

Chapter 1

Basic classical statistical mechanics of lattice spin systems

1.1 Classical spin systems

The topic of this chapter is classical spin systems on the lattice. Classical spin systems are idealized versions of magnets. Although many magnetic phenomena in materials are inherently quantum mechanical, a many properties are well described at least qualitatively by classical spin systems. In fact, we will see in the next chapter that one can obtain quantum spin systems by taking a limit of classical ones! This is not a typo; the catch is that for this to work, the classical system must be in one dimension higher.

Each degree of freedom of a spin system models a magnetic moment at a fixed location. A classical spin can be thought of simply as the angular momentum vector of a spinning particle, and so can be represented as an arrow. Even though classically this can in principle be of any magnitude, one typically requires that the magnitude of the angular momentum be a constant, so the arrow is of fixed length. (This requirement is of course natural in the quantum context, as will be discussed in the next chapter.) A given classical magnetic model the requires specifying:

1. Any constraints on where the arrow is allowed to point. Two famous examples are the “XY” model, where the arrow is constrained to point in a two-dimensional plane, and the “Ising” model, where the arrow is allowed to point in only two directions. More complicated examples are called sigma models, where instead of an arrow, the degrees of freedom take values on some manifold. In this language, the spins in the XY model take values on a circle.
2. Where the spins are. For example, they may be at each point of some lattice in some dimension, or they could be distributed continuously (i.e. an arrow for every point in space). In the latter case, the model would be called a “field theory”.
3. How the spin interact, i.e. the energy associated with each particular configuration

of spins. Two major types of interaction are called “ferromagnetic”, where the energy is lowered when two spins are aligned, and “antiferromagnetic”, where the energy is lowered when they point in opposite directions. A given model can include interactions of both types; much recent research has been devoted to studying this situation.

One important thing to note now is that although not all physical systems can be approximated well by a spin system or sigma model, their range of applicability extends much further than it might appear at first glance. It is fairly obvious that most (all?) bosonic systems can be modeled by some spin system. However, in one spatial dimension it is possible even to study fermionic systems using spin systems, a topic I will discuss in depth. In fact, this correspondence can sometimes be extended in some special cases beyond one dimension.

1.2 Why the lattice?

A central role in this book is played by spin systems defined on some lattice. It thus seems a good idea to say a little about why this is.

One reason is many many-body physical systems of interest can effectively be treated as being on a lattice. For example, many systems are often well treated by utilizing a “tight-binding” model, where each electrons are treated as being located at a fixed nucleus, and so live at particular points in space. They may hop from one model, or in some cases, they are at fixed locations. In the latter case, we can then ignore their charge and study only their spin. This situation then reduces to a (quantum) spin system on a lattice.

In fact, in many cases, defining models on a lattice opens up more possibilities for studying the physics. Models where the degrees of freedom take on discrete values in general can only be defined on the lattice. In some situations, the physics itself arises from the interplay between the degrees of freedom and the particular lattice they live on. A famous example is that of geometrical frustration, as exemplified by the two-dimensional antiferromagnetic Ising model on a triangular lattice. Calling the two directions allowed in the Ising model “up” and “down”, antiferromagnetic interactions make adjacent spins different. On the square lattice, it is possible to make a low-energy state by alternating up and down spins, but on a triangular lattice it is not: around a triangle, there must be at least two like spins adjacent to each other. Such a bond is said to be “unsatisfied”, and so the spins are “frustrated”. The resulting physics is completely different from the unfrustrated model.

One obvious reason why working on the lattice is that it makes the models easier to formulate, since probabilities and other physical quantities are defined in terms of sums instead of integrals. However, much of the point of statistical mechanics is show that physical quantities can be rewritten into integrals that can then be done in some fashion (either analytically, by a series, or by a computer). So a fair question is then:

even if it's easier to formulate a given model on the lattice, why not pass to a continuum formulation as quickly as possible, and then take advantage of the vast amount of tools of field theory?

In some situations, the answer is simply that one indeed can analyze the problem thoroughly using field theory. That's great, but the point of this book is to study situations where this is not so instant. I will discuss situations that can be analyzed thoroughly using field theory, but where the correct field theory to use is not immediately apparent. For example, in the two-dimensional (unfrustrated) Ising model, the most useful field theory is that of free fermions, a fact hardly obvious at first glance. Even subtler cases are those of frustrated models, where it is not obvious what their correct effective degrees of freedom are.

In fact, the study of lattice models has value well beyond just ease of formulation; in many cases lattice models a *precise* method of formulating a problem that is singular in the continuum. A more formal way of stating this in field theory language is that putting a model on the lattice provides a “regulator” for various infinities that appear in certain field-theory formulations.

To understand this point, however, does not require knowing any field theory, but can be nicely illustrated in the XY model. The XY model appears in many different contexts, but the simplest formulation is to think of each degree of freedom as a magnet that can point in any direction in the plane, i.e. an arrow. Equivalently, we can say each degree of freedom takes values on a circle, each point on the circle corresponding to one direction of the arrow. Depending on the physical situation of interest, there may be a degree of freedom for every point in continuous space, or on a lattice. One of the interesting features of the XY model in two dimensions is that it has vortices. A vortex is a configuration of arrows such that as one moves around a circle in real space, the arrow rotates around by an angle $2n\pi$, for n some non-zero integer. If real space is a continuum instead of a lattice, the a vortex configuration necessarily contains a point where the magnitude is singular (the arrows at a single point cannot point in all directions). To resolve this, one must complicate the problem further by some method, for example by allowing the length of the vector to vary. Putting the model on a lattice is for most purposes the simplest way of resolving this problem.

Another situation I will discuss is a special class of classical models in two dimensions (or quantum models in one dimension) called integrable models. Here there are powerful methods to do exact computations even when interactions are very strong. Some of the methods work on the lattice, some in the continuum. Both enable the derivation of results that the other can't.

I saved the simplest reason for last. As has been known to condensed-matter physicists for nearly a century, if one Fourier transforms variables defined on a lattice into momentum space, the momenta take on continuous values. The range of allowed values is called the “Brillouin zone”. Thus we end up with a model defined on a continuous space, just like in field theory. Here the boundaries of the space are determined by which lattice the original model is defined on instead of which box your system is in, but it is

still a continuous space. Of course, if one then does quantum mechanics in a finite box on the lattice the allowed momenta are then quantized, and we end up back on a lattice!

1.3 The partition function and correlators

The basic object in classical statistical mechanics is the *Boltzmann weight* $P(n)$. In thermal equilibrium, it gives the probability as the system is in a given configuration labeled by n as

$$P(n) = \frac{e^{-\beta E_n}}{Z} \quad (1.1)$$

where E_n is the energy of this configuration, $\beta = 1/k_B T$ is the inverse temperature (times the Boltzmann constant) and Z is an important object called the *partition function*. It is defined by the requirement that the probabilities sum to 1, so

$$Z = \sum_n e^{-\beta E_n} \quad (1.2)$$

If the degrees of freedom take on continuous values, or the model is defined in the continuum, then this sum is replaced by an integral. Note that indeed the probability of a given configuration increases as the energy gets lower, and that as the temperature gets higher and higher, the energies involved must get larger and larger to make a difference. Note also that if we shift all the energies by some constant E_0 , the probabilities are left unchanged: both the numerator and denominator are multiplied by the same constant $e^{-\beta E_0}$.

The point of this book is to study the statistical mechanics of *interacting* many-body systems, where E_n depends on mutual properties of the degrees of freedom. An important and simple example of such a system is the *Ising model*. This is a spin system where the spin is constrained to point in one of two directions. These directions are described by a variable $\sigma_i = 1$ for the "up" direction and $\sigma_i = -1$ for the down direction, where i labels the site of the lattice on which sits the spin. For a lattice of N sites, there are therefore 2^N different configurations in the model. The simplest interaction one can imagine is to assign one energy if neighboring spins are the same, and another if neighboring spins are different. This can be compactly described as

$$E(\{\sigma_i\}) = -J \sum_{\langle ij \rangle} \sigma_i \sigma_j . \quad (1.3)$$

The notation is as follows. Each configuration of spins (i.e. a set of N variables ± 1) is labeled by $\{\sigma_i\}$. Nearest-neighbor sites i and j are labeled $\langle ij \rangle$, so this sum is not over all sites, but rather over all "bonds", pairs of nearest-neighbor sites. (Amusing fact: on the square lattice, there are two bonds for every site.) The parameter J is the strength of the coupling, and is usually called the "coupling" for short. With the minus sign in

front, $J > 0$ corresponds to a ferromagnetic coupling, while $J < 0$ is antiferromagnetic. The partition function is then

$$Z = \sum_{\{\sigma_i = \pm 1\}} e^{-\beta E(\{\sigma_i\})}, \quad (1.4)$$

where this sum is over the 2^N different configurations. Despite this simple definition, the statistical mechanics of the Ising model has a wealth of interesting properties that will be explored in the pages to come.

From the partition function one obtains *expectation values* of physical quantities. An expectation value is an average over all the configuration of the system. For example, the average of the energy (sometimes called the “internal energy”, sometimes confusingly just called the “energy”) is

$$\langle E \rangle = \sum_n E_n P(n) = \frac{1}{Z} \sum_n E_n e^{-\beta E_n}. \quad (1.5)$$

Here I have used the standard physics notation of brackets to represent an expectation value even in non-quantum systems. This can be rewritten if desired as

$$\langle E \rangle = -\frac{\partial}{\partial \beta} \ln(Z).$$

The specific heat C is defined to be the *response* of the internal energy to a small change in temperature. In other words, it is the derivative

$$C = \frac{\partial}{\partial T} \langle E \rangle = k_B \frac{\partial}{\partial T} T^2 \frac{\partial \ln Z}{\partial T}. \quad (1.6)$$

It is also (and sometimes even more) illuminating to study a system by looking at its correlation functions, or *correlators* for short. A correlator, roughly speaking, describes how the degrees of freedom at one point are affected by the degrees of freedom at another point or points. For example, the basic correlator in the Ising model is the two-point function

$$\langle \sigma_a \sigma_b \rangle \equiv \sum_{\{\sigma_i\}} \sigma_a \sigma_b P(\{\sigma_i\}) = \frac{1}{Z} \sum_{\{\sigma_i\}} \sigma_a \sigma_b e^{-\beta E(\{\sigma_i\})} \quad (1.7)$$

for any two fixed points a and b . The product $\sigma_a \sigma_b = 1$ when the two spins are the same and $= -1$ when the two are different, so this correlator gives a notion of how the two spins are correlated; if $\langle \sigma_a \sigma_b \rangle = 1$, then the two spins are perfectly aligned; in every configuration the spins point in the same direction. Likewise, if $\langle \sigma_a \sigma_b \rangle = -1$ they are perfectly oppositely aligned. If this correlator is zero, then they are uncorrelated.

1.4 Three kinds of phases

Even in simple interacting systems, the two point function will not have such simple values as 0 or ± 1 except perhaps in some extreme limits. How it does depend on the separation is valuable information: the behavior of correlators is one of the best ways of understanding and characterizing different phases of matter. Consider a correlator defined so that in the non-interacting limit (or equivalently, the $T \rightarrow \infty$ limit) it vanishes. Letting L be the system size, there are three main types of behavior of such a two-point function as $L \gg |a - b| \gg 1$:

- **ordered:**

$$\langle \sigma_a \sigma_b \rangle \sim \text{non zero constant}$$

- **disordered:**

$$\langle \sigma_a \sigma_b \rangle \sim e^{-|a-b|/\xi}$$

- **critical:**

$$\langle \sigma_a \sigma_b \rangle \sim \frac{1}{|a - b|^{2x}}$$

Of course, the picture here is oversimplified: different types of correlators can occur in the same system. Nevertheless, for most systems this simple picture is exceptionally useful, so a few words here about each of the different types of phases is appropriate.

Order In an ordered spin system, the spins are correlated arbitrarily far from each other: if we know one spin is pointing in some direction, then we know that even far away, how another spin is most likely going to point. One can think of this as saying the spins “line up”, but this is a little imprecise. For example, in an antiferromagnet, order typically means that spins an odd number of sites apart are anti-aligned, while those an even number of sites apart are aligned.

One slight subtlety is worth noting. Often people characterize order in terms of a one-point function (of course of some variable whose thermodynamic average vanishes in the non-interacting limit.) In the Ising model, this is simply $\langle \sigma_a \rangle$, the probability that a given spin is either up or down. This sounds simpler to study than the two-point function, but the catch is that in many cases, including the Ising model, it vanishes for any value of a , b and J . This is easy to prove, and follows from the fact that $E(\{\sigma_i\}) = E(-\{\sigma_i\})$, i.e. the energy remains invariant if we flip all the spins. As long as the boundary conditions are invariant under spin flip, we have

$$-\langle \sigma_a \rangle = - \sum_{\{\sigma_i = \pm 1\}} \sigma_a e^{-\beta E(\{\sigma_i\})} = \sum_{\{s_i = \mp 1\}} s_a e^{-\beta E(\{s_i\})}$$

where $s_i = -\sigma_i$. However, σ_i and s_i are variables that we sum over, so we are welcome to rename this variable to σ_i . Thus the last expression is simply $\langle \sigma_a \rangle$, so that we have $-\langle \sigma_a \rangle = \langle \sigma_a \rangle = 0$.

Figure 1.1: A simple phase diagram.

A more general way of characterizing order therefore is to use the above definition: a non-vanishing value of a two-point function at arbitrarily large separation. An *order parameter* then can be defined as the square root of this absolute value.

Disorder In a disordered phase the two-point function class off to zero exponentially. Thus a convenient measure is the *correlation length* ξ . A useful heuristic is then to think of degrees of freedom closer than ξ as being correlated, and those further apart as being uncorrelated. One thing to bear in mind is that not all disordered systems are the same. We will discuss in the later parts of this book systems that are disordered by this definition, but which possess topological order, where an expectation value of some non-local quantities has a non-vanishing expectation value.

Criticality In many ways, this is the most difficult and most interesting of these cases. The correlator is said to be *algebraically decaying*. Notice in particular that in algebraic decay, the spatial dependence does not involve in any way a dimensional parameter, but only the dimensionless parameter x . Thus this behavior is characteristic of a critical phase, where the system is scale invariant: there is no parameter like a correlation length. The reason for the factor of two in the definition of x stems from this picture; in the continuum limit x can be thought of as the “scaling dimension” of the spin field.

A simple but common phase diagram in statistical mechanics is illustrated in figure 1.1. At temperatures much larger than any parameter with dimensions of energy (like J in the Ising model), one can effectively neglect the energy altogether. The partition function is then a sum over all allowed configurations with the same weight for each. The two-point function of any local operators vanishes quickly as they are brought apart (as long as non-local constraints are not present). The system is thus in a disordered phase at sufficiently large temperature. At low enough temperatures, the opposite occurs: since the Boltzmann weight is proportional to $e^{-E/T}$, the configurations with very low energy are much more likely. This can lead to order: the configurations that have the lowest energy have some regular pattern, and so there is a non-vanishing local order parameter. When there is a disordered phase at high temperature and an ordered phase at high temperature, there must be a phase transition in between. There are two categories of phase transitions, often referred to as *first-order*, and *continuous*. In the latter case, correlators become algebraically decaying, and the system is critical. In the former case, the system remains disordered at the phase transition, and abruptly changes its behavior

This book will discuss many examples of the phase diagram, and many examples where other behavior occurs. These are most prevalent in low dimensions, where it turns

out to be more difficult to have conventionally ordered phases, but where also occurs many more possibilities for interesting (and intricate) phase diagrams.

1.5 The transfer matrix, and solving the one-dimensional Ising model

To illustrate in a particular example some of the general ideas discussed so far, an excellent starting point is the one-dimensional Ising model. This is probably the simplest many-body system that can be solved easily and completely.

The way of solving the one-dimensional Ising model is to use the *transfer matrix*, a very convenient object utilized frequently when studying classical lattice models. The idea is very elegant: one singles out a single spatial direction, say the x direction. Then the partition function is built up by “evolving” the system from one value of $x = x_0$ to the next value $x = x_1$. Repeating this process yields the full partition function after all values of x are summed over. This method is not only useful for studying classical statistical mechanics, but by replacing this spatial evolution with evolution in real time, is fundamental in making the connection between quantum and classical statistical mechanics, as will be discussed in the next chapter.

In the one-dimensional Ising model, there is only a single site at a fixed value of x , so the transfer matrix therefore “evolves” the system from one spin to the next. Therefore, consider two nearest-neighbor sites in the Ising model. There are four configurations on these two sites: $++$, $--$, $+-$, and $-+$. The Boltzmann weights $w(\sigma_1\sigma_2)$ for these configurations with the energy in (1.3) are:

$$w(++) = w(--) = e^{\beta J}, \quad w(+-) = w(-+) = e^{-\beta J} .$$

Now consider the one-dimensional Ising model on three sites, and take the simplest boundary condition, where both spins σ_1 and σ_3 are both “fixed” to specific values. Summing over all values of the middle spin $\sigma_2 = \pm 1$, the partition function with both end spins fixed to be up is

$$\sum_{\sigma_2=\pm 1} w(+\sigma_2)w(\sigma_2+) = e^{2J} + e^{-2J} .$$

The general expression for the partition function with fixed boundary conditions on three sites is likewise

$$Z_3(\text{fixed}) = \sum_{\sigma_2=\pm 1} w(\sigma_1\sigma_2)w(\sigma_2\sigma_3) .$$

Likewise, in a four-site one-dimensional system, the partition function with fixed boundary conditions is

$$Z_4(\text{fixed}) = \sum_{\sigma_2=\pm 1, \sigma_3=\pm 1} w(\sigma_1\sigma_2)w(\sigma_2\sigma_3)w(\sigma_3\sigma_4) .$$

These expressions look exactly like matrix multiplication written out in terms of the matrix elements. The matrix being multiplied is called the transfer matrix and includes all the interactions between nearest-neighbor spins. For the one-dimensional Ising model, it is

$$T = \begin{pmatrix} w(++) & w(+-) \\ w(-+) & w(--)\end{pmatrix} = \begin{pmatrix} e^J & e^{-J} \\ e^{-J} & e^J \end{pmatrix} \quad (1.8)$$

To compute the partition function with fixed boundary conditions from the transfer matrix, the boundary conditions are treated as vectors on which the transfer matrix acts. (Treating boundary conditions as an element of a vector space turns out to have profound consequences that we will explore down the road.) Here, the basis elements of this vector space are the two spin values $+$ and $-$, which in the vector space correspond to

$$v_+ = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad v_- = \begin{pmatrix} 0 \\ 1 \end{pmatrix} .$$

The partition function of a system with just two sites and fixed boundary conditions (i.e. both spins fixed!) is then

$$Z_2(\text{fixed}) = (v_{\sigma_1})^T T v_{\sigma_2} = w(\sigma_1 \sigma_2) ,$$

where the T in the exponent is not the transfer matrix, but the transpose! Likewise for three sites and fixed boundary conditions, the earlier expressions are recovered by

$$Z_3(\text{fixed}) = (v_{\sigma_1})^T T^2 v_{\sigma_3} ,$$

while in general for N spins

$$Z_N(\text{fixed}) = (v_{\sigma_1})^T T^{N-1} v_{\sigma_N} , \quad (1.9)$$

Each time the transfer matrix is applied corresponds to adding one *bond* to the system.

The transfer matrix now can easily be used to find analogous expressions for other boundary conditions. *Free* boundary conditions allow the spins on the end to take on all allowed values, so

$$Z_N(\text{free}) = \sum_{\sigma_1=\pm 1, \sigma_N=\pm 1} (v_{\sigma_1})^T T^{N-1} v_{\sigma_N} = \begin{pmatrix} 1 & 1 \end{pmatrix} T^{N-1} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad (1.10)$$

A very useful boundary condition to take is *periodic*, where an extra interaction $-J\sigma_N\sigma_1$ is added to the energy, including a bond between the two end spins. In this case, the model is translation invariant: shifting all the spins by one site mod N does not change the energy. The transfer matrix must therefore “evolve” the first spin back to the N th one, so N powers appear. This is equivalent to having fixed boundary conditions on

$N + 1$ sites where $\sigma_{N+1} = \sigma_1$, and then summing over both possible values. The final expression is particularly elegant:

$$\begin{aligned} Z_N(\text{periodic}) &= \sum_{\sigma_1=\pm 1} (v_{\sigma_1})^T T^N v_{\sigma_1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} T^N \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} T^N \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &= \text{tr } T^N . \end{aligned} \tag{1.11}$$

The partition functions with different boundary conditions are simply related; e.g. $Z_N(\text{periodic}) = Z_N(++) + Z_N(--)$.

The beauty of using a transfer matrix in a one-dimensional system is that the partition function now can be computed simply by diagonalizing the transfer matrix. For the one-dimensional Ising model, the eigenvalues are $\lambda_1 = 2 \cosh(\beta J)$ and $\lambda_2 = 2 \sinh(\beta J)$ with respective eigenvectors

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 \\ -1 \end{pmatrix} .$$

Since the transfer matrix here is hermitian, the eigenvalues are real, and the left eigenvectors are the the transpose of the right ones. This means that for free boundary conditions,

$$Z_N(\text{free}) = 2(2 \cosh(\beta J))^{N-1}$$

while for periodic,

$$\begin{aligned} Z_N(\text{periodic}) &= (2 \cosh(\beta J))^N + (2 \sinh(\beta J))^N \\ &= (2 \cosh(\beta J))^N (1 + (2 \tanh(\beta J))^N) \\ &\approx (2 \cosh(\beta J))^N \quad \text{for } N \gg 1 . \end{aligned}$$

The approximate expression in the last line follows from the fact that $\tanh(\beta J) < 1$ for any finite value of βJ (i.e. at any non-zero temperature and finite coupling). In the free case, decompose the vectors v_{σ_j} in terms of the eigenvectors:

$$v_{\pm} = \frac{1}{2} \left(\begin{pmatrix} 1 \\ 1 \end{pmatrix} \pm \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right) .$$

to get

$$\begin{aligned} Z_N(++) = Z_N(--) &= \frac{1}{2} (2 \cosh(\beta J))^{N-1} (1 + (\tanh(\beta J))^{N-1}) , \\ Z_N(+ -) = Z_N(- +) &= \frac{1}{2} (2 \cosh(\beta J))^{N-1} (1 - (\tanh(\beta J))^{N-1}) . \end{aligned}$$

One useful thing to note is that the free energy per site $\propto \ln(Z)/N$ is independent of boundary conditions when N is large.

It is straightforward to compute the two point function (1.7) by writing them in terms of these partition functions. An interesting and simple one is the correlator between the

two end spins with free boundary conditions. There are four possibilities for the spin configurations on the ends. The free boundary conditions mean that all four need to be summed over, but the $\sigma_1\sigma_N$ in the correlator means that when these spins are opposite, there is a minus sign in front of these terms. This yields

$$\langle\sigma_1\sigma_N\rangle_{\text{free}} = \frac{Z_N(++)+Z_N(--)-Z_N(+)-Z_N(-)}{Z_N(\text{free})} = (\tanh(\beta J))^{N-1} .$$

This correlator falls off *exponentially* with distance, so even for small systems, the end spins are effectively uncorrelated. The general two-point correlator is similar. For periodic boundary conditions, breaking the sum into four parts in a similar fashion gives

$$\langle\sigma_a\sigma_b\rangle_{\text{periodic}} = \frac{2Z_{|a-b|}(++)Z_{N-|a-b|}(++)+2Z_{|a-b|}(+-)Z_{N-|a-b|}(-)}{Z_N(\text{periodic})}$$

where the products arise factors of 2 arise from the fact that $Z(++) = Z(--)$ and $Z(-) = Z(+)$. Plugging in the explicit expressions and letting $1 \ll |a-b| \ll N$ gives

$$\langle\sigma_a\sigma_b\rangle \approx (\tanh(\beta J))^{|a-b|} .$$

This expression is independent of boundary conditions in this limit, as long as a and b are far from the boundaries.

The two-point correlator is therefore

$$\langle\sigma_a\sigma_b\rangle \approx e^{-|a-b|/\xi} . \tag{1.12}$$

where the correlation length ξ is

$$\xi = -\frac{1}{\ln(\tanh(\beta J))} . \tag{1.13}$$

The correlation between two spin falls off *exponentially* with distance, so the one-dimensional Ising model is in a *disordered* phase for any non-zero temperature. The zero-temperature case is somewhat pathological, since there is no meaning to equilibrium in a zero-temperature classical system. In the next section, I describe how essentially all one-dimensional classical systems with local interactions are disordered.

1.6 The free energy

On very general grounds, a one-dimensional classical system cannot order. This follows from a simple and elegant argument, one version originally given by Peierls. To give this argument, it is useful to discuss a quantity fundamental to thermodynamics, the *free energy*. The definition of the free energy in statistical mechanics is quite simple:

$$F = -k_B T \ln(Z) . \tag{1.14}$$

The relation of the free energy to thermodynamics arises in the following fashion. First define the *density of states* as

$$\tilde{\rho}(E) = \sum_n \delta(E - E_n) \quad (1.15)$$

where $\delta(E)$ is the Dirac delta function. The density of states is therefore highly discontinuous, but the standard assumption, valid most (but not all!) of the time, is that there is a energy scale ϵ much smaller than the interaction energies (e.g. J in Ising), but much larger than the typical spacing between energy levels. If this assumption holds, then the density of states can be “smoothed” simply by averaging its values over all the regions of size ϵ . Namely, this smoothed density of states is defined so that $\rho(E)\epsilon$ gives the number of states in a region of size ϵ around the energy E , i.e.

$$\rho(E) = \frac{1}{\epsilon} \int_{E-\epsilon/2}^{E+\epsilon/2} d\tilde{E} \tilde{\rho}(\tilde{E}) .$$

The partition function can then be rewritten as

$$\begin{aligned} Z &= \sum_n e^{-\beta E_n} = \int dE \tilde{\rho}(E) e^{-\beta E} \\ &\approx \int dE \rho(E) e^{-\beta E} = \int dE e^{-\beta(E-TS(E))} . \end{aligned} \quad (1.16)$$

where the *entropy* is defined as¹

$$S(E) \equiv k_B \ln(\rho(E)) . \quad (1.17)$$

The next assumption to make is that the integrand in (1.14) is sharply peaked around some value of E . With this assumption, the expectation value of the energy is precisely this value of E , and the quantity

$$\mathcal{F}(E) \equiv E - TS(E) \quad (1.18)$$

has a *minimum* at $E = \langle E \rangle$. Thus for E near $\langle E \rangle$,

$$\mathcal{F}(E) = \mathcal{F}(\langle E \rangle) + T\alpha(E - \langle E \rangle)^2 + \dots$$

where $\alpha = -\partial^2 S / (\partial E)^2 / 2|_{E=\langle E \rangle}$. Plugging this into (1.14) gives

$$Z \approx e^{-\beta \mathcal{F}(\langle E \rangle)} \sqrt{\frac{\pi k_B}{\alpha}}$$

¹I am being slightly sloppy, since ρ as defined here is dimensionful and one really shouldn't take the log of a dimensionful quantity. One can make everything nice by picking some arbitrary energy scale and rescaling the energy and hence the density of states into a dimensionless quantity; this arbitrary energy scale of course amounts to a choice of units for the energy.

Finally, using the statistical-mechanical definition of the free energy (1.12) gives

$$F \approx \mathcal{F}(\langle E \rangle) ;$$

the constant in front of Z generally can be neglected because it is independent of the size of the system, whereas F and \mathcal{F} will be linear in it. (If one is interested in boundary phenomena, it can *not* necessarily be neglected.) Those familiar with thermodynamics will recognize (1.16) as the standard thermodynamical definition of the free energy. The purpose of this derivation is to explicitly display the assumptions that make statistical mechanical and the thermodynamical definitions equivalent. In traditional thermodynamic language, the equilibrium configuration is found by *minimizing* the free energy. Statistical mechanics, remarkably, builds this into the computation of the partition function.

1.7 The Peierls argument: why is difficult to order in one dimension and possible in two

Thinking in terms of the free energy is useful in understanding whether or not a system can be ordered. The Peierls argument², in a version due to Landau and Lifshitz³, is a way of making this precise. With the explicit expression for the energy that goes into the definition of the Boltzmann weights, it typically is easy to find a configuration (or configurations) that minimize the *energy*. Such configurations often are ordered; for example, in the Ising model with $J > 0$, the configuration with all spins up and that with all spins down each minimize the energy (in any dimension). If the energy of such an ordered configuration also is near a minimum of the *free energy*, then the configurations dominating the partition function are also ordered, and thus the expectation value of the order parameter will be non-zero. The system will be in an ordered phase. However, a minimum of the energy does not guarantee a minimum of the free energy; the entropy at some other value of E can be higher, and so for large enough T , this will yield a lower value of the free energy. In fact, for obvious reasons, there are typically more disordered configurations than ordered configurations, so $S(E)$ is rarely at a minimum for an ordered configuration. Thus for high enough T , a system will always be disordered. The question is if there are values of T low enough so that a system can order.

This question often is rephrased in terms of competition between energy and entropy. Changing an ordered configuration will increase the energy by some amount ΔE , but there are many ways of changing it, and so the entropy will increase by some amount ΔS as well. If $\Delta E > T\Delta S$, then the ordered configuration wins; the system remains ordered. If $\Delta E < T\Delta S$, then the system is disordered: the free energy is lowered by destroying the order. If there is a value of T where $\Delta E = T\Delta S$, then there is a phase

²R. Peierls, Math. Proc. Cambridge Phil. Soc. **32**, 477 (1936)

³*Statistical Physics*

Figure 1.2: Clusters in the two dimensional Ising model.

transition at this point. The presence of order therefore can be tested by computing the energy and entropy of configurations “near” to the ordered configuration.

First consider the ferromagnetic ($J > 0$) nearest-neighbor Ising model in one dimension with N spins. The *magnetization* M of a configuration is defined as

$$M(\{\sigma_i\}) \equiv \sum_i \sigma_i = \# \text{ up spins} - \# \text{ down spins} .$$

Consider a given configuration with $M = \mu N$ for some $\mu > 0$. Such a configuration is *ordered*; it is more likely that a spin is up rather than down. To investigate whether it is a minimum of the free energy, flip with the first L spins, giving a configuration with $M \approx \mu(N - 2L)$. The energy of the new configuration only differs from the original by (at most) $4J$ for any L , because only (at most) two bonds are changed. The entropy change, however, is much larger, because there are at least μN choices of L that have the same energy shift of $+2$; recall the entropy is a function of the energy, so all configurations of the same energy contribute to the entropy. Thus

$$\Delta \mathcal{F} = \Delta E - T \Delta S \approx 4J - k_B T \ln(\mu N) .$$

As as long as J is finite and $T \neq 0$, for large enough N , this must be negative: there are configurations with smaller magnetization that have lower free energy. Repeating this lowers the free energy further and further, showing that the minimum of the free energy occurs for zero magnetization. The Ising model in one dimension is not ordered! This was shown in the previous section, but by extending this argument, Landau and Lifshitz explain how *any* one-dimensional classical system with *local* interactions cannot order.

Essentially the only way of avoiding this conclusion is if there are long-range interactions (or an infinite number of degrees of freedom per site). A simple argument due to Thouless⁴ shows that a one-dimensional Ising model with interactions between spins σ_i and σ_j that falls off with distance as $|i - j|^{-b}$ is disordered if $b > 2$, but is ordered if $b < 2$. The reason is that flipping all the spins within a region of length $N \gg L \gg 1$ changes the energy by an amount proportional to L^{2-b} . The most interesting case is $b = 2$, where flipping all these spins gives an energy proportional to $\ln L$. Thus this energy increase is of the same order as the increase in entropy. Thus for $b = 2$ there can be (and in fact is) a transition between ordered and disordered phases.

The result in two and higher dimensions is the opposite: not only is order possible in the Ising model, but can be proved to exist at sufficiently low temperatures. Since for βJ small enough there cannot be order, there must be a phase transition. Again

⁴Phys. Rev. **187**, 732 (1969)

one studies the dependence of the energy and entropy of each “cluster” of flipped spins in an ordered background, as drawn in two dimensions in figure 1.2. The boundary of a cluster is often called a domain wall (even though in two dimensions, the “wall” is one-dimensional). In the Ising model with nearest-neighbor interactions and $J > 0$, the energy difference between the configurations with and without the cluster is $2JB$, where B is the number of “broken” bonds, i.e. those links on the lattice that have a $+$ on one end and $-$ on the other. B depends on the size of the *wall*, not the size of the cluster. Thus in two dimensions, a cluster surrounded by a contour of length L has energy $2JL$.

Computing the precise entropy of the clusters is a difficult and often impossible problem, but the Peierls argument requires only understanding how it depends on the wall “area”. For simplicity, consider a two-dimensional square lattice of N sites with the boundary fixed to be up spins all around. The domain walls surround clusters of down spins; they must be closed contours because of the boundary conditions. They can be chosen so that they never intersect. The number $n(L)$ of contours of length L is bounded from above by

$$n(L) \leq \frac{2N}{L} 3^L . \quad (1.19)$$

This arises from the following counting. There are at most 3^L contours of length L going through a given point, because as one creates the contour by laying down L links of length 1 each, the path can go straight, turn left, or turn right. One can fit at most $2N/L$ contours of length L in the system, as follows from noting that the maximum length of the contours in the system is $2N$. This results in the above bound can be vastly improved upon⁵, but this is unimportant for the Peierls argument: the key part of this bound is that $n(L)$ increases at most exponentially in L , so the entropy increases at most linearly. The upper bound on $n(L)$ and hence on the entropy results in a *lower* bound on its contribution to the free energy. Rerunning the energy vs. entropy argument gives

$$\Delta\mathcal{F} \leq 2JL - k_B T (L \ln(3) - \ln L + \ln(2N)) .$$

For sufficiently low T , the first term is larger than the second for any L , since $L \geq 4$. The ordered configuration can be stable! This argument is straightforward to generalize to higher dimensions, and the conclusion is the same: a phase transition in an Ising system can occur in any dimension from two on up.

Making the Peierls argument rigorous is not difficult.⁶ This puts an upper bound on the number of down spins N_- , and so a lower bound on the magnetization $\langle M \rangle = \langle N_+ - N_- \rangle = N - 2\langle N_- \rangle$. The magnetization here need not vanish because the symmetry under spin flip is broken by the boundary conditions; with different boundary conditions one can consider for example $\langle |M| \rangle$. Consider a particular configuration of spins. The maximum number of down spins inside a cluster surrounded by a contour of length L is

⁵J. L. Lebowitz and A. Mazel, J. Stat. Phys. **90**, 1051 (1998)

⁶R.B. Griffiths, Phys. Rev. **136**, A437 (1964)

$(L/4)^2$, so the number of total down spins in this configuration is bounded by

$$N_-(\{\sigma_i\}) \leq \sum_{L=4,6,8,\dots} \frac{L^2}{16} \sum_{i=1}^{n(L)} X_L^{(i)}$$

where $X_L^{(i)} = 1$ if the i th contour of length L is in this configuration, and is 0 otherwise. Single out a particular cluster of length L and label it C . The Boltzmann weight of any configuration containing C is suppressed by a factor of $e^{-2L\beta J}$ relative to the configuration obtained by removing C (i.e. flipping all the spins inside). Define \sum'_n to be the sum over all configurations containing C . The probability of having C is bounded from above by

$$\langle X_L^{(i)} \rangle = \frac{\sum'_n e^{-\beta E_n}}{\sum_n e^{-\beta E_n}} \leq e^{-2L\beta J} ,$$

since the denominator is a sum of positive terms, and includes all the configurations with C removed. Putting these two upper bounds together with the bound on the total number of clusters in (1.17) means that the expectation value of the number of down spins N_- is bounded from above by

$$\langle N_- \rangle \leq \frac{N}{8} \sum_{L=4,6,8,\dots} L 3^L e^{-2L\beta J} .$$

Letting $\kappa = 9e^{-4\beta J}$ and doing the sum for $\kappa < 1$ gives

$$\langle M \rangle \geq N - \frac{N\kappa^2}{2} \frac{2 - \kappa}{(1 - \kappa)^2} .$$

For κ small enough ($< .539$) the coefficient in front of N greater than zero, so the magnetization is macroscopically large. A non-vanishing magnetization means that the two-dimensional Ising model is ordered at low enough temperature.

With a little more work this can be generalized to any boundary condition and other lattices, as long as $J > 0$. For $J < 0$, the interesting counterexamples are lattices with geometric frustration: these typically will not order in the traditional sense.