5. Phase Transitions

A phase transition is an abrupt, discontinuous change in the properties of a system. We've already seen one example of a phase transition in our discussion of Bose-Einstein condensation. In that case, we had to look fairly closely to see the discontinuity: it was lurking in the derivative of the heat capacity. In other phase transitions — many of them already familiar — the discontinuity is more manifest. Examples include steam condensing to water and water freezing to ice.

In this section we'll explore a couple of phase transitions in some detail and extract some lessons that are common to all transitions.

5.1 Liquid-Gas Transition

Recall that we derived the van der Waals equation of state for a gas (2.31) in Section 2.5. We can write the van der Waals equation as

$$p = \frac{k_B T}{v - b} - \frac{a}{v^2} \tag{5.1}$$

where v = V/N is the volume per particle. In the literature, you will also see this equation written in terms of the particle density $\rho = 1/v$.

On the right we fix T at different values and sketch the graph of p vs. V determined by the van der Waals equation. These curves are *isotherms* — line of constant temperature. As we can see from the diagram, the isotherms take three different shapes depending on the value of T. The top curve shows the isotherm for large values of T. Here we can effectively ignore the $-a/v^2$ term. (Recall that v cannot take values smaller than b, reflecting the fact that atoms cannot approach to arbitrar-



Figure 34:

ily closely). The result is a monotonically decreasing function, essentially the same as we would get for an ideal gas. In contrast, when T is low enough, the second term in (5.1) can compete with the first term. Roughly speaking, this happens when $k_B T \sim a/v$ is in the allowed region v > b. For these low value of the temperature, the isotherm has a wiggle. At some intermediate temperature, the wiggle must flatten out so that the bottom curve looks like the top one. This happens when the maximum and minimum meet to form an inflection point. Mathematically, we are looking for a solution to dp/dv = $d^2p/dv^2 = 0$. It is simple to check that these two equations only have a solution at the critical temperature $T = T_c$ given by

$$k_B T_c = \frac{8a}{27b} \tag{5.2}$$

Let's look in more detail at the $T < T_c$ curve. For a range of pressures, the system can have three different choices of volume. A typical, albeit somewhat exagerated, example of this curve is shown in the figure below. What's going on? How should we interpret the fact that the system can seemingly live at three different densities $\rho = 1/v$?

First look at the middle solution. This has some fairly weird properties. We can see from the graph that the gradient is positive: $dp/dv|_T > 0$. This means that if we apply a force to the container to squeeze the gas, the pressure decreases. The gas doesn't push back; it just relents. But if we expand the gas, the pressure increases and the gas pushes harder. Both of these properties are telling us that the gas in that state is unstable.





If we were able to create such a state, it wouldn't hang around for long because any tiny perturbation would lead to a rapid, explosive change in its density. If we want to find states which we are likely to observe in Nature then we should look at the other two solutions.

The solution to the left on the graph has v slightly bigger than b. But, recall from our discussion of Section 2.5 that b is the closest that the atoms can get. If we have $v \sim b$, then the atoms are very densely packed. Moreover, we can also see from the graph that |dp/dv| is very large for this solution which means that the state is very difficult to compress: we need to add a great deal of pressure to change the volume only slightly. We have a name for this state: it is a *liquid*.

You may recall that our original derivation of the van der Waals equation was valid only for densities much lower than the liquid state. This means that we don't really trust (5.1) on this solution. Nonetheless, it is interesting that the equation predicts the existence of liquids and our plan is to gratefully accept this gift and push ahead to explore what the van der Waals tells us about the liquid-gas transition. We will see that it captures many of the qualitative features of the phase transition. The last of the three solutions is the one on the right in the figure. This solution has $v \gg b$ and small |dp/dv|. It is the *gas state*. Our goal is to understand what happens in between the liquid and gas state. We know that the naive, middle, solution given to us by the van der Waals equation is unstable. What replaces it?

5.1.1 Phase Equilibrium

Throughout our derivation of the van der Waals equation in Section 2.5, we assumed that the system was at a fixed density. But the presence of two solutions — the liquid and gas state — allows us to consider more general configurations: part of the system could be a liquid and part could be a gas.

How do we figure out if this indeed happens? Just because both liquid and gas states can exist, doesn't mean that they can cohabit. It might be that one is preferred over the other. We already saw some conditions that must be satisfied in order for two systems to sit in equilibrium back in Section 1. Mechanical and thermal equilibrium are guaranteed if two systems have the same pressure and temperature respectively. But both of these are already guaranteed by construction for our two liquid and gas solutions: the two solutions sit on the same isotherm and at the same value of p. We're left with only one further requirement that we must satisfy which arises because the two systems can exchange particles. This is the requirement of chemical equilibrium,

$$\mu_{\text{liquid}} = \mu_{\text{gas}} \tag{5.3}$$

Because of the relationship (4.22) between the chemical potential and the Gibbs free energy, this is often expressed as

$$g_{\text{liquid}} = g_{\text{gas}} \tag{5.4}$$

where g = G/N is the Gibbs free energy per particle.

Notice that all the equilibrium conditions involve only intensive quantities: p, T and μ . This means that if we have a situation where liquid and gas are in equilibrium, then we can have any number N_{liquid} of atoms in the liquid state and any number N_{gas} in the gas state. But how can we make sure that chemical equilibrium (5.3) is satisfied?

Maxwell Construction

We want to solve $\mu_{\text{liquid}} = \mu_{\text{gas}}$. We will think of the chemical potential as a function of p and T: $\mu = \mu(p, T)$. Importantly, we won't assume that $\mu(p, T)$ is single valued since that would be assuming the result we're trying to prove! Instead we will show that if we fix T, the condition (5.3) can only be solved for a very particular value of pressure

p. To see this, start in the liquid state at some fixed value of p and T and travel along the isotherm. The infinitesimal change in the chemical potential is

$$d\mu = \left. \frac{\partial \mu}{\partial p} \right|_T dp$$

However, we can get an expression for $\partial \mu / \partial p$ by recalling that arguments involving extensive and intensive variables tell us that the chemical potential is proportional to the Gibbs free energy: $G(p, T, N) = \mu(p, T)N$ (4.22). Looking back at the variation of the Gibbs free energy (4.21) then tells us that

$$\left. \frac{\partial G}{\partial p} \right|_{N,T} = \left. \frac{\partial \mu}{\partial p} \right|_T N = V \tag{5.5}$$

Integrating along the isotherm then tells us the chemical potential of any point on the curve,

$$\mu(p,T) = \mu_{\text{liquid}} + \int_{p_{\text{liquid}}}^{p} dp' \ \frac{V(p',T)}{N}$$

When we get to gas state at the same pressure $p = p_{\text{liquid}}$ that we started from, the condition for equilibrium is $\mu = \mu_{\text{liquid}}$. Which means that the integral has to vanish. Graphically this is very simple to de-



Figure 36:

scribe: the two shaded areas in the graph must have equal area. This condition, known as the *Maxwell construction*, tells us the pressure at which gas and liquid can co-exist.

I should confess that there's something slightly dodgy about the Maxwell construction. We already argued that the part of the isotherm with dp/dv > 0 suffers an instability and is unphysical. But we needed to trek along that part of the curve to derive our result. There are more rigorous arguments that give the same answer.

For each isotherm, we can determine the pressure at which the liquid and gas states are in equilibrium. The gives us the *co-existence curve*, shown by the dotted line in Figure 37. Inside this region, liquid and gas can both exist at the same temperature and pressure. But there is nothing that tells us how much gas there should be and how much liquid: atoms can happily move from the liquid state to the gas state. This means that while the density of gas and liquid is fixed, the *average* density of the system is not. It can vary between the gas density and the liquid density simply by changing the amount of liquid. The upshot of this argument is that inside the co-existence curves, the isotherms simply become flat lines, reflecting the fact that the density can take any value. This is shown in graph on the right of Figure 37.



Figure 37: The co-existence curve in red, resulting in constant pressure regions consisting of a harmonious mixture of vapour and liquid.

To illustrate the physics of this situation, suppose that we sit at some fixed density $\rho = 1/v$ and cool the system down from a high temperature to $T < T_c$ at a point inside the co-existence curve so that we're now sitting on one of the flat lines. Here, the system is neither entirely liquid, nor entirely gas. Instead it will split into gas, with density $1/v_{\text{gas}}$, and liquid, with density $1/v_{\text{liquid}}$ so that the average density remains 1/v. The system undergoes *phase separation*. The minimum energy configuration will typically be a single phase of



Figure 38:

liquid and one of gas because the interface between the two costs energy. (We will derive an expression for this energy in Section 5.5). The end result is shown on the right. In the presence of gravity, the higher density liquid will indeed sink to the bottom.

Meta-Stable States

We've understood what replaces the unstable region of the van der Waals phase diagram. But we seem to have removed more states than anticipated: parts of the Van der Waals isotherm that had dp/dv < 0 are contained in the co-existence region and replaced by the flat pressure lines. This is the region of the *p*-*V* phase diagram that is contained between the two dotted lines in the figure to the right. The outer dotted line is the co-existence curve. The inner dotted curve is constructed to pass through the stationary points of the van der Waals isotherms. It is called the *spinodal curve*.



Figure 39:

The van der Waals states which lie between the spinodal curve and the co-existence curve are good states. But they are meta-stable. One can show that their Gibbs free energy is higher than that of the liquid-gas equilibrium at the same p and T. However, if we compress the gas very slowly we can coax the system into this state. It is known as a supercooled vapour. It is delicate. Any small disturbance will cause some amount of the gas to condense into the liquid. Similarly, expanding a liquid beyond the coexistence curve results in an meta-stable, superheated liquid.

5.1.2 The Clausius-Clapeyron Equation

We can also choose to plot the liquid-gas phase diagram on the p-T plane. Here the co-existence region is squeezed into a line: if we're sitting in the gas phase and increase the pressure just a little bit at fixed $T < T_c$ then we jump immediately to the liquid phase. This appears as a discontinuity in the volume. Such discontinuities are the sign of a phase transition. The end result is sketched in the figure to the right; the thick solid line denotes the presence of a phase transition.



Figure 40:

Either side of the line, all particles are either in the gas or liquid phase. We know from (5.4) that the Gibbs free energies (per particle) of these two states are equal,

$$g_{\text{liquid}} = g_{\text{gas}}$$

So G is continuous as we move across the line of phase transitions. Suppose that we sit on the line itself and move along it. How does g change? We can easily compute this from (4.21),

$$dG_{\text{liquid}} = -S_{\text{liquid}}dT + V_{\text{liquid}}dp \quad \Rightarrow \quad dg_{\text{liquid}} = -s_{\text{liquid}}dT + v_{\text{liquid}}dp$$

where, just as g = G/N an v = V/N, the entropy density is s = S/N. Equating this with the free energy in the gaseous phase gives

$$dg_{\text{liquid}} = -s_{\text{liquid}}dT + v_{\text{liquid}}dp = dg_{\text{gas}} = -s_{\text{gas}}dT + v_{\text{gas}}dp$$

This can be rearranged to gives us a nice expression for the slope of the line of phase transitions in the p - T plane. It is

$$\frac{dp}{dT} = \frac{s_{\text{gas}} - s_{\text{liquid}}}{v_{\text{gas}} - v_{\text{liquid}}}$$

We usually define the specific latent heat

$$L = T(s_{\rm gas} - s_{\rm liquid})$$

This is the energy released per particle as we pass through the phase transition. We see that the slope of the line in the p-T plane is determined by the ratio of latent heat released in the phase transition and the discontinuity in volume. The result is known as the *Clausius-Clapeyron* equation,

$$\frac{dp}{dT} = \frac{L}{T(v_{\rm gas} - v_{\rm liquid})} \tag{5.6}$$

There is a classification of phase transitions, due originally to Ehrenfest. When the n^{th} derivative of a thermodynamic potential (either F or G usually) is discontinuous, we say we have an n^{th} order phase transition. In practice, we nearly always deal with first, second and (very rarely) third order transitions. The liquid-gas transition releases latent heat, which means that $S = -\partial F/\partial T$ is discontinuous. Alternatively, we can say that $V = \partial G/\partial p$ is discontinuous. Either way, it is a *first order phase transition*. The Clausius-Clapeyron equation (5.6) applies to any first order transition.

As we approach $T \to T_c$, the discontinuity diminishes and $S_{\text{liquid}} \to S_{\text{gas}}$. At the critical point $T = T_c$ we have a second order phase transition. Above the critical point, there is no sharp distinction between the gas phase and liquid phase.

For most simple materials, the phase diagram above is part of a larger phase diagram which includes solids at smaller temperatures or higher pressures. A generic version of such a phase diagram is shown to the right. The van der Waals equation is missing the physics of solidification and includes only the liquid-gas line.



Figure 41:

An Approximate Solution to the Clausius-Clapeyron Equation

We can solve the Clausius-Clapeyron solution if we make the following assumptions:

- The latent heat L is constant.
- $v_{\rm gas} \gg v_{\rm liquid}$, so $v_{\rm gas} v_{\rm liquid} \approx v_{\rm gas}$. For water, this is an error of less than 0.1%
- Although we derived the phase transition using the van der Waals equation, now we've got equation (5.6) we'll pretend the gas obeys the ideal gas law $pv = k_B T$.

With these assumptions, it is simple to solve (5.6). It reduces to

$$\frac{dp}{dT} = \frac{Lp}{k_B T^2} \quad \Rightarrow \quad p = p_0 e^{-L/k_B T}$$

5.1.3 The Critical Point

Let's now return to discuss some aspects of life at the critical point. We previously worked out the critical temperature (5.2) by looking for solutions to simultaneous equations $\partial p/\partial v = \partial^2 p/\partial v^2 = 0$. There's a slightly more elegant way to find the critical point which also quickly gives us p_c and v_c as well. We rearrange the van der Waals equation (5.1) to get a cubic,

$$pv^3 - (pb + k_BT)v^2 + av - ab = 0$$

For $T < T_c$, this equation has three real roots. For $T > T_c$ there is just one. Precisely at $T = T_c$, the three roots must therefore coincide (before two move off onto the complex plane). At the critical point, this curve can be written as

$$p_c(v - v_c)^3 = 0$$

Comparing the coefficients tells us the values at the critical point,

$$k_B T_c = \frac{8a}{27b}$$
 , $v_c = 3b$, $p_c = \frac{a}{27b^2}$ (5.7)

The Law of Corresponding States

We can invert the relations (5.7) to express the parameters a and b in terms of the critical values, which we then substitute back into the van der Waals equation. To this end, we define the *reduced* variables,

$$\bar{T} = \frac{T}{T_c}$$
 , $\bar{v} = \frac{v}{v_c}$ $\bar{p} = \frac{p}{p_c}$

The advantage of working with \overline{T} , \overline{v} and \overline{p} is that it allows us to write the van der Waals equation (5.1) in a form that is universal to all gases, usually referred to as the law of corresponding states

$$\bar{p} = \frac{8}{3} \frac{\bar{T}}{\bar{v} - 1/3} - \frac{3}{\bar{v}^2}$$

Moreover, because the three variables T_c , p_c and v_c at the critical point are expressed in terms of just two variables, a and b (5.7), we can construct a combination of them



Figure 42: The co-existence curve for gases. Data is plotted for Ne, Ar, Kr, Xe, N_2 , O_2 , CO and CH_4 .

which is independent of a and b and therefore supposedly the same for all gases. This is the universal compressibility ratio,

$$\frac{p_c v_c}{k_B T_c} = \frac{3}{8} = 0.375 \tag{5.8}$$

Comparing to real gases, this number is a little high. Values range from around 0.28 to 0.3. We shouldn't be too discouraged by this; after all, we knew from the beginning that the van der Waals equation is unlikely to be accurate in the liquid regime. Moreover, the fact that gases have a critical point (defined by three variables T_c , p_c and v_c) guarantees that a similar relationship would hold for any equation of state which includes just two parameters (such as a and b) but would most likely fail to hold for equations of state that included more than two parameters.

Dubious as its theoretical foundation is, the law of corresponding states is the first suggestion that something remarkable happens if we describe a gas in terms of its reduced variables. More importantly, there is striking experimental evidence to back this up! Figure 42 shows the Guggenheim plot, constructed in 1945. The co-existence curve for 8 different gases in plotted in reduced variables: \overline{T} along the vertical axis; $\overline{\rho} = 1/\overline{v}$ along the horizontal. The gases vary in complexity from the simple monatomic gas Ne to the molecule CH_4 . As you can see, the co-existence curve for all gases is essentially the same, with the chemical make-up largely forgotten. There is clearly something interesting going on. How to understand it?

Critical Exponents

We will focus attention on physics close to the critical point. It is not immediately

obvious what are the right questions to ask. It turns out that the questions which have the most interesting answer are concerned with how various quantities change as we approach the critical point. There are lots of ways to ask questions of this type since there are many quantities of interest and, for each of them, we could approach the critical point from different directions. Here we'll look at the behaviour of three quantities to get a feel for what happens.

First, we can ask what happens to the difference in (inverse) densities $v_{\text{gas}} - v_{\text{liquid}}$ as we approach the critical point along the co-existence curve. For $T < T_c$, or equivalently $\bar{T} < 1$, the reduced van der Waals equation (5.8) has two stable solutions,

$$\bar{p} = \frac{8\bar{T}}{3\bar{v}_{\text{liquid}} - 1} - \frac{3}{\bar{v}_{\text{liquid}}^2} = \frac{8\bar{T}}{3\bar{v}_{\text{gas}} - 1} - \frac{3}{\bar{v}_{\text{gas}}^2}$$

If we solve this for \overline{T} , we have

$$\bar{T} = \frac{(3\bar{v}_{\text{liquid}} - 1)(3\bar{v}_{\text{gas}} - 1)(\bar{v}_{\text{liquid}} + \bar{v}_{\text{gas}})}{8\bar{v}_{\text{gas}}^2\bar{v}_{\text{liquid}}^2}$$

Notice that as we approach the critical point, $\bar{v}_{gas}, \bar{v}_{liquid} \rightarrow 1$ and the equation above tells us that $\bar{T} \rightarrow 1$ as expected. We can see exactly how we approach $\bar{T} = 1$ by expanding the right right-hand side for small $\epsilon \equiv \bar{v}_{gas} - \bar{v}_{liquid}$. To do this quickly, it's best to notice that the equation is symmetric in \bar{v}_{gas} and \bar{v}_{liquid} , so close to the critical point we can write $\bar{v}_{gas} = 1 + \epsilon/2$ and $\bar{v}_{liquid} = 1 - \epsilon/2$. Substituting this into the equation above and keeping just the leading order term, we find

$$\bar{T} \approx 1 - \frac{1}{16} (\bar{v}_{\text{gas}} - \bar{v}_{\text{liquid}})^2$$

Or, re-arranging, as we approach T_c along the co-existence curve,

$$v_{\rm gas} - v_{\rm liquid} \sim (T_c - T)^{1/2}$$
 (5.9)

This is the answer to our first question.

Our second variant of the question is: how does the volume change with pressure as we move along the critical isotherm. It turns out that we can answer this question without doing any work. Notice that at $T = T_c$, there is a unique pressure for a given volume $p(v, T_c)$. But we know that $\partial p/\partial v = \partial^2 p/\partial v^2 = 0$ at the critical point. So a Taylor expansion around the critical point must start with the cubic term,

$$p - p_c \sim (v - v_c)^3$$
 (5.10)

This is the answer to our second question.

Our third and final variant of the question concerns the compressibility, defined as

$$\kappa = -\frac{1}{v} \left. \frac{\partial v}{\partial p} \right|_T \tag{5.11}$$

We want to understand how κ changes as we approach $T \to T_c$ from above. In fact, we met the compressibility before: it was the feature that first made us nervous about the van der Waals equation since κ is negative in the unstable region. We already know that at the critical point $\partial p/\partial v|_{T_c} = 0$. So expanding for temperatures close to T_c , we expect

$$\left. \frac{\partial p}{\partial v} \right|_{T;v=v_c} = -a(T - T_c) + \dots$$

This tells us that the compressibility should diverge at the critical point, scaling as

$$\kappa \sim (T - T_c)^{-1} \tag{5.12}$$

We now have three answers to three questions: (5.9), (5.10) and (5.12). Are they right?! By which I mean: do they agree with experiment? Remember that we're not sure that we can trust the van der Waals equation at the critical point so we should be nervous. However, there is also reason for some confidence. Notice, in particular, that in order to compute (5.10) and (5.12), we didn't actually need any details of the van der Waals equation. We simply needed to assume the existence of the critical point and an analytic Taylor expansion of various quantities in the neighbourhood. Given that the answers follow from such general grounds, one may hope that they provide the correct answers for a gas in the neighbourhood of the critical point even though we know that the approximations that went into the van der Waals equation aren't valid there. Fortunately, that isn't the case: physics is much more interesting than that!

The experimental results for a gas in the neighbourhood of the critical point do share one feature in common with the discussion above: they are completely independent of the atomic make-up of the gas. However, the scaling that we computed using the van der Waals equation is not fully accurate. The correct results are as follows. As we approach the critical point along the co-existence curve, the densities scale as

$$v_{\rm gas} - v_{\rm liquid} \sim (T_c - T)^{\beta}$$
 with $\beta \approx 0.32$

(Note that the exponent β has nothing to do with inverse temperature. We're just near the end of the course and running out of letters and β is the canonical name for this exponent). As we approach along an isotherm,

$$p - p_c \sim (v - v_c)^{\delta}$$
 with $\delta \approx 4.8$

Finally, as we approach T_c from above, the compressibility scales as

$$\kappa \sim (T - T_c)^{-\gamma}$$
 with $\gamma \approx 1.2$

The quantities β , γ and δ are examples of *critical exponents*. We will see more of them shortly. The van der Waals equation provides only a crude first approximation to the critical exponents.

Fluctuations

We see that the van der Waals equation didn't do too badly in capturing the dynamics of an interacting gas. It gets the qualitative behaviour right, but fails on precise quantitative tests. So what went wrong? We mentioned during the derivation of the van der Waals equation that we made certain approximations that are valid only at low density. So perhaps it is not surprising that it fails to get the numbers right near the critical point v = 3b. But there's actually a deeper reason that the van der Waals equation fails: fluctuations.

This is simplest to see in the grand canonical ensemble. Recall that back in Section 1 that we argued that $\Delta N/N \sim 1/\sqrt{N}$, which allowed us to happily work in the grand canonical ensemble even when we actually had fixed particle number. In the context of the liquid-gas transition, fluctuating particle number is the same thing as fluctuating density $\rho = N/V$. Let's revisit the calculation of ΔN near the critical point. Using (1.45) and (1.48), the grand canonical particle number (1.42) is

$$\langle N \rangle = V \left. \frac{\partial p}{\partial \mu} \right|_{T,V}$$

We already have an expression for the variance in the particle number in (1.43),

$$\Delta N^2 = \frac{1}{\beta} \left. \frac{\partial \langle N \rangle}{\partial \mu} \right|_{T,V}$$

Dividing these two expressions, we have

$$\frac{\Delta N^2}{N} = \frac{1}{V\beta} \left. \frac{\partial \langle N \rangle}{\partial \mu} \right|_{T,V} \left. \frac{\partial \mu}{\partial p} \right|_{T,V} = \frac{1}{V\beta} \left. \frac{\partial \langle N \rangle}{\partial p} \right|_{T,V}$$

But we can re-write this expression using the general relationship between partial derivatives $\partial x/\partial y|_z \partial y/\partial z|_x \partial z/\partial x|_y = -1$. We then have

$$\frac{\Delta N^2}{N} = -\frac{1}{\beta} \left. \frac{\partial \langle N \rangle}{\partial V} \right|_{p,T} \left. \frac{1}{V} \left. \frac{\partial V}{\partial p} \right|_{N,T}$$

This final expression relates the fluctuations in the particle number to the compressibility (5.11). But the compressibility is diverging at the critical point and this means that there are large fluctuations in the density of the fluid at this point. The result is that any simple equation of state, like the van der Waals equation, which works only with the average volume, pressure and density will miss this key aspect of the physics.

Understanding how to correctly account for these fluctuations is the subject of *critical* phenomena. It has close links with the renormalization group and conformal field theory which also arise in particle physics and string theory. You will meet some of these ideas in next year's Statistical Field Theory course. Here we will turn to a different phase transition which will allow us to highlight some of the key ideas.

5.2 The Ising Model

The Ising model is one of the touchstones of modern physics; a simple system that exhibits non-trivial and interesting behaviour.

The Ising model consists of N sites in a d-dimensional lattice. On each lattice site lives a quantum spin that can sit in one of two states: spin up or spin down. We'll call the eigenvalue of the spin on the i^{th} lattice site s_i . If the spin is up, $s_i = +1$; if the spin is down, $s_i = -1$.

The spins sit in a magnetic field that endows an energy advantage to those which point up,

$$E_B = -B\sum_{i=1}^N s_i$$

(A comment on notation: B should be properly denoted H. We're sticking with B to avoid confusion with the Hamiltonian. There is also a factor of the magnetic moment which has been absorbed into the definition of B). The lattice system with energy E_B is equivalent to the two-state system that we first met when learning the techniques of statistical mechanics back in Section 1.2.3. However, the Ising model contains an additional complication that makes the system much more interesting: this is an interaction between neighbouring spins. The full energy of the system is therefore,

$$E = -J\sum_{\langle ij\rangle} s_i s_j - B\sum_i s_i \tag{5.13}$$

The notation $\langle ij \rangle$ means that we sum over all "nearest neighbour" pairs in the lattice. The number of such pairs depends both on the dimension d and the type of lattice. We'll denote the number of nearest neighbours as q. For example, in d = 1 a lattice has q = 2; in d = 2, a square lattice has q = 4. A square lattice in d dimensions has q = 2d.

If J > 0, neighbouring spins prefer to be aligned ($\uparrow\uparrow$ or $\downarrow\downarrow$). In the context of magnetism, such a system is called a *ferromagnet*. If J < 0, the spins want to antialign ($\uparrow\downarrow$). This is an *anti-ferromagnet*. In the following, we'll choose J > 0 although for the level of discussion needed for this course, the differences are minor.

We work in the canonical ensemble and introduce the partition function

$$Z = \sum_{\{s_i\}} e^{-\beta E[s_i]}$$
(5.14)

While the effect of both J > 0 and $B \neq 0$ is to make it energetically preferable for the spins to align, the effect of temperature will be to randomize the spins, with entropy winning over energy. Our interest is in the average spin, or average *magnetization*

$$m = \frac{1}{N} \sum_{i} \langle s_i \rangle = \frac{1}{N\beta} \frac{\partial \log Z}{\partial B}$$
(5.15)

The Ising Model as a Lattice Gas

Before we develop techniques to compute the partition function (5.14), it's worth pointing out that we can drape slightly different words around the mathematics of the Ising model. It need not be interpreted as a system of spins; it can also be thought of as a lattice description of a gas.

To see this, consider the same *d*-dimensional lattice as before, but now with particles hopping between lattice sites. These particles have hard cores, so no more than one can sit on a single lattice site. We introduce the variable $n_i \in \{0, 1\}$ to specify whether a given lattice site, labelled by *i*, is empty $(n_i = 0)$ or filled $(n_i = 1)$. We can also introduce an attractive force between atoms by offering them an energetic reward if they sit on neighbouring sites. The Hamiltonian of such a lattice gas is given by

$$E = -4J\sum_{\langle ij\rangle} n_i n_j - \mu \sum_i n_i$$

where μ is the chemical potential which determines the overall particle number. But this Hamiltonian is trivially the same as the Ising model (5.13) if we make the identification

$$s_i = 2n_i - 1 \in \{-1, 1\}$$

The chemical potential μ in the lattice gas plays the role of magnetic field in the spin system while the magnetization of the system (5.15) measures the average density of particles away from half-filling.

5.2.1 Mean Field Theory

For general lattices, in arbitrary dimension d, the sum (5.14) cannot be performed. An exact solution exists in d = 1 and, when B = 0, in d = 2. (The d = 2 solution is originally due to Onsager and is famously complicated! Simpler solutions using more modern techniques have since been discovered).

Here we'll develop an approximate method to evaluate Z known as *mean field theory*. We write the interactions between neighbouring spins in term of their deviation from the average spin m,

$$s_i s_j = [(s_i - m) + m][(s_j - m) + m]$$

= $(s_i - m)(s_j - m) + m(s_j - m) + m(s_i - m) + m^2$

The mean field approximation means that we assume that the fluctuations of spins away from the average are small which allows us to neglect the first term above. Notice that this isn't the statement that the variance of an individual spin is small; that can never be true because s_i takes values +1 or -1 so $\langle s_i^2 \rangle = 1$ and the variance $\langle (s_i - m)^2 \rangle$ is always large. Instead, the mean field approximation is a statement about fluctuations between spins on neighbouring sites, so the first term above can be neglected when summing over $\sum_{\langle ij \rangle}$. We can then write the energy (5.13) as

$$E_{\rm mf} = -J \sum_{\langle ij \rangle} [m(s_i + s_j) - m^2] - B \sum_i s_i$$

= $\frac{1}{2} JNqm^2 - (Jqm + B) \sum_i s_i$ (5.16)

where the factor of Nq/2 in the first term is simply the number of nearest neighbour pairs $\sum_{\langle ij \rangle}$. The factor or 1/2 is there because $\sum_{\langle ij \rangle}$ is a sum over pairs rather than a sum of individual sites. (If you're worried about this formula, you should check it for a simple square lattice in d = 1 and d = 2 dimensions). A similar factor in the second term cancelled the factor of 2 due to $(s_i + s_j)$.

We see that the mean field approximation has removed the interactions. The Ising model reduces to the two state system that we saw way back in Section 1. The result of the interactions is that a given spin feels the average effect of its neighbour's spins through a contribution to the effective magnetic field,

$$B_{\text{eff}} = B + Jqm$$

Once we've taken into account this extra contribution to B_{eff} , each spin acts indepen-



Figure 43: $tanh(Jqm\beta)$ for $Jq\beta < 1$

Figure 44: $tanh(Jqm\beta)$ for $Jq\beta > 1$

dently and it is easy to write the partition function. It is

$$Z = e^{-\frac{1}{2}\beta JNqm^2} \left(e^{-\beta B_{\text{eff}}} + e^{\beta B_{\text{eff}}} \right)^N$$
$$= e^{-\frac{1}{2}\beta JNqm^2} 2^N \cosh^N \beta B_{\text{eff}}$$
(5.17)

However, we're not quite done. Our result for the partition function Z depends on B_{eff} which depends on m which we don't yet know. However, we can use our expression for Z to self-consistently determine the magnetization (5.15). We find,

$$m = \tanh(\beta B + \beta Jqm) \tag{5.18}$$

We can now solve this equation to find the magnetization for various values of T and B: m = m(T, B). It is simple to see the nature of the solutions using graphical methods.

B=0

Let's first consider the situation with vanishing magnetic field, B = 0. The figures above show the graph linear in m compared with the tanh function. Since $\tanh x \approx x - \frac{1}{3}x^3 + \ldots$, the slope of the graph near the origin is given by βJq . This then determines the nature of the solution.

- The first graph depicts the situation for $\beta Jq < 1$. The only solution is m = 0. This means that at high temperatures $k_BT > Jq$, there is no average magnetization of the system. The entropy associated to the random temperature flucutations wins over the energetically preferred ordered state in which the spins align.
- The second graph depicts the situation for $\beta Jq > 1$. Now there are three solutions: $m = \pm m_0$ and m = 0. It will turn out that the middle solution, m = 0, is unstable. (This solution is entirely analogous to the unstable solution of the van der Waals equation. We will see this below when we compute the free energy).

For the other two possible solutions, $m = \pm m_0$, the magnetization is non-zero. Here we see the effects of the interactions begin to win over temperature. Notice that in the limit of vanishing temperature, $\beta \to \infty$, $m_0 \to 1$. This means that all the spins are pointing in the same direction (either up or down) as expected.

• The critical temperature separating these two cases is

$$k_B T_c = Jq \tag{5.19}$$

The results described above are perhaps rather surprising. Based on the intuition that things in physics always happen smoothly, one might have thought that the magnetization would drop slowly to zero as $T \to \infty$. But that doesn't happen. Instead the magnetization turns off abruptly at some finite value of the temperature $T = T_c$, with no magnetization at all for higher temperatures. This is the characteristic behaviour of a phase transition.

$\mathbf{B}\neq\mathbf{0}$

For $B \neq 0$, we can solve the consistency equation (5.18) in a similar fashion. There are a couple of key differences to the $B \neq 0$ case. Firstly, there is now no phase transition at fixed B as we vary temperature T. Instead, for very large temperatures $k_B T \gg Jq$, the magnetization goes smoothly to zero as

$$m \to \frac{B}{k_B T}$$
 as $T \to \infty$

At low temperatures, the magnetization again asymptotes to the state $m \to \pm 1$ which minimizes the energy. Except this time, there is no ambiguity as to whether the system chooses m = +1 or m = -1. This is entirely determined by the sign of the magnetic field B.

In fact the low temperature behaviour requires slightly more explanation. For small values of B and T, there are again three solutions to (5.18). This follows simply from continuity: there are three solutions for $T < T_c$ and B = 0 shown in Figure 44 and these must survive in some neighbourhood of B = 0. One of these solutions is again unstable. However, of the remaining two only one is now stable: that with $\operatorname{sign}(m) = \operatorname{sign}(B)$. The other is meta-stable. We will see why this is the case shortly when we come to discuss the free energy.





Figure 45: Magnetization with B = 0 Figure 46: Magnetization at $B \neq 0$. and the phase transition

The net result of our discussion is depicted in the figures above. When B = 0there is a phase transition at $T = T_c$. For $T < T_c$, the system can sit in one of two magnetized states with $m = \pm m_0$. In contrast, for $B \neq 0$, there is no phase transition as we vary temperature and the system has at all times a preferred magnetization whose sign is determined by that of B. Notice however, we do have a phase transition if we fix temperature at $T < T_c$ and vary B from negative to positive. Then the magnetization jumps discontinuously from a negative value to a positive value. Since the magnetization is a first derivative of the free energy (5.15), this is a first order phase transition. In contrast, moving along the temperature axis at B = 0 results in a second order phase transition at $T = T_c$.

5.2.2 Critical Exponents

It is interesting to compare the phase transition of the Ising model with that of the liquid-gas phase transition. The two are sketched in the Figure 47 above. In both cases, we have a first order phase transition and a quantity jumps discontinuously at $T < T_c$. In the case of the liquid-gas, it is the density $\rho = 1/v$ that jumps as we vary pressure; in the case of the Ising model it is the magnetization m that jumps as we vary the magnetic field. Moreover, in both cases the discontinuity disappears as we approach $T = T_c$.

We can calculate critical exponents for the Ising model. To compare with our discussion for the liquid-gas critical point, we will compute three quantities. First, consider the magnetization at B = 0. We can ask how this magnetization decreases as we tend towards the critical point. Just below $T = T_c$, m is small and we can Taylor expand (5.18) to get

$$m \approx \beta Jqm - \frac{1}{3}(\beta Jqm)^3 + \dots$$



Figure 47: A comparison of the phase diagram for the liquid-gas system and the Ising model.

The magnetization therefore scales as

$$m_0 \sim \pm (T_c - T)^{1/2}$$
 (5.20)

This is to be compared with the analogous result (5.9) from the van der Waals equation. We see that the values of the exponents are the same in both cases. Notice that the derivative dm/dT becomes infinite as we approach the critical point. In fact, we had already anticipated this when we drew the plot of the magnetization in Figure 45.

Secondly, we can sit at $T = T_c$ and ask how the magnetization changes as we approach B = 0. We can read this off from (5.18). At $T = T_c$ we have $\beta Jq = 1$ and the consistency condition becomes $m = \tanh(B/Jq + m)$. Expanding for small B gives

$$m \approx \frac{B}{Jq} + m - \frac{1}{3} \left(\frac{B}{Jq} + m \right)^3 + \ldots \approx \frac{B}{Jq} + m - \frac{1}{3} m^3 + \mathcal{O}(B^2)$$

So we find that the magnetization scales as

$$m \sim B^{1/3}$$
 (5.21)

Notice that this power of 1/3 is again familiar from the liquid-gas transition (5.10) where the van der Waals equation gave $v_{\text{gas}} - v_{\text{liquid}} \sim (p - p_c)^{1/3}$.

Finally, we can look at the magnetic susceptibility $\chi,$ defined as

$$\chi = N \left. \frac{\partial m}{\partial B} \right|_T$$

This is analogous to the compressibility κ of the gas. We will ask how χ changes as we approach $T \to T_c$ from above at B = 0. We differentiate (5.18) with respect to B to get

$$\chi = \frac{N\beta}{\cosh^2\beta Jqm} \left(1 + \frac{Jq}{N}\chi\right)$$

We now evaluate this at B = 0. Since we want to approach $T \to T_c$ from above, we can also set m = 0 in the above expression. Evaluating this at B = 0 gives us the scaling

$$\chi = \frac{N\beta}{1 - Jq\beta} \sim (T - T_c)^{-1} \tag{5.22}$$

Once again, we see that same critical exponent that the van der Waals equation gave us for the gas (5.12).

5.2.3 Validity of Mean Field Theory

The phase diagram and critical exponents above were all derived using the mean field approximation. But this was an unjustified approximation. Just as for the van der Waals equation, we can ask the all-important question: are our results right?

There is actually a version of the Ising model for which the mean field theory is exact: it is the $d = \infty$ dimensional lattice. This is unphysical (even for a string theorist). Roughly speaking, mean field theory works for large d because each spin has a large number of neighbours and so indeed sees something close to the average spin.

But what about dimensions of interest? Mean field theory gets things most dramatically wrong in d = 1. In that case, no phase transition occurs. We will derive this result below where we briefly describe the exact solution to the d = 1 Ising model. There is a general lesson here: in low dimensions, both thermal and quantum fluctuations are more important and invariably stop systems forming ordered phases.

In higher dimensions, $d \geq 2$, the crude features of the phase diagram, including the existence of a phase transition, given by mean field theory are essentially correct. In fact, the very existence of a phase transition is already worthy of comment. The defining feature of a phase transition is behaviour that jumps discontinuously as we vary β or B. Mathematically, the functions must be non-analytic. Yet all properties of the theory can be extracted from the partition function Z which is a sum of smooth, analytic functions (5.14). How can we get a phase transition? The loophole is that Zis only necessarily analytic if the sum is finite. But there is no such guarantee when the number of lattice sites $N \to \infty$. We reach a similar conclusion to that of Bose-Einstein condensation: phase transitions only strictly happen in the thermodynamic limit. There are no phase transitions in finite systems.

What about the critical exponents that we computed in (5.20), (5.21) and (5.22)? It turns out that these are correct for the Ising model defined in $d \ge 4$. (We will briefly sketch why this is true at the end of this Chapter). But for d = 2 and d = 3, the critical exponents predicted by mean field theory are only first approximations to the true answers. For d = 2, the exact solution (which goes quite substantially past this course) gives the critical exponents to be,

$$m_0 \sim (T_c - T)^{\beta}$$
 with $\beta = \frac{1}{8}$
 $m \sim B^{1/\delta}$ with $\delta = 15$
 $\chi \sim (T - T_c)^{-\gamma}$ with $\gamma = \frac{7}{4}$

The biggest surprise is in d = 3 dimensions. Here the critical exponents are not known exactly. However, there has been a great deal of numerical work to determine them. They are given by

$$\beta \approx 0.32$$
 , $\delta \approx 4.8$, $\gamma \approx 1.2$

But these are exactly the same critical exponents that are seen in the liquid-gas phase transition. That's remarkable! We saw above that the mean field approach to the Ising model gave the same critical exponents as the van der Waals equation. But they are both wrong. And they are both wrong in the same, complicated, way! Why on earth would a system of spins on a lattice have anything to do with the phase transition between a liquid and gas? It is as if all memory of the microscopic physics — the type of particles, the nature of the interactions — has been lost at the critical point. And that's exactly what happens.

What we're seeing here is evidence for *universality*. There is a single theory which describes the physics at the critical point of the liquid gas transition, the 3d Ising model and many other systems. This is a theoretical physicist's dream! We spend a great deal of time trying to throw away the messy details of a system to focus on the elegant essentials. But, at a critical point, Nature does this for us! Although critical points in two dimensions are well understood, there is still much that we don't know about critical points in three dimensions. This, however, is a story that will have to wait for another day.

5.3 Some Exact Results for the Ising Model

This subsection is something of a diversion from our main interest. In later subsections, we will develop the idea of mean field theory. But first we pause to describe some exact results for the Ising model using techniques that do not rely on the mean field approximation. Many of the results that we derive have broader implications for systems beyond the Ising model. As we mentioned above, there is an exact solution for the Ising model in d = 1 dimension and, when B = 0, in d = 2 dimensions. Here we will describe the d = 1 solution but not the full d = 2 solution. We will, however, derive a number of results for the d = 2 Ising model which, while falling short of the full solution, nonetheless provide important insights into the physics.

5.3.1 The Ising Model in d = 1 Dimensions

We start with the *Ising chain*, the Ising model on a one dimensional line. Here we will see that the mean field approximation fails miserably, giving qualitatively incorrect results: the exact results shows that there are no phase transitions in the Ising chain.

The energy of the system (5.13) can be trivially rewritten as

$$E = -J\sum_{i=1}^{N} s_i s_{i+1} - \frac{B}{2}\sum_{i=1}^{N} (s_i + s_{i+1})$$

We will impose periodic boundary conditions, so the spins live on a circular lattice with $s_{N+1} \equiv s_1$. The partition function is then

$$Z = \sum_{s_1=\pm 1} \dots \sum_{s_N=\pm 1} \prod_{i=1}^{N} \exp\left(\beta J s_i s_{i+1} + \frac{\beta B}{2} (s_i + s_{i+1})\right)$$
(5.23)

The crucial observation that allows us to solve the problem is that this partition function can be written as a product of matrices. We adopt notation from quantum mechanics and define the 2×2 matrix,

$$\langle s_i | T | s_{i+1} \rangle \equiv \exp\left(\beta J s_i s_{i+1} + \frac{\beta B}{2} (s_i + s_{i+1})\right)$$
(5.24)

The row of the matrix is specified by the value of $s_i = \pm 1$ and the column by $s_{i+1} = \pm 1$. T is known as the *transfer matrix* and, in more conventional notation, is given by

$$T = \begin{pmatrix} e^{\beta J + \beta B} & e^{-\beta J} \\ e^{-\beta J} & e^{\beta J - \beta B} \end{pmatrix}$$

The sums over the spins \sum_{s_i} and product over lattice sites \prod_i in (5.23) simply tell us to multiply the matrices defined in (5.24) and the partition function becomes

$$Z = \operatorname{Tr} \left(\langle s_1 | T | s_2 \rangle \langle s_2 | T | s_3 \rangle \dots \langle s_N | T | s_1 \rangle \right) = \operatorname{Tr} T^N$$
(5.25)

where the trace arises because we have imposed periodic boundary conditions. To complete the story, we need only compute the eigenvalues of T to determine the partition function. A quick calculation shows that the two eigenvalues of T are

$$\lambda_{\pm} = e^{\beta J} \cosh\beta B \pm \sqrt{e^{2\beta J} \cosh^2\beta B - 2\sinh 2\beta J}$$
(5.26)

where, clearly, $\lambda_{-} < \lambda_{+}$. The partition function is then

$$Z = \lambda_{+}^{N} + \lambda_{-}^{N} = \lambda_{+}^{N} \left(1 + \frac{\lambda_{-}^{N}}{\lambda_{+}^{N}} \right) \approx \lambda_{+}^{N}$$
(5.27)

where, in the last step, we've used the simple fact that if λ_+ is the largest eigenvalue then $\lambda_-^N/\lambda_+^N \approx 0$ for very large N.

The partition function Z contains many quantities of interest. In particular, we can use it to compute the magnetisation as a function of temperature when B = 0. This, recall, is the quantity which is predicted to undergo a phase transition in the mean field approximation, going abruptly to zero at some critical temperature. In the d = 1Ising model, the magnetisation is given by

$$m = \frac{1}{N\beta} \left. \frac{\partial \log Z}{\partial B} \right|_{B=0} = \frac{1}{\lambda_+\beta} \left. \frac{\partial \lambda_+}{\partial B} \right|_{B=0} = 0$$

We see that the true physics for d = 1 is very different than that suggested by the mean field approximation. When B = 0, there is no magnetisation! While the J term in the energy encourages the spins to align, this is completely overwhelmed by thermal fluctuations for any value of the temperature.

There is a general lesson in this calculation: thermal fluctuations always win in one dimensional systems. They never exhibit ordered phases and, for this reason, never exhibit phase transitions. The mean field approximation is bad in one dimension.

5.3.2 2d Ising Model: Low Temperatures and Peierls Droplets

Let's now turn to the Ising model in d = 2 dimensions. We'll work on a square lattice and set B = 0. Rather than trying to solve the model exactly, we'll have more modest goals. We will compute the partition function in two different limits: high temperature and low temperature. We start here with the low temperature expansion.

The partition function is given by the sum over all states, weighted by $e^{-\beta E}$. At low temperatures, this is always dominated by the lowest lying states. For the Ising model,

we have

$$Z = \sum_{\{s_i\}} \exp\left(\beta J \sum_{\langle ij \rangle} s_i s_j\right)$$

The low temperature limit is $\beta J \to \infty$, where the partition function can be approximated by the sum over the first few lowest energy states. All we need to do is list these states.

The ground states are easy. There are two of them: spins all up or spins all down. For example, the ground state with spins all up looks like



Each of these ground states has energy $E = E_0 = -2NJ$.

The first excited states arise by flipping a single spin. Each spin has q = 4 nearest neighbours – denoted by red lines in the example below – each of which leads to an energy cost of 2J. The energy of each first excited state is therefore $E_1 = E_0 + 8J$.



There are, of course, N different spins that we we can flip and, correspondingly, the first energy level has a degeneracy of N.

To proceed, we introduce a diagrammatic method to list the different states. We draw only the "broken" bonds which connect two spins with opposite orientation and, as in the diagram above, denote these by red lines. We further draw the flipped spins as red dots, the unflipped spins as blue dots. The energy of the state is determined simply by the number of red lines in the diagram. Pictorially, we write the first excited state as



The next lowest state has six broken bonds. It takes the form



where the extra factor of 2 in the degeneracy comes from the two possible orientations (vertical and horizontal) of the graph.

Things are more interesting for the states which sit at the third excited level. These have 8 broken bonds. The simplest configuration consists of two, disconnected, flipped spins



The factor of N in the degeneracy comes from placing the first graph; the factor of N-5 arises because the flipped spin in the second graph can sit anywhere apart from on the five vertices used in the first graph. Finally, the factor of 1/2 arises from the interchange of the two graphs.

There are also three further graphs with the same energy E_3 . These are



and

where the degeneracy comes from the two orientations (vertical and horizontal). And, finally,



where the degeneracy comes from the four orientations (rotating the graph by 90°).

Adding all the graphs above together gives us an expansion of the partition function in power of $e^{-\beta J} \ll 1$. This is

$$Z = 2e^{2N\beta J} \left(1 + Ne^{-8\beta J} + 2Ne^{-12\beta J} + \frac{1}{2}(N^2 + 9N)e^{-16\beta J} + \dots \right)$$
(5.29)

where the overall factor of 2 originates from the two ground states of the system. We'll make use of the specific coefficients in this expansion in Section 5.3.4. Before we focus on the physics hiding in the low temperature expansion, it's worth making a quick comment that something quite nice happens if we take the log of the partition function,

$$\log Z = \log 2 + 2N\beta J + Ne^{-8\beta J} + 2Ne^{-12\beta J} + \frac{9}{2}Ne^{-16\beta J} + \dots$$

The thing to notice is that the N^2 term in the partition function (5.29) has cancelled out and log Z is proportional to N, which is to be expected since the free energy of the system is extensive. Looking back, we see that the N^2 term was associated to the disconnected diagrams in (5.28). There is actually a general lesson hiding here: the partition function can be written as the exponential of the sum of connected diagrams. We saw exactly the same issue arise in the cluster expansion in (2.37).

Peierls Droplets

Continuing the low temperature expansion provides a heuristic, but physically intuitive, explanation for why phase transitions happen in $d \ge 2$ dimensions but not in d = 1. As we flip more and more spins, the low energy states become *droplets*, consisting of a region of space in which all the spins are flipped, surrounded by a larger sea in which the spins have their original alignment. The energy cost of such a droplet is roughly

$$E \sim 2JL$$

where L is the perimeter of the droplet. Notice that the energy does not scale as the area of the droplet since all spins inside are aligned with their neighbours. It is only those on the edge which are misaligned and this is the reason for the perimeter scaling. To understand how these droplets contribute to the partition function, we also need to know their degeneracy. We will now argue that the degeneracy of droplets scales as

Degeneracy
$$\sim e^{\alpha L}$$

for some value of α . To see this, consider firstly the problem of a random walk on a 2d square lattice. At each step, we can move in one of four directions. So the number of paths of length L is

$$\#$$
paths $\sim 4^L = e^{L \log 4}$

Of course, the perimeter of a droplet is more constrained than a random walk. Firstly, the perimeter can't go back on itself, so it really only has three directions that it can move in at each step. Secondly, the perimeter must return to its starting point after L steps. And, finally, the perimeter cannot self-intersect. One can show that the number of paths that obey these conditions is

#paths
$$\sim e^{\alpha L}$$

where $\log 2 < \alpha < \log 3$. Since the degeneracy scales as $e^{\alpha L}$, the entropy of the droplets is proportional to L.

The fact that both energy and entropy scale with L means that there is an interesting competition between them. At temperatures where the droplets are important, the partition function is schematically of the form

$$Z \sim \sum_L e^{\alpha L} e^{-2\beta JL}$$

For large β (i.e. low temperature) the partition function converges. However, as the temperature increases, one reaches the critical temperature

$$k_B T_c \approx \frac{2J}{\alpha} \tag{5.30}$$

where the partition function no longer converges. At this point, the entropy wins over the energy cost and it is favourable to populate the system with droplets of arbitrary sizes. This is the how one sees the phase transition in the partition function. For temperature above T_c , the low-temperature expansion breaks down and the ordered magnetic phase is destroyed. We can also use the droplet argument to see why phase transitions don't occur in d = 1 dimension. On a line, the boundary of any droplet always consists of just two points. This means that the energy cost to forming a droplet is always E = 2J, regardless of the size of the droplet. But, since the droplet can exist anywhere along the line, its degeneracy is N. The net result is that the free energy associated to creating a droplet scales as

$$F \sim 2J - k_B T \log N$$

and, as $N \to \infty$, the free energy is negative for any T > 0. This means that the system will prefer to create droplets of arbitrary length, randomizing the spins. This is the intuitive reason why there is no magnetic ordered phase in the d = 1 Ising model.

5.3.3 2d Ising Model: High Temperatures

We now turn to the 2d Ising model in the opposite limit of high temperature. Here we expect the partition function to be dominated by the completely random, disordered configurations of maximum entropy. Our goal is to find a way to expand the partition function in $\beta J \ll 1$.

We again work with zero magnetic field, B = 0 and write the partition function as

$$Z = \sum_{\{s_i\}} \exp\left(\beta J \sum_{\langle ij \rangle} s_i s_j\right) = \sum_{\{s_i\}} \prod_{\langle ij \rangle} e^{\beta J s_i s_j}$$

There is a useful way to rewrite $e^{\beta J s_i s_j}$ which relies on the fact that the product $s_i s_j$ only takes ± 1 . It doesn't take long to check the following identity:

$$e^{\beta J s_i s_j} = \cosh \beta J + s_i s_j \sinh \beta J$$
$$= \cosh \beta J \left(1 + s_i s_j \tanh \beta J\right)$$

Using this, the partition function becomes

$$Z = \sum_{\{s_i\}} \prod_{\langle ij \rangle} \cosh \beta J \left(1 + s_i s_j \tanh \beta J \right)$$

= $(\cosh \beta J)^{qN/2} \sum_{\{s_i\}} \prod_{\langle ij \rangle} \left(1 + s_i s_j \tanh \beta J \right)$ (5.31)

where the number of nearest neighbours is q = 4 for the 2d square lattice.

With the partition function in this form, there is a natural expansion which suggests itself. At high temperatures $\beta J \ll 1$ which, of course, means that $\tanh \beta J \ll 1$. But the partition function is now naturally a product of powers of $\tanh \beta J$. This is somewhat analogous to the cluster expansion for the interacting gas that we met in Section 2.5.3. As in the cluster expansion, we will represent the expansion graphically.

We need no graphics for the leading order term. It has no factors of $\tanh\beta J$ and is simply

$$Z \approx (\cosh \beta J)^{2N} \sum_{\{s_i\}} 1 = 2^N (\cosh \beta J)^{2N}$$

That's simple.

Let's now turn to the leading correction. Expanding the partition function (5.31), each power of $\tanh \beta J$ is associated to a nearest neighbour pair $\langle ij \rangle$. We'll represent this by drawing a line on the lattice:

$$_{i} \circ \circ \circ_{j} = s_{i} s_{j} \tanh \beta J$$

But there's a problem: each factor of $\tanh \beta J$ in (5.31) also comes with a sum over all spins s_i and s_j . And these are +1 and -1 which means that they simply sum to zero,

$$\sum_{s_i, s_j} s_i s_j = +1 - 1 - 1 + 1 = 0$$

How can we avoid this? The only way is to make sure that we're summing over an even number of spins on each site, since then we get factors of $s_i^2 = 1$ and no cancellations. Graphically, this means that every site must have an even number of lines attached to it. The first correction is then of the form

$$\begin{array}{c}
1 \\
3 \\
3 \\
4
\end{array} = (\tanh \beta J)^4 \sum_{\{s_i\}} s_1 s_2 s_2 s_3 s_3 s_4 s_4 s_1 = 2^4 (\tanh \beta J)^4$$

There are N such terms since the upper left corner of the square can be on any one of the N lattice sites. (Assuming periodic boundary conditions for the lattice). So including the leading term and first correction, we have

$$Z = 2^N (\cosh\beta J)^{2N} \left(1 + N (\tanh\beta J)^4 + \ldots \right)$$

We can go further. The next terms arise from graphs of length 6 and the only possibilities are rectangles, oriented as either landscape or portrait. Each of them can sit on one of N sites, giving a contribution

$$\begin{array}{c} & & \\ & &$$

Things get more interesting when we look at graphs of length 8. We have four different types of graphs. Firstly, there are the trivial, disconnected pair of squares

$$\underbrace{\bigcap_{n=1}^{\infty}}_{n=1} \underbrace{\bigcap_{n=1}^{\infty}}_{n=1} = \frac{1}{2}N(N-5)(\tanh\beta J)^8$$

Here the first factor of N is the possible positions of the first square; the factor of N-5 arises because the possible location of the upper corner of the second square can't be on any of the vertices of the first, but nor can it be on the square one to the left of the upper corner of the first since that would give a graph that looks like \square which has three lines coming off the middle site and therefore vanishes when we sum over spins. Finally, the factor of 1/2 comes because the two squares are identical.

The other graphs of length 8 are a large square, a rectangle and a corner. The large square gives a contribution

There are two orientations for the rectangle. Including these gives a factor of 2,

Finally, the corner graph has four orientations, giving

Adding all contributions together gives us the first few terms in high temperature expansion of the partition function

$$Z = 2^{N} (\cosh\beta J)^{2N} \left(1 + N (\tanh\beta J)^{4} + 2N (\tanh\beta J)^{6} + \frac{1}{2} (N^{2} + 9N) (\tanh\beta J)^{8} + \dots \right)$$
(5.32)

There's some magic hiding in this expansion which we'll turn to in Section 5.3.4. First, let's just see how the high energy expansion plays out in the d = 1 dimensional Ising model.

The Ising Chain Revisited

Let's do the high temperature expansion for the d = 1 Ising chain with periodic boundary conditions and B = 0. We have the same partition function (5.31) and the same issue that only graphs with an even number of lines attached to each vertex contribute. But, for the Ising chain, there is only one such term: it is the closed loop. This means that the partition function is



$$Z = 2^{N} (\cosh \beta J)^{N} \left(1 + (\tanh \beta J)^{N} \right)$$
 Figure 48:

In the limit $N \to \infty$, $(\tanh \beta J)^N \to 0$ at high temperatures and even the contribution from the closed loop vanishes. We're left with

$$Z = (2\cosh\beta J)^N$$

This agrees with our exact result for the Ising chain given in (5.27), which can be seen by setting B = 0 in (5.26) so that $\lambda_{+} = 2 \cosh \beta J$.

5.3.4 Kramers-Wannier Duality

In the previous sections we computed the partition function perturbatively in two extreme regimes of low temperature and high temperature. The physics in the two cases is, of course, very different. At low temperatures, the partition function is dominated by the lowest energy states; at high temperatures it is dominated by maximally disordered states. Yet comparing the partition functions at low temperature (5.29) and high temperature (5.32) reveals an extraordinary fact: the expansions are the same! More concretely, the two series agree if we exchange

$$e^{-2\beta J} \longleftrightarrow \tanh \beta J$$
 (5.33)

Of course, we've only checked the agreement to the first few orders in perturbation theory. Below we shall prove that this miracle continues to all orders in perturbation theory. The symmetry of the partition function under the interchange (5.33) is known as *Kramers-Wannier duality*. Before we prove this duality, we will first just assume that it is true and extract some consequences.

We can express the statement of the duality more clearly. The Ising model at temperature β is related to the same model at temperature $\tilde{\beta}$, defined as

$$e^{-2\tilde{\beta}J} = \tanh\beta J \tag{5.34}$$

This way of writing things hides the symmetry of the transformation. A little algebra shows that this is equivalent to

$$\sinh 2\tilde{\beta}J = \frac{1}{\sinh 2\beta J}$$

Notice that this is a hot/cold duality. When βJ is large, $\tilde{\beta} J$ is small. Kramers-Wannier duality is the statement that, when B = 0, the partition functions of the Ising model at two temperatures are related by

$$Z[\beta] = \frac{2^{N} (\cosh \beta J)^{2N}}{2e^{2N\tilde{\beta}J}} Z[\tilde{\beta}]$$

= 2^{N-1} (\cosh \beta J \sinh \beta J)^{N} Z[\tilde{\beta}] (5.35)

This means that if you know the thermodynamics of the Ising model at one temperature, then you also know the thermodynamics at the other temperature. Notice however, that it does not say that all the physics of the two models is equivalent. In particular, when one system is in the ordered phase, the other typically lies in the disordered phase.

One immediate consequence of the duality is that we can use it to compute the exact critical temperature T_c . This is the temperature at which the partition function is singular in the $N \to \infty$ limit. (We'll discuss a more refined criterion in Section 5.4.3). If we further assume that there is just a single phase transition as we vary the temperature, then it must happen at the special self-dual point $\beta = \tilde{\beta}$. This is

$$k_B T = \frac{2J}{\log(\sqrt{2}+1)} \approx 2.269 \, J$$

The exact solution of Onsager confirms that this is indeed the transition temperature. It's also worth noting that it's fully consistent with the more heuristic Peierls droplet argument (5.30) since $\log 2 < \log(\sqrt{2} + 1) < \log 3$.

Proving the Duality

So far our evidence for the duality (5.35) lies in the agreement of the first few terms in the low and high temperature expansions (5.29) and (5.32). Of course, we could keep computing further and further terms and checking that they agree, but it would be nicer to simply prove the equality between the partition functions. We shall do so here. The key idea that we need can actually be found by staring hard at the various graphs that arise in the two expansions. Eventually, you will realise that they are the same, albeit drawn differently. For example, consider the two "corner" diagrams



The two graphs are *dual*. The red lines in the first graph intersect the black lines in the second as can be seen by placing them on top of each other:



The same pattern occurs more generally: the graphs appearing in the low temperature expansion are in one-to-one correspondence with the dual graphs of the high temperature expansion. Here we will show how this occurs and how one can map the partition functions onto each other.

Let's start by writing the partition function in the form (5.31) that we met in the high temperature expansion and presenting it in a slightly different way,

$$Z[\beta] = \sum_{\{s_i\}} \prod_{\langle ij \rangle} (\cosh \beta J + s_i s_j \sinh \beta J)$$
$$= \sum_{\{s_i\}} \prod_{\langle ij \rangle} \sum_{k_{ij}=0,1} C_{k_{ij}}[\beta J] (s_i s_j)^{k_{ij}}$$

where we have introduced the rather strange variable k_{ij} associated to each nearest neighbour pair that takes values 0 and 1, together with the functions.

$$C_0[\beta J] = \cosh \beta J$$
 and $C_1[\beta J] = \sinh \beta J$

The variables in the original Ising model were spins on the lattice sites. The observation that the graphs which appear in the two expansions are dual suggests that it might be profitable to focus attention on the links between lattice sites. Clearly, we have one link for every nearest neighbour pair. If we label these links by l, we can trivially rewrite the partition function as

$$Z = \sum_{k_l=0,1} \prod_{l} \sum_{\{s_i\}} C_{k_l} [\beta J] (s_i s_j)^{k_l}$$

Notice that the strange label k_{ij} has now become a variable that lives on the links l rather than the original lattice sites i.

At this stage, we do the sum over the spins s_i . We've already seen that if a given spin, say s_i , appears in a term an odd number of times, then that term will vanish when we sum over the spin. Alternatively, if the spin s_i appears an even number of times, then the sum will give 2. We'll say that a given link l is turned on in configurations with $k_l = 1$ and turned off when $k_l = 0$. In this language, a term in the sum over spin s_i contributes only if an even number of links attached to site i are turned on. The partition function then becomes

$$Z = 2^{N} \sum_{k_{l}} \prod_{l} C_{k_{l}}[\beta J] \bigg|_{\text{Constrained}}$$
(5.36)

Now we have something interesting. Rather than summing over spins on lattice sites, we're now summing over the new variables k_l living on links. This looks like the partition function of a totally different physical system, where the degrees of freedom live on the links of the original lattice. But there's a catch – that big "Constrained" label on the sum. This is there to remind us that we don't sum over all k_l configurations; only those for which an even number of links are turned on for every lattice site. And that's annoying. It's telling us that the k_l aren't really independent variables. There are some constraints that must be imposed.

Fortunately, for the 2d square lattice, there is a simple way to solve the constraint. We introduce yet more variables, \tilde{s}_i which, like the original spin variables, take values ± 1 . However, the \tilde{s}_i do not live on the original lattice sites. Instead, they live on the vertices of the dual lattice. For the 2d square lattice, the dual vertices are drawn in the figure. The original lattice sites are in white; the dual lattice sites in black.



Figure 49:

The link variables k_l are related to the two nearest spin variables \tilde{s}_i as follows:

$$k_{12} = \frac{1}{2}(1 - \tilde{s}_1 \tilde{s}_2)$$

$$k_{13} = \frac{1}{2}(1 - \tilde{s}_2 \tilde{s}_3)$$

$$k_{14} = \frac{1}{2}(1 - \tilde{s}_3 \tilde{s}_4)$$

$$k_{15} = \frac{1}{2}(1 - \tilde{s}_1 \tilde{s}_4)$$

Notice that we've replaced four variables k_l taking values 0, 1 with four variables \tilde{s}_i taking values ± 1 . Each set of variables gives 2^4 possibilities. However, the map is not one-to-one. It is not possible to construct for all values of k_l using the parameterization in terms of \tilde{s}_i . To see this, we need only look at

$$k_{12} + k_{13} + k_{14} + k_{15} = 2 - \frac{1}{2} (\tilde{s}_1 \tilde{s}_2 + \tilde{s}_2 \tilde{s}_3 + \tilde{s}_3 \tilde{s}_4 + \tilde{s}_1 \tilde{s}_4)$$

= $2 - \frac{1}{2} (\tilde{s}_1 + \tilde{s}_3) (\tilde{s}_2 + \tilde{s}_4)$
= 0, 2, or 4

In other words, the number of links that are turned on must be even. But that's exactly what we want! Writing the k_l in terms of the auxiliary spins \tilde{s}_i automatically solves the constraint that is imposed on the sum in (5.36). Moreover, it is simple to check that for every configuration $\{k_l\}$ obeying the constraint, there are two configurations of $\{\tilde{s}_i\}$. This means that we can replace the constrained sum over $\{k_l\}$ with an unconstrained sum over $\{\tilde{s}_i\}$. The only price we pay is an additional factor of 1/2.

$$Z[\beta] = \frac{1}{2} 2^N \sum_{\{\tilde{s}_i\}} \prod_{\langle ij \rangle} C_{\frac{1}{2}(1-\tilde{s}_i \tilde{s}_j)}[\beta j]$$

Finally, we'd like to find a simple expression for C_0 and C_1 in terms of \tilde{s}_i . That's easy enough. We can write

$$C_k[\beta J] = \cosh\beta J \exp(k\log\tanh\beta J)$$

= $(\sinh\beta J \cosh\beta J)^{1/2} \exp\left(-\frac{1}{2}\tilde{s}_i\tilde{s}_j\log\tanh\beta J\right)$

Substituting this into our newly re-written partition function gives

$$Z[\beta] = 2^{N-1} \sum_{\{\tilde{s}_i\}} \prod_{\langle ij \rangle} (\sinh \beta J \cosh \beta J)^{1/2} \exp\left(-\frac{1}{2} \tilde{s}_i \tilde{s}_j \log \tanh \beta J\right)$$
$$= 2^{N-1} (\sinh \beta J \cosh \beta J)^N \sum_{\{\tilde{s}_i\}} \exp\left(-\frac{1}{2} \log \tanh \beta J \sum_{\langle ij \rangle} \tilde{s}_i \tilde{s}_j\right)$$

But this final form of the partition function in terms of the dual spins \tilde{s}_i has exactly the same functional form as the original partition function in terms of the spins s_i . More precisely, we can write

$$Z[\beta] = 2^{N-1} (\sinh 2\beta J)^N Z[\tilde{\beta}]$$

where

$$e^{-2\tilde{\beta}J} = \tanh\beta J$$

as advertised previously in (5.34). This completes the proof of Kramers-Wannier duality in the 2d Ising model on a square lattice.

The concept of *duality* of this kind is a major feature in much of modern theoretical physics. The key idea is that when the temperature gets large there may be a different set of variables in which a theory can be written where it appears to live at low temperature. The same idea often holds in quantum theories, where duality maps strong coupling problems to weak coupling problems.

The duality in the Ising model is special for two reasons: firstly, the new variables \tilde{s}_i are governed by the same Hamiltonian as the original variables s_i . We say that the Ising model is self-dual. In general, this need not be the case — the high temperature limit of one system could look like the low-temperature limit of a very different system. Secondly, the duality in the Ising model can be proven explicitly. For most systems, we have no such luck. Nonetheless, the idea that there may be dual variables in other, more difficult theories, is compelling. Commonly studied examples include the exchange particles and vortices in two dimensions, and electrons and magnetic monopoles in three dimensions.

5.4 Landau Theory

We saw in Sections 5.1 and 5.2 that the van der Waals equation and mean field Ising model gave the same (sometimes wrong!) answers for the critical exponents. This suggests that there should be a unified way to look at phase transitions. Such a method was developed by Landau. It is worth stressing that, as we saw above, the Landau approach to phase transitions often only gives qualitatively correct results. However, its advantage is that it is extremely straightforward and easy. (Certainly much easier than the more elaborate methods needed to compute critical exponents more accurately). The Landau theory of phase transitions is based around the free energy. We will illustrate the theory using the Ising model and then explain how to extend it to different systems. The free energy of the Ising model in the mean field approximation is readily attainable from the partition function (5.17),

$$F = -\frac{1}{\beta}\log Z = \frac{1}{2}JNqm^2 - \frac{N}{\beta}\log\left(2\cosh\beta B_{\text{eff}}\right)$$
(5.37)

So far in this course, we've considered only systems in equilibrium. The free energy, like all other thermodynamic potentials, has only been defined on equilibrium states. Yet the equation above can be thought of as an expression for F as a function of m. Of course, we could substitute in the equilibrium value of m given by solving (5.18), but it seems a shame to throw out F(m) when it is such a nice function. Surely we can put it to some use!

The key step in Landau theory is to treat the function F = F(T, V; m) seriously. This means that we are extending our viewpoint away from equilibrium states to a whole class of states which have a constant average value of m. If you want some words to drape around this, you could imagine some external magical power that holds mfixed. The free energy F(T, V; m) is then telling us the equilibrium properties in the presence of this magical power. Perhaps more convincing is what we do with the free energy in the absence of any magical constraint. We saw in Section 4 that equilibrium is guaranteed if we sit at the minimum of F. Looking at extrema of F, we have the condition

$$\frac{\partial F}{\partial m} = 0 \quad \Rightarrow \quad m = \tanh\beta B_{\rm eff}$$

But that's precisely the condition (5.18) that we saw previously. Isn't that nice!

In the context of Landau theory, m is called an *order parameter*. When it takes nonzero values, the system has some degree of order (the spins have a preferred direction in which they point) while when m = 0 the spins are randomised and happily point in any direction.

For any system of interest, Landau theory starts by identifying a suitable order parameter. This should be taken to be a quantity which vanishes above the critical temperature at which the phase transition occurs, but is non-zero below the critical temperature. Sometimes it is obvious what to take as the order parameter; other times less so. For the liquid-gas transition, the relevant order parameter is the difference in densities between the two phases, $v_{gas} - v_{liquid}$. For magnetic or electric systems, the order parameter is typically some form of magnetization (as for the Ising model) or the polarization. For the Bose-Einstein condensate, superfluids and superconductors, the order parameter is more subtle and is related to off-diagonal long-range order in the one-particle density matrix¹¹, although this is usually rather lazily simplified to say that the order parameter can be thought of as the macroscopic wavefunction $|\psi|^2$.

Starting from the existence of a suitable order parameter, the next step in the Landau programme is to write down the free energy. But that looks tricky. The free energy for the Ising model (5.37) is a rather complicated function and clearly contains some detailed information about the physics of the spins. How do we just write down the free energy in the general case? The trick is to assume that we can expand the free energy in an analytic power series in the order parameter. For this to be true, the order parameter must be small which is guaranteed if we are close to a critical point (since m = 0 for $T > T_c$). The nature of the phase transition is determined by the kind of terms that appear in the expansion of the free energy. Let's look at a couple of simple examples.

5.4.1 Second Order Phase Transitions

We'll consider a general system (Ising model; liquid-gas; BEC; whatever) and denote the order parameter as m. Suppose that the expansion of the free energy takes the general form

$$F(T;m) = F_0(T) + a(T)m^2 + b(T)m^4 + \dots$$
(5.38)

One common reason why the free energy has this form is because the theory has a symmetry under $m \to -m$, forbidding terms with odd powers of m in the expansion. For example, this is the situation in the Ising model when B = 0. Indeed, if we expand out the free energy (5.37) for the Ising model for small m using $\cosh x \approx 1 + \frac{1}{2}x^2 + \frac{1}{4!}x^4 + \ldots$ and $\log(1+y) \approx y - \frac{1}{2}y^2 + \ldots$ we get the general form above with explicit expressions for $F_0(T)$, a(T) and b(T),

$$F_{Ising}(T;m) = -Nk_B T \log 2 + \left(\frac{NJq}{2}(1 - Jq\beta)\right)m^2 + \left(\frac{N\beta^3 J^4 q^4}{24}\right)m^4 + \dots$$

The leading term $F_0(T)$ is unimportant for our story. We are interested in how the free energy changes with m. The condition for equilibrium is given by

$$\frac{\partial F}{\partial m} = 0 \tag{5.39}$$

But the solutions to this equation depend on the sign of the coefficients a(T) and

 $^{^{11}\}mathrm{See},$ for example, the book "Quantum Liquids" by Anthony Leggett



Figure 50: Free energy when a(T) > 0 Figure 51: Free energy when a(T) < 0

b(T). Moreover, this sign can change with temperature. This is the essence of the phase transitions. In the following discussion, we will assume that b(T) > 0 for all T. (If we relax this condition, we have to also consider the m^6 term in the free energy which leads to interesting results concerning so-called tri-critical points).

The two figures above show sketches of the free energy in the case where a(T) > 0and a(T) < 0. Comparing to the explicit free energy of the Ising model, a(T) > 0when $T > T_c = Jq/k_B$ and a(T) < 0 when $T < T_c$. When a(T) > 0, we have just a single equilibrium solution to (5.39) at m = 0. This is typically the situation at high temperatures. In contrast, at a(T) < 0, there are three solutions to (5.39). The solution m = 0 clearly has higher free energy: this is now the unstable solution. The two stable solutions sit at $m = \pm m_0$. For example, if we choose to truncate the free energy (5.38) at quartic order, we have

$$m_0 = \sqrt{\frac{-a}{2b}} \qquad T < T_c$$

If a(T) is a smooth function then the equilibrium value of m changes continuously from m = 0 when a(T) > 0 to $m \neq 0$ at a(T) < 0. This describes a second order phase transition occurring at T_c , defined by $a(T_c) = 0$.

Once we know the equilibrium value of m, we can then substitute this back into the free energy F(T;m) in (5.38). This gives the thermodynamic free energy F(T) of the system in equilibrium that we have been studying throughout this course. For the quartic free energy, we have

$$F(T) = \begin{cases} F_0(T) & T > T_c \\ F_0(T) - a^2/4b & T < T_c \end{cases}$$
(5.40)

Because $a(T_c) = 0$, the equilibrium free energy F(T) is continuous at $T = T_c$. Moreover, the entropy $S = -\partial F/\partial T$ is also continuous at $T = T_c$. However, if you differentiate the

equilibrium free energy twice, you will get a term a'^2/b which is generically not vanishing at $T = T_c$. This means that the heat capacity $C = T\partial S/\partial T$ changes discontinuously at $T = T_c$, as befits a second order phase transition. A word of warning: if you want to compute equilibrium quantities such as the heat capacity, it's important that you first substitute in the equilibrium value of m and work with (5.40) rather than with (5.38). If you don't, you miss the fact that the magnetization also changes with T.

We can easily compute critical exponents within the context of Landau theory. We need only make further assumptions about the behaviour of a(T) and b(T) in the vicinity of T_c . If we assume that near $T = T_c$, we can write

$$b(T) \approx b_0$$
 , $a(T) \approx a_0(T - T_c)$ (5.41)

then we have

$$m_0 \approx \pm \sqrt{\frac{a_0}{2b_0}} (T_c - T)^{1/2} \qquad T < T_c$$

which reproduces the critical exponent (5.9) and (5.20) that we derived for the van der Waals equation and Ising model respectively.

Notice that we didn't put any discontinuities into the free energy. Everything in F(T; m) was nice and smooth. When Taylor expanded, it has only integer powers of m and T as shown in (5.38) and (5.41). But the minima of F behave in a non-analytic fashion as seen in the expression for m_0 above.

Landau's theory of phase transitions predicts this same critical exponent for all values of the dimension d of the system. But we've already mentioned in previous contexts that the critical exponent is in fact only correct for $d \ge 4$. We will understand how to derive this criterion from Landau theory in the next section.

Spontaneous Symmetry Breaking

As we approach the end of the course, we're touching upon a number of ideas that become increasingly important in subsequent developments in physics. We already briefly met the idea of universality and critical phenomena. Here I would like to point out another very important idea: *spontaneous symmetry breaking*.

The free energy (5.38) is invariant under the \mathbb{Z}_2 symmetry $m \to -m$. Indeed, we said that one common reason that we can expand the free energy only in even powers of m is that the underlying theory also enjoys this symmetry. But below T_c , the system must pick one of the two ground states $m = +m_0$ or $m = -m_0$. Whichever choice it makes breaks the \mathbb{Z}_2 symmetry. We say that the symmetry is spontaneously broken by the choice of ground state of the theory.

Spontaneous symmetry breaking has particularly dramatic consequences when the symmetry in question is continuous rather than discrete. For example, consider a situation where the order parameter is a complex number ψ and the free energy is given by (5.38) with $m = |\psi|^2$. (This is effectively what happens for BECs, superfluids and superconductors). Then we should only look at the m > 0 solutions so that the ground state has $|\psi|^2 = +m_0$. But this leaves the phase of ψ completely undetermined. So there is now a continuous choice of ground states: we get to sit anywhere on the circle parameterised by the phase of ψ . Any choice that the system makes spontaneously breaks the U(1) rotational symmetry which acts on the phase of ψ . Some beautiful results due to Nambu and Goldstone show that the much of the physics of these systems can be understood simply as a consequence of this symmetry breaking. The ideas of spontaneous symmetry breaking are crucial in both condensed matter physics and particle physics. In the latter context, it is intimately tied with the Higgs mechanism.

5.4.2 First Order Phase Transitions

Let us now consider a situation where the expansion of the free energy also includes odd powers of the order parameter

$$F(T;m) = F_0(T) + \alpha(T)m + a(T)m^2 + \gamma(T)m^3 + b(T)m^4 + \dots$$

For example, this is the kind of expansion that we get for the Ising model free energy (5.37) when $B \neq 0$, which reads

$$F_{\text{Ising}}(T;m) = -Nk_BT\log 2 + \frac{JNq}{2}m^2 - \frac{N}{2k_BT}(B + Jqm)^2 + \frac{N}{24(k_BT)^3}(B + Jqm)^4 + \dots$$

Notice that there is no longer a symmetry relating $m \to -m$: the *B* field has a preference for one sign over the other.

If we again assume that b(T) > 0 for all temperatures, the crude shape of the free energy graph again has two choices: there is a single minimum, or two minima and a local maximum.

Let's start at suitably low temperatures for which the situation is depicted in Figure 52. The free energy once again has a double well, except now slightly skewed. The local maximum is still an unstable point. But this time around, the minima with the lower free energy is preferred over the other one. This is the true ground state of the system. In contrast, the point which is locally, but not globally, a minimum corresponds to a meta-stable state of the system. In order for the system to leave this state, it must first fluctuate up and over the energy barrier separating the two.



Figure 52: The free energy of the Ising model for B < 0, B = 0 and B > 0.

In this set-up, we can initiate a first order phase transition. This occurs when the coefficient of the odd terms, $\alpha(T)$ and $\gamma(T)$ change sign and the true ground state changes discontinuously from m < 0 to m > 0. In some systems this behaviour occurs when changing temperature; in others it could occur by changing some external parameter. For example, in the Ising model the first order phase transition is induced by changing B.

At very high temperature, the double well potential is lost in favour of a single minimum as depicted in the figure to the right. There is a unique ground state, albeit shifted from m = 0 by the presence of the $\alpha(T)$ term above (which translates into the magnetic field *B* in the Ising model). The temperature at which the meta-stable ground state of the system is lost corresponds to the spinodal point in our discussion of the liquid-gas transition.



Figure 53:

One can play further games in Landau theory, looking at how the shape of the free energy can change as we vary temperature or other parameters. One can also use this framework to give a simple explanation of the concept of hysteresis. You can learn more about these from the links on the course webpage.

5.4.3 Lee-Yang Zeros

You may have noticed that the flavour of our discussion of phase transitions is a little different from the rest of the course. Until now, our philosophy was to derive everything from the partition function. But in this section, we dumped the partition function as soon as we could, preferring instead to work with the macroscopic variables such as the free energy. Why didn't we just stick with the partition function and examine phase transitions directly?

The reason, of course, is that the approach using the partition function is hard! In this short section, which is somewhat tangential to our main discussion, we will describe how phase transitions manifest themselves in the partition function.

For concreteness, let's go back to the classical interacting gas of Section 2.5, although the results we derive will be more general. We'll work in the grand canonical ensemble, with the partition function

$$\mathcal{Z}(z,V,T) = \sum_{N} z^{N} Z(N,V,T) = \sum_{N} \frac{z^{N}}{N! \lambda^{3N}} \int \prod_{i} d^{3} r_{i} \ e^{-\beta \sum_{j < k} U(r_{jk})}$$
(5.42)

To regulate any potential difficulties with short distances, it is useful to assume that the particles have hard-cores so that they cannot approach to a distance less than r_0 . We model this by requiring that the potential satisfies

$$U(r_{jk}) = 0 \quad \text{for} \quad r_{jk} < r_0$$

But this has an obvious consequence: if the particles have finite size, then there is a maximum number of particles, N_V , that we can fit into a finite volume V. (Roughly this number is $N_V \sim V/r_0^3$). But that, in turn, means that the canonical partition function Z(N, V, T) = 0 for $N > N_V$, and the grand partition function \mathcal{Z} is therefore a finite polynomial in the fugacity z, of order N_V . But if the partition function is a finite polynomial, there can't be any discontinuous behaviour associated with a phase transition. In particular, we can calculate

$$pV = k_B T \log \mathcal{Z} \tag{5.43}$$

which gives us pV as a smooth function of z. We can also calculate

$$N = z \frac{\partial}{\partial z} \log \mathcal{Z} \tag{5.44}$$

which gives us N as a function of z. Eliminating z between these two functions (as we did for both bosons and fermions in Section 3) tells us that pressure p is a smooth function of density N/V. We're never going to get the behaviour that we derived from the Maxwell construction in which the plot of pressure vs density shown in Figure 37 exhibits a discontinuous derivative.

The discussion above is just re-iterating a statement that we've alluded to several times already: there are no phase transitions in a finite system. To see the discontinuous behaviour, we need to take the limit $V \to \infty$. A theorem due to Lee and Yang¹² gives us a handle on the analytic properties of the partition function in this limit.

¹²This theorem was first proven for the Ising model in 1952. Soon afterwards, the same Lee and Yang proposed a model of parity violation in the weak interaction for which they won the 1957 Nobel prize.

The surprising insight of Lee and Yang is that if you're interested in phase transitions, you should look at the zeros of \mathcal{Z} in the complex z-plane. Let's firstly look at these when V is finite. Importantly, at finite V there can be no zeros on the positive real axis, z > 0. This follows from the definition of \mathcal{Z} given in (5.42) where it is a sum of positive quantities. Moreover, from (5.44), we can see that \mathcal{Z} is a monotonically increasing function of z because we necessarily have N > 0. Nonetheless, \mathcal{Z} is a polynomial in z of order N_V so it certainly has N_V zeros somewhere in the complex z-plane. Since $\mathcal{Z}^*(z) = \mathcal{Z}(z^*)$, these zeros must either sit on the real negative axis or come in complex pairs.

However, the statements above rely on the fact that \mathcal{Z} is a finite polynomial. As we take the limit $V \to \infty$, the maximum number of particles that we can fit in the system diverges, $N_V \to \infty$, and \mathcal{Z} is now defined as an infinite series. But infinite series can do things that finite ones can't. The Lee-Yang theorem says that as long as the zeros of \mathcal{Z} continue to stay away from the positive real axis as $V \to \infty$, then no phase transitions can happen. But if one or more zeros happen to touch the positive real axis, life gets more interesting.

More concretely, the Lee-Yang theorem states:

• Lee-Yang Theorem: The quantity

$$\Theta = \lim_{V \to \infty} \left(\frac{1}{V} \log \mathcal{Z}(z, V, T) \right)$$

exists for all z > 0. The result is a continuous, non-decreasing function of z which is independent of the shape of the box (up to some sensible assumptions such as Surface Area/ $V \sim V^{-1/3}$ which ensures that the box isn't some stupid fractal shape).

Moreover, let R be a fixed, volume independent, region in the complex z plane which contains part of the real, positive axis. If R contains no zero of $\mathcal{Z}(z, V, T)$ for all $z \in R$ then Θ is a an analytic function of z for all $z \in R$. In particular, all derivatives of Θ are continuous.

In other words, there can be no phase transitions in the region R even in the $V \to \infty$ limit. The last result means that, as long as we are safely in a region R, taking derivatives of with respect to z commutes with the limit $V \to \infty$. In other words, we are allowed to use (5.44) to write the particle density n = N/V as

$$\lim_{V \to \infty} n = \lim_{V \to \infty} z \frac{\partial}{\partial z} \left(\frac{p}{k_B T} \right) = z \frac{\partial \Theta}{\partial z}$$

However, if we look at points z where zeros appear on the positive real axis, then Θ will generally not be analytic. If $d\Theta/dz$ is discontinuous, then the system is said to undergo a first order phase transition. More generally, if $d^m\Theta/dz^m$ is discontinuous for m = n, but continuous for all m < n, then the system undergoes an n^{th} order phase transition. We won't offer a proof of the Lee-Yang theorem. Instead we just illustrate the general idea with an example.

A Made-Up Example

Ideally, we would like to start with a Hamiltonian which exhibits a first order phase transition, compute the associated grand partition function \mathcal{Z} and then follow its zeros as $V \to \infty$. However, as we mentioned above, that's hard! Instead we will simply make up a partition function \mathcal{Z} which has the appropriate properties. Our choice is somewhat artificial,

$$\mathcal{Z}(z,V) = (1+z)^{[\alpha V]}(1+z^{[\alpha V]})$$

Here α is a constant which will typically depend on temperature, although we'll suppress this dependence in what follows. Also,

$$[x] =$$
 Integer part of x

Although we just made up the form of \mathcal{Z} , it does have the behaviour that one would expect of a partition function. In particular, for finite V, the zeros sit at

$$z = -1$$
 and $z = e^{\pi i (2n+1)/[\alpha V]}$ $n = 0, 1, \dots, [\alpha V] - 1$

As promised, none of the zeros sit on the positive real axis. However, as we increase V, the zeros become denser and denser on the unit circle. From the Lee-Yang theorem, we expect that no phase transition will occur for $z \neq 1$ but that something interesting could happen at z = 1.

Let's look at what happens as we send $V \to \infty$. We have

$$\Theta = \lim_{V \to \infty} \frac{1}{V} \log \mathcal{Z}(z, V)$$

=
$$\lim_{V \to \infty} \frac{1}{V} \left([\alpha V] \log(1+z) + \log(1+z^{[\alpha V]}) \right)$$

=
$$\begin{cases} \alpha \log(1+z) & |z| < 1\\ \alpha \log(1+z) + \alpha \log z & |z| > 1 \end{cases}$$

We see that Θ is continuous for all z as promised. But it is only analytic for $|z| \neq 1$.

We can extract the physics by using (5.43) and (5.44) to eliminate the dependence on z. This gives us the equation of state, with pressure p as a function of n = V/N. For |z| < 1, we have

$$p = \alpha k_B T \log\left(\frac{\alpha}{\alpha - n}\right)$$
 $n \in [0, \alpha/2)$, $p < k_B T \log 2$

While for |z| > 1, we have

$$p = \alpha k_B T \log\left(\frac{2\alpha n}{(2\alpha - n)^2}\right)$$
 $n \in (3\alpha/2, 2\alpha)$, $p > k_B T \log 2$

The key point is that there is a jump in particle density of $\Delta n = \alpha$ at $p = \alpha k_B T \log 2$. Plotting this as a function of p vs v = 1/n, we find that we have a curve that is qualitatively identical to the pressure-volume plot of the liquid-gas phase diagram under the co-existence curve. (See, for example, figure 37). This is a first order phase transition.

5.5 Landau-Ginzburg Theory

Landau's theory of phase transition focusses only on the average quantity, the order parameter. It ignores the fluctuations of the system, assuming that they are negligible. Here we sketch a generalisation which attempts to account for these fluctuations. It is known as Landau-Ginzburg theory.

The idea is to stick with the concept of the order parameter, m. But now we allow the order parameter to vary in space so it becomes a function $m(\vec{r})$. Let's restrict ourselves to the situation where there is a symmetry of the theory $m \to -m$ so we need only consider even powers in the expansion of the free energy. We add to these a gradient term whose role is to captures the fact that there is some stiffness in the system, so it costs energy to vary the order parameter from one point to another. (For the example of the Ising model, this is simply the statement that nearby spins want to be aligned). The free energy is then given by

$$F[m(\vec{r})] = \int d^d r \ \left[a(T)m^2 + b(T)m^4 + c(T)(\nabla m)^2 \right]$$
(5.45)

where we have dropped the constant $F_0(T)$ piece which doesn't depend on the order parameter and hence plays no role in the story. Notice that we start with terms quadratic in the gradient: a term linear in the gradient would violate the rotational symmetry of the system. We again require that the free energy is minimised. But now F is a functional – it is a function of the function $m(\vec{r})$. To find the stationary points of such objects we need to use the same kind of variational methods that we use in Lagrangian mechanics. We write the variation of the free energy as

$$\delta F = \int d^d r \left[2am \,\delta m + 4bm^3 \,\delta m + 2c\nabla m \cdot \nabla \delta m \right]$$
$$= \int d^d r \left[2am + 4bm^3 - 2c\nabla^2 m \right] \delta m$$

where to go from the first line to the second we have integrated by parts. (We need to remember that c(T) is a function of temperature but does not vary in space so that ∇ doesn't act on it). The minimum of the free energy is then determined by setting $\delta F = 0$ which means that we have to solve the Euler-Lagrange equations for the function $m(\vec{r})$,

$$c\nabla^2 m = am + 2bm^3 \tag{5.46}$$

The simplest solutions to this equation have m constant, reducing us back to Landau theory. We'll assume once again that a(T) > 0 for $T > T_c$ and a(T) < 0 for $T < T_c$. Then the constant solutions are m = 0 for $T > T_c$ and $m = \pm m_0 = \pm \sqrt{-a/2b}$ for $T < T_c$. However, allowing for the possibility of spatial variation in the order parameter also opens up the possibility for us to search for more interesting solutions.

Domain Walls

Suppose that we have $T < T_c$ so there exist two degenerate ground states, $m = \pm m_0$. We could cook up a situation in which one half of space, say x < 0, lives in the ground state $m = -m_0$ while the other half of space, x > 0 lives in $m = +m_0$. This is exactly the situation that we already met in the liquid-gas transition and is depicted in Figure 38. It is also easy to cook up the analogous configuration in the Ising model. The two regions in which the spins point up or down are called domains. The place where these regions meet is called the *domain wall*.

We would like to understand the structure of the domain wall. How does the system interpolate between these two states? The transition can't happen instantaneously because that would result in the gradient term $(\nabla m)^2$ giving an infinite contribution to the free energy. But neither can the transition linger too much because any point at which $m(\vec{r})$ differs significantly from the value m_0 costs free energy from the m^2 and m^4 terms. There must be a happy medium between these two. To describe the system with two domains, $m(\vec{r})$ must vary but it need only change in one direction: m = m(x). Equation (5.46) then becomes an ordinary differential equation,

$$\frac{d^2m}{dx^2} = \frac{am}{c} + \frac{2bm^3}{c}$$

This equation is easily solved. We should remember that in order to have two vacua, $T < T_c$ which means that a < 0. We then have

$$m = m_0 \tanh\left(\sqrt{\frac{-a}{2c}}x\right)$$

where $m_0 = \sqrt{-a/2b}$ is the constant ground state solution for the spin. As $x \to \pm \infty$, the tanh function tends towards ± 1 which means that $m \to \pm m_0$. So this solution indeed interpolates between the two domains as required. We learn that the width of the domain wall is given by $\sqrt{-2c/a}$. Outside of this region, the magnetisation relaxes exponentially quickly back to the ground state values.

We can also compute the cost in free energy due to the presence of the domain wall. To do this, we substitute the solution back into the expression for the free energy (5.45). The cost is not proportional to the volume of the system, but instead proportional to the area of the domain wall. This means that if the system has linear size L then the free energy of the ground state scales as L^d while the free energy required by the wall scales only as L^{d-1} . It is simple to find the parametric dependence of this domain wall energy without doing any integrals; the energy per unit area scales as $\sqrt{-ca^3/b}$. Notice that as we approach the critical point, and $a \to 0$, the two vacua are closer, the width of the domain wall increases and its energy decreases.

5.5.1 Correlations

One of the most important applications of Landau-Ginzburg theory is to understand the correlations between fluctuations of the system at different points in space. Suppose that we know that the system has an unusually high fluctuation away from the average at some point in space, let's say the origin $\vec{r} = 0$. What is the effect of this on nearby points?

There is a simple way to answer this question that requires us only to solve the differential equation (5.46). However, there is also a more complicated way to derive the same result which has the advantage of stressing the underlying physics and the role played by fluctuations. Below we'll start by deriving the correlations in the simple manner. We'll then see how it can also be derived using more technical machinery.

We assume that the system sits in a given ground state, say $m = +m_0$, and imagine small perturbations around this. We write the magnetisation as

$$m(\vec{r}) = m_0 + \delta m(\vec{r}) \tag{5.47}$$

If we substitute this into equation (5.46) and keep only terms linear in δm , we find

$$c\nabla^2 \delta m + \frac{2a}{c} \delta m = 0$$

where we have substituted $m_0^2 = -a/2b$ to get this result. (Recall that a < 0 in the ordered phase). We now perturb the system. This can be modelled by putting a delta-function source at the origin, so that the above equation becomes

$$c\nabla^2 \delta m + \frac{2a}{c} \delta m = \frac{1}{2c} \,\delta^d(0)$$

where the strength of the delta function has been chosen merely to make the equation somewhat nicer. It is straightforward to solve the asymptotic behaviour of this equation. Indeed, it is the same kind of equation that we already solved when discussing the Debye-Hückel model of screening. Neglecting constant factors, it is

$$\delta m(\vec{r}) \sim \frac{e^{-r/\xi}}{r^{(d-1)/2}}$$
 (5.48)

This tells us how the perturbation decays as we move away from the origin. This equation has several names, reflecting the fact that it arises in many contexts. In liquids, it is usually called the *Ornstein-Zernicke* correlation. It also arises in particle physics as the Yukawa potential. The length scale ξ is called the *correlation length*

$$\xi = \sqrt{\frac{-c}{2a}} \tag{5.49}$$

The correlation length provides a measure of the distance it takes correlations to decay. Notice that as we approach a critical point, $a \rightarrow 0$ and the correlation length diverges. This provides yet another hint that we need more powerful tools to understand the physics at the critical point. We will now take the first baby step towards developing these tools.

5.5.2 Fluctuations

The main motivation to allow the order parameter to depend on space is to take into the account the effect of fluctuations. To see how we can do this, we first need to think a little more about the meaning of the quantity $F[m(\vec{r})]$ and what we can use it for. To understand this point, it's best if we go back to basics. We know that the true free energy of the system can be equated with the log of the partition function (1.36). We'd like to call the true free energy of the system F because that's the notation that we've been using throughout the course. But we've now called the Landau-Ginzburg functional $F[m(\vec{r})]$ and, while it's closely related to the true free energy, it's not quite the same thing as we shall shortly see. So to save some confusion, we're going to change notation at this late stage and call the true free energy A. Equation (1.36) then reads $A = -k_BT \log Z$, which we write this as

$$e^{-\beta A} = Z = \sum_{n} e^{-\beta E_n}$$

We would like to understand the right way to view the functional $F[m(\vec{r})]$ in this framework. Here we give a heuristic and fairly handwaving argument. A fuller treatment involves the ideas of the renormalisation group.

The idea is that each microstate $|n\rangle$ of the system can be associated to some specific function of the spatially varying order parameter $m(\vec{r})$. To illustrate this, we'll talk in the language of the Ising model although the discussion generalises to any system. For the Ising model, we could associate a magnetisation $m(\vec{r})$ to each lattice site by simply averaging over all the spins within some distance of that point. Clearly, this will only lead to functions that take values on lattice sites rather than in the continuum. But if the functions are suitably well behaved it should be possible to smooth them out into continuous functions $m(\vec{r})$ which are essentially constant on distance scales smaller than the lattice spacing. In this way, we get a map from the space of microstates to the magnetisation, $|n\rangle \mapsto m(\vec{r})$. But this map is not one-to-one. For example, if the averaging procedure is performed over enough sites, flipping the spin on just a single site is unlikely to have much effect on the average. In this way, many microstates map onto the same average magnetisation. Summing over just these microstates provides a first principles construction of the $F[m(\vec{r})]$,

$$e^{-\beta F[m(\vec{r})]} = \sum_{n|m(\vec{r})} e^{-\beta E_n}$$
 (5.50)

Of course, we didn't actually perform this procedure to get to (5.45): we simply wrote it down the most general form in the vicinity of a critical point with a bunch of unknown coefficients a(T), b(T) and c(T). But if we were up for a challenge, the above procedure tells us how we could go about figuring out those functions from first principles. More importantly, it also tells us what we should do with the Landau-Ginzburg free energy. Because in (5.50) we have only summed over those states that correspond to a particular value of $m(\vec{r})$. To compute the full partition function, we need to sum over all states. But we can do that by summing over all possible values of $m(\vec{r})$. In other words,

$$Z = \int Dm(\vec{r}) \ e^{-\beta F[m(\vec{r})]} \tag{5.51}$$

This is a tricky beast: it is a *functional integral*. We are integrating over all possible function $m(\vec{r})$, which is the same thing as performing an infinite number of integrations. (Actually, because the order parameters $m(\vec{r})$ arose from an underlying lattice and are suitably smooth on short distance scales, the problem is somewhat mitigated).

The result (5.51) is physically very nice, albeit mathematically somewhat daunting. It means that we should view the Landau-Ginzburg free energy as a new effective Hamiltonian for a continuous variable $m(\vec{r})$. It arises from performing the partition function sum over much of the microscopic information, but still leaves us with a final sum, or integral, over fluctuations in an averaged quantity, namely the order parameter.

To complete the problem, we need to perform the function integral (5.51). This is *hard*. Here "hard" means that the majority of unsolved problems in theoretical physics can be boiled down to performing integrals of this type. Yet the fact it's hard shouldn't dissuade us, since there is a wealth of rich and beautiful physics hiding in the path integral, including the deep reason behind the magic of universality. We will start to explore some of these ideas in next year's course on Statistical Field Theory.