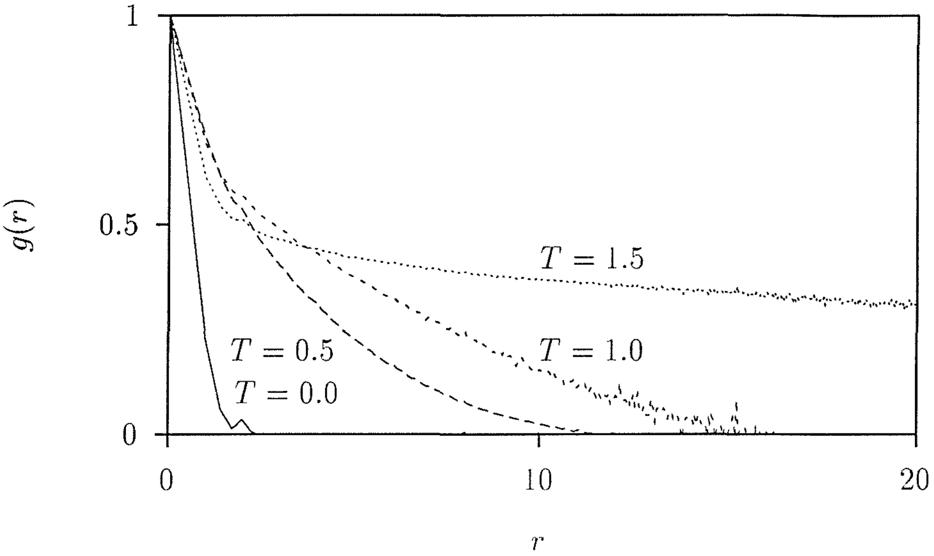

(a) Correlations

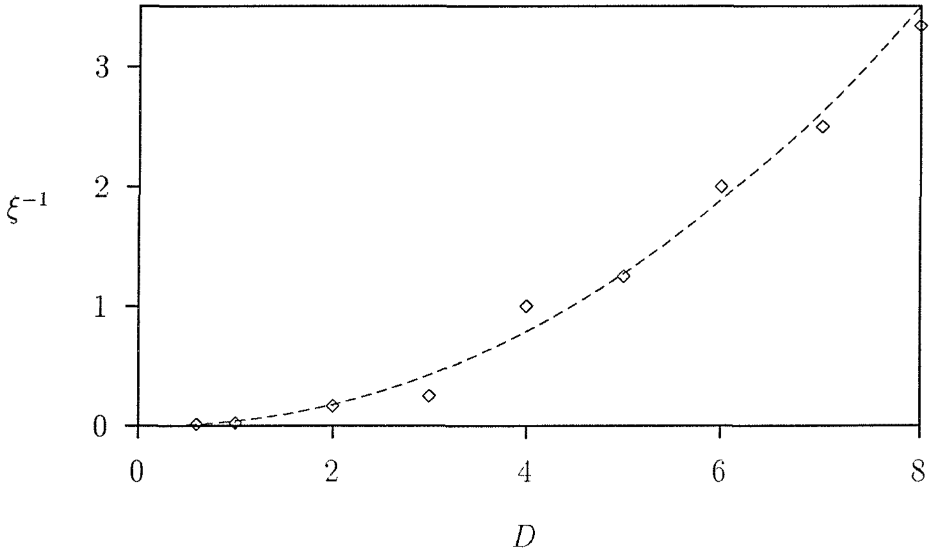

(b) Domain size

Figure 6.14: Low-temperature spin correlations (a) for quenched systems. As before, we force-fit an exponential to these data, to give approximate domain sizes (b). The dotted line shows $D^{2.15}$.

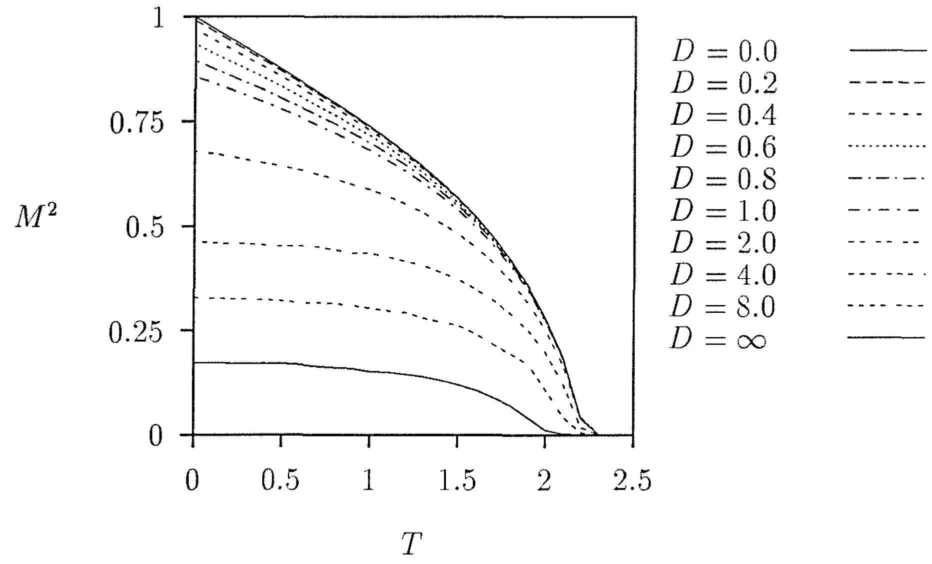

Figure 6.15: Spontaneous Magnetisation as a function of temperature for various anisotropy strengths.

6.4.1 Magnetisation of a Finite System

We find that the addition of random anisotropy suppresses the magnetisation of the $3D$ XY model at a given system size, as shown in figure 6.15. However, we find the reduction for a given anisotropy strength to be significantly smaller than for the two-dimensional system: the magnetisation for $D = 1$ in a 31^3 system is about the same as for only $D = 0.2$ in a 62^2 system (figure 5.24). This may be due to the smaller linear dimension of the system (we use more spins in three dimensions than in two, but the linear dimension L is smaller, and it is the ratio between correlation length ξ and L that influences finite size behaviour), or it may be due to the fact that the pure three-dimensional system is 'more' magnetic in the first place (section 2.2). Even at $D = \infty$, the 31^3 system shows a non-zero magnetisation which is in good quantitative agreement with [46].

The plot of zero-temperature magnetisation against anisotropy strength D in figure 6.16 bears remarkable resemblance to the predictions of mean field theory [29]. This again suggests a finite size effect — in a small system, the magnetisation does not get a chance to change direction across the sample before reaching the other edge. The difference from full alignment is shown in figure 6.16(b), and for small D it is in reasonable agreement with the $\delta M \sim D^2$ predicted by mean field theory [29].

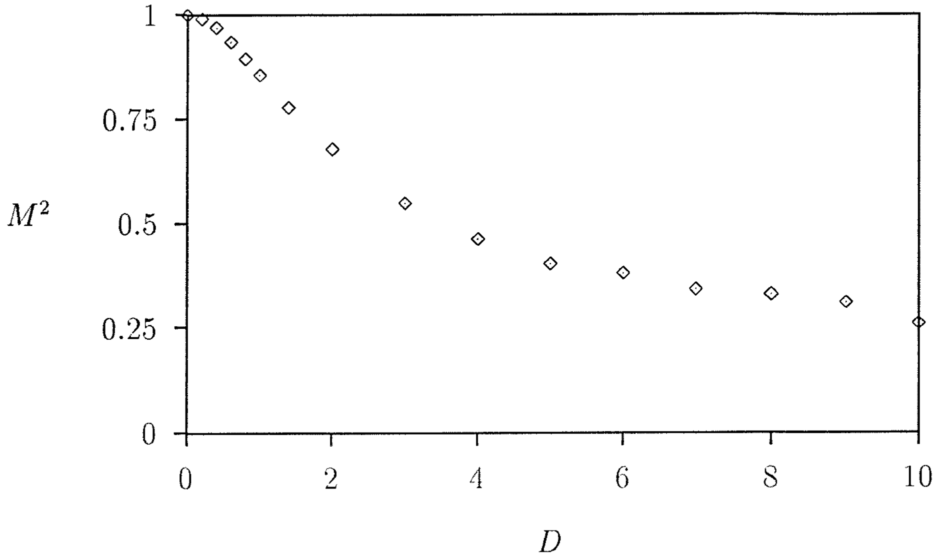

(a) Magnetisation

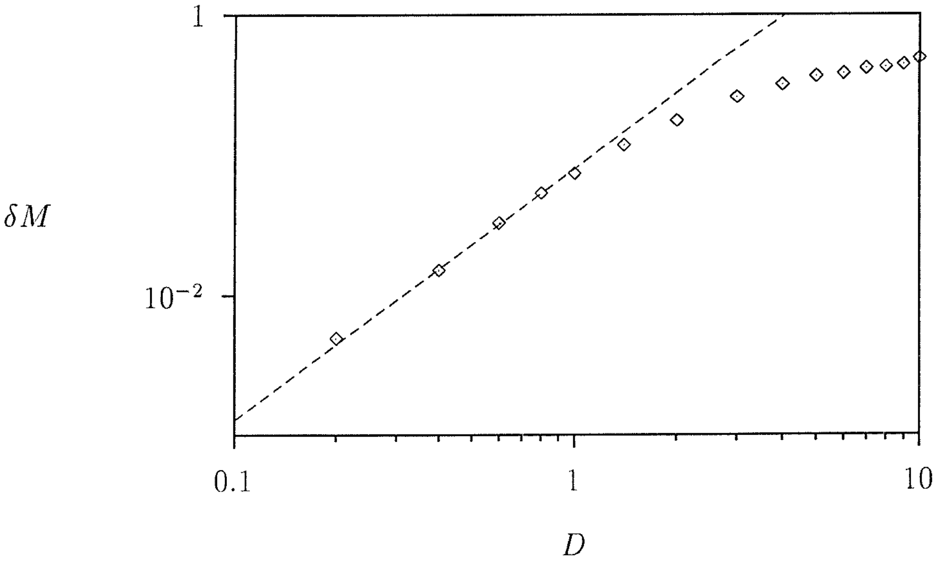

(b) Departure from full alignment

Figure 6.16: Ground-state magnetisation as a function of anisotropy strength. The line in (b) shows $\delta M \sim D^{1.8}$ in reasonable agreement with mean field theory [29]

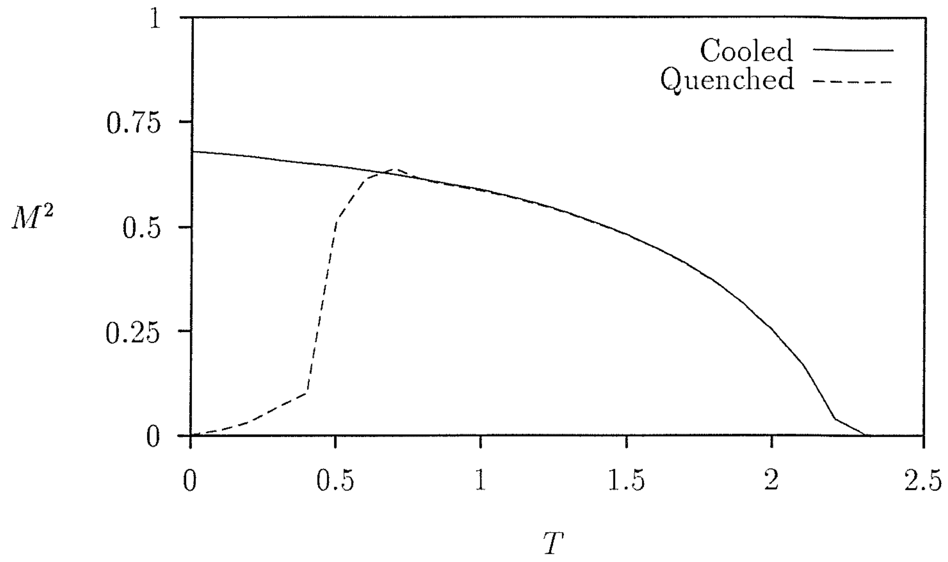

Figure 6.17: Comparison of cooled and quenched magnetisation for 31^3 system with $D = 2$

Quenched Systems

As in the two-dimensional system, the system must be warmed to a temperature $T^*(D)$ before the magnetisation achieves the same value as for the cooled system. As would be expected from the preceeding results, T^* is much larger for a given anisotropy strength than for the two-dimensional system; a typical result is shown in figure 6.17. T^* does grow rapidly with D (figure 6.26), however, so this may be a feature of the small system-size — as discussed in section 2.6, glassiness is a dynamic feature, and the size of the energy barriers depends on the system size.

6.4.2 Finite Size Scaling

In the the study on the two-dimensional system, the scaling of the net magnetisation with system size seemed to give the most reliable (qualitative) indication of a change in 'static' behaviour (on the timescales of our simulations). We have made a similar study of the scaling of the three-dimensional system, but we are hampered somewhat by the limited system dimension which we have been able to study. As noted at the beginning of this chapter, we were able to use a little time on another machine to work with a slightly larger system.

Figure 6.18 shows the scaling behaviour of the non-random system, with a close-up on the low-temperature behaviour in figure 6.18(b). Because we expect

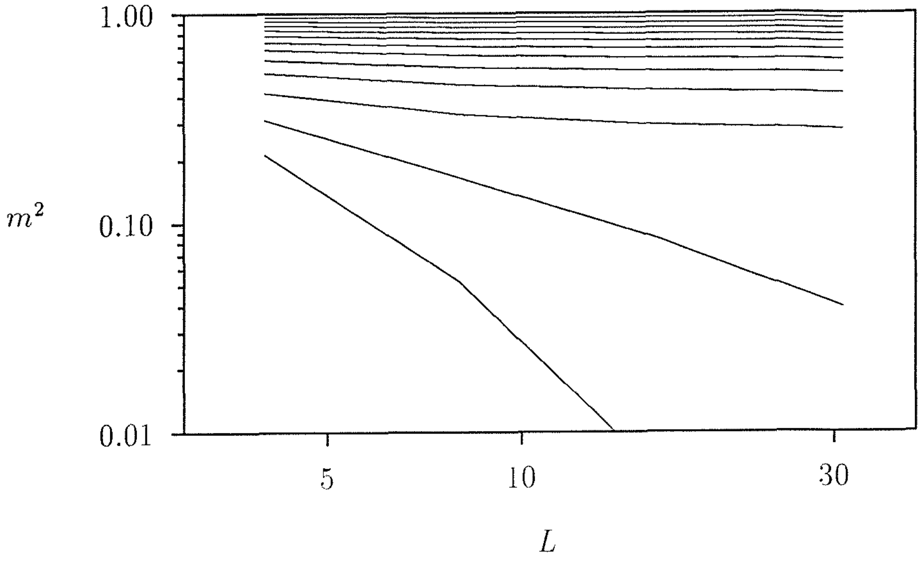

(a) High Temperatures

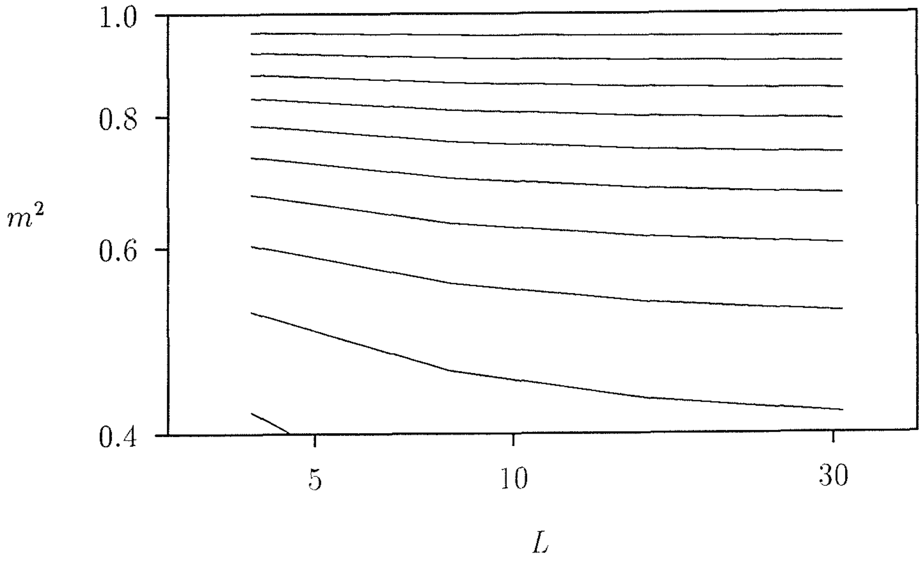

(b) Low Temperatures

Figure 6.18: Scaling of magnetisation of nonrandom system with system size. For clarity, figure (b) shows a close-up of the low-temperature data. Each line is a different temperature, from 0 at the top, increasing in steps of $0.2J$ down the graph, with $T = 2.4J$ at the bottom.

long-range order at low temperature, the magnetisation should not change with system-size. Putting $g(r) = a + b/r$ in equation 5.14, we find $m(L)^2 \sim a' + b'L^{-1}$, so that for sufficiently large system, m is independent of L, as required. The scaling curves can be shown to fit this form exactly, at low temperature.

At high temperatures ($T > 2.2$), the scaling behaviour changes over to $m^2 \sim L^{-3}$, as expected for a paramagnetic régime.

The corresponding scaling behaviour for our random system is shown on figure 6.19. We were able to calculate the magnetisation at only one point for a larger system; this single point is invaluable, however, in demonstrating that the magnetisation is falling as some power of system size L (straight-line on log-log plot).

At low temperatures, the curves appears to be bending down with increasing L, in a similar way to the two-dimensional system (figure 5.27); it is not clear from this figure whether they are heading towards a different power law from the intermediate temperature results, but when results for different D's are taken into account, it seems likely that the power is the same. At higher temperatures, the magnetisation begins with with a more rapid decay at small L, but head towards the same power law for larger system sizes. These may be effects due to the $a_L(T)$.

At high temperatures, the curves are heading towards L^{-3} as before. The crossover is not sufficiently clear-cut for the transition temperature to be identified using this approach.

A power law of magnetisation with system size is consistent with the explicit measurements on the spin correlations in section 6.3, where we found the correlations $g(r) \sim a_L(T) + br^{-\lambda}$: putting this form into equation 5.14, we should expect to find $m^2 \sim a_L(T) + L^{-\lambda}$. We measure the gradient of the magnetisation scaling curves, and present the results in figure 6.20.

The powers from the finite-size scaling are somewhat smaller than those from the direct measurement on the correlation function, though the general trends are the same. [63] found similar behaviour, but extrapolation of the correlation results to large L brought the results into agreement.

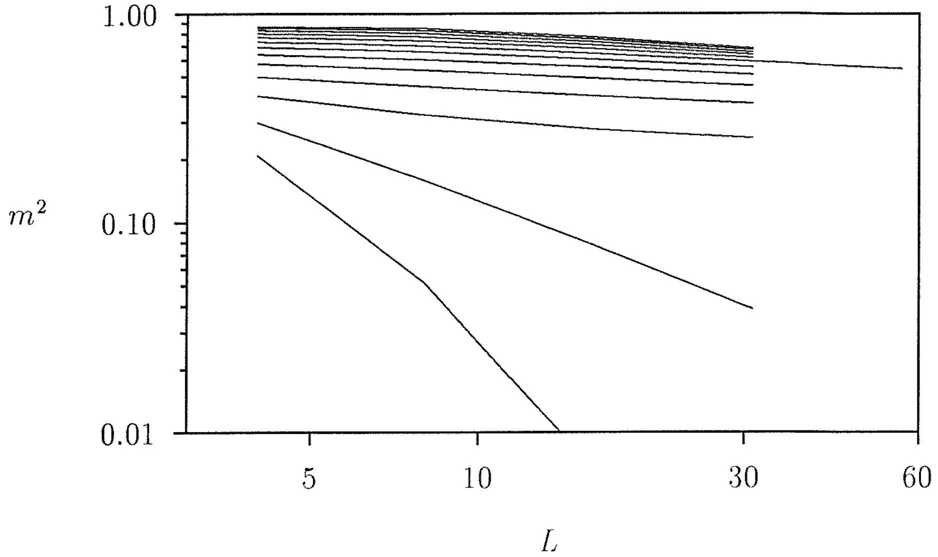

(a) High Temperatures

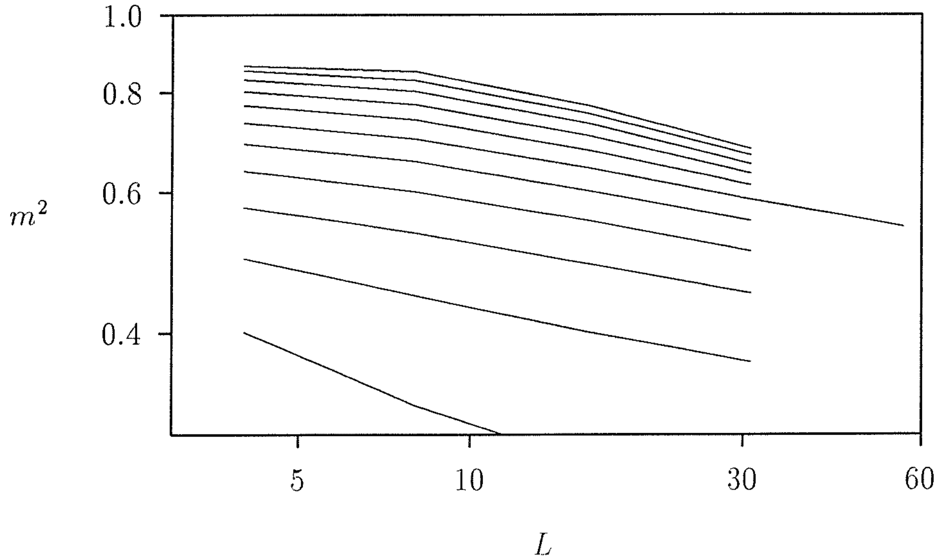

(b) Low Temperatures

Figure 6.19: Scaling of magnetisation of $D = 2$ system with system size. Each line is a different temperature, from 0 at the top, increasing in steps of $0.2J$ down the graph, with $T = 2.4J$ at the bottom. The curves at low temperature show a similar 'bunching' to the two-dimensional system. We were able to use time on another computer system to make one measurement at an even larger system size, to confirm the power-law decay.

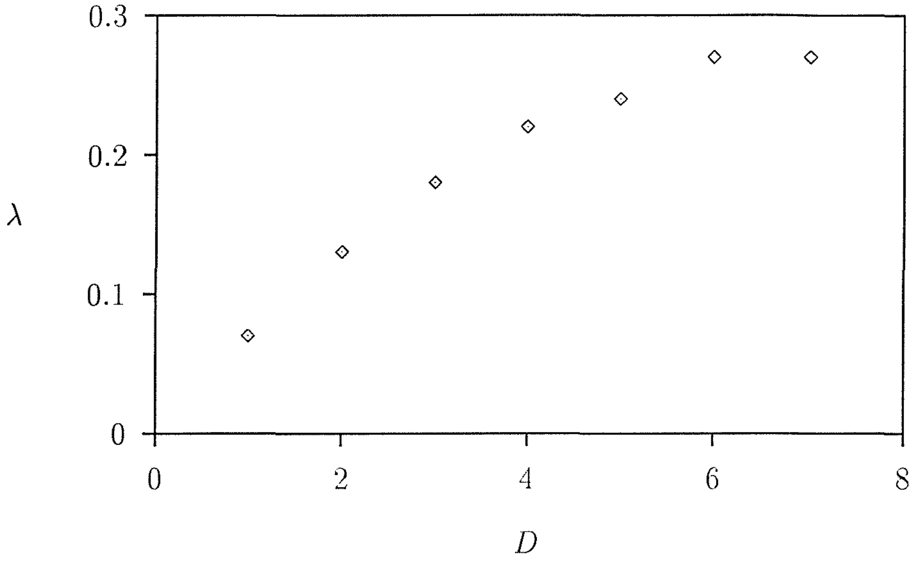

Figure 6.20: Gradient of magnetisation curves

6.4.3 Approach to Saturation

In order to investigate the CIM predictions for the approach to saturation (summarised in section 2.8.10), we have performed studies with large magnetic fields. As was found in the two dimensional case (figure 5.29, the system size is relatively unimportant, so the following is based on results from a 16^3 system, to minimise the computational load.

Figure 6.21 demonstrates thermal desaturation (section 2.8.9) on a system with weak anisotropy ($D = 0.25$), with a dotted line indicating the H^{-1} predicated by equation 2.94.

In strong anisotropy, at low temperature, the desaturation in large field is predicted to have take the form $\delta m \sim D^2/H^2$ (section 2.8). As before (section 5.4.4), we find an extended crossover region with the H dependence changing, until it saturates at H^{-2} for $D > 4$: for $D = 0.25$, we find H^{-1}, which we interpreted as thermal desaturation above; for $D = 1$, we find $H^{-1.2}$; by $D = 2$ it has become $H^{-1.5}$, and it has reached $H^{-1.8}$ by $D = 4$. This is shown on figure 6.22.

We have been unable to identify the FWA régime with $H^{-\frac{1}{2}}$, as predicted by equation 2.90. The curves in the figures do have to pass through $H^{-\frac{1}{2}}$ in order to reach H^{-1}, but they do not spend sufficiently long with this behaviour for the FWA régime to be formally identified. As noted in section 2.8.7, there is some confusion over the existence of the FWA phase.

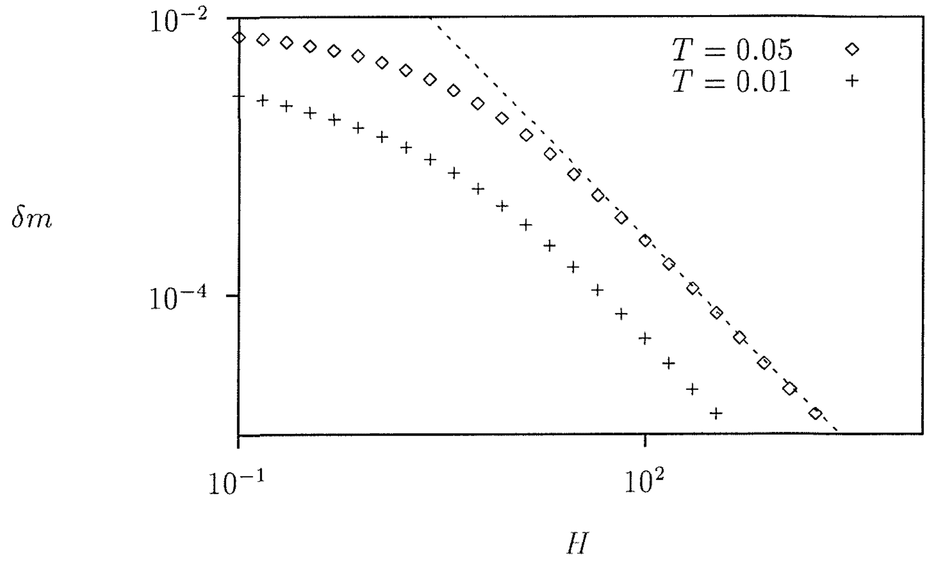

Figure 6.21: Thermal Demagnetisation in weakly anisotropic system ($D = 0.25$) with H^{-1} behaviour predicated by equation 2.94.

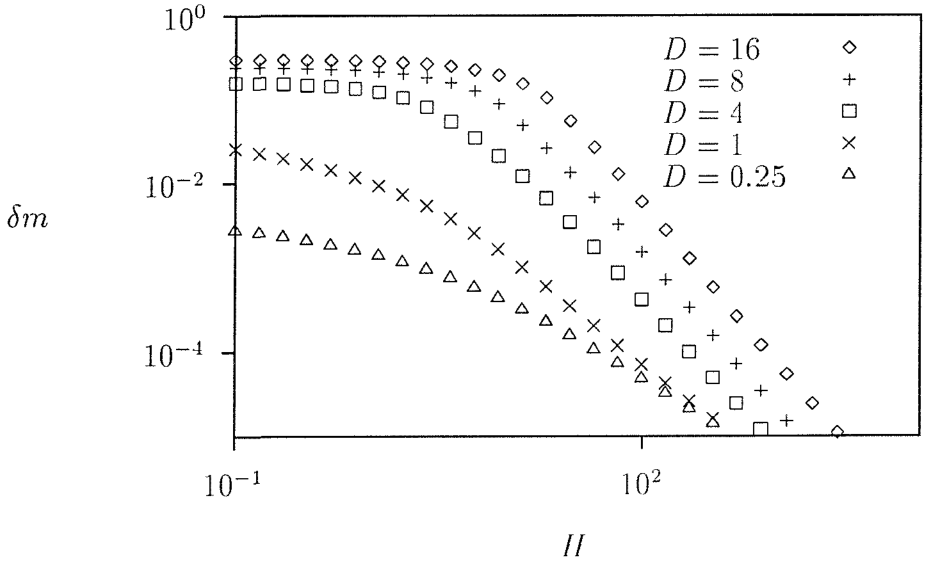

Figure 6.22: Approach to saturation at low temperature for various D, tending towards the predicted H^{-2} for $D > 4$

6.5 Dynamic Effects

In the two-dimensional system, we found that the details of the quenched behaviour depended on the system size and the timescale of the experiments, as expected. We have run some very long simulations on the 31^3 system in order to discover the time dependence of the system. Figure 6.23(a) shows the magnetisation of the system as a function of time over 5×10^6 steps per spin, from a random start with anisotropy strength 4 at a temperature of $0.5J$.

The jumps in magnetisation appear to be correlated to the jumps in energy, as expected. Even after 10^6 sweeps through the lattice, the system appears to have some way to go before the various parameters reach the values from the cooled simulations, which are chosen to be the extremes of the scales on the figures.

Figure 6.24 compares the relaxation of the energy towards that of the cooled system for various temperatures for the $D = 4$ 31^3 system. The relaxation time clearly falls abruptly at some temperature dependent on D, and the apparent glassiness disappears at temperatures above this $T^*(D)$. This will be discussed further in chapter 7.

6.5.1 Hysteresis

Because hysteresis is an important criterion for the existence of magnetism in a given material, we have tried to measure an $M - H$ characteristic for the three dimensional system. However, while we found that the curves converged in the two-dimensional system as we reduced the field sweep-rate, we were not able to find a stable curve in three dimensions for sweep rates accessible with the current implementation and with the current hardware.

Figure 6.25 shows a set of curves for a 16^3 system with $D = 4$ at $T = 0.1J$. We perform 7,500 sweeps at each field point, and vary the field in steps of between $0.1J$ and $0.001J$. For each reduction in sweep rate by a factor of 2, the coercive field appears to be reduced by a constant amount. For a given sweep rate, the coercive field increases with anisotropy strength and decreases with increasing temperature, but without a more careful analysis of the dynamics of the system, it is impossible to make a quantitative analysis.

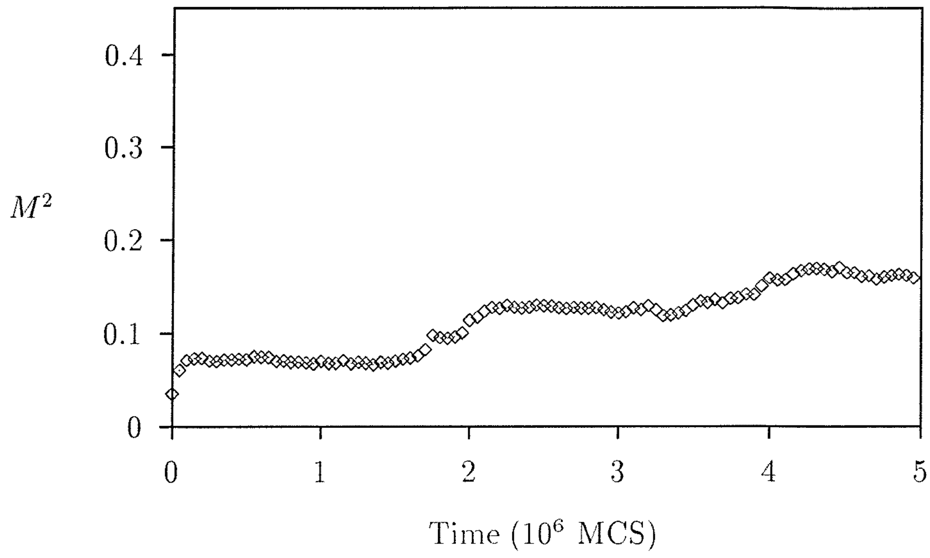

(a) Magnetisation

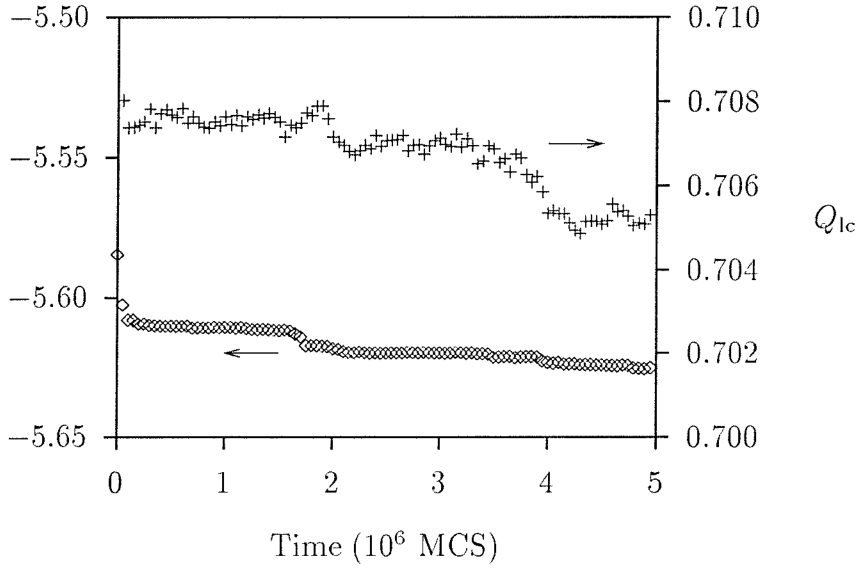

(b) Energy

Figure 6.23: Magnetisation (a) and energy and liquid crystal order parameter (b) as a function of time for 31^3 system with $D = 4$ and $T = 0.5$. The upper limit of the magnetisation scale (0.45) is the value of the magnetisation for a system cooled to $T = 0.5$, while the lower limit of the scales in (b) are the equivalent targets for energy and $Q_{\rm lc}$. Though the changes look rather large on these figures, the jumps are extremely small in practice — the excess energy above the cooled system (lower limit of scale) is about the same as the separation of the cooled and quenched curves on figure 6.7.

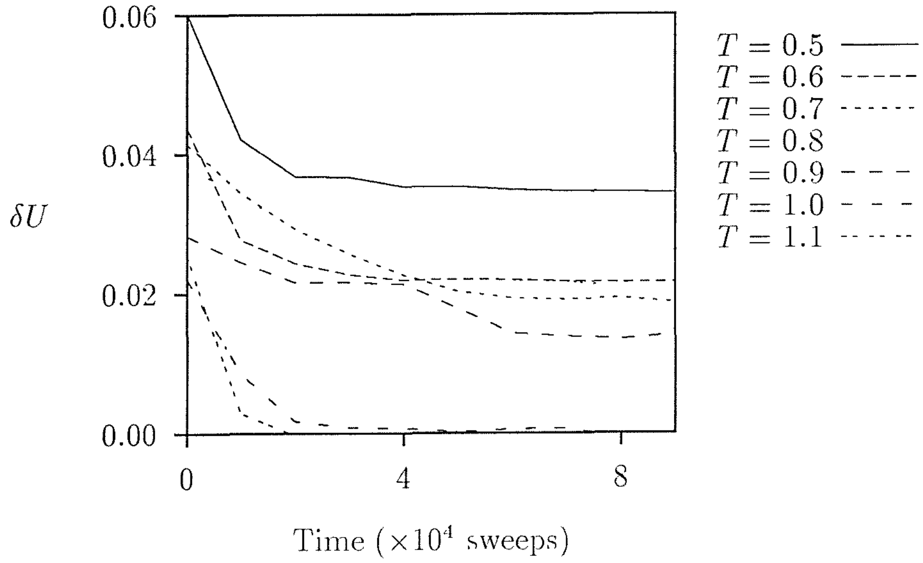

Figure 6.24: Comparison of rate of relaxation of energy against temperature] The relaxation time appears to fall dramatically at $T \sim 1$ for this system ($D = 4$ and $N = 31^3$).

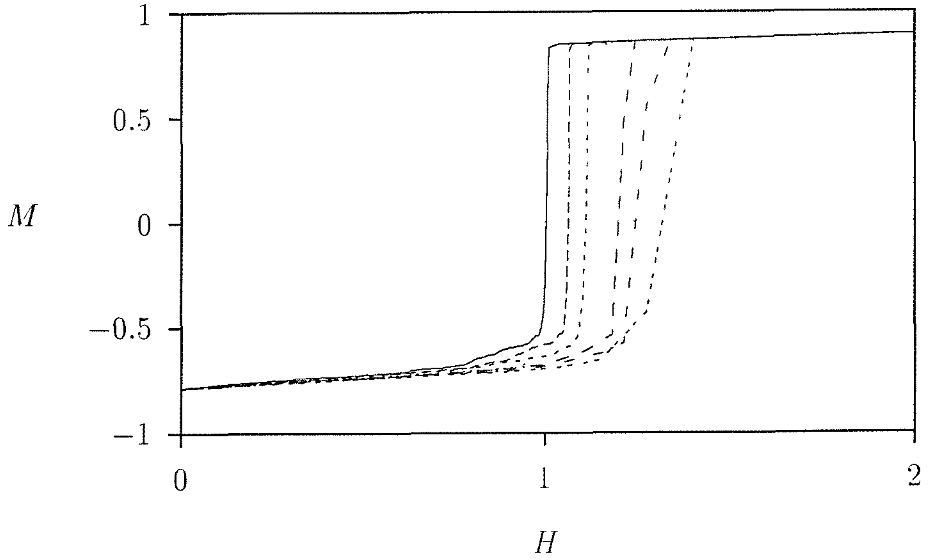

Figure 6.25: Dependence of $H_{\rm co}$ on field steps of between $0.1J$ and $0.001J$.

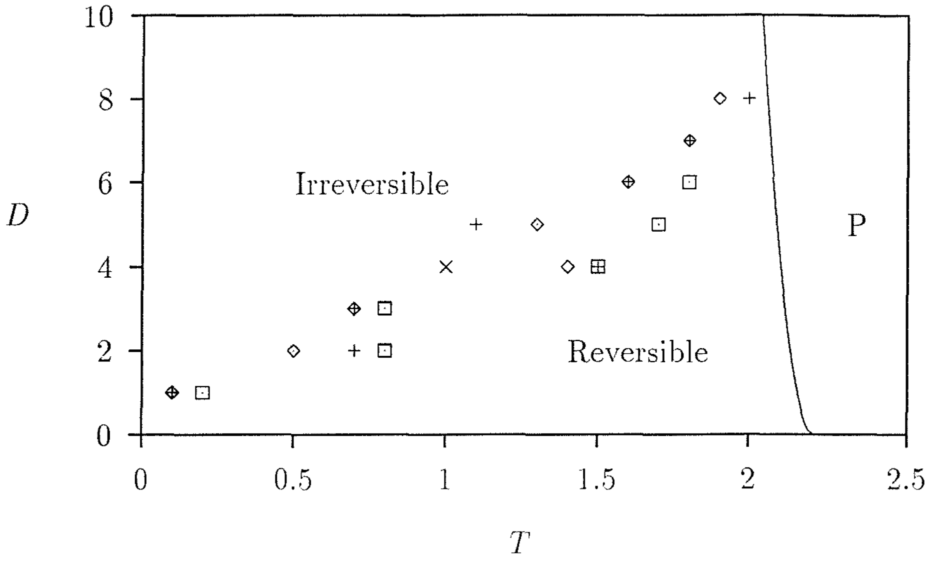

Figure 6.26: Tentative Phase Diagram for three-dimensional system. The line delineating the high-temperature paramagnetic phase (P) is the specific-heat scaling curve, extrapolated to small D. We compare energy ($\diamond$), magnetisation ($+$) and correlations ($\Box$) for cooled and quenched systems, and include the explicit observation of the relaxation time ($\times$), in order to separate reversible from irreversible (glassy) behaviour.

6.6 Summary

As before, we collect together the various pieces of evidence to produce a tentative phase diagram (figure 6.26):

- The specific heat scaling curve provides a clear delineation for the transition to paramagnetic behaviour at high temperatures. We extrapolate the curve towards $D = 0$, where the finite size effects obscured the effect.
- We use the differences in the energy, magnetisation and correlations between the cooled and quenched systems to separate reversible from irreversible behaviour.

Like the two-dimensional system, there is no thermodynamic indication of the onset of irreversibility. In the two-dimensional system, we did find a qualitative change in the correlation function to support the irreversibility data, but here we have no such change; the only evidence is an increased relaxation time. This will be discussed further in chapter 7.

Chapter 7

Discussion and Conclusions

In any Monte Carlo study, results are valid only if precautions are taken to ensure finite size effects are controlled, and that adequate equilibration and simulation time has been allowed, particularly in the vicinity of phase transitions. Multispin and cluster update algorithms have been advocated as techniques for overcoming critical slowing down. Despite criticism in the literature over the use of a single-spin update algorithm, we have used such an algorithm in our study the XY model with random anisotropy. We have performed simulations on a wide variety of timescales in order to establish the validity of our results.

We feel that if our simulations show time-dependence, then real systems may also show similar behaviour. In particular, for example, simulations using a cluster update algorithm cannot exhibit hysteresis, since that algorithm permits spin flips to be applied globally, and the symmetry of the Hamiltonian guarantees that the system will flip to follow the field. On the other hand, it is true that simulations on a finite system using any valid update algorithm should in practice eventually flip all the spins; on reasonable simulation timescales, however, local update algorithms do show hysteresis, which is also seen in experiments on real systems. Thus, algorithms which improve the ergodicity of the simulation do not necessarily give the 'correct' or desired answers.

We have not explicitly performed each simulation over a large number of different configurations of the random axis; this should be done on at least a small part of the phase space to confirm that these results are free from sample to sample variations. That said, we have not strictly kept the same sets of axes across all simulations at a given system size, so that there has been some variation. Our results are in good quantitative agreement with previously published results in the various limits ($T \to 0$, $D \to 0$, $D \to \infty$) so we are fairly confident of the validity of our results.

7.1 Two dimensional system

7.1.1 Chudnovsky-Imry-Ma

The most important theoretical prediction which we have sought to confirm is the zero-temperature phenomenological Chudnovsky-Imry-Ma (CIM) theory reviewed in section 2.8. For small anisotropy strength, this predicts a breakup of the system into domains, and makes a variety of predictions about changes in various system parameters with anisotropy strength.

Our results for systems cooled to zero temperature are in good agreement with these domain predictions:

- For small anisotropy strengths, the energy contains a component proportional to $D^{1.75}$ in reasonable agreement with D^2 in equation 2.75;
- In an external magnetic field, the approach to saturation is as D^2/H as predicted for the ferromagnet with wandering axes.

In addition, the domain size from quenched correlations agrees with predicted D^{-1} behaviour. We have been unable to verify the predictions for the magnetic susceptibility quantitatively, as noted in section 5.4.2. The susceptibility does appear to be strongly suppressed at low temperature.

For larger anisotropy strengths, the domain size reaches a lower limit of one lattice spacing, and a different régime is predicted. In this régime we find:

- a component of energy inversely proportional to D, as predicted (equation 2.77);
- Approach to magnetic saturation in an external field like $\left(\frac{D}{H}\right)^2$ (equation 2.92)

In addition, at non-zero temperatures, the approach to saturation in a magnetic field is observed to go like T/H, as predicted by equation 2.94.

7.1.2 Cardy and Ostlund

The renormalisation results of Cardy and Ostlund [35] predict that the low-temperature Kosterlitz-Thouless phase will survive the addition of the randomness at intermediate temperatures, *provided that the number of equivalent easy directions is at least 3.* Other theoretical analyses concur with this picture.

Our simulations with only two easy axes fall outwith the remit for this prediction, but we *do* find that the spin-correlations decay with the same power η as the pure system at intermediate temperatures, suggesting the survival of the Kosterlitz-Thouless phase above a temperature proportional to D.

The notion of thermal depinning section 2.8.4 is appealing, since it uses simple arguments to predict the overcoming of the anisotropy barriers due to thermal excitations. Our results show a depinning temperature proportional to D, in quantitative dispute with the D^2 prediction in two dimensions (equation 2.78).

7.1.3 The Paramagnetic Transition

We have not paid particular attention to the effect of the anisotropy on the transition to paramagnetism around T_{KT}; there remains dispute in the literature on the exact signature of the transition in the pure system, and we did not attempt to compound the difficulties with the addition of randomness.

The bump in the specific heat for the pure system is an anomalous feature, occurring *above* the phase transition. We find that the bump does not move for small anisotropy strengths, and we interpret this evidence as support for the survival of the Kosterlitz-Thouless phase at intermediate temperatures. For larger anisotropy strengths, the peak does broaden and move to higher temperatures. We have not investigated this behaviour in detail.

7.1.4 The Low Temperature Phase

We have tended to refer to the low temperature phase as 'glassy' without really attempting to justify the use. The data from the quenched simulations implies some irreversibility at low temperature, and the magnetic order has been reduced by the anisotropy.

A true spin glass shows many ground states separated by energy barriers whose magnitude increases with system size, so that the dynamics of the system depend on the timescales used. Our results certainly imply energy barriers, but we do not have sufficient data to form conclusions on the size dependence. But the measurements on vortex lifetimes certainly indicate the presence of energy barriers, which can be overcome with sufficient thermal energy.

We also see static changes in behaviour at temperatures which appear to coincide with the glassiness. In particular, the spin correlations and the finite-size scaling of the magnetisation indicate a qualitative change in behaviour. There are no pinned vortices in systems cooled to these temperatures, which precludes an effect solely due to topological defects due to dimensionality of the system.

7.2 Three Dimensional System

In three dimensions we find that the long-range order found in the pure system at low temperatures is destroyed by the addition of the random anisotropy, as predicted by most workers in this field. We deduce that there remains a phase transition, separating a paramagnetic régime at high temperature from a low-temperature state with algebraic correlations, as predicted by Aharony and Pytte [37]. We find, however, non-universal behaviour, with the exponent dependent on D. Other work has been done in the infinite-anisotropy limit with the justification that system parameters can be renormalised.

We have been able to fit a scaling curve to the specific heat peak, in order to calculate a crossover scaling exponent for the transition temperature, though for small system sizes the rounding is obscured by finite-size effects — this will have adversely affected Reed's work on this system around T_c [51].

Algebraic decay of spin correlations corresponds to a phase with infinite magnetic susceptibility, as was predicted by Aharony and Pytte [37]. This result was also found by Fisch in simulations in the Ising limit [45, 47].

It appears that the equilibrium low-temperature phase does *not* agree with the predictions of CIM for small anisotropy:

- We find evidence for an infinite-susceptibility phase, which conflicts with the prediction of a finite anisotropy by CIM and others.
- The energy is found to contain a component proportional to D^2 for $D < 5$,

whereas equation 2.75 predicts D^4 in three dimensions. The anisotropy component of the energy from the liquid crystal order parameter measurements confirm the D^2 behaviour.

- In a large external field, we are unable to observe the FWA state predicted in section 2.8.7

For larger anisotropy strengths, our results do agree with predictions made for the limit where the cluster size falls to the lattice spacing (sections 2.8.3 and 2.8.8).

Simulations on systems quenched from high temperature *do* show some CIM properties. In particular, the spin are correlated over length scales characteristic of the predictions in section 2.8.

7.2.1 The Low Temperature Phase

As in the two-dimensional system, we find significant irreversible effects at low temperatures: the properties of the system cooled to low temperature are very different to those of systems quenched as if from high temperature. In particular, the magnetisation of the quenched (finite) system is strongly suppressed below that of the cooled system, and the system dynamics are extremely slow. Only after the system is warmed above a D-dependent *depinning temperature* do the dynamics speed up, and the cooled and quenched systems agree.

The simple argument in section 2.8.4 predicts that $T^* \sim D^4$ in three dimensions, but the separation line in figure 6.26 is much closer to linear. However, because we find that the D-dependence of energy is not consistent with the Chudnovsky-Imry-Ma prediction of D^4, but does agree with the two-dimensional observation D^2, we should not be surprised to find a similar irreversibility curve dependence.

In contrast to the two-dimensional system, however, the glassiness is not coupled with a qualitative change in behaviour. The onset of glassiness in the two-dimensional system was associated with a change in the spin correlations and finite-size analysis, but there is no such change here — the irreversibility in the three-dimensional system appears to be entirely a dynamic effect.

In the two-dimensional system, we found a qualitative change in the nature of the spin correlation function at low temperature. Here, we find that the correlations are algebraic both above and below T^*, with the same power of decay. The only indication of the apparent freezing is that the amplitude of the correlation function stops changing with temperature around T^*, giving a 'bunching' of the size-scaling curves very similar to the two-dimensional system.

7.3 Possible Implementation Enhancements

In section 4.2, it was explained that some effort was expended to host the transputer system from the group's microVAX, since the pc 'operating system' was just not sufficiently flexible or stable to host simulations lasting days or weeks. The code was subsequently modified so that some of the program actually ran on the vax host, allowing spin correlations to be calculated on a machine with sufficient memory to hold all spins at once, while the transputers continued to generate new configurations.

It has already been noted that one possible enhancement to the program would be to use one of the more recent cluster-update algorithms: since the random anisotropy magnet is expected to form clusters, a cluster update algorithm seems a natural choice. As noted in section 4.4.5, it might be simpler to send the spins to the (sequential) host for the cluster update part of the simulation, particularly since the group's latest alpha workstation gives an order of magnitude in computing power over the vax, and can be used as the transputer host with nothing more than a recompile of the software. However, if large quantities of data need to be transferred between transputers and host, it is likely that the bridge interface would turn out to be a bottleneck, and the communications overhead would be become significant.

Recently, Linux, a free unix implementation for the PC, has become available, and the group has found this sufficiently stable that a 486 *could* be trusted to host the transputers for the days or weeks required for a simulation. So it would be worth investigating putting the transputer interface card into a pc and recompiling the code – a 486 is intermediate in performance between microVAX and DEC alpha, while the direct connection reduces the communications bottleneck between transputers and host. The group's linux PC's are also on the internet, so that accessibility to the system is not lost. A little more effort would allow several PCs and even the alpha to cooperate on some aspects of the problem while the transputers continue to generate spin configurations.

7.3.1 Graphics

While the real-time graphics were very illuminating for the two-dimensional system, visualising the three-dimensional system as two-dimensional slices was not particularly useful, since the simulations generally use an order of magnitude more spins, and therefore took an order of magnitude longer. It would probably be better to dump the current configuration every so often, then recreate the graphics 'offline' — it is possible that there are packages around which would allow sophisticated navigation through the system as it changes.

7.4 Parallel Computers

It was stated at the beginning of this thesis that one of the goals of the work was to investigate the use of parallel computers for monte-carlo simulations. We were very fortunate to inherit a 32-node transputer system of which we had sole use, putting a significant amount of computing power at our disposal — there is no doubt that we could not have performed simulations over such a wide region of phase space without this. However, inevitably, the hardware manufacturers have continued to increase the performance of their CPUs, so that the group's (sequential) DEC alpha workstation now performs with a similar performance, but is of course much easier to program, and doesn't carry the internal communications overheads. The manufacturers of dedicated parallel computers have not really kept up, and the next generation of transputers (the T9000) has still not been released. Specialist companies are, however, now looking into making parallel computers out of conventional chips — the mass market means that the price of the Intel 486/Pentium and the DEC alpha is much lower than the parallel CPUs, and there are parallel computers based on these. Standards are slowly emerging for message passing, etc. — it will make it much easier to justify investing time and resources in the development of parallel programs if there is some guarantee that such a program can be ported to other manufacturer's parallel computers. As it is, the existing OCCAM code for this simulation is unlikely to be of use on any other contemporary parallel computers, though the (sequential) front end code which parses and actions the simulation control files can almost certainly be salvaged and reused.

Parallel computers will always give an unbeatable price/performance ratio. But the effort required to make the most of them remains prohibitive, and since sequential computers increase in performance by an order of magnitude every few years, it is probable that the effort is not justified. The most efficient way of using a parallel computer is to run a separate job on each node, since this avoids communication overheads. (This work could not be organised this way because if insufficient memory per node.) And if each processor is running independently, there is no need for specialist parallel hardware; VMS is a particularly good operating system from this respect since batch queues can be distributed over a cluster of machines, and there are several batch systems emerging for unix. It is a matter of personal frustration that considerable computer power lies untapped in the physics dept. at the University of Southampton; a large number of 486's in the various offices spend almost all day and certainly all night lying idle, when they could so easily be doing useful simulations in the background.

Appendix A

Simulation Control Language

The details of the simulation control language are presented here, in the hope that this work can be continued with the existing software. The code will be made available by anonymous ftp from sotona.phys.soton.ac.uk [152.78.192.42] though it is possible that lack of disk space will require that the files be moved to tape — instructions will be left there for how to proceed.

A.1 Introduction

The control language is a description of a simulation in terms of low level primitive operations. The controlling process, which may be running on either a transputer or the host system, reads and parses a control file, and sends out a stream of instructions to the workers to implement the simulation. In this implementation, the primitives are messages sent to the transputer network for action, but on a sequential machine the same primitives could be actioned as subroutine calls. In this way, the details of the parallel nature of the program are hidden from the user.

The code was developed in this way because a parallel program is implemented as a number of processes executing simultaneously – changes to the flow of the simulation required changes to several files, which was tedious and error prone, and required long recompile times. Because the compiler runs on the transputers, the code could not be recompiled while a simulation was running.

Another advantage of the textual control file is that it can be written to the results file in its entirety, so that each data file contains full details of the

Variable	Use
B	External field strengh
C	External field direction (0–2π)
D	Anisotropy
T	Temperature
K	System depth (in three-dimensional simulations)
L	System width
M	System height
N	System size $= L \times M \times J$
W	Number of workers
Z	Number of quantisation axes
U	System energy
X	Magnetisation in x direction
Y	Magnetisation in y direction
Q	Liquid crystal order parameter
V	Vortex density
W	Half-vortex density

Table A.1: variables used by the simulation control language

simulation which generated it, minimising danger of mixing up the large number of data files produced.

A.2 Variables

The control language has 52 variables, A to Z and a to z. The upper case letters are reserved for reading or setting system parameters, and changes made to them are conveyed to the rest of the simulation code, while the lower case letters are used for loop control, and as temporary workspace in the controlling process only. Table A.2 shows the system variables currently implemented. The variables in the first section are used to control the simulation. Those in the second section are set up by the code-initialisation section, so that the script can customise itself – for example, it may do correlations less frequently on a large system. In this implementation, these are supplied to the controlling process by the transputer system during initialisation, since the transputer part of the code has the size compiled in, for speed. The last section shows the variables which are updated by the *getstate* primitive, and which are written to the file during a *study*.

A.3 Primitive operations

The primitive operations are described here, in the order they might typically be used in a simulation – the first few are typically done once to configure the simulation, and the rest might make up the body of a loop over temperature or field. *expression* denotes a reasonably simple expression involving variables and constants, with operators +,−,∗, / and ∧ (power), and with functions exp, log and sqrt. *direction* can be 'x', 'y', 'h', 'd' or 'random' to set spins or anisotropy axes to x direction, y direction, aligned along field, aligned along anisotropy axes or random, respectively – an arbitrary direction can be selected by setting the field in that direction, then aligning with field. $i \ldots n$ denote integer constants. f denotes a file number. *string* denotes a comma-seperated list of elements "constant string", *nl* for end-of-line, or variable, with optional suffix for specifying required accuracy. Suffix is :n or _n where n is number of decimal places, and _n is used to replace the decimal point with _, which is required for numbers in directory names under VMS, for example.

seed *expression* sets the seed for the pseudorandom number generator

axes *direction* sets the direction of the anisotropy axes. (Cannot specify direction of 'd' here.)

spins *direction* sets the direction of the spins

cycle m n defines a cycle to be m sweeps through the lattice, attempting a rotation at each site, followed by n sweeps attempting a spin flip at each site. A cycle is the unit timestep of the simulation. n is typically 0 or 1 when D is small, and $n \approx m$ for larger D.

graphics n sets the number of cycles between graphics updates

correlations m n enables correlations to be made every m cycles during the simulation. If $n \neq 0$, only n sites are chosen at random to be compared with all the others in each configuration – for large systems, this allows the same computational effort to sample correlations from a larger number of configurations than studying each configuration exhaustively.

order f i directs that spin glass order parameter (equation 2.47) should be computed every cycle, with reference to the configuration stored in store number i, and written to file f

histogram r enables histogramming (section 3.7.2) with a resolution of r in the (single-site) energy

plot T,i,... selects which variable(s) are written to file during a study

open f *string* opens file number f with given name

tofile f *string* writes a line to file f

write *string* writes line to screen

wait n performs n cycles without measurement for equilibriation

store i stores current configuration to store number i

restore i restores the current configuration from store i

study f m n performs n sequential studies each of m cycles, writing n averaged values to file f – making several measurements in this way allows some check on any drift in parameters over time. This could be done using combinations of *wait*, *getstate* and *tofile*, but since this where most time is spent, it was implemented as a primitive operation.

getstate performs a measurement on the current configuration, storing results in the system variables indicated in table A.1

writecorr f writes the accumulated correlations to file f

writehist f writes accumulated histogram results to file f

close f closes file f

writespins *string* write current spin configuration to named file.

writeaxes *string* write current anisotropy axes to named file.

repeat ...**until** *expression* simple flow control

if *expression*[**then** *action*] simple decision making

:*label* ...**goto** *label* usually used with *if* to build more sophisticated flow control

It should be noted that if the system is being brought to equilibrium without any measurements (**repeat**, **wait** n, **until**...), then a **getstate** should be put into the loop, to prevent a build-up of **wait** commands in the buffers (since **repeat**...**until** is handled entirely within the front end, and **wait** does not required the front end to wait until the transputers have finished. See section A.7.

A.4 File format

A standard file format evolved, where data was output in blocks, with each group optionally preceeded by a single number on a line as a header; this header was usually temperature, for cooled or quenched simulations, but could be field or time, depending on the simulation. If no header was given, the filename was parsed for a number; data was stored in files with names such as 'MX_D0.400' for M_x at $D = 0.400$.

is used in files to mark comments. Each file opened by the **open** primitive has the entire control file written to it as a block of comments, so that data files are self-documenting, in order to avoid confusion when a large number of data files are produced.

The **study** primitive outputs one complete set of parameters on each line, in the order: energy; specific heat; liquid crystal order parameter; magnetisation in the x and y directions, followed by mean, mean square and mean of fourth power; magnetic susceptibility in x and y directions, and net susceptibility; count of vortices and half-vortices. See section 4.5 for more details. These data were typically written to a file with a name like '_D0.400', and a simple program averaged the parameters over the different runs and wrote to files such as 'U_D0.400', 'C_D0.400', etc.

The spin glass order parameter are written simply as lines of order parameter against time. Initially, the correlations were written as blocks of M lines of L columns of ASCII numbers, but this was taking up too much filespace, especially when it came to three dimensions. Additionally, by writing only a few decimal places, information was lost. Laterally, the data was written as a block of binary data. One snag is that the VAX and the transputer use different internal representations for floating point numbers, so the block of binary data was preceeded by a comment line detailing whether it was IEEE or VAX floating point data.

Most of our simulations used the VAX as the host (section 4.4.4) so this has not been a problem in practice, but this may become relevent if the code is used further. (The group's DEC Alpha can be told to use IEEE floats as a compile-time option, so moving from VAX to Alpha may eliminate this problem.)

The **writehist** command writes each block of histogram data as lines of increasing energy, with each line containing: the (unnormaslised) count; the energy of the bin; the mean, mean square and mean of the fourth power of the magnetisation; and the liquid crystal order parameter.

The **writespins** and **writeaxes** primitives do not write the entire simulation control file as comments because these output files were intended for plotting by a postscript program, and parsing is not a strongpoint of this language.

A.5 Data Analysis

A suite of simple analysis programs was written to allow various manipulations to the data to be made. Because of the standard file format, a library of routines was implemented to read the data file, parse file name for temperature, parse the command line for standard options, and so on, so that simple utilities could be put together with little effort. These programs will be made available by anonymous ftp with the rest of the code — see section A. Each of the following utilities provides a usage summary when given -? as an argument:

analyse.c : reads the output from the **study** primitive, writing each averaged parameter and standard deviation to a file whose default name is the concatenation of the parameter to the input file name. Typically, the input name was of the form '_D0.400' and output went to 'U_D0.400', 'C_D0.400', etc..

docorr.c : performs a similar role for the correlations; reads the raw binary data, takes a fourier transform, performs a circular/spherical average and writes the output as three columns r, $g(r)$ and $G(k)$.

analhist.c : analyses histogram data output, allowing extrapolation of data to nearby temperatures, and outputting either shifted histograms or thermal averages at the new temperature. Care is taken to avoid numerical overflows by only ever scaling numbers down; silent truncation to zero is permitted.

It should also be noted that the full system energy must be used, not the energy per spin. [This caused much confusion initially.]

chngevar.c : reads a set of files representing different values of a parameter, changing variable. For example, given files such as M_D0.000, M_D0.100, M_D0.200. . . , each containing data for various temperatures, this program will output M as a function of D, with each different temperature as a different data set.

seperate.c : typically used on the output from chngevar, this program takes an input file comprising several data sets, and writes each set to a different output file. The program takes as a parameter a template for the output filenames.

manip.c : allows arbitrary manipulations to be applied to the data files in a general way. For example, this was used to calculate the direction of magnetisation from the x and y components, for section 5.5.3.

A.6 User Guide

Once an account on the group's VMScluster has been allocated, it is a simple matter to perform simulations using the existing software. Access to the system is available from anywhere on the internet; the VAX is accessed through the DNS name **sotona.phys.soton.ac.uk** [152.78.192.42].

A local command **LEARN** is provided for setting up system extensions on a per-process basis. Lines of the form

```
$ LEARN BATCH MAKE TOOLSET
$ EXEC == "$SYS$DISK:[]'"
```

should be put into your LOGIN.COM file, so that in each session the local commands **exec**, **make** and **batch** are available, as well as the compilers for the transputer. Online help is available:

HELP MEIKO describes the transputer toolset, with some local implementation details which, if nothing else, should contain sufficient technical information for the University's High Performance Computing Group (formerly the Parallel Computing Support Group) to give assistance. Manuals and courses on parallel programming and the transputer toolset are available from them.

HELP MAKE gives information on the local version of the unix 'make' utility.

HELP BATCH gives information on the local 'batch' command.

exec is the way we usually invoke c programs to give them parameters. (It's a bit of DCL magic that is beyond the scope of this manual !)

The program comes in two parts: an OCCAM part which runs on the Meiko; and a C front end which usually runs on the vax, but can also be run on a transputer by making a small change to the configuration file 'xy.cfs'.

A.6.1 Organisation of the Source Code

Because there is a lot of overlap between the two- and three-dimensional code, the common code is placed in one directory, with subdirectories with the dimension-specific code.

The code is in two distinct parts; the C front end compiles and runs the simulation control language, sending messages to the OCCAM program which actually performs the simulation. These two parts are compiled separately.

The front end consists of the files

header.h a shared header file, defining structures, the message protocol, etc.

main.c the main program, which co-ordinates activity and runs the simulation.

parse.c contains the routines required to compile the simulation control file into an internal p-code.

actions.c contains the subroutines to action the simulation control primitives. Each primitive corresponds to a subroutine in actions.c, for the most part.

The front end is compiled for two or three dimensions, selected by a macro in header.h. It should be possible in principle to reuse this language compiler by replacing actions.c with subroutines appropriate to different simulations —

the actual simulation work can be put in actions.c, rather than simply sending messages to another program.

The simulation code is compiled from one of the subdirectories. For several files. the main code is in a file in the parent directory, which is included into the compiled code. Frequently, the code in the parent directory does an include for code in the subdirectory — this will pick up the version from the appropriate subdirectory.

The OCCAM source is split into the following files, some of whos names are historical. The reader is referred to figure 4.1.

xy.cfs contains the transputer configuration description, mapping the processes in the source to the physical transputers.

mstrschd.occ is the 'master' scheduler, which communicates with the front end, forwarding commands to the workers and performing averages and histogramming. This file includes mstrschd.occ in the parent directory, which in turn sources mstrschd.inc and mstr2.inc

sched.occ implements the routing harness illustrated on figure 4.1. It routes commands to and results from the workers

worker.occ actually performs the simulations. Files tables.inc, random.occ, settable.occ are read from the parent directory. In addition, the inner loops which are handcoded in assembler are in seperate files state_mc.occ, flip_mc.occ and mc_mc.occ, in order to make it easier to substitute occam versions state_occ.occ, flip_occ.occ and mc_mc.occ for testing and development. The machine code versions also optionally include the occam versions, to verify that the machine code and occam versions produce the same answers.

graphics.occ drivers the graphics hardware

A.6.2 Recompiling

The code must be recompiled if the system size is to be changed. It is necessary to set the system size in file consts.inc, and set the number of transputers in xy.cfs. A version of the unix utility *make* is used to recompile the program; simply typing

```
$ MAKE
```

at the DCL command line will recompile the program.

A.6.3 Preparing to Simulate

Any text editor can be used to prepare simulation control language files — editors available include **fue**, a folding editor derived from micro-emacs, as well as the standard VMS editors EDT and EVE. See the example control files in section A.7. Failing that, the files can be prepared on any other machine and sent to the VAX by ftp.

A.6.4 Running the Simulation

Two steps are used to run the program in this implementation:

- The transputer part of the code is *booted* onto the Meiko using a modified version of the standard **iserver** transputer server:

  ```
  $ ISERVER /SR /SG WIRING.DAT /SC XY.BTL
  ```

  which first of all configures the transputer links according to the description in the file 'wiring.dat'. There are files 'wire4.wir', 'wire8.wir', 'wire16.wir' and 'wire31.wir' which contain the setup for the usual system sizes. Further details are available with HELP MEIKO CONFIGURING.

- The VAX end is contained in an executable XY.EXE — this can be run interactively using

  ```
  $ EXEC XY file.scf [T=<temp>] [D=<anis>] ...
  ```

  while developing a simulation. Variables can optionally be set on the command line using variable=value, as indicated. Because DCL is case insensitive, only upper-case variables can be set in this way. Production runs are submited to the VMS queue MEIKO$BATCH by simply prepending BATCH/MEIKO to the command line:

```
$ BATCH/MEIKO EXEC XY file.scf ...
```

though normal DCL command procedures can be SUBMIT-ed as usual if that is preferred.

It is not necessary to boot the transputer code for each run; provided each simulation ends without error, the code will automatically restart itself in preparation for the next run, so that several runs can be submitted, with different anisotropy strengths for example. If there are several users using the MEIKO$BATCH queue, then it would of course be necessary to reboot each time, since it must be assumed that other jobs will run between your jobs; there is a local command **MBATCH** which allows the boot and run to be controlled from one line:

```
$ LEARN MBATCH
$ MBATCH "iserver/sr/sg wire31.wir /sc xy.btl" "exec xy cool.scf"
```

which will do both commands in batch. Alternatively, a DCL file containing the lines

```
$ SET DEF [directory]
$ ISERVER /SR /SG ...
$ EXEC XY file.scf ...
$ EXIT
```

may be submitted to the queue in the usual VMS way (HELP SUBMIT). SET DEF is the VMS equaivalent of the DOS/unix cd command. (There is a local CD command available —HELP CD for details). The BATCH command automatically arranges for the batch job to run in the directory from which the command was submitted. It is even possible for the simulation control file to be sent inline to the front end:

```
$ ISERVER ...
$ EXEC XY SYS$INPUT ...
# control file goes here
set T=1.4
repeat
```

```
  wait 100
  getstate
  set T=T-0.1
until T<0.1
end
$ EXIT
```

and DCL will send all the lines from the EXEC line to the next line beginning with $ to XY.EXE as the control file (SYS$INPUT). As always, there are many ways of doing things...

A.7 Example Programs

Two example simulation control programs are presented. Figure A.1 shows the usual sort of simulation performed, where the system starts at high temperature, and is cooled, performing simulations at each successive temperature. Figure A.2 shows a more novel use of the language, where the lifetime of the vortices in an initially quenched system is measured. Ten independent measurements are made, and the mean and standard deviation are output. This is the program used for the results in section 5.5.1.

```
# Example program 1

# randomise the axes and spins
seed 91
axes random
spins random

#define each cycle as four rotation sweeps and 1 flip sweep.
cycle 4 1

# D defaults to 0.4 unless it was set on command line
if D==0 then set D=0.4

# write results to a file who's name depends
# on anisotropy strength, in subdirectory
# [.cool_<L>] (vms syntax) (eg [.cool_62]_D0.400)

open 0 "[.cool_", L:0, "]_D", D:3

#start at high temperature
set T=1.4

#loop over temperature
repeat
  write "T=",T,nl  # inform user how far we've got
  wait 1000        # equilbriate
  study 0 2000 5   # 5 averages over 2000 cycles each, to file 0
  set T=T-0.1
until T<0.05       # allow for rounding errors

#close file
close 0

end
```

Figure A.1: Typical simulation control program, where the system is progressively cooled, with simulations run at each temperature.

```
# Example program 2
open 0 "vortex_life.dat"
repeat
  set T=0.5
  repeat
    set s=0       # sum of life
    set q=0       # sum of squares
    set i=0       # number of measurements made
    repeat
      set l=0     # lifetime of these vortices
      axes random
      spins random
      repeat
        wait 50
        set l=l+50
        getstate
        v=V*N
        if l>=99999 then set v=0  # dont wait any longer
      until v==0
      write " annihilated after ", l:0, nl
      # add this result into the averages
      set s=s+l
      set q=q+l*l
      set i=i+1
    until i==10   # ten measurements for each pair of parameters
    # work out the statistics...
    set m=s/10
    set v=sqrt((q/10)-(m*m))
    # .. and write to file 0
    tofile 0 D:1," ",T:1," ",m," ",v, nl
    set T=T-0.1
  until T<0.05
  tofile 0 nl
  set D=D+0.1
until D>1.05
end
```

Figure A.2: Control program for a more novel simulation: average the vortex lifetime over ten sets of axes and initial configurations for each T and D.

Bibliography

[1] D. R. Denholm and T. J. Sluckin, Phys. Rev. B **48**, 901 (1993).

[2] D. R. Denholm, T. J. Sluckin, and B. D. Rainford, J. Mag. Mag. Mat. **104**, 103 (1991).

[3] D. R. Denholm, T. J. Sluckin, and B. D. Rainford, Acta Physica Polonica **25**, 219 (1992).

[4] B. I. Bleaney and B. Bleaney, *Electricity and Magnetism* (Clarendon Press, Oxford, 1957).

[5] N. Ashcroft and N. D. Mermin, *Solid State Physics* (Saunders College, Philadelphia, 1988).

[6] C. Kittel, *Introduction to Solid State Physics* (Wiley, New York, 1986).

[7] K. H. Fischer and J. A. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, 1991).

[8] R. Harris, M. Plischke, and M. J. Zuckermann, Phys. Rev. Lett. **31**, 160 (1973).

[9] K. M. Lee and M. J. O'Shea, J. Appl. Phys. **63**, 3740 (1988).

[10] B. Dieny and B. Barbara, Phys. Rev. Lett. **57**, 1169 (1986).

[11] D. J. Sellmyer and S. Nafis, Phys. Rev. Lett. **57**, 1173 (1986).

[12] S. F. Edwards and P. W. Anderson, J. Phys. F **5**, 965 (1975).

[13] L. Onsager, Phys. Rev. **65**, 117 (1944).

[14] Ma, *Statistical Mechanics* (World Scientific, Singapore, 1985).

[15] V. L. Berezinsky, Zh. Exsp. Teor. Fiz. **59**, 907 (1970).

[16] J. M. Kosterlitz and D. J. Thouless, J. Phys. C: Solid State Phys. **6**, 1181 (1973).

[17] J. Tobochnik and G. V. Chester, Phys. Rev. B **20**, 3761 (1979).

[18] J. M. Kosterlitz, J. Phys. C: Solid State Phys. **7**, 1046 (1974).

[19] L. P. Kadanoff *et al.*, Rev. Mod. Phys. **39**, 395 (1971).

[20] P. L. Doussal and T. Giamarchi (unpublished).

[21] Y. Imry and S. Ma, Phys. Rev. Lett. **35**, 1399 (1975).

[22] E. M. Chudnovsky and R. A. Serota, Phys. Rev. B **26**, 2697 (1982).

[23] E. M. Chudnovsky and R. A. Serota, J. Phys. C: Solid State Phys. **16**, 4181 (1983).

[24] E. M. Chudnovsky, W. M. Saslow, and R. A. Serota, Phys. Rev. B **33**, 251 (1986).

[25] E. M. Chudnovsky, J. Mag. Mag. Mat. **79**, 127 (1989).

[26] E. M. Chudnovsky, J. Appl. Phys. **64**, 5770 (1988).

[27] J. M. D. Coey, J. Appl. Phys. **49**. 1646 (1978).

[28] R. Harris and D. Zobin, J. Phys. F **7**, 7 (1977).

[29] E. Callen, Y. J. Liu, and J. R. Cullen, Phys. Rev. B **16**, 263 (1977).

[30] J. D. Patterson, G. R. Gruzalski, and D. J. Sellmyer, Phys. Rev. B **18**, 1377 (1978).

[31] B. Derrida and J. Vannimenus, J. Phys. C: Solid State Phys. **13**, 3261 (1980).

[32] J. H. Chen and T. C. Lubensky, Phys. Rev. B **16**, 2106 (1977).

[33] V. S. Dotsenko and M. V. Feigelman, J. Phys. C: Solid State Phys. **14**, L823 (1981).

[34] V. S. Dotsenko and M. V. Feigelman, J. Phys. C: Solid State Phys. **16**, L803 (1983).

[35] J. L. Cardy and S. Ostlund, Phys. Rev. B **25**, 6899 (1982).

[36] A. J. Bray and M. A. Moore, J. Phys. C: Solid State Phys. **18**, L139 (1985).

[37] A. Aharony and E. Pytte, Phys. Rev. Lett. **45**, 1583 (1980).

[38] A. Aharony and E. Pytte, Phys. Rev. B **27**, 5872 (1983).

[39] M. C. Chi and R. Alben, J. Appl. Phys. **48**, 2987 (1977).

[40] M. C. Chi and T. Egami, J. Appl. Phys. **50**, 1651 (1979).

[41] R. Harris and S. H. Sung, J. Phys. F **8**, L299 (1978).

[42] R. Harris, J. Phys. F **10**, 2545 (1980).

[43] C. Jayaprakash and S. Kirkpatrick, Phys. Rev. B **21**, 4072 (1980).

[44] A. Chakrabarti, J. Appl. Phys. **63**, 3735 (1988).

[45] R. Fisch, Phys. Rev. B **39**, 873 (1989).

[46] R. Fisch, Phys. Rev. B **42**, 540 (1990).

[47] R. Fisch, Phys. Rev. Lett. **66**, 2041 (1991).

[48] R. Fisch, Phys. Rev. B **46**, 242 (1992).

[49] P. Reed, J. Phys.: Condens. Matter **1**, 3037 (1989).

[50] P. Reed, J. Phys. A: Math. Gen. **24**, L1299 (1991).

[51] J. Reed, J. Phys. A: Math. Gen. **24**, L117 (1991).

[52] B. Dieny and B. Barbara, Phys. Rev. B **41**, 11549 (1990).

[53] W. M. Saslow and N. C. Koon, Phys. Rev. B **49**, 3386 (1994).

[54] N. Metropolis *et al.*, J. Chem. Phys. **21**, 1087 (1953).

[55] M. E. Fisher, in *Critical Phenomena*, edited by M. S. Green (Academic Press, New York, 1971).

[56] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett. **58**, 86 (1987).

[57] U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).

[58] A. Ferrenberg and R. Swendsen, Phys. Rev. Lett. **61**, 2653 (1988).

[59] C. R. Askew, Ph.D. thesis, University of Southampton, Department of Physics, 1989.

[60] C. A. R. Hoare, *Communicating sequential processes* (Prentice-Hall International, Hemel Hempstead, 1985).

[61] M. Flanigan and P. Tamayo, Int. Jour. of Modern Physics C **3**, 1235 (1992).

[62] R. Dickman and E. M. Chudnovsky, Phys. Rev. B **44**, 4397 (1991).

[63] D. J. Cleaver, S. Kralj, T. J. Sluckin, and M. P. Allen, Random Anisotropy Nematics in *Liquid crystals in finite geometries formed by polymer and porous networks*, edited eds. G. P. Crawford and S. Žumer (Wiley, 1995) to be published.