Preprint typeset in JHEP style - HYPER VERSION

Michaelmas Term, 2017

Statistical Field Theory

University of Cambridge Part III Mathematical Tripos

David Tong

Department of Applied Mathematics and Theoretical Physics,
Centre for Mathematical Sciences,
Wilberforce Road,
Cambridge, CB3 OBA, UK

http://www.damtp.cam.ac.uk/user/tong/sft.html
d.tong@damtp.cam.ac.uk

Recommended Books and Resources

There are a large number of books which cover the material in these lectures, although often from very different perspectives. They have titles like "Critical Phenomena", "Phase Transitions", "Renormalisation Group" or, less helpfully, "Advanced Statistical Mechanics". Here are some that I particularly like

- Nigel Goldenfeld, *Phase Transitions and the Renormalization Group*

A great book, covering the basic material that we'll need and delving deeper in places.

- Mehran Kardar, *Statistical Physics of Fields*

The second of two volumes on statistical mechanics. It cuts a concise path through the subject, at the expense of being a little telegraphic in places. It is based on lecture notes which you can find on the web; a link is given on the course website.

- John Cardy, *Scaling and Renormalisation in Statistical Physics*

A beautiful little book from one of the masters of conformal field theory. It covers the material from a slightly different perspective than these lectures, with more focus on renormalisation in real space.

- Chaikin and Lubensky, *Principles of Condensed Matter Physics*

- Shankar, *Quantum Field Theory and Condensed Matter*

Both of these are more all-round condensed matter books, but with substantial sections on critical phenomena and the renormalisation group. Chaikin and Lubensky is more traditional, and packed full of content. Shankar covers modern methods of QFT, with an easygoing style suitable for bedtime reading.

A number of excellent lecture notes are available on the web. Links can be found on the course webpage: http://www.damtp.cam.ac.uk/user/tong/sft.html.

Contents

Acknowledgements

These lectures are aimed at beginning graduate students. They assume a background in statistical mechanics and thermodynamics. No quantum field theory is required, although the techniques that we develop here are useful in that context.

The course is built on the foundation of previous courses given in Cambridge by Ron Horgan and Matt Wingate. I'm grateful to the students for their questions and to Carl Turner for comments. I'm supported by the Royal Society and Alex Considine Tong.

Conventions

In a previous course on Statistical Physics, we diligently kept factors of the Boltzmann constant, k_B, in every equation. Now it is time to grow up. We set $k_B = 1$. This means that we measure temperature in units of energy.

0. Introduction

Superficially, this course is about *phase transitions*. This is the name given to the abrupt, discontinuous changes that occur when matter is deformed in some way, whether through heating or squeezing or something else. The familiar example is the violent shaking of a pot on a stove as water approaches its boiling point, and bubbles of steam erupt from within.

Despite their familiarity, phase transitions are striking, and even a little disconcerting. Usually in physics, things happens gradually. This fact is sewn into the heart of classical physics where the positions and momenta of particles are described by smooth, differentiable functions. Indeed, historically, the idea that change happens only infinitesimally resulted in the discovery of calculus. Yet, somehow, what holds on the micro level fails at the macro. Phase transitions tell us that a large number of particles can behave collectively in a way that any individual particle cannot, with the macroscopic properties of a system changing discontinuously.

A closer look at what happens at phase transitions – in particular at so-called critical points – reveals something startling. Many different substances, regardless of their microscopic composition, exhibit identical behaviour at a phase transition. This is not just a qualitative statement, but a quantitative one. For example, as a liquid changes into a gas at the critical temperature T_c, the heat capacity diverges as

$$c \sim \frac{1}{|T - T_c|^{0.11008\ldots}}$$

The exponent is not known precisely. It is thought not to be a rational number, but should instead be viewed as a universal mathematical constant, similar to π or e, but more subtle. Remarkably, the same exponent occurs for all gases. It also occurs in other systems, including a certain class of magnets. It's as if all knowledge of the microscopic physics has been washed away, leaving us with something pure, that carries only a vague memory of what lies underneath. This phenomenon is known as *universality*.

All of this makes phase transitions interesting. They involve violence, universal truths and competition between rival states. The story of phase transitions is, quite literally, the song of fire and ice.

And yet these are not the only reasons to study phase transitions. In our attempt to understand what happens as water boils, we will need to develop new tools and a new way of thinking about the world. This leads us to a paradigm which now underlies huge swathes of physics, far removed from its humble origin of a pot on a stove. This

paradigm revolves around two deep facts about the Universe we inhabit: Nature is organised by symmetry. And Nature is organised by scale.

Nature is Organised by Symmetry

When I was a kid, I was told that there are three phases of matter: solid, liquid and gas. (Actually, this isn't quite true. Knowing that I was interested in this kind of stuff, the teacher conspiratorially let on that there was a fourth phase of matter, "plasma". To this day, I have no idea why. My best guess is that this fitted better with some old view of the basic elements as earth, water, air and fire.)

It won't be any surprise to learn that the real world is much more interesting than the one we're introduced to as kids. There are not three phases of matter, nor four: there are many. A key insight, due to Landau, is that these different phases are characterised by symmetry.

In this scheme, a solid differs from a liquid because its crystal structure breaks the translational and rotational symmetries of space. Moreover, solids with different crystal structures should be viewed as different phases of matter because they break these symmetries in different ways. Perhaps more surprisingly, liquids and gases break no such symmetries and so should be viewed as the same phase. When you include further symmetries, such as rotations of spins in a magnet or more subtle quantum counterparts, this classification opens up a wide range of possibilities that allows us to understand almost all the known forms of matter.

This characterisation has its advantages. First, we can be sure that any attempt to change a material from one symmetry class to another will necessarily involve a violent phase transition. Second, it turns out that understanding the symmetries of a system will immediately determine many of its properties, especially at low temperature.

Moreover, the classification of matter in terms of symmetry has a power that goes far beyond its initial regime of application. The vacuum of space is, in many ways, like a complicated material, with quantum effects playing the role of thermal fluctuations. The vacuum can sit in different phases and is thought to have undergone several phase transitions as the Universe cooled after the big bang, each of which can be understood in terms of symmetries. All the ideas that we will develop here carry directly to theories of particle physics, cosmology and beyond.

Nature is Organised by Scale

There is an order to the Universe we live in. Roughly speaking, little things affect big things. Not the other way round.

This is something you already know: particle physics underlies nuclear and atomic physics; atomic physics underlies condensed matter and chemistry; and so on up the chain. It's certainly true that it can be difficult to make the leap from one level to the next, and new creative ideas are needed at each step, but this doesn't change the fact that there is an ordering. Big things don't affect little things. This is the reason there are no astrology departments in universities.

But there is another aspect to this story, one which is not often stressed. Little things affect big things, but they rarely affect very big things. Instead, little things affect slightly bigger things. And these, in turn, affect slightly bigger things too. But as you go up the chain, you lose the information about what came long before.

This again is something that you know. A zoologist who is interested in the way that starlings flock has little reason to study the dynamics of the Higgs boson. It's also the reason that science is possible in the first place: neither Newton nor Einstein needed to understand how quantum gravity works on microscopic distance scales to write down theories that work extraordinarily well on larger scales.

In the 1970s a mathematical formalism was developed that makes these ideas concrete. This formalism is called the *renormalisation group* and provides a framework to describe physics at different scales. The renormalisation group gets little coverage in popular science articles, yet is arguably the single most important advance in theoretical physics in the past 50 years. While zoologists may have little need to talk to particle physicists, the right way to understand both the Higgs boson and the flocking of starlings is through the language of the renormalisation group.

These two ideas – symmetry and scale – now dominate the way we think about physics. Yet both have their origins in the simple question: what happens when you boil water? The purpose of this course is to find out.

1. From Spins to Fields

The ideas that we will introduce in these lectures have wide applicability across many areas of physics. However, we will spend much of these lectures studying just a single example. This example is known as the *Ising model* and it provides a particularly simple model of a magnet. Despite its simplicity, the Ising model already contains many of the concepts that will keep us occupied for the rest of the course. Having understood the Ising model in some detail, we see how these ideas play out in many other phases of matter in Section 4.

1.1 The Ising Model

The Ising model is easy to state, yet hard to solve. We have a lattice in d spatial dimensions, containing N sites. On each lattice site $i = 1, \ldots, N$, there lives a discrete variable which can take one of two states: $s_i = +1$ or $s_i = -1$.

It is useful to refer to these states as *spins*, with $s_i = +1$ corresponding to spin up, and $s_i = -1$ to spin down. However, we won't be using any machinery from quantum mechanics: there are no non-commuting operators, nor quantum superpositions of states. In this sense, the Ising model, while discrete, is purely classical.

The collection of spins $\{s_i\}$ has energy

$$E = -B\sum_i s_i - J\sum_{\langle ij\rangle} s_i s_j \tag{1.1}$$

The first term arises due to an external magnetic field, B that we impose on the system. It has the effect that, for $B > 0$, the spins want to be up because that will lower their energy. (Properly this should be the magnetising field H, but we're using B to avoid confusion with the Hamiltonian. There is also a factor of the magnetic moment that we've absorbed into B.)

With the first term alone, the individual spins don't talk to each other and the model is easy to solve. It is the second term that makes life more interesting. This is an interaction between neighbouring spins. 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.

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 anti-align ($\uparrow\downarrow$). This is an *anti-ferromagnet*. In the following, we choose $J > 0$ although, for our purposes, the differences are minor.

(For what it's worth, the anti-ferromagnetic case, with $J < 0$, becomes more subtle on certain lattices where it's not possible to arrange the spins so that they are opposite to all their neighbours. A 2d triangular lattice provides a simple example. The resulting physics is interesting and is known as (geometrical) *frustration*. We will not discuss it here.)

We are interested in the physics of the Ising model at a finite temperature T. We can already get some intuition for what will happen. The interaction energy encourages the spins to align in the same way. The magnetic field encourages the spins to align in a particular way. Meanwhile, the temperature encourages the spins to ignore both the interactions and magnetic field because, at finite temperature, energy is no longer at a premium. Instead, entropy becomes more important. Since there are many more random configurations than aligned configurations, the temperature will tend to mess up the nice ordered states that the interactions and magnetic field have so carefully prepared. Already we can see that, like any good story, we're starting with a healthy dose of narrative tension between our main characters.

In the canonical ensemble, the probability of sitting in a configuration of spins $\{s_i\}$ is given by

$$p[s_i] = \frac{e^{-\beta E[s_i]}}{Z} \tag{1.2}$$

where $\beta = 1/T$ and the normalisation factor Z is called the *partition function*, and is given by

$$Z(T, J, B) = \sum_{\{s_i\}} e^{-\beta E[s_i]} \tag{1.3}$$

The beauty of statistical mechanics is that even though Z is first introduced as an innocuous normalisation factor, it actually contains everything we want to known about the system. If we're able to perform the sum to compute Z, we can extract any information we want[1]. For example, the interplay between energy and entropy S is captured by the thermodynamic free energy

$$F_{\text{thermo}}(T, B) = \langle E \rangle - TS = -T \log Z \tag{1.4}$$

As the notation shows, this is a function of thermodynamic variables, T and B. (We don't usually add the subscript "thermo" to the free energy but for much of these lectures we'll be working with a more refined version of the free energy which we'll shortly introduce.)

[1]The basic machinery of partition functions was described in the first course on Statistical Physics.

Of particular interest to us will be the average spin of the configuration, which we refer to as the equilibrium *magnetisation*.

$$m = \frac{1}{N}\langle \sum_i s_i \rangle \tag{1.5}$$

This takes values in the range $m \in [-1, +1]$. From our discussion above, we would expect that, for $B > 0$, $m \to +1$ at low temperatures where the spins are ordered, and $m \to 0$ at high temperatures where the spins are arranged randomly. We'd like to make this intuition more precise. Using the probability (1.2), it is straightforward to check that the magnetisation can be written as

$$m = \frac{1}{N}\sum_{\{s_i\}} \frac{e^{-\beta E[s_i]}}{Z} \sum_i s_i = \frac{1}{N\beta}\frac{\partial \log Z}{\partial B} \tag{1.6}$$

Taking further derivatives allows us to compute higher moments of the distribution, and so learn about fluctuations around the average. We will see this in action later in these lectures.

Our task is now clear: we should compute the sum (1.3). This is easier said than done. It turns out that the sum is straightforward in a $d = 1$ dimensional lattice, and you will do this on an example sheet. An exact solution also exists in $d = 2$ when $B = 0$, originally due to Onsager. It is not straightforward. In higher dimensions, no exact solutions are available, although various expansions and tricks have been invented to manipulate the sum (1.3) to extract some interesting information.

In these lectures, we will not attempt to directly sum the microscopic spins in the partition function. Instead, we will rephrase the question. We will massage the partition function into a somewhat different form, one that has much broader application.

1.1.1 The Effective Free Energy

The key idea was explained in the introduction: Nature is organised by scale. The energy of the Ising model (1.1) provides the microscopic description of our system in terms of individual spins. Computing the partition function exactly gives us a macroscopic description of the system in terms of thermodynamic variables like temperature T and magnetic field B. What we're looking for is something in between.

We're going to get to this "something in between" in two steps: we'll do something rather crude here, and then do a better job in Section 1.3. For our crude attempt, we rewrite the partition function (1.3) in the following manner:

$$Z = \sum_m \sum_{\{s_i\}|m} e^{-\beta E[s_i]} := \sum_m e^{-\beta F(m)} \tag{1.7}$$

where the notation $\{s_i\}|m$ means all configurations of spins such that $\frac{1}{N}\sum_i s_i = m$. In other words, we first sum over all configurations with fixed average magnetisation $m = \frac{1}{N}\sum s_i$, and subsequently sum over all possible m.

Note that we're using m here in a subtly different way to the original definition (1.5). In the sum (1.7), the magnetisation refers to the average magnetisation of a given configuration and can take any value. In contrast, in (1.5) we are talking about the equilibrium value of magnetisation, averaged over all configurations in the canonical ensemble. We will see shortly how to find this equilibrium value.

The average magnetisation lies in the range $-1 \leq m \leq 1$. Strictly speaking, this takes only discrete values, quantised in units of $2/N$. However, we are ultimately interested in the limit of large N, so we don't lose anything by writing this as an integral,

$$Z = \frac{N}{2}\int_{-1}^{+1} dm\; e^{-\beta F(m)} \tag{1.8}$$

where the factor of $N/2$ is the (inverse) width between the allowed m values. Such overall coefficients in the partition function are unimportant for the physics, and we will not be careful in keeping them below.

This way of writing things has allowed us to define something new: an effective free energy, $F(m)$, which depends on the magnetisation m of the system, in addition to both T and B. This goes beyond the usual, thermodynamic idea of free energy $F_{\rm thermo}$ given in (1.4) which is defined only in equilibrium, where the magnetisation m takes a specific value, determined by (1.6).

Note that (1.8) looks very much like a standard path integral, with $F(m)$ playing the role of the energy for m. But there is a difference because, unlike in the microscopic theory, $F(m)$ can depend on temperature. This means that the β dependence in the exponent can be more complicated than we're used to. We'll see this explicitly below.

The effective free energy $F(m)$ contains more information than the thermodynamic free energy $F_{\rm thermo}$. In particular, $F(m)$ can tell us the correct, equilibrium value of the magnetisation m. To see this, we define the free energy per unit spin,

$$f(m) = \frac{F(m)}{N}$$

Our partition function becomes

$$Z = \int dm\; e^{-\beta N f(m)}$$

Here N is a very large number (think $N \sim 10^{23}$) while $\beta f(m) \sim 1$. Integrals of this kind are very well approximated by the value of m which minimises $f(m)$, an approximation known as the *saddle point* or *steepest descent*,

$$\left.\frac{\partial f}{\partial m}\right|_{m=m_{\rm min}} = 0$$

The minimum $m_{\rm min}$ is the equilibrium value of the magnetisation that we previously computed in (1.6). Substituting this into the the partition function, we have

$$Z \approx e^{-\beta N f(m_{\rm min})} \quad \Rightarrow \quad F_{\rm thermo} \approx F(m_{\rm min}) \tag{1.9}$$

In this way, we can reconstruct $F_{\rm thermo}$ from knowledge of $F(m)$.

1.1.2 Mean Field Theory

To proceed, we want to find a way to compute $F(m) = Nf(m)$, defined in (1.7). But this is tantamount to performing the sum in the path integral and, as we stressed above, this is not easy. We need to find another way.

We will use a method called the *mean field approximation*. Here the word "approximation" is somewhat generous; a better name would be the mean field "guess" since there is little justification for what we're about to do. Instead, the purpose is to provide a starting point for our discussion. Not all the consequences that we derive from this guess will be accurate, but as the lectures progress we'll get a better understanding about what we can trust, what we can't, and how to do better.

We wish to sum over configurations $\{s_i\}$ with $\sum_i s_i = Nm$. We can get an estimate for the energy of such configurations by replacing each spin s in (1.1) with its expectation (i.e. *mean*) value $\langle s\rangle = m$,

$$E = -B\sum_i m - J\sum_{\langle ij\rangle} m^2 \quad \Rightarrow \quad \frac{E}{N} = -Bm - \frac{1}{2}Jqm^2 \tag{1.10}$$

Here q denotes the number of nearest neighbours of each spin. 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 $q = 2d$. The factor or $1/2$ is there because $\sum_{\langle ij\rangle}$ is a sum over pairs rather than a sum of individual sites.

Among other things, this means that we're assuming (falsely!) that the energy depends only on the magnetisation m of a configuration, rather than any more finely grained details. The sole reason for doing this is that it makes the resulting sum over

configurations $\{s_i\}|m$ very easy to do: we simply need to count the number of configurations with magnetisation m. A configuration with $N_\uparrow$ up spins and $N_\downarrow = N - N_\uparrow$ down spins has magnetisation

$$m = \frac{N_\uparrow - N_\downarrow}{N} = \frac{2N_\uparrow - N}{N}$$

The number of such configurations is

$$\Omega = \frac{N!}{N_\uparrow!(N - N_\uparrow)!}$$

and we can use Stirling's formula to evaluate

$$\begin{aligned} &\log\Omega \approx N\log N - N_\uparrow \log N_\uparrow - (N - N_\uparrow)\log(N - N_\uparrow) \\ \Rightarrow \quad &\frac{\log\Omega}{N} \approx \log 2 - \frac{1}{2}(1+m)\log(1+m) - \frac{1}{2}(1-m)\log(1-m) \end{aligned}$$

In our mean field approximation, the effective free energy defined in (1.7) is then given by

$$e^{-\beta N f(m)} \approx \Omega(m)\, e^{-\beta E(m)}$$

Substituting the energy (1.10), and taking logs of both sides, we find ourselves with the following expression:

$$f(m) \approx -Bm - \frac{1}{2}Jqm^2 - T\left(\log 2 - \frac{1}{2}(1+m)\log(1+m) - \frac{1}{2}(1-m)\log(1-m)\right)$$

From this, we can compute the equilibrium value for m. As explained above, this occurs at the minimum

$$\frac{\partial f}{\partial m} = 0 \quad \Rightarrow \quad \beta(B + Jqm) = \frac{1}{2}\log\left(\frac{1+m}{1-m}\right)$$

A little algebra shows us that the equilibrium magnetisation obeys the self-consistency condition

$$m = \tanh\left(\beta B + \beta Jqm\right) \tag{1.11}$$

There's a nice intuition behind this equation. It can be derived by assuming that each spin experiences an effective magnetic field given by $B_{\text{eff}} = B + Jqm$, which includes an extra contribution from the spins around it. In this way, the tricky interactions in the Ising model have been replaced by an averaged effective field B_{eff}. This is sometimes called the *mean field* and gives its name to this technique.

There are a number of ways forward at this point. We could, for example, analyse the properties of the Ising model by looking at solutions to the self-consistency condition (1.11); this was the strategy taken in the Statistical Physics lectures. However, instead we're going to focus on the free energy $f(m)$, since this will prove to be the more useful tool moving forward.

1.2 Landau Approach to Phase Transitions

A phase transition occurs when some quantity changes in a discontinuous fashion. For us, this quantity is m which, as we will see, will exhibit non-analytic behaviour as we vary different parameters.

Landau theory is a simple, yet effective way to understand the qualitative manner in which these phase transitions occur. It is based on two ideas: the free energy, and symmetry. Here we will apply Landau's theory to the Ising model, focussing first on the free energy. We will comment on the role of symmetry at the end of this section, although the full power of this idea will only become apparent in Section 4 where we describe a whole raft of further examples.

In the previous section, we introduced the effective free energy $F(m) = Nf(m)$, which is a function of the magnetisation m. This kind of object is the starting point for Landau's theory. In this context, the magnetisation is known as an *order parameter*. Using the mean field approximation, we computed this to be

$$f(m) \approx -Bm - \frac{1}{2}Jqm^2 - T\left(\log 2 - \frac{1}{2}(1+m)\log(1+m) - \frac{1}{2}(1-m)\log(1-m)\right)$$

Landau's approach focusses on situations where the order parameter is small. Here we can Taylor expand,

$$f(m) \approx -T\log 2 - Bm + \frac{1}{2}(T-Jq)m^2 + \frac{1}{12}Tm^4 + \ldots \tag{1.12}$$

As we've seen above, the preferred value of m is simply the minimum of $f(m)$. The idea of Landau theory is simply to observe how the function $f(m)$, and in particular its minima, change as we vary different parameters. Crucially, the story is rather different depending on whether $B \neq 0$ or $B = 0$. We will deal with each of these in turn.

1.2.1 $B = 0$: A Continuous Phase Transitions

Let's first consider the situation with vanishing magnetic field, $B = 0$, so that there is no term in (1.12) that is linear in m. Since our interest lies in the m dependence of

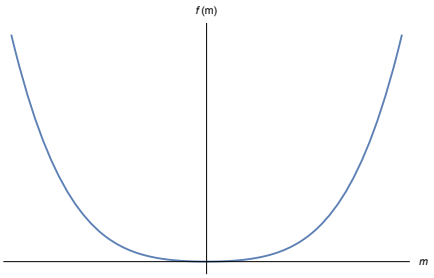

Figure 1: Free energy when $T > T_c$

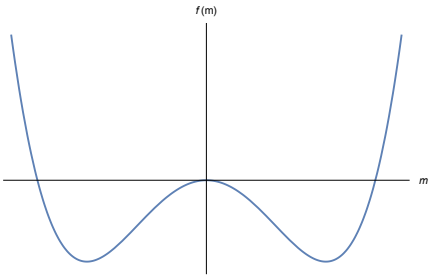

Figure 2: Free energy when $T < T_c$

$f(m)$, we won't lose anything if we drop the constant $-T\log 2$ term. We're left with the free energy

$$f(m) \approx \frac{1}{2}(T - Jq)m^2 + \frac{1}{12}Tm^4 + \dots \tag{1.13}$$

The behaviour of the free energy depends on whether the quadratic term is positive or negative. To distinguish between these two, we define the *critical temperature*

$$T_c = Jq$$

The free energy is sketched in the case where $T > T_c$ (on the left) and $T < T_c$ (on the right). We see that at high temperatures, the magnetisation vanishes at the minimum: $m = 0$. This is in agreement with our expectation that temperature will randomise the spins. However, as the temperature is reduced below T_c, the point $m = 0$ becomes a maximum of the free energy and the minima now lie at $m = \pm m_0$ which, if we chose to truncate the free energy (1.13) at order m^4, is given by

$$m_0 = \sqrt{\frac{3(T_c - T)}{T}} \tag{1.14}$$

This form is valid only when T is close to T_c, so that m is small and higher order terms can be ignored. As the temperature is lowered further, the minimum m_0 grows. We're then obliged to turn to the full form of the free energy $f(m)$ which, among other things, knows that the magnetisation lies in the range $m \in [-1, +1]$.

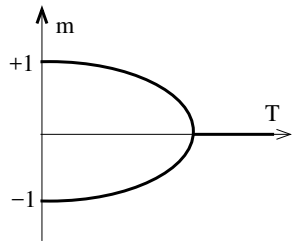

Figure 3:

The upshot is that as we vary the temperature, the magnetisation takes the form shown on the right. This is perhaps somewhat surprising. Usually in physics, things turn on gradually. But here the magnetisation turns off

abruptly at $T = T_c$, and remains zero for all higher temperatures. This kind of sharp change is characteristic of a phase transition. When $m = 0$, we say that the system sits in the *disordered phase*; when $m \neq 0$, it is in the *ordered* phase.

The magnetisation itself is continuous and, for this reason, it is referred to as a *continuous phase transition* or, sometimes, a *second order phase transition*.

As an aside: phase transitions can be classified by looking at the thermodynamic free energy $F_{\text{thermo}} = Nf(m_{\text{min}})$ and taking derivatives with respect to some thermodynamic variable, like T. If the discontinuity first shows up in the n^{th} derivative, it is said to be an n^{th} order phase transition. However, in practice we very rarely have to deal with anything other than first order transitions (which we will see below) and second order transitions. In the present case, we'll see shortly that the heat capacity is discontinuous, confirming that it is indeed a second order transition.

The essence of a phase transition is that some quantity is discontinuous. Yet, this should make us nervous. In principle, everything is determined by the partition function Z, defined in (1.3), which is a sum of smooth, analytic functions. How is it possible, then, to get the kind of non-analytic behaviour characteristic of a phase transition? The loophole is that Z is only necessarily analytic if the sum is finite. But there is no such guarantee that it remains analytic when the number of lattice sites $N \to \infty$. This means that genuine, discontinuous phase transitions only occur in infinite systems. In reality, we have around $N \approx 10^{23}$ atoms. This gives rise to functions which are, strictly speaking, smooth, but which change so rapidly that they act, to all intents and purposes, as if they were discontinuous.

These kind of subtleties are brushed under the carpet in the mean field approach that we're taking here. However, it's worth pointing out that the free energy, $f(m)$ is an analytic function which, when Taylor expanded, gives terms with integer powers of m. Nonetheless, the minima of f behave in a non-analytic fashion.

For future purposes, it will be useful to see how the heat capacity $C = \partial\langle E\rangle/\partial T$ changes as we approach T_c. In the canonical ensemble, the average energy is given by $\langle E\rangle = -\partial \log Z/\partial\beta$. From this, we find that we can write the heat capacity as

$$C = \beta^2 \frac{\partial^2}{\partial\beta^2} \log Z \tag{1.15}$$

To proceed, we need to compute the partition function Z, by evaluating $f(m)$ at the minimum value m_{min} as in (1.9). When $T > T_c$, this is simple: we have $m_{\text{min}} = 0$ and $f(m_{\text{min}}) = 0$ (still neglecting the constant $T \log 2$ term which doesn't contribute to the

heat capacity). In contrast, when $T < T_c$ the minimum lies at $m = m_0$ given in (1.14), and the free energy is $f(m_0) = -\frac{3}{4}(T_c - T)^2/T$.

Now we simply need to differentiate to get the heat capacity, C. The leading contribution as $T \to T_c$ comes from differentiating the $(T_c - T)$ piece, rather than the $1/T$ piece. We have

$$c = \frac{C}{N} \to \begin{cases} 0 & \text{as } T \to T_c^+ \\ 3/2 & \text{as } T \to T_c^- \end{cases} \tag{1.16}$$

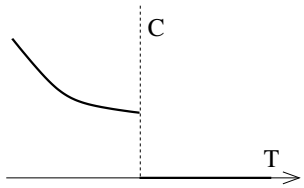

Figure 4:

We learn that the heat capacity jumps discontinuously. The part of the heat capacity that comes from differentiating $(T_c - T)$ terms is often called the *singular piece*. We'll be seeing more of this down the line.

Spontaneous Symmetry Breaking

There is one further aspect of the continuous phase transition that is worth highlighting. The free energy (1.13) is invariant under the $\mathbf{Z}_2$ symmetry $m \to -m$. This is no coincidence: it follows because our microscopic definition of the Ising model (1.1) also enjoys this symmetry when $B = 0$.

However, 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 $\mathbf{Z}_2$ symmetry. When a symmetry of a system is not respected by the ground state we say that the symmetry is *spontaneously broken*. This will become an important theme for us as we move through the course. It is also an idea which plays an important role in particle physics.

Strictly speaking, spontaneous symmetry breaking can only occur in infinite systems, with the magnetisation defined by taking the limit

$$m = \lim_{B \to 0} \lim_{N \to \infty} \frac{1}{N} \langle \sum_i s_i \rangle$$

It's important that we take the limit $N \to \infty$, before we take the limit $B \to 0$. If we do it the other way round, we find $\langle \sum s_i \rangle \to 0$ as $B \to 0$ for any finite N.

1.2.2 $B \neq 0$: First Order Phase Transitions

Let's now study what happens when $B \neq 0$. Once again, we ignore the constant $-T \log 2$ term in the free energy (1.12). We're left with the free energy

$$f(m) \approx -Bm + \frac{1}{2}(T - Jq)m^2 + \frac{1}{12}Tm^4 + \ldots \tag{1.17}$$

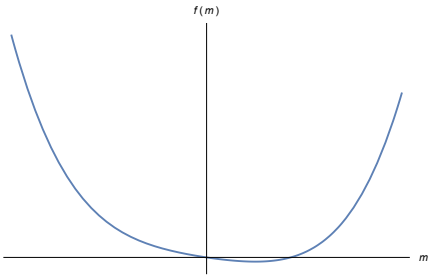

Figure 5: Free energy at high temperatures

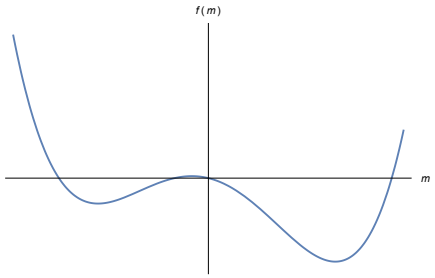

Figure 6: Free energy at low temperatures.

There are again two, qualitatively different forms of this function at low and high temperatures, shown in the figures above for $B > 0$.

At low temperatures there are two minima, but one is always lower than the other. The global minima is the true ground state of the system. The other minima is a *meta-stable state*. The system can exit the meta-stable state by fluctuating up, and over the energy barrier separating it from the ground state, and so has a finite lifetime. As we increase the temperature, there is a temperature (which depends on B) beyond which the meta-stable state disappears. This temperature is referred to as the *spinodal point*. It will not play any further role in these lectures.

For us, the most important issue is that the ground state of the system – the global minimum of $f(m)$ – does not qualitatively change as we vary the temperature. At high temperatures, the magnetisation asymptotes smoothly to zero as

$$m \to \frac{B}{T} \qquad \text{as } T \to \infty$$

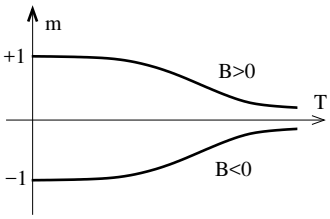

Figure 7:

At low temperatures, the magnetisation again asymptotes to the state $m \to \pm 1$ which minimises 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. A sketch of the magnetisation as a function of temperature is shown on the right. The upshot is that, for $B \neq 0$, there is no phase transition as a function of the temperature.

However, we do not have to look too hard to find a phase transition: we just need to move along a different path in the phase diagram. Suppose that we keep T fixed at a

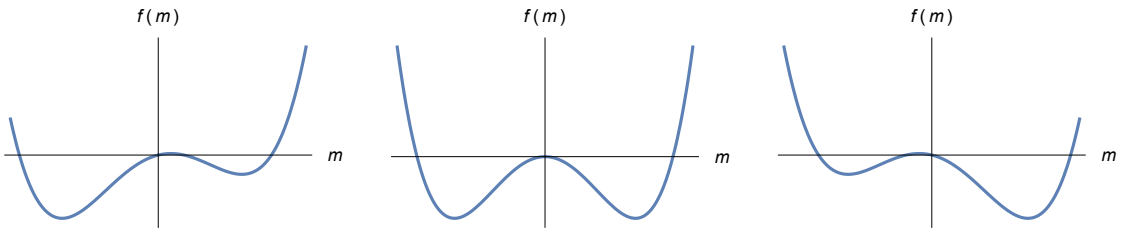

Figure 8: $T < T_c$: The free energy of the Ising model for $B < 0$ on the left, $B = 0$ in the middle, and $B > 0$ on the right.

value below T_c. We then vary the magnetic field from $B < 0$ to $B > 0$. The resulting free energy is shown in Figure 8.

We see that the magnetisation jumps discontinuously from $-m_0$ to $+m_0$ as B flips from negative to positive. This is an example of a *first order phase transition*.

Our analysis above has left us with the following picture of the phase diagram for the Ising model: if we vary B from positive to negative then we cross the red line in the figure and the system suffers a first order phase transition. Note, however, that if we first raise the temperature then it's always possible to move from any point $B \neq 0$ to any other point without suffering a phase transition.

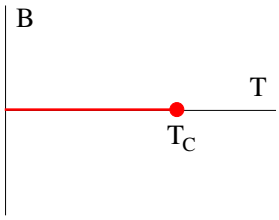

Figure 9:

This line of first order phase transitions ends at a second order phase transition at $T = T_c$. This is referred to as the *critical point*.

Close to the Critical Point

It will prove interesting to explore what happens when we sit close to the critical temperature $T = T_c$. There are a bunch of different questions that we can ask. Suppose, for example, that we sit at $T = T_c$ and vary the magnetic field: how does the magnetisation change? Here the free energy (1.17) becomes

$$f(m) \approx -Bm + \frac{1}{12}T_c m^4 + \ldots$$

For m small, where we can neglect the higher order terms, minimising the free energy gives $m^3 \sim B$. So, when $B > 0$ we have

$$m \sim B^{1/3} \tag{1.18}$$

and when $B < 0$ we have $m \sim -|B|^{1/3}$.

Here is another question: the *magnetic susceptibility* χ is defined as

$$\chi = \left.\frac{\partial m}{\partial B}\right|_T \tag{1.19}$$

We will compute this at $B = 0$, and as $T \to T_c$ from both above and below. First, from above: when $T \gtrsim T_c$ we can keep just the linear and quadratic terms in the free energy

$$f(m) \approx -Bm + \frac{1}{2}(T - T_c)m^2 \quad \Rightarrow \quad m \approx \frac{B}{T - T_c} \quad \Rightarrow \quad \chi \sim \frac{1}{T - T_c}$$

When $T < T_c$, we need to work a little harder. For small B, we can write the minimum of (1.17) as $m = m_0 + \delta m$ where m_0 is given by (1.14). Working to leading order in B/T, we find

$$m \approx m_0 + \frac{B}{2(T_c - T)} \quad \Rightarrow \quad \chi \sim \frac{1}{T_c - T}$$

We can combine these two results by writing

$$\chi \sim \frac{1}{|T - T_c|} \tag{1.20}$$

We'll see the relevance of this shortly.

1.2.3 Validity of Mean Field Theory

The first thing that we should ask ourselves is: are the results above right?! We have reason to be nervous because they were all derived using the mean field approximation, for which we offered no justification. On the other hand, there is reason to be optimistic because, at the end of the day, the structure of the phase diagram followed from some fairly basic properties of the Taylor expansion of the free energy.

In this section and the next, we will give some spoilers. What follows is a list of facts. In large part, the rest of these lectures will be devoted to explaining where these facts come from. It turns out that the validity of mean field theory depends strongly on the spatial dimension d of the theory. We will explain this in detail shortly but here is the take-home message:

- In $d = 1$ mean field theory fails completely. There are no phase transitions.
- In $d = 2$ and $d = 3$ the basic structure of the phase diagram is correct, but detailed predictions at $T \approx T_c$ are wrong.

- In $d \geq 4$, mean field theory gives the right answers.

This basic pattern holds for all other models that we will look at too. Mean field theory always fails completely for $d \leq d_l$, where d_l known as the *lower critical dimension*. For the Ising model, $d_l = 1$, but we will later meet examples where this approach fails in $d_l = 2$.

In contrast, mean field theory always works for $d \geq d_c$, where d_c is known as the *upper critical dimension*. For the Ising model, mean field theory works because, as d increases, each spin has a larger number of neighbours and so indeed experiences something close to the average spin.

Critical Exponents

What about the intermediate dimensions, $d_l < d < d_c$? These are very often the dimensions of interest: for the Ising model it is $d = 2$ and $d = 3$. Here the crude structure of the phase diagram predicted by mean field theory is correct, but it gives misleading results near the critical point $T = T_c$.

To explain this, recall that we computed the behaviour of four different quantities as we approach the critical point. For three of these, we fixed $B = 0$ and dialled the temperature towards the critical point. We found that, for $T < T_c$, the magnetisation (1.14) varies as

$$m \sim (T_c - T)^\beta \qquad \text{with } \beta = \frac{1}{2} \tag{1.21}$$

The heat capacity (1.16) varies as

$$c \sim c_\pm |T - T_c|^{-\alpha} \qquad \text{with } \alpha = 0 \tag{1.22}$$

where the $c_\pm$ is there to remind us that there is a discontinuity as we approach T_c from above or below. The magnetic susceptibility (1.20) varies as

$$\chi \sim \frac{1}{|T - T_c|^\gamma} \qquad \text{with } \gamma = 1$$

The fourth quantity requires us to take a different path in the phase diagram. This time we fix T and dial the magnetic field B towards zero, in which case the magnetisation (1.18) varies as

$$m \sim B^{1/\delta} \qquad \text{with } \delta = 3 \tag{1.23}$$

The coefficients α, β, γ and δ are known as *critical exponents*. (The Greek letters are standard notation; more generally, one can define a whole slew of these kind of objects.)

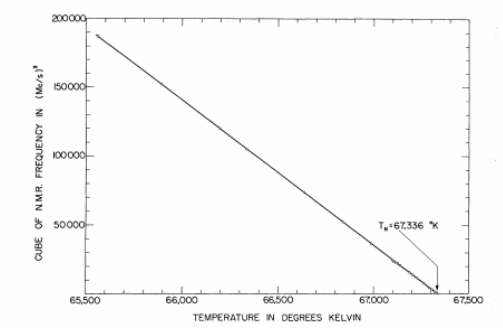

Figure 10: Magnetisation m^3, plotted against temperature for an Ising magnet, suggests that $1/\beta \approx 3$.

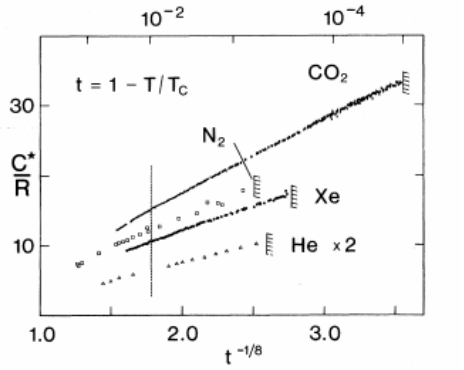

Figure 11: Heat capacity $c \sim |T - T_c|^{-\alpha}$ for various gases shows good agreement with $\alpha \approx 1/8$.

When the Ising model is treated correctly, one finds that these quantities do indeed scale to zero or infinity near T_c, but with different exponents. Here is a table of our mean field (MF) values, together with the true results for $d = 2$ and $d = 3$,

	MF	$d = 2$	$d = 3$
α	0 (disc.)	0 (log)	0.1101
β	$\frac{1}{2}$	$\frac{1}{8}$	0.3264
γ	1	$\frac{7}{4}$	1.2371
δ	3	15	4.7898

Note that the heat capacity has critical exponent $\alpha = 0$ in both mean field and in $d = 2$, but the discontinuity seen in mean field is replaced by a log divergence in $d = 2$.

The $d = 2$ results are known analytically, while the $d = 3$ results are known only numerically (to about 5 or 6 significant figures; I truncated early in the table above). Both the $d = 2$ and $d = 3$ results are also in fairly good agreement with experiment, which confirm that the observed exponents do not take their mean field values. For example, the left-hand figure above shows the magnetisation $m^{1/\beta} \sim (T_c - T)$ taken from MnF_2, a magnet with uniaxial anisotropy which is thought to be described by the Ising model[2]. The data shows a good fit to $1/\beta = 3$; as shown in the table, it is now thought that the exponent is not a rational number, but $1/\beta \approx 3.064$.

[2]This data is taken from P. Heller and G. Benedek, *Nuclear Magnetic Resonance in MnF_2 Near the Critical Point*, Phys. Rev. Lett. 8, 428 (1962).

This kind of behaviour is very surprising. It's rare that we see any kind of non-analytic behaviour in physics, but rarer still to find exponents that are not integers or simple fractions. What is going on here? This is one of the questions we will answer as these lectures progress.

1.2.4 A First Look at Universality

Before we go digging further into the Ising model, there is one other important aspect that deserves a mention at this point. The phase diagram for the Ising model is rather similar to the phase diagram for the transition between a liquid and gas, now drawn in the pressure-temperature plane. This is shown in the figure. In both cases, there is a line of first order transitions, ending at a critical point. For water, the critical point lies at a temperature $T_c \approx 374$ °C and a pressure $p_c \approx 218$ atm.

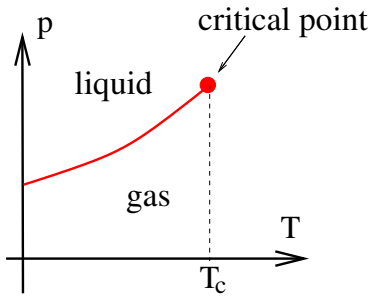

Figure 12:

The similarities are not just qualitative. One can use an appropriate equation of state for an interacting gas – say, the van der Waals equation – to compute how various quantities behave as we approach the critical point. (See the lectures on Statistical Physics for details of these calculations.) As we cross the first order phase transition, keeping the pressure constant, the volume V of the liquid/gas jumps discontinuously. This suggests that the rescaled volume $v = V/N$, where N is the number of atoms in the gas, is analogous to m in the Ising model. We can then ask how the jump in v changes as we approach the critical point. One finds,

$$v_{\rm gas} - v_{\rm liquid} \sim (T_c - T)^\beta \qquad \text{where } \beta = \frac{1}{2}$$

From the phase diagram, we see that the pressure p is analogous to the magnetic field B. We could then ask how the volume changes with pressure as we approach the critical point keeping $T = T_c$ fixed. We find

$$v_{\rm gas} - v_{\rm liquid} \sim (p - p_c)^{1/\delta} \qquad \text{where } \delta = 3$$

Finally, we want the analog of the magnetic susceptibility. For a gas, this is the compressibility, $\kappa = -\frac{1}{v}\frac{\partial v}{\partial p}\big|_T$. As we approach the critical point, we find

$$\kappa \sim \frac{1}{|T - T_c|^\gamma} \qquad \text{where } \gamma = 1$$

It has probably not escaped your attention that these critical exponents are exactly the same as we saw for the Ising model when treated in mean field. The same is also

true for the heat capacity C_V which approach different constant values as the critical point is approached from opposite sides.

However, the coincidence doesn't stop there. Because, it turns out, the critical exponents above are also wrong! The true critical exponents for the liquid-gas transitions in $d = 2$ and $d = 3$ dimensions are the same as those of the Ising model, listed previously in the table. For example, experimental data for the critical exponent α of a number of different gases was plotted two pages back[3] showing that it is approximately $\alpha \approx 0.1$, and certainly not consistent with mean field expectations.

This is an astonishing fact. It's telling us that at a second order phase transition, all memory of the underlying microscopic physics is washed away. Instead, there is a single theory which describes the physics at the critical point of the liquid-gas transition, the 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. Whatever drastic "spherical cow" approximation you make doesn't matter: if you capture the correct physics, you will get the exact answer! The fact that many different systems are described by the same critical point is called *universality*.

We might ask: does every second order phase transition have the same critical exponents as the Ising model? The answer is no! Instead, in each dimension d there is a set of critical points. Any system that undergoes a second order phase transition is governed by one member of this set. If two systems are governed by the same critical point, we say that they lie in the same *universality class*. The choice of universality class is, in large part, dictated by the symmetries of the problem; we will see some examples in Section 4.

The Ising Model as a Lattice Gas

It is, at first sight, surprising that a magnet and gas lie in the same universality class. However, there is a different interpretation of the Ising model that makes it look a little more gassy.

[3]The data is taken from J. Lipa, C. Edwards, and M. Buckingham *Precision Measurement of the Specific Heat of CO_2 Near the Critical Point*", Phys. Rev. Lett. 25, 1086 (1970). This was before the theory of critical phenomena was well understood. The data shows a good fit with $\alpha \approx 1/8$, not too far from the now accepted value $\alpha \approx 0.11$. Notice that the data stops around $t = 1 - T/T_c \approx 10^{-4}$. This is apparently because the effect of gravity becomes important as the critical point is approached, making experiments increasingly difficult.

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 (1.1) if we make the identification

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

The chemical potenial μ in the lattice gas plays the role of magnetic field in the spin system while the magnetisation of the system (1.6) measures the average density of particles away from half-filling.

There's no a priori reason to think that the Ising model is a particular good description of a gas. Nonetheless, this interpretation may make it a little less surprising that the Ising model and a gas share the same critical point.

1.3 Landau-Ginzburg Theory

The idea of universality – that many different systems enjoy the same critical point – is a powerful one. It means that if we want to accurately describe the critical point, we needn't concern ourselves with the messy details of any specific system. Instead, we should just search for the simplest model which gives the correct physics and work with that.

What is the simplest model? The Landau approach – in which the configuration of the Ising model is reduced to a single number m – is too coarse. This is because it misses any spatial variation in the system. And, as we will see shortly, the critical point is all about spatial variations. Here we describe a simple generalisation of Landau's ideas which allows the system to move away from a homogeneous state. This generalisation is known as *Landau-Ginzburg theory*[4].

[4]It is also known as *Ginzburg-Landau* theory. The original paper, dating from 1950, has the authors in alphabetical order and constructs the free energy of a superconductor. Here we use the term Landau-Ginzburg theory to reflect the fact that, following Landau's earlier work, these ideas apply more broadly to any system.

The key idea is simple: we take the order parameter from Landau theory – which, for the Ising model, is the magnetisation m – and promote it to a field which can vary in space, $m(\mathbf{x})$. This is referred to as a *local order parameter*.

How do we do this? We started with a lattice. We divide this lattice up many into boxes, with sides of length a. Each of these boxes contains many lattice sites – call it N' – but the box size a is smaller than any other length scale in the game. (In particular, a should be smaller than something called the *correlation length* that we will encounter below.) For each box, we define the average magnetisation $m(\mathbf{x}) = \frac{1}{N'}\sum_i s_i$, where $\mathbf{x}$ denotes the centre of the box. This kind of procedure is known as *coarse-graining*.

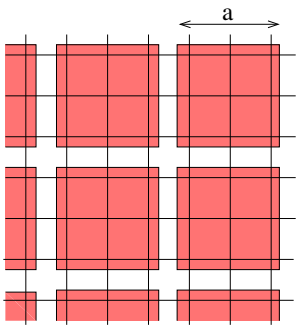

Figure 13:

There are all sorts of subtleties involved in this construction, and we will do our best to sweep them firmly under the rug. First, $m(\mathbf{x})$ takes values only at certain values of $\mathbf{x}$ which label the positions of boxes. To mitigate this, we take the number of boxes N/N' to be big enough so that $\mathbf{x}$ can be treated as a continuous variable. Second, at each $\mathbf{x}$, $m(\mathbf{x})$ is quantised in units of $1/N'$. We will take N' big enough so that $m(\mathbf{x})$ can take any value in $[-1, +1]$. The upshot of this is that we will treat $m(\mathbf{x})$ as a continuous function. However, we will at some point need to remember that it makes no sense for $m(\mathbf{x})$ to vary on very short distance scales – shorter than the separation between boxes.

You may have noticed that we've not been particularly careful about defining $m(\mathbf{x})$ and this cavalier attitude will continue below. You might reasonably ask: does the physics depend on the details of how we coarse grain? The answer is no. This is the beauty of universality. We'll see this more clearly as we proceed.

Our next step is to repeat the ideas that we saw in Section 1.1.1. Following (1.7), we write the partition function as

$$Z = \sum_{m(\mathbf{x})} \sum_{\{s_i\}|m(\mathbf{x})} e^{-\beta E[s_i]} := \sum_{m(\mathbf{x})} e^{-\beta F[m(\mathbf{x})]} \tag{1.24}$$

Here, the notation $\{s_i\}|m(\mathbf{x})$ means that we sum over all configurations of spins such that the coarse graining yields $m(\mathbf{x})$. In general, there will be many spin configurations for each $m(\mathbf{x})$. We then sum over all possible values of $m(\mathbf{x})$.

This procedure has allowed us to define a free energy $F[m(\mathbf{x})]$. This is a *functional*, meaning that you give it a function $m(\mathbf{x})$ and it spits back a number, $F[m(\mathbf{x})]$. This is known as the *Landau-Ginzburg free energy*.

We will invoke one last notational flourish. We're left in (1.24) with a sum over all possible configurations $m(\mathbf{x})$. With the assumption that $m(\mathbf{x})$ is a continuous function, this is usually written as

$$Z = \int \mathcal{D}m(\mathbf{x})\ e^{-\beta F[m(\mathbf{x})]} \tag{1.25}$$

This is a *functional integral*, also known as a *path integral*. The notation $\mathcal{D}m(\mathbf{x})$ – which is usually shortened to simply $\mathcal{D}m$ – means that we should sum over all field configurations $m(\mathbf{x})$.

Path integrals may be somewhat daunting at first sight. But it's worth remembering where it comes from: an integration over $m(\mathbf{x})$ at each point $\mathbf{x}$ labelling a box. In other words, it's nothing more than an infinite number of normal integrals. We will start to learn how to play with these objects in Section 2.

The path integral looks very much like a usual partition function, but with the Landau-Ginzburg free energy $F[m(\mathbf{x})]$ playing the role of an effective Hamiltonian for the continuous variable $m(\mathbf{x})$. There is some nice intuition behind this. In the thermal ensemble, a given field configuration $m(\mathbf{x})$ arises with probability

$$p[m(\mathbf{x})] = \frac{e^{-\beta F[m(\mathbf{x})]}}{Z} \tag{1.26}$$

The path integral (1.25), is nothing but the usual partition function, now for a field $m(\mathbf{x})$ rather than the more familiar variables of classical physics like position and momentum. In other words, we're doing the statistical mechanics of fields, rather than particles. The clue was in the title of these lectures.

1.3.1 The Landau-Ginzburg Free Energy

The next step is to ask: how do we calculate $F[m(\mathbf{x})]$? This seems tricky: already in our earlier discussion of Landau theory, we had to resort to an unjustified mean field approximation. What are we going to do here?

The answer to this question is wonderfully simple, as becomes clear if we express the question in a slightly different way: what could the free energy $F[m(\mathbf{x})]$ possibly be? There are a number of constraints on $F[m(\mathbf{x})]$ that arise from its microscopic origin:

- Locality: The nearest neighbour interactions of the Ising model mean that a spin on one site does not directly affect a spin on a far flung site. It only does so

through the intermediate spins. The same should be true of the magnetisation field $m(\mathbf{x})$. The result is that the free energy should take the form

$$F[m(\mathbf{x})] = \int d^d x \; f[m(\mathbf{x})]$$

where $f[m(\mathbf{x})]$ is a local function. It can depend on $m(\mathbf{x})$, but also on $\nabla m(\mathbf{x})$ and higher derivatives. These gradient terms control how the field $m(\mathbf{x})$ at one point affects the field at neighbouring points.

- Translation and Rotation Invariance: The original lattice has a discrete translation symmetry. For certain lattices (e.g. a square lattice) there can be discrete rotation symmetries. At distances much larger than the lattice scale, we expect that the continuum version of both these symmetries emerges, and our free energy should be invariant under them.

- $\mathbf{Z}_2$ symmetry. When $B = 0$, the original Ising model (1.1) is invariant under the symmetry $s_i \to -s_i$, acting simultaneously on all sites. This symmetry is inherited in our coarse-grained description which should be invariant under

$$m(\mathbf{x}) \to -m(\mathbf{x}) \tag{1.27}$$

  When $B \neq 0$, the Ising model is invariant under $m(\mathbf{x}) \to -m(\mathbf{x})$, together with $B \to -B$. Again, our free energy should inherit this symmetry.

- Analyticity: We will make one last assumption: that the free energy density is an analytic function of $m(\mathbf{x})$ and its derivatives. Our primary interest lies in the critical point where m first becomes non-zero, although we will also use this formalism to describe first order phase transitions where $m(\mathbf{x})$ is small. In both cases, we can Taylor expand the free energy in m and restrict attention to low powers of m.

  Furthermore, we will restrict attention to situations where $m(\mathbf{x})$ varies rather slowly in space. In particular, we will assume that $m(\mathbf{x})$ varies appreciably only over distances that are much larger than the distance a between boxes. This means that we can also consider a gradient expansion of $f[m(\mathbf{x})]$, in the dimensionless combination $a\nabla$. This means that ∇m terms are more important than $a\nabla^2 m$ terms and so on.

With these considerations in place, we can now simply write down the general form of the free energy. When $B = 0$, the symmetry (1.27) means that the free energy can

depend only on even powers of m. The first few terms in the expansion are

$$F[m(\mathbf{x})] = \int d^d x \ \left[\frac{1}{2}\alpha_2(T)m^2 + \frac{1}{4}\alpha_4(T)m^4 + \frac{1}{2}\gamma(T)(\nabla m)^2 + \ldots\right] \tag{1.28}$$

There can also be a $F_0(T)$ piece – like the $T\log 2$ that appeared in the Landau free energy – which doesn't depend on the order parameter and so, for the most part, will play 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.

When $B \neq 0$, we can have further terms in the free energy that are odd in m, but also odd in B, such as Bm and Bm^3. Each of these comes with a coefficient which is, in general, a function of T.

The arguments that led us to (1.28) are very general and powerful; we will see many similar arguments in Section 4. The downside is that we are left with a bunch of unknown coefficients $\alpha_2(T)$, $\alpha_4(T)$ and $\gamma(T)$. These are typically hard to compute from first principles. One way forward is to import our results from Landau mean field approach. Indeed, for constant $m(\mathbf{x}) = m$, the free energy (1.28) coincides with our earlier result (1.13) in Landau theory, with

$$\alpha_2(T) \sim (T - T_c) \quad \text{and} \quad \alpha_4(T) \sim \frac{1}{3}T \tag{1.29}$$

Happily, however, the exact form of these functions will not be important. All we will need is that these functions are analytic in T, and that $\alpha_4(T) > 0$ and $\gamma(T) > 0$, while $\alpha_2(T)$ flips sign at the second order phase transition.

Looking ahead, there is both good news and bad. The good news is that the path integral (1.25), with Landau-Ginzburg free energy (1.28), does give a correct description of the critical point, with the surprising d-dependent critical exponents described in Section 1.2.3. The bad news is that this path integral is hard to do! Here "hard" means that many of the unsolved problems in theoretical physics can be phrased in terms of these kinds of path integrals. Do not fear. We will tread lightly.

1.3.2 The Saddle Point and Domain Walls

We are going to build up slowly to understand how we can perform the functional integral (1.25). As a first guess, we'll resort to the saddle point method and assume that the path integral is dominated by the configurations which minimise $F[m(\mathbf{x})]$. In subsequent sections, we'll treat the integral more seriously and do a better job.

To find the minima of functionals like $F[m(\mathbf{x})]$, we use the same kind of variational methods that we met when working with Lagrangians in the lectures on Classical Dynamics. We take some fixed configuration $m(\mathbf{x})$ and consider a nearby configuration $m(\mathbf{x}) + \delta m(\mathbf{x})$. The change in the free energy is then

$$\begin{aligned}\delta F &= \int d^d x \ \left[\alpha_2 m\, \delta m + \alpha_4 m^3\, \delta m + \gamma \nabla m \cdot \nabla \delta m\right] \\ &= \int d^d x \ \left[\alpha_2 m + \alpha_4 m^3 - \gamma \nabla^2 m\right] \delta m\end{aligned}$$

where, to go from the first line to the second, we have integrated by parts. This encourages us to introduce the *functional derivative*,

$$\frac{\delta F}{\delta m(\mathbf{x})} = \alpha_2 m(\mathbf{x}) + \alpha_4 m^3(\mathbf{x}) - \gamma \nabla^2 m(\mathbf{x})$$

Note that I've put back the $\mathbf{x}$ dependence to stress that, in contrast to F, $\delta F/\delta m(\mathbf{x})$ is evaluated at some specific position $\mathbf{x}$.

If the original field configuration $m(\mathbf{x})$ was a minimum of the free energy it satisfies the Euler-Lagrange equations,

$$\left.\frac{\delta F}{\delta m}\right|_{m(\mathbf{x})} = 0 \quad \Rightarrow \quad \gamma \nabla^2 m = \alpha_2 m + \alpha_4 m^3 \tag{1.30}$$

The simplest solutions to this equation have m constant. This recovers our earlier results from Landau theory. When $T > T_c$, we have $\alpha_2 > 0$ and the ground state has $m = 0$. In contrast, when $T < T_c$, $\alpha_2 < 0$ and there is a degenerate ground state $m = \pm m_0$ with

$$m_0 = \sqrt{-\frac{\alpha_2}{\alpha_4}} \tag{1.31}$$

This is the same as our previous expression (1.14), where we replaced α_2 and α_4 with the specific functions (1.29). We see that what we previously called the mean field approximation, is simply the saddle point approximation in Landau-Ginzburg theory. For this reason, the term "mean field theory" is given to any situation where we write down the free energy, typically focussing on a Taylor expansion around $m = 0$, and then work with the resulting Euler-Lagrange equations.

Domain Walls

The Landau-Ginzburg theory contains more information than our earlier Landau approach. It also tells us how the magnetisation changes in space.

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 $m = +m_0$. 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 $m(\mathbf{x})$ 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, $m(\mathbf{x})$ must vary but it need only change in one direction: $m = m(x)$. Equation (1.30) then becomes an ordinary differential equation,

$$\gamma \frac{d^2 m}{dx^2} = \alpha_2 m + \alpha_4 m^3$$

This equation is easily solved. If we insist that the field interpolate between the two different ground states, $m \to \mp m_0$ as $x \to \mp\infty$, then the solution is given by

$$m = m_0 \tanh\left(\frac{x - X}{W}\right) \tag{1.32}$$

This is plotted in the figure. Here X is the position of the domain wall and

$$W = \sqrt{-\frac{2\gamma}{\alpha_2}}$$

is its width. For $|x - X| > W$, the magnetisation relaxes exponentially quickly back to the ground state values $\pm m_0$.

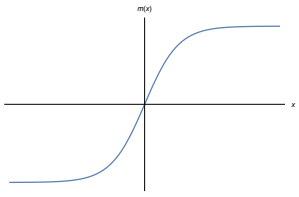

Figure 14:

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 (1.28). 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 additional free energy required by the wall scales only as L^{d-1}. It is simple to see that the excess cost in free energy of the domain wall has parametric dependence

$$F_{\text{wall}} \sim L^{d-1} \sqrt{-\frac{\gamma \alpha_2^3}{\alpha_4^2}} \tag{1.33}$$

Here's the quick way to see this: from the solution (1.32), the field m changes by $\Delta m \sim m_0$ over a distance $\Delta x \sim W$. This means that $dm/dx \sim \Delta m/\Delta x \sim m_0/W$ and the gradient term in the free energy contributes $F \sim L^{d-1}\int dx\ \gamma(dm/dx)^2 \sim L^{d-1}W\gamma(m_0/W)^2$, where the second expression comes because the support of the integral is only non-zero over the width W where the domain wall is changing. This gives the parametric dependence (1.33). There are further contributions from the m^2 and m^4 terms in the potential, but our domain wall solves the equation of motion whose purpose is to balance the gradient and potential terms. This means that the potential terms contribute with the same parametric dependence as the gradient terms, a fact you can check by hand by plugging in the solution.

Notice that as we approach the critical point, and $\alpha_2 \to 0$, the two vacua are closer, the width of the domain wall increases and its energy decreases.

1.3.3 The Lower Critical Dimension

We stated above that the Ising model has no phase transition in dimension $d = 1$. The reason behind this can be traced to the existence of domain walls, which destroy any attempt to sit in the ordered phase.

Let's set $\alpha_2(T) < 0$ where we would expect two, ordered ground states $m = \pm m_0$. To set up the problem, we start by considering the system on a finite interval of length L. At one end of this interval – say the left-hand edge $x = -L/2$ – we'll fix the magnetisation to be in its ground state $m = +m_0$. One might think that the preferred state of the system is then to remain in $m = +m_0$ everywhere. But what actually happens?

There is always a probability that a domain wall will appear in the thermal ensemble and push us over to the other ground state $m = -m_0$. The probability for a wall to appear at some point $x = X$ is given by (1.26)

$$p[\text{wall at } x = X] = \frac{e^{-\beta F_{\text{wall}}}}{Z}$$

This looks like it's exponentially suppressed compared to the probability of staying put. However, we get an enhancement because the domain wall can sit anywhere on the line $-L/2 \leq x \leq L/2$. This means that the probability becomes

$$p[\text{wall anywhere}] = \frac{e^{-\beta F_{\text{wall}}}}{Z}\frac{L}{W}$$

For a large enough system, the factor of L/W will overwhelm any exponential suppression. This is an example of the entropy of a configuration – which, in this context is $\log(L/W)$ – outweighing the energetic cost.

Of course, once we've got one domain wall there's nothing to stop us having another, flipping us back to $m = +m_0$. If we have an even number of domain walls along the line, we will be back at $m = m_0$ by the time we get to the right-hand edge at $x = +L/2$; an odd number and we'll sit in the other vacuum $m = -m_0$. Which of these is most likely?

The probability to have n walls, placed anywhere, is

$$p[n \text{ walls}] = \frac{e^{-n\beta F_{\rm wall}}}{Z}\frac{1}{W^n}\int_{-L/2}^{L/2} dx_1 \int_{x_1}^{L/2} dx_2 \ldots \int_{x_{n-1}}^{L/2} dx_n = \frac{1}{Zn!}\left(\frac{L\, e^{-\beta F_{\rm wall}}}{W}\right)^n$$

This means that the probability that we start at $m = m_0$ and end up at $m = m_0$ is

$$p[m_0 \to m_0] = \frac{1}{Z}\sum_{n \text{ even}} \frac{1}{n!}\left(\frac{L\, e^{-\beta F_{\rm wall}}}{W}\right)^n = \frac{1}{Z}\cosh\left(\frac{L\, e^{-\beta F_{\rm wall}}}{W}\right)$$

Meanwhile, the probability that the spin flips is

$$p[m_0 \to -m_0] = \frac{1}{Z}\sum_{n \text{ odd}} \frac{1}{n!}\left(\frac{L\, e^{-\beta F_{\rm wall}}}{W}\right)^n = \frac{1}{Z}\sinh\left(\frac{L\, e^{-\beta F_{\rm wall}}}{W}\right)$$

We see that, at finite L, there is some residual memory of the boundary condition that we imposed on the left-hand edge. However, as $L \to \infty$, this gets washed away. You're just as likely to find the spins up as down on the right-hand edge.

Although we phrased this calculation in terms of pinning a choice of ground state on a boundary, the same basic physics holds on an infinite line. Indeed, this is a general principle: whenever we have a Landau-Ginzburg theory characterised by a discrete symmetry – like the $\mathbf{Z}_2$ of the Ising model – then the ordered phase will have a number of degenerate, disconnected ground states which spontaneously break the symmetry. In all such cases, the lower critical dimension is $d_l = 1$ and in all cases the underlying reason is the same: fluctuations of domain walls will take us from one ground state to another and destroy the ordered phase. It is said that the domain walls *proliferate*.

We could try to run the arguments above in dimensions $d \geq 2$. The first obvious place that it fails is that the free energy cost of the domain wall (1.33) now scales with the system size, L. This means that as L increases, we pay an exponentially large cost from $e^{-F_{\rm wall}}$. Nonetheless, one could envisage a curved domain wall, which encloses some finite region of the other phase. It turns out that this is not sufficient to disorder the system. However, in $d = 2$, the fluctuations of the domain wall become important as one approaches the critical point.

1.3.4 Lev Landau: 1908-1968

Lev Landau was one of the great physicists of the 20th century. He made important contributions to nearly all areas of physics, including the theory of magnetism, phase transitions, superfluids and superconductors, Fermi liquids and quantum field theory. He founded a school of Soviet physics whose impact lasts to this day.

Landau was, by all accounts, boorish and did not suffer fools gladly. This did not sit well with the authorities and, in 1938, he was arrested in one of the great Soviet purges and sentenced to 10 years in prison. He was rescued by Kapitza who wrote a personal letter to Stalin arguing, correctly as it turned out, that Landau should be released as he was the only person who could understand superfluid helium. Thus, for his work on superfluids, Landau was awarded both his freedom and, later, the Nobel prize. His friend, the physicist Yuri Rumer, was not so lucky, spending 10 years in prison before exile to Siberia.

Legend has it that Landau hated to write. The extraordinarily ambitious, multi-volume *Course of Theoretical Physics* was entirely written by his co-author Evgeny Lifshitz, usually dictated by Landau. As Landau himself put it: "Evgeny is a great writer: he cannot write what he does not understand".

Figure 15:

Here is a story about a visit by Niels Bohr to Landau's Moscow Institute. Bohr was asked by a journalist how he succeeded in creating such a vibrant atmosphere in Copenhagen. He replied "I guess the thing is, I've never been embarrassed to admit to my students that I'm a fool". This was translated by Lifshitz as: "I guess the thing is, I've never been embarrassed to admit to my students that they're fools". According to Kapista, the translation was no mistake: it simply reflected the difference between Landau's school and Bohr's.

In 1962, Landau suffered a car accident. He survived but never did physics again.

2. My First Path Integral

It's now time to understand a little better how to deal with the path integral

$$Z = \int \mathcal{D}m(\mathbf{x})\ e^{-\beta F[m(\mathbf{x})]} \tag{2.1}$$

Our strategy – at least for now – will be to work in situations where the saddle point dominates, with the integral giving small corrections to this result. In this regime, we can think of the integral as describing the thermal fluctuations of the order parameter $m(\mathbf{x})$ around the equilibrium configuration determined by the saddle point. As we will see, this approach fails to work at the critical point, which is the regime we're most interested in. We will then have to search for other techniques, which we will describe in Section 3.

Preparing the Scene

Before we get going, we're going to change notation. First, we will change the name of our main character and write the magnetisation as

$$m(\mathbf{x}) \ \rightarrow \ \phi(\mathbf{x})$$

If you want a reason for this, I could tell you that the change of name is in deference to universality and the fact that the field could describe many things, not just magnetisation. But the real reason is simply that fields in path integrals should have names like ϕ. (This is especially true in quantum field theory where m is reserved for the mass of the particle.)

We start by setting $B = 0$; we'll turn B back on in Section 2.2. The free energy is then

$$F[\phi(\mathbf{x})] = \int d^d x\ \left[\frac{1}{2}\alpha_2(T)\phi^2 + \frac{1}{4}\alpha_4(T)\phi^4 + \frac{1}{2}\gamma(T)(\nabla\phi)^2 + \ldots\right]$$

Roughly speaking, path integrals are trivial to do if $F[\phi(\mathbf{x})]$ is quadratic in ϕ, and possible to do if the higher order terms in $F[\phi(\mathbf{x})]$ give small corrections. If the higher order terms in $F[\phi(\mathbf{x})]$ are important, then the path integral is typically impossible without the use of numerics. Here we'll start with the trivial, building up to the "possible" in later chapters.

To this end, throughout the rest of this chapter we will restrict attention to a free energy that contains no terms higher than quadratic in $\phi(\mathbf{x})$. We have to be a little careful about whether we sit above or below the critical temperature. When $T > T_c$, things are easy and we simply throw away all higher terms and work with

$$F[\phi(\mathbf{x})] = \frac{1}{2}\int d^d x \ \left(\gamma \nabla\phi \cdot \nabla\phi + \mu^2\phi^2\right) \tag{2.2}$$

where $\mu^2 = \alpha_2(T)$ is positive.

A word of caution. We are ignoring the quartic terms purely on grounds of expediency: this makes our life simple. However, these terms become increasingly important as $\mu^2 \sim \alpha_2(T) \to 0$ and we approach the critical point. This means that nothing we are about to do can be trusted near the the critical point. Nonetheless, we will be able to extract some useful intuition from what follows. We will then be well armed to study critical phenomena in Section 3.

When $T < T_c$, and $\alpha_2(T) < 0$, we need to do a little more work. Now the saddle point does not lie at $\phi = 0$, but rather at $\langle\phi\rangle = \pm m_0$ given in (1.31). In particular, it's crucial that we keep the quartic term because this rescues the field from the upturned potential it feels at the origin.

However, it's straightforward to take this into account. We simply compute the path integral about the appropriate minimum by writing

$$\tilde{\phi}(\mathbf{x}) = \phi(\mathbf{x}) - \langle\phi\rangle \tag{2.3}$$

Substituting this into the free energy gives

$$F[\tilde{\phi}(\mathbf{x})] = F[m_0] + \frac{1}{2}\int d^d x \left[\alpha_2'(T)\tilde{\phi}^2 + \gamma(T)(\nabla\tilde{\phi})^2 + \ldots\right] \tag{2.4}$$

where now the $\ldots$ include terms of order $\tilde{\phi}^3$ and $\tilde{\phi}^4$ and higher, all of which we've truncated. Importantly, there are no terms linear in $\tilde{\phi}$. In fact, this was guaranteed to happen: the vanishing of the linear terms is equivalent to the requirement that the equation of motion (1.30) is obeyed. The new quadratic coefficient is

$$\alpha_2'(T) = \alpha_2(T) + 3m_0^2\alpha_4(T) = -2\alpha_2(T) \tag{2.5}$$

In particular, $\alpha_2'(T) > 0$.

Practically, this means that we take the calculations that we do at $T > T_c$ with the free energy (2.2) and trivially translate them into calculations at $T < T_c$. We just need to remember that we're working with a shifted ϕ field, and that we should take $\mu^2 = \alpha_2'(T) = |2\alpha_2(T)|$. Once again, the same caveats hold: our calculations should not be trusted near the critical point $\mu^2 = 0$.

2.1 The Thermodynamic Free Energy Revisited

For our first application of the path integral, we will compute something straightforward and a little bit boring: the thermodynamic free energy. There will be a little insight to be had from the result of this, although the main purpose of going through these steps is to prepare us for what's to come.

We've already found some contributions to the thermodynamic free energy. There is the constant term $F_0(T)$ and, if we're working at $T < T_c$, the additional contribution $F[m_0]$ in (2.4). Here we are interested in further contributions to $F_{\rm thermo}$, coming from fluctuations of the field. To keep the formulae simple, we will ignore these two earlier contributions; you can always put them back in if you please.

Throughout this calculation, we'll set $B = 0$ so we're working with the free energy (2.2). There is a simple trick to compute the partition function when the free energy is quadratic: we work in Fourier space. We write the Fourier transform of the magnetisation field as

$$\phi_{\mathbf{k}} = \int d^d x \; e^{-i\mathbf{k}\cdot\mathbf{x}} \, \phi(\mathbf{x})$$

Since our original field $\phi(\mathbf{x})$ is real, the Fourier modes obeys $\phi_{\mathbf{k}}^\star = \phi_{-\mathbf{k}}$.

The $\mathbf{k}$ are wavevectors. Following the terminology of quantum mechanics, we will refer to $\mathbf{k}$ as the *momentum*. At this stage, we should remember something about our roots. Although we're thinking of $\phi(\mathbf{x})$ as a continuous field, ultimately it arose from a lattice and so can't vary on very small distance scales. This means that the Fourier modes must all vanish for suitably high momentum,

$$\phi_{\mathbf{k}} = 0 \quad \text{for } |\mathbf{k}| > \Lambda$$

Here Λ is called the *ultra-violet (UV) cut-off*. In the present case, we can think of $\Lambda = \pi/a$, with a the distance between the boxes that we used to coarse grain when first defining $\phi(\mathbf{x})$.

We can recover our original field by the inverse Fourier transform. It's useful to have two different scenarios in mind. In the first, we place the system in a finite spatial volume $V = L^d$ with periodic boundary conditions. In this case, the momenta take discrete values,

$$\mathbf{k} = \frac{2\pi}{L}\mathbf{n} \qquad \mathbf{n} \in \mathbf{Z}^d \tag{2.6}$$

and the inverse Fourier transform is

$$\phi(\mathbf{x}) = \frac{1}{V}\sum_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{x}}\,\phi_{\mathbf{k}} \tag{2.7}$$

Alternatively, if we send $V \to \infty$, the sum over $\mathbf{k}$ modes becomes an integral and we have

$$\phi(\mathbf{x}) = \int \frac{d^dk}{(2\pi)^d} e^{i\mathbf{k}\cdot\mathbf{x}}\,\phi_{\mathbf{k}} \tag{2.8}$$

In what follows, we'll jump backwards and forwards between these two forms. Ultimately, we will favour the integral form. But there are times when it will be simpler to think of the system in a finite volume as it will make us less queasy about some of the formulae we'll encounter.

For now, let's work with the form (2.8). We substitute this into the free energy to find

$$F[\phi_{\mathbf{k}}] = \frac{1}{2}\int \frac{d^dk_1}{(2\pi)^d}\frac{d^dk_2}{(2\pi)^d}\int d^dx\ \left(-\gamma\mathbf{k}_1\cdot\mathbf{k}_2 + \mu^2\right)\phi_{\mathbf{k}_1}\phi_{\mathbf{k}_2}\,e^{i(\mathbf{k}_1+\mathbf{k}_2)\cdot\mathbf{x}}$$

The integral over $\mathbf{x}$ is now straightforward and gives us a delta function

$$\int d^dx\ e^{i(\mathbf{k}_1+\mathbf{k}_2)\cdot\mathbf{x}} = (2\pi)^d\delta^d(\mathbf{k}_1+\mathbf{k}_2)$$

and the free energy takes the simple form

$$F[\phi_{\mathbf{k}}] = \frac{1}{2}\int \frac{d^dk}{(2\pi)^d}\ \left(\gamma k^2+\mu^2\right)\phi_{\mathbf{k}}\phi_{-\mathbf{k}} = \frac{1}{2}\int \frac{d^dk}{(2\pi)^d}\ \left(\gamma k^2+\mu^2\right)\phi_{\mathbf{k}}\phi^\star_{\mathbf{k}} \tag{2.9}$$

Now we can see the benefit of working in Fourier space: at quadratic order, the free energy decomposes into individual $\phi_{\mathbf{k}}$ modes. This is because the Fourier basis is the eigenbasis of $-\gamma\nabla^2+\mu^2$, allowing us to diagonalise this operator.

To perform the functional integral, we also need to change the measure. Recall that the path integral was originally an integral over $\phi(\mathbf{x})$ for each value of $\mathbf{x}$ labelling the position of a box. Now it is an integral over all Fourier modes, which we write as

$$\int \mathcal{D}\phi(\mathbf{x}) = \prod_{\mathbf{k}} \left[\mathcal{N} \int d\phi_{\mathbf{k}} d\phi^\star_{\mathbf{k}} \right] \tag{2.10}$$

I should warn you that this equation, as written, isn't quite right because we should remember that $\phi(x)$ is real, which means that $\phi^\star_{\mathbf{k}} = \phi_{-\mathbf{k}}$. So the measure above double counts the number of integrations we're doing. We could try to amend our notation to take care of this, but it's better to just remember what the measure above means. We won't have to remember for long: the issue is resolved by the time we get to (2.11) below.

I've included a normalisation constant $\mathcal{N}$ in the measure. I'll make no attempt to calculate this and, ultimately, it won't play any role because, having computed the partition function, the first thing we do is take the log and differentiate. At this point, $\mathcal{N}$ will drop out. In later formulae, we'll simply ignore it. But it's worth keeping it in for now.

Our path integral is now

$$Z = \prod_{\mathbf{k}} \mathcal{N} \int d\phi_{\mathbf{k}} d\phi^\star_{\mathbf{k}} \ \exp\left(-\frac{\beta}{2} \int \frac{d^d k}{(2\pi)^d} \ \left(\gamma k^2 + \mu^2\right) |\phi_{\mathbf{k}}|^2 \right)$$

If this still looks daunting, we just need to recall that in finite volume, the integral in the exponent is really a sum over discrete momentum values,

$$\begin{aligned} Z &= \prod_{\mathbf{k}} \mathcal{N} \int d\phi_{\mathbf{k}} d\phi^\star_{\mathbf{k}} \ \exp\left(-\frac{\beta}{2V} \sum_{\mathbf{k}} \ \left(\gamma k^2 + \mu^2\right) |\phi_{\mathbf{k}}|^2 \right) \\ &= \prod_{\mathbf{k}} \left[\mathcal{N} \int d\phi_{\mathbf{k}} d\phi^\star_{\mathbf{k}} \ \exp\left(-\frac{\beta}{2V} \left(\gamma k^2 + \mu^2\right) |\phi_{\mathbf{k}}|^2 \right) \right] \end{aligned}$$

Note that the placement of brackets shifted in the two lines, because the sum in the exponent got absorbed into the overall product. If this is confusing, it might be worth comparing what we've done above to a simple integral of the form $\int dx\, dy\ e^{-x^2-y^2} = (\int dx\ e^{-x^2})(\int dy\ e^{-y^2})$.

We're left with something very straightforward: it's simply a bunch of decoupled Gaussian integrals, one for each value of $\mathbf{k}$. Recall that a Gaussian integral over a

single variable is given by

$$\int_{-\infty}^{+\infty} dx\ e^{-x^2/2a} = \sqrt{2\pi a} \tag{2.11}$$

Applying this for each $\mathbf{k}$, we have our expression for the path integral

$$Z = \prod_{\mathbf{k}} \mathcal{N} \sqrt{\frac{2\pi TV}{\gamma k^2 + \mu^2}}$$

where the square root is there, despite the fact that we're integrating over complex $\phi_{\mathbf{k}}$, because $\phi^\star_{\mathbf{k}} = \phi_{-\mathbf{k}}$ is not an independent variable. Note that we have a product over all $\mathbf{k}$. In finite volume, where the possible $\mathbf{k}$ are discrete, there's nothing fishy going on. But as we go to infinite volume, this will become a product over a continuous variable $\mathbf{k}$.

We can now compute the contribution of these thermal fluctuations to the thermodynamic free energy. The free energy per unit volume is given by $Z = e^{-\beta F_{\rm thermo}}$ or,

$$\frac{F_{\rm thermo}}{V} = -\frac{T}{V} \log Z = -\frac{T}{2V} \sum_{\mathbf{k}} \log \left(\frac{2\pi TV\mathcal{N}^2}{\gamma k^2 + \mu^2} \right)$$

We can now revert back to the integral over $\mathbf{k}$, rather than the sum by writing

$$\frac{F_{\rm thermo}}{V} = -\frac{T}{2} \int \frac{d^d k}{(2\pi)^d} \log \left(\frac{2\pi TV\mathcal{N}^2}{\gamma k^2 + \mu^2} \right)$$

This final equation might make you uneasy since an explicit factor of the volume V remains in the argument, but we've sent $V \to \infty$ to convert from $\sum_{\mathbf{k}}$ to $\int d^d k$. At this point, the normalisation factor $\mathcal{N}$ will ride to the rescue. However, as advertised previously, none of these issues are particularly important since they drop out when we compute physical quantities. Let's look at the simplest example.

2.1.1 The Heat Capacity

Our real interest lies in the heat capacity per unit volume, $c = C/V$. Specifically, we would like to understand the temperature dependence of the heat capacity. This is given by (1.15),

$$c = \frac{\beta^2}{V} \frac{\partial^2}{\partial \beta^2} \log Z = \frac{1}{2} \left(T^2 \frac{\partial^2}{\partial T^2} + 2T \frac{\partial}{\partial T} \right) \int \frac{d^d k}{(2\pi)^d} \log \left(\frac{2\pi TV\mathcal{N}^2}{\gamma k^2 + \mu^2} \right)$$

The derivatives hit both the factor of T in the numerator, and any T dependence in the coefficients γ and μ^2. For simplicity, let's work at $T > T_c$. We'll take γ constant and $\mu^2 = T - T_c$. A little bit of algebra shows that the contribution to the heat capacity from the fluctuations is given by

$$c = \frac{1}{2}\int \frac{d^dk}{(2\pi)^d}\left[1 - \frac{2T}{\gamma k^2 + \mu^2} + \frac{T^2}{(\gamma k^2 + \mu^2)^2}\right] \tag{2.12}$$

The first of these terms has a straightforward interpretation: it is the usual "$\frac{1}{2}k_B$" per degree of freedom that we expect from equipartition, albeit with $k_B = 1$. (This can be traced to the original β in $e^{-\beta F}$.)

The other two terms come from the temperature dependence in $F[\phi(\mathbf{x})]$. What happens next depends on the dimension d. Let's look at the middle term, proportional to

$$\int_0^\Lambda dk\ \frac{k^{d-1}}{\gamma k^2 + \mu^2}$$

For $d \geq 2$, this integral diverges as we remove the UV cut-off Λ. In contrast, when $d = 1$ it is finite as $\Lambda \to \infty$. When it is finite, we can easily determine the leading order temperature dependence of the integral by rescaling variables. We learn that

$$\int_0^\Lambda dk\ \frac{k^{d-1}}{\gamma k^2 + \mu^2} \sim \begin{cases} \Lambda^{d-2} & \text{when } d \geq 2 \\ 1/\mu & \text{when } d = 1 \end{cases} \tag{2.13}$$

When $d = 2$, the term Λ^0 should be replaced by a logarithm. Similarly, the final term in (2.12) is proportional to

$$\int_0^\Lambda dk\ \frac{k^{d-1}}{(\gamma k^2 + \mu^2)^2} \sim \begin{cases} \Lambda^{d-4} & \text{when } d \geq 4 \\ \mu^{d-4} & \text{when } d < 4 \end{cases}$$

again, with a logarithm when $d = 4$.

What should we take from this? When $d \geq 4$, the leading contribution to the heat capacity involves a temperature independent constant, Λ, albeit a large one. This constant will be the same on both sides of the transition. (The heat capacity itself is not quite temperature independent as it comes with the factor of T^2 from the numerator of (2.12), but this doesn't do anything particularly dramatic.) In contrast, when $d < 4$, the leading order contribution to the heat capacity is proportional to μ^{d-4}. And, this leads to something more interesting.

To see this interesting behaviour, we have to do something naughty. Remember that our calculation above isn't valid near the critical point, $\mu^2 = 0$, because we've ignored the quartic term in the free energy. Suppose, however, that we throw caution to the wind and apply our result here anyway. We learn that, for $d < 4$, the heat capacity diverges at the critical point. The leading order behaviour is

$$c \sim |T - T_c|^{-\alpha} \quad \text{with } \alpha = 2 - \frac{d}{2} \tag{2.14}$$

This is to be contrasted with our mean field result which gives $\alpha = 0$.

As we've stressed, we can't trust the result (2.14). And, indeed, this is not the right answer for the critical exponent. But it does give us some sense for how the mean field results can be changed by the path integral. It also gives a hint for why the critical exponents are not affected when $d \geq 4$, which is the upper critical dimension.

2.2 Correlation Functions

The essential ingredient of Landau-Ginzburg theory – one that was lacking in the earlier Landau approach – is the existence of spatial structure. With the local order parameter $\phi(\mathbf{x})$, we can start to answer questions about how the magnetisation varies from point to point.

Such spatial variations exist even in the ground state of the system. Mean field theory – which is synonymous with the saddle point of the path integral – tells us that the expectation value of the magnetisation is constant in the ground state

$$\langle \phi(\mathbf{x}) \rangle = \begin{cases} 0 & T > T_c \\ \pm m_0 & T < T_c \end{cases} \tag{2.15}$$

This makes us think of the ground state as a calm fluid, like the Cambridge mill pond when the tourists are out of town. This is misleading. The ground state is not a single field configuration but, as always in statistical mechanics, a sum over many possible configurations in the thermal ensemble. This is what the path integral does for us. The importance of these other configurations will determine whether the ground state is likely to contain only gentle ripples around the background (2.15), or fluctuations so wild that it makes little sense to talk about an underlying background at all.

These kind of spatial fluctuations of the ground state are captured by *correlation functions*. The simplest is the two-point function $\langle \phi(\mathbf{x})\phi(\mathbf{y}) \rangle$, computed using the probability distribution (1.26). This tells us how the magnetisation at point $\mathbf{x}$ is correlated with the magnetisation at $\mathbf{y}$. If, for example, there is an unusually large fluctuation at $\mathbf{y}$, what will the magnitude of the field most likely be at $\mathbf{x}$?

Because $\langle\phi(\mathbf{x})\rangle$ takes different values above and below the transition, it is often more useful to compute the *connected correlation function*,

$$\langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle_c = \langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle - \langle\phi\rangle^2 \tag{2.16}$$

If you're statistically inclined, this is sometimes called a *cumulant* of the random variable $\phi(\mathbf{x})$.

The path integral provides a particularly nice way to compute connected correlation functions of this kind. We consider the system in the presence of a magnetic field B, but now allow $B(\mathbf{x})$ to also vary in space. We take the free energy to be

$$F[\phi(\mathbf{x})] = \int d^dx \ \left[\frac{\gamma}{2}(\nabla\phi)^2 + \frac{\mu^2}{2}\phi^2(\mathbf{x}) - B(\mathbf{x})\phi(\mathbf{x})\right] \tag{2.17}$$

We can now think of the partition function as a functional of $B(\mathbf{x})$.

$$Z[B(\mathbf{x})] = \int \mathcal{D}\phi \ e^{-\beta F}$$

For what it's worth, $Z[B(\mathbf{x})]$ is related to the Legendre transform of $F[\phi(\mathbf{x})]$.

Now that Z depends on the function $B(\mathbf{x})$ it is a much richer and more complicated object. Indeed, it encodes all the information about the fluctuations of the theory. Consider, for example, the functional derivative of $\log Z$,

$$\frac{1}{\beta}\frac{\delta \log Z}{\delta B(\mathbf{x})} = \frac{1}{\beta Z}\frac{\delta Z}{\delta B(\mathbf{x})} = \frac{1}{Z}\int \mathcal{D}\phi \ \phi(\mathbf{x})\, e^{-\beta F} = \langle\phi(\mathbf{x})\rangle_B$$

Here I've put a subscript B on $\langle\cdot\rangle_B$ to remind us that this is the expectation value computed in the presence of the magnetic field $B(\mathbf{x})$. If our real interest is in what happens as we approach the critical point, we can simply set $B = 0$.

Similarly, we can take two derivatives of $\log Z$. Now when the second derivative hits, it can either act on the exponent $e^{-\beta F}$, or on the $1/Z$ factor in front. The upshot is that we get

$$\frac{1}{\beta^2}\frac{\delta^2 \log Z}{\delta B(\mathbf{x})\delta B(\mathbf{y})} = \frac{1}{\beta^2 Z}\frac{\delta^2 Z}{\delta B(\mathbf{x})\delta B(\mathbf{y})} - \frac{1}{\beta^2 Z^2}\frac{\delta Z}{\delta B(\mathbf{x})}\frac{\delta Z}{\delta B(\mathbf{y})}$$

or

$$\frac{1}{\beta^2}\frac{\delta^2 \log Z}{\delta B(\mathbf{x})\delta B(\mathbf{y})} = \langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle_B - \langle\phi(\mathbf{x})\rangle_B\langle\phi(\mathbf{y})\rangle_B$$

which is precisely the connected correlation function (2.16). In what follows, we'll mostly work above the critical temperature so that $\langle\phi\rangle_{B=0} = 0$. In this case, we set $B = 0$ to find

$$\frac{1}{\beta^2}\left.\frac{\delta^2 \log Z}{\delta B(\mathbf{x})\delta B(\mathbf{y})}\right|_{B=0} = \langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle \tag{2.18}$$

All that's left is for us to compute the path integral $Z[B(\mathbf{x})]$.

2.2.1 The Gaussian Path Integral

As in our calculation of the thermodynamic free energy, we work in Fourier space. The free energy is now a generalisation of (2.9),

$$F[\phi_{\mathbf{k}}] = \int \frac{d^d k}{(2\pi)^d}\left[\frac{1}{2}\left(\gamma k^2 + \mu^2\right)\phi_{\mathbf{k}}\phi_{-\mathbf{k}} - B_{-\mathbf{k}}\phi_{\mathbf{k}}\right]$$

where $B_{\mathbf{k}}$ are the Fourier modes of $B(\mathbf{x})$. To proceed, we complete the square, and define the shifted magnetisation

$$\hat{\phi}_{\mathbf{k}} = \phi_{\mathbf{k}} - \frac{B_{\mathbf{k}}}{\gamma k^2 + \mu^2}$$

We can then write the free energy as

$$F[\hat{\phi}_{\mathbf{k}}] = \int \frac{d^d k}{(2\pi)^d}\left[\frac{1}{2}\left(\gamma k^2 + \mu^2\right)|\hat{\phi}_{\mathbf{k}}|^2 - \frac{1}{2}\frac{|B_{\mathbf{k}}|^2}{\gamma k^2 + \mu^2}\right]$$

Our path integral is

$$Z = \prod_{\mathbf{k}} \int d\hat{\phi}_{\mathbf{k}} d\hat{\phi}^\star_{\mathbf{k}}\; e^{-\beta F[\hat{\phi}_{\mathbf{k}}]}$$

where we've shifted the integration variable from $\phi_{\mathbf{k}}$ to $\hat{\phi}_{\mathbf{k}}$; there is no Jacobian penalty for doing this. We've also dropped the normalisation constant $\mathcal{N}$ that we included in our previous measure (2.10) on the grounds that it clutters equations and does nothing useful.

The path integral now gives

$$Z[B(\mathbf{x})] = e^{-\beta F_{\rm thermo}}\,\exp\left(\frac{\beta}{2}\int \frac{d^d k}{(2\pi)^d}\,\frac{|B_{\mathbf{k}}|^2}{\gamma k^2 + \mu^2}\right)$$

The first term $e^{-\beta F_{\rm thermo}}$ is just the contribution we saw before. It does not depend on the magnetic field $B(\mathbf{x})$ and won't contribute to the correlation function. (Specifically,

it will drop out when we differentiate $\log Z$.) The interesting piece is the dependence on the Fourier modes $B_{\mathbf{k}}$. To get back to real space $B(\mathbf{x})$, we simply need to do an inverse Fourier transform. We have

$$Z[B(\mathbf{x})] = e^{-\beta F_{\rm thermo}} \exp\left(\frac{\beta}{2}\int d^dx d^dy\ B(\mathbf{x})G(\mathbf{x}-\mathbf{y})B(\mathbf{y})\right) \tag{2.19}$$

where

$$G(\mathbf{x}) = \int \frac{d^dk}{(2\pi)^d}\frac{e^{-i\mathbf{k}\cdot\mathbf{x}}}{\gamma k^2+\mu^2} \tag{2.20}$$

We're getting there. Differentiating the partition function as in (2.18), we learn that the connected two-point function is

$$\langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle = \frac{1}{\beta}G(\mathbf{x}-\mathbf{y}) \tag{2.21}$$

We just need to do the integral (2.20).

Computing the Fourier Integral

To start, note that the integral $G(\mathbf{x})$ is rotationally invariant, and so $G(\mathbf{x}) = G(r)$ with $r = |\mathbf{x}|$. We write the integral as

$$G(r) = \frac{1}{\gamma}\int \frac{d^dk}{(2\pi)^d}\frac{e^{-i\mathbf{k}\cdot\mathbf{x}}}{k^2+1/\xi^2}$$

where we've introduced a length scale

$$\xi^2 = \frac{\gamma}{\mu^2} \tag{2.22}$$

This is called the *correlation length* and it will prove to be important as we move forwards. We'll discuss it more in Section 2.2.3.

To proceed, we use a trick. We can write

$$\frac{1}{k^2+1/\xi^2} = \int_0^\infty dt\ e^{-t(k^2+1/\xi^2)}$$

Using this, we have

$$\begin{aligned} G(r) &= \frac{1}{\gamma}\int\frac{d^dk}{(2\pi)^d}\int_0^\infty dt\ e^{-i\mathbf{k}\cdot\mathbf{x}-t(k^2+1/\xi^2)} \\ &= \frac{1}{\gamma}\int\frac{d^dk}{(2\pi)^d}\int_0^\infty dt\ e^{-t(\mathbf{k}+i\mathbf{x}/2t)^2}e^{-r^2/4t-t/\xi^2} \\ &= \frac{1}{\gamma(4\pi)^{d/2}}\int_0^\infty dt\ t^{-d/2}e^{-r^2/4t-t/\xi^2} \end{aligned} \tag{2.23}$$

where, in going to the last line, we've simply done the d Gaussian integrals over $\mathbf{k}$. At this point there are a number of different routes. We could invoke some special-functionology and note that we can massage the integral into the form of a Bessel function

$$K_\nu(z) = \frac{1}{2}\int_0^\infty dt\ t^{\nu-1}e^{-z(t+1/t)/2}$$

whose properties you can find in some dog-eared mathematical methods textbook. However, our interest is only in the behaviour of $G(r)$ in various limits, and for this purpose it will suffice to perform the integral (2.23) using a saddle point. We ignore overall constants, and write the integral as

$$G(r) \sim \int_0^\infty dt\ e^{-S(t)} \quad \text{with} \quad S(t) = \frac{r^2}{4t} + \frac{t}{\xi^2} + \frac{d}{2}\log t$$

The saddle point $t = t_\star$ sits at $S'(t_\star) = 0$. We then approximate the integral as

$$G(r) \sim \int_0^\infty dt\ e^{-S(t_\star)+S''(t_\star)t^2/2} = \sqrt{\frac{\pi}{2S''(t_\star)}}\, e^{-S(t_\star)}$$

For us the saddle lies at

$$S'(t_\star) = 0 \quad \Rightarrow \quad t_\star = \frac{\xi^2}{2}\left(-\frac{d}{2} + \sqrt{\frac{d^2}{4} + \frac{r^2}{\xi^2}}\right)$$

There are two different limits that we are interested in: $r \gg \xi$ and $r \ll \xi$. We'll deal with them in turn:

$\underline{r \gg \xi}$: In this regime, we have $t_\star \approx r\xi/2$. And so $S(t_\star) \approx r/\xi + (d/2)\log(r\xi/2)$. One can also check that $S''(t_\star) \approx 4/r\xi^3$. The upshot is that the asymptotic form of the integral scales as

$$G(r) \sim \frac{1}{\xi^{d/2-3/2}}\,\frac{e^{-r/\xi}}{r^{d/2-1/2}} \qquad r \gg \xi$$

At large distance scales, the correlation function falls off exponentially.

$\underline{r \ll \xi}$: In the other regime, the saddle point lies at $t_\star \approx r^2/2d$, giving $S(t_\star) \approx d/2 + (d/2)\log(r^2/2d)$ and $S''(t_\star) \approx 2d^3/r^4$. Putting this together, we see that for $r \ll \xi$, the fall-off is only power law at short distances,

$$G(r) \sim \frac{1}{r^{d-2}} \qquad r \ll \xi$$

We learn that the correlation function changes its form at the distance scale $r \sim \xi$, with the limiting form

$$\langle \phi(\mathbf{x})\phi(\mathbf{y})\rangle \sim \begin{cases} \dfrac{1}{r^{d-2}} & r \ll \xi \\ \dfrac{e^{-r/\xi}}{r^{(d-1)/2}} & r \gg \xi \end{cases} \tag{2.24}$$

This is known as the *Ornstein-Zernicke* correlation.

2.2.2 The Correlation Function is a Green's Function

The result (2.24) is important and we we'll delve a little deeper into it shortly. But first, it will prove useful to redo the calculation above in real space, rather than Fourier space, to highlight some of the machinery hiding behind our path integral.

To set some foundations, we start with a multi-dimensional integral over n variables. Suppose that $\mathbf{y}$ is an n-dimensional vector. The simple Gaussian integral now involves an invertible $n \times n$ matrix G,

$$\int_{-\infty}^{+\infty} d^n y \; e^{-\frac{1}{2}\mathbf{y}\cdot G^{-1}\mathbf{y}} = \det{}^{1/2}(2\pi G)$$

This result follows straighforwardly from the single-variable Gaussian integral (2.11), by using a basis that diagonalises G. Similarly, if we introduce an n-dimensional vector $\mathbf{B}$, we can complete the square to find

$$\int_{-\infty}^{+\infty} d^n y \; e^{-\frac{1}{2}\mathbf{y}\cdot G^{-1}\mathbf{y}+\mathbf{B}\cdot\mathbf{y}} = \det{}^{1/2}(2\pi G)\, e^{\frac{1}{2}\mathbf{B}\cdot G\mathbf{B}} \tag{2.25}$$

Now let's jump to the infinite dimensional, path integral version of this. Throughout this section, we've been working with a quadratic free energy

$$F[\phi(\mathbf{x})] = \int d^d x \;\; \left[\frac{1}{2}\gamma(\nabla\phi)^2 + \frac{1}{2}\mu^2\phi^2(\mathbf{x}) - B(\mathbf{x})\phi(\mathbf{x})\right] \tag{2.26}$$

We can massage this into the form of the exponent in (2.25) by writing

$$F[\phi(\mathbf{x})] = \int d^d x \int d^d y \; \frac{1}{2}\phi(\mathbf{x})G^{-1}(\mathbf{x},\mathbf{y})\phi(\mathbf{y}) - \int d^d x \; B(\mathbf{x})\phi(\mathbf{x})$$

where we've introduced the "infinite dimensional matrix", more commonly known as an operator

$$G^{-1}(\mathbf{x},\mathbf{y}) = \delta^d(\mathbf{x}-\mathbf{y})\left(-\gamma\nabla_y^2 + \mu^2\right) \tag{2.27}$$

Note that this is the operator that appears in the saddle point evaluation of the free energy, as we saw earlier in (1.30).

Given the operator G^{-1}, what is the inverse operator $G(\mathbf{x}, \mathbf{y})$? We have another name for the inverse of an operator: it is called a *Green's function*. In the present case, $G(\mathbf{x}, \mathbf{y})$ obeys the equation

$$(-\gamma\nabla_x^2 + \mu^2)G(\mathbf{x}, \mathbf{y}) = \delta^d(\mathbf{x} - \mathbf{y})$$

By translational symmetry, we have $G(\mathbf{x}, \mathbf{y}) = G(\mathbf{x} - \mathbf{y})$. You can simply check that the Green's function is indeed given in Fourier space by our previous result (2.20)

$$G(\mathbf{x}) = \int \frac{d^d k}{(2\pi)^d} \frac{e^{-i\mathbf{k}\cdot\mathbf{x}}}{\gamma k^2 + \mu^2}$$

This route led us to the same result we had previously. Except we learn something new: the correlation function is the same thing as the Green's function, $\langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle = \beta^{-1}G(\mathbf{x}, \mathbf{y})$, and hence solves,

$$(-\gamma\nabla^2 + \mu^2)\,\langle\phi(\mathbf{x})\phi(0)\rangle = \frac{1}{\beta}\delta^d(\mathbf{x})$$

This is telling us that if we perturb the system at the origin then, for a free energy of the quadratic form (2.26), the correlator $\langle\phi(\mathbf{x})\phi(0)\rangle$ responds by solving the original saddle point equation.

There is one further avatar of the correlation function that is worth mentioning: it is related to the *susceptibility*. Recall that previously we defined the susceptibility in (1.19) as $\chi = \partial m/\partial B$. Now, we have a more refined version of susceptibility which knows about the spatial structure,

$$\chi(\mathbf{x}, \mathbf{y}) = \frac{\delta\langle\phi(\mathbf{x})\rangle}{\delta B(\mathbf{y})}$$

But, from our discussion above, this is exactly the correlation function $\chi(\mathbf{x}, \mathbf{y}) = \beta\langle\phi(\mathbf{x})\phi(\mathbf{y})\rangle$. We can recover our original, coarse grained susceptibility as

$$\chi = \int d^d x\ \chi(\mathbf{x}, 0) = \beta \int d^d x\ \langle\phi(\mathbf{x})\phi(0)\rangle \tag{2.28}$$

The two point correlation function will play an increasingly important role in later calculations. For this reason it is given its own name: it is called the *propagator*. Propagators of this kind also arose in the lectures on Quantum Field Theory. In that case, the propagator was defined for a theory in Minkowski space, which led to an ambiguity (of integration contour) and a choice of different propagators: advanced, retarded or Feynman. In the context of Statistical Field Theory, we are working in Euclidean space and there is no such ambiguity.

2.2.3 The Correlation Length

Let's now look a little more closely at the expression (2.24) for the correlation function which, in two different regimes, scales as

$$\langle \phi(\mathbf{x})\phi(\mathbf{y})\rangle \sim \begin{cases} \dfrac{1}{r^{d-2}} & r \ll \xi \\ \\ \dfrac{e^{-r/\xi}}{r^{(d-1)/2}} & r \gg \xi \end{cases} \tag{2.29}$$

where $r = |\mathbf{x}-\mathbf{y}|$. The exponent contains a length scale, ξ that we previously defined as the *correlation length*, given in terms of the parameters in the free energy as $\xi^2 = \gamma/\mu^2$.

We see from (2.29) that all correlations die off quickly at distances $r \gg \xi$. In contrast, for $r \ll \xi$ there is only a much slower, power-law fall-off. In this sense, ξ provides a characteristic length scale for the fluctuations. In a given thermal ensemble, there will be patches where the magnetisation is slightly higher, or slightly lower than the average $\langle m\rangle$. The size of these patches will be no larger than ξ.

Recall that, close to the critical point, $\mu^2 \sim |T-T_c|$. This means that as we approach $T = T_c$, the correlation length diverges as

$$\xi \sim \frac{1}{|T-T_c|^{1/2}} \tag{2.30}$$

This is telling us that system will undergo fluctuations of arbitrarily large size. This is the essence of a second order phase transition, and as we move forward we will try to better understand these fluctuations.

Numerical Simulations of the Ising Model

It's useful to get a sense for what these fluctuations look like. We start in the disordered phase with $T > T_c$. In the figures you can see two typical configurations that contribute to the partition function of the Ising model[5]. The up spins are shown in yellow, the down spins in blue.

On the left, the temperature is $T_1 > T_c$, while on the right the temperature is $T_2 > T_1 > T_c$. In both pictures, the spins look random. And yet, you can see by eye that there is something different between the two pictures; on the right, when the

[5]These images were generated by the Metropolis algorithm using a mathematica programme created by Daniel Schroeder. It's well worth playing with to get a feel for what's going on. Ising simulations, in various formats, can be found on his webpage http://physics.weber.edu/thermal/computer.html.

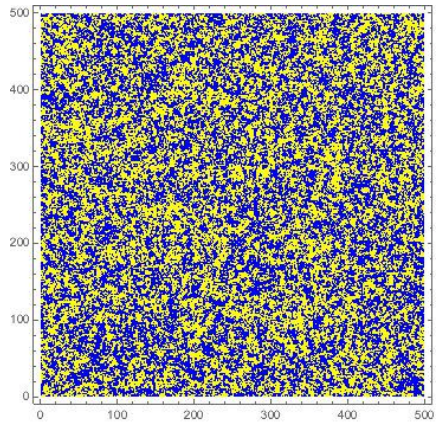

Figure 16: Spins with when $T > T_c$

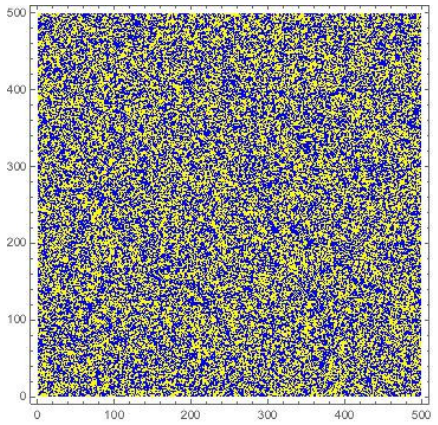

Figure 17: Spins when $T \gg T_c$

temperature is higher, the spins are more finely intertwined, with a yellow spin likely to have a blue dot sitting right next to it. Meanwhile, on the left, the randomness is coarser.

What you're seeing here is the correlation length at work. In each picture, ξ sets the typical length scale of fluctuations. In the right-hand picture, where the temperature is higher, the correlation length is smaller.

There is a similar story in the ordered phase, with $T < T_c$. Once again, we show two configurations below. Now the system must choose between one of the two ground states; here the choice is that the yellow up spins are dominant. The left-hand configuration has temperature $T_1' < T_c$, and the right-hand configuration temperature $T_2' < T_1' < T_c$. We see that sizes of the fluctuations around the ordered phase become smaller the further we sit from the critical point.

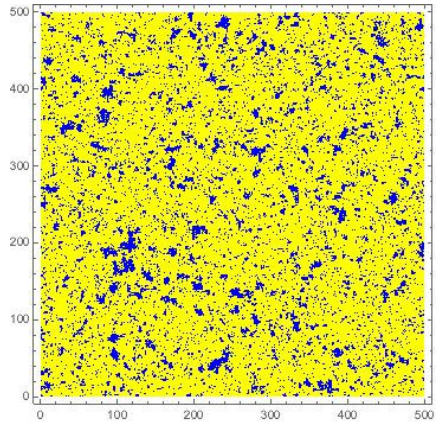

Figure 18: Spins with when $T < T_c$

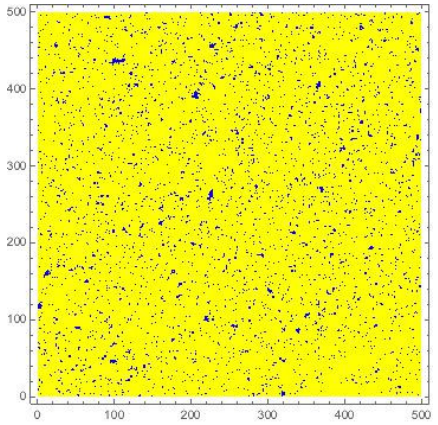

Figure 19: Spins when $T \ll T_c$

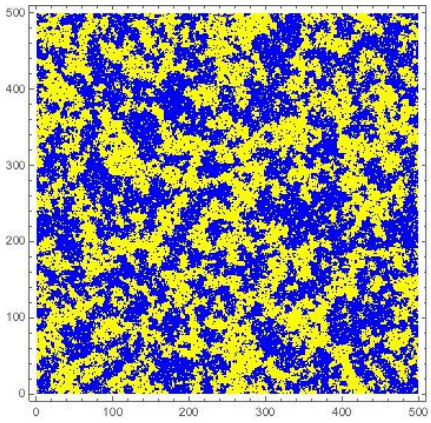

Figure 20: $T = T_c$

Finally, we can ask what happens when we sit at the critical point $T = T_c$. A typical configuration is shown in Figure 20. Although it may not be obvious, there is now no characteristic length scale in the picture. Instead, fluctuations occur on all length scales, big and small[6]. This is the meaning of the diverging correlation length $\xi \to \infty$.

Critical Opalescence

There is a nice experimental realisation of these large fluctuations, which can be seen in liquid-gas transitions or mixing transitions between two different fluids. (Both of these lie in the same universality class as the Ising model.) As we approach the second order phase transition, transparent fluids become cloudy, an effect known as *critical opalescence*. What's happening is that the size of the density fluctuations is becoming larger and larger, until they begin to scatter visible light.

More Critical Exponents

We saw in previous sections that we get a number of power-laws at critical points, each of which is related to a critical exponent. The results above give us two further exponents to add to our list. First, we have a correlation length ξ which diverges at the critical point with power (2.30)

$$\xi \sim \frac{1}{|T - T_c|^\nu} \qquad \text{where } \nu = \frac{1}{2}$$

[6]The best demonstration that I've seen of this scale invariance at the critical point is this Youtube video by Douglas Ashton.

Similarly, we know that the correlation function itself is a power law at the critical point, with exponent

$$\langle \phi(\mathbf{x})\phi(\mathbf{y})\rangle \sim \frac{1}{r^{d-2+\eta}} \quad \text{where } \eta = 0$$

Each of these can be thought of as a mean field prediction, in the sense that we are treating the path integral only in quadratic order, which neglects important effects near the critical point. Given our previous discussion, it may not come as a surprise to learn that these critical exponents are correct when $d \geq 4$. However, they are not correct in lower dimensions. Instead one finds

	MF	$d=2$	$d=3$
η	0	$\frac{1}{4}$	0.0363
ν	$\frac{1}{2}$	1	0.6300

This gives us another challenge, one we will rise to in Section 3.

2.2.4 The Upper Critical Dimension

We're finally in a position to understand why the mean field results hold in high dimensions, but are incorrect in low dimensions. Recall our story so far: when $T < T_c$, the saddle point suggests that

$$\langle \phi(\mathbf{x})\rangle = \pm m_0$$

Meanwhile, there are fluctuations around this mean field value described, at long distances, by the correlation function (2.29). In order to trust our calculations, these fluctuations should be smaller than the background around which they're fluctuating. In other words, we require $\langle \phi^2\rangle \ll \langle \phi\rangle^2$.

It's straightforward to get an estimate for this. We know that the fluctuations decay after a distance $r \gg \xi$. We can gain a measure of their importance if we integrate over a ball of radius ξ. We're then interested in the ratio

$$R = \frac{\int_0^\xi d^dx\ \langle \phi(\mathbf{x})\phi(0)\rangle}{\int_0^\xi d^dx\ m_0^2} \sim \frac{1}{m_0^2\,\xi^d}\int_0^\xi dr\ \frac{r^{d-1}}{r^{d-2}} \sim \frac{\xi^{2-d}}{m_0^2}$$

In order to trust mean field theory, we require that this ratio is much less than one. This is the *Ginzburg criterion*. We can anticipate trouble as we approach a critical point, for here ξ diverges and m_0 vanishes. According to mean field theory, these two quantities scale as

$$m_0 \sim |T - T_c|^{1/2} \quad \text{and} \quad \xi \sim |T - T_c|^{-1/2}$$

results which can be found, respectively, in (1.21) and (2.30). This means that the ratio R scales as

$$R \sim |T - T_c|^{(d-4)/2}$$

We learn that, as we approach the critical point, mean field – which, in this context, means computing small fluctuations around the saddle point – appears trustworthy only if

$$d \geq d_c = 4$$

This is the *upper critical dimension* for the Ising model. Actually, at the critical dimension $d = 4$ there is a logarithmic divergence in R and so we have to treat this case separately; we'll be more careful about this in the Section 3.

For dimensions $d < 4$, mean field theory predicts its own demise. We'll see how to make progress in Section 3.

2.3 The Analogy with Quantum Field Theory

There is a very close analogy between the kinds of field theories we're looking at here, and those that arise in quantum field theory. This analogy is seen most clearly in Feynman's path integral approach to quantum mechanics[7]. Correlation functions in both statistical and quantum field theories are captured by partition functions

$$\text{Statistical Field Theory:} \qquad Z = \int \mathcal{D}\phi \; e^{-\beta \int d^d x \; \mathcal{F}(\phi)}$$

$$\text{Quantum Field Theory:} \qquad Z = \int \mathcal{D}\phi \; e^{\frac{i}{\hbar} \int d^d x \; \mathcal{L}(\phi)}$$

You don't need to be a visionary to see the similarities. But there are also some differences: the statistical path integral describes a system in d spatial dimensions, while the quantum path integral describes a system in d *spacetime* dimensions, or $d-1$ spatial dimensions.

The factor of i in the exponent of the quantum path integral can be traced to the fact that it describes a system evolving in time, and means that it has more subtle convergence properties than its statistical counterpart. In practice, to compute anything in the quantum partition function, one tends to rotate the integration contour and work with Euclidean time,

$$\tau = it \tag{2.31}$$

[7]You will see this in next term's *Advanced Quantum Field Theory* course.

This is known as a *Wick rotation*. After this, the quantum and statistical partition functions are mathematically the same kind of objects

$$Z = \int \mathcal{D}\phi \; e^{\frac{i}{\hbar}S[\phi]} \quad \rightarrow \quad Z = \int \mathcal{D}\phi \; e^{-S_E[\phi]/\hbar}$$

where $S_E[\phi]$ is the Euclidean action, and is analogous to the free energy in statistical mechanics. If the original action $S[\phi]$ was Lorentz invariant, then the Euclidean action $S_E[\phi]$ will be rotationally invariant. Suddenly, the $d = 4$ dimensional field theories, which seemed so unrealistic in the statistical mechanics context, take on a new significance.

By this stage, the only difference between the two theories is the words we drape around them. In statistical mechanics, the path integral captures the thermal fluctuations of the local order parameter, with a strength controlled by the temperature β; in quantum field theory the path integral captures the quantum fluctuations of the field ϕ, with a strength controlled by $\hbar$. This means that many of the techniques we will develop in this course can be translated directly to quantum field theory and high energy physics. Moreover, as we will see in the next section, much of the terminology has its roots in the applications to quantum field theory.

Note that the second order phase transition occurs in our theory when the coefficient of the quadratic term, vanishes: $\mu^2 = 0$. From the perspective of quantum field theory, a second order phase transition describes massless particles.

Given that the similarities are so striking, one could ask if there are any differences between statistical and quantum field theories. The answer is yes: there are some quantum field theories which, upon Wick rotation, do not have real Euclidean actions. Perhaps the simplest example is Maxwell (or Yang-Mills) theory, with the addition of a "theta term", proportional to $\epsilon^{\mu\nu\rho\sigma}F_{\mu\nu}F_{\rho\sigma}$. This gives rise to subtle effects in the quantum theory. However, because it contains a single time derivative, it becomes imaginary in the τ variable (2.31) and, correspondingly, there is no interpretation of $e^{-S_E[\phi]/\hbar}$ as probabilities in a thermal ensemble.

A Different Perspective on the Lower Critical Dimension

A statistical field theory in $d = 1$ spatial dimensions is related to quantum field theory in $d = 0$ spatial dimensions. But we have a name for this: we call it quantum mechanics.

Viewed in this way, the lower critical dimension becomes something very familiar. Consider the quartic potential $V(x)$ shown in the figure. Classical considerations suggest that there are two ground states, one for each of the minima. But we know that this is not the way things work in quantum mechanics. Instead, there is a unique ground state in which the wavefunction has support in both minima, but with the expectation value $\langle x \rangle = 0$. Indeed, the domain wall calculation that we described in Section 1.3.3 is the same calculation that one uses to describe quantum tunnelling using the path integral.

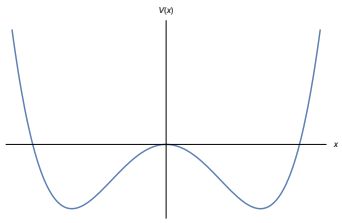

Figure 21:

Dressed in fancy language, we could say that quantum tunnelling means that the $\mathbf{Z}_2$ symmetry cannot be spontaneously broken in quantum mechanics. This translates to the statement that there are no phase transitions in $d = 1$ statistical field theories.

3. The Renormalisation Group

We've built up the technology of field theory and path integrals, and I've promised you that this is sufficient to understand what happens at a second order phase transition. But so far, we've made little headway. All we've seen is that as we approach the critical point, fluctuations dominate and the Gaussian path integral is no longer a good starting point. We need to take the interactions into account.

Sometimes in physics, you can understand a phenomenon just by jumping in and doing the right calculation. And we will shortly do this, using perturbation theory to understand how the ϕ^4 terms change the critical exponents. However, to really understand second order phase transitions requires something more: we will need to set up the right framework in which to think of physics at various length scales. This set of ideas was developed in the 1960s and 1970s, by people like Leo Kadanoff, Michael Fisher and, most importantly, Kenneth Wilson. It goes by the name of the *Renormalisation Group*.

3.1 What's the Big Idea?

Let's start by painting the big picture. As in the previous section, we're going to consider a class of theories based around a single scalar field $\phi(\mathbf{x})$ in d dimensions. (We will consider more general set-ups in Section 4.) The free energy takes the now familiar form,

$$F[\phi] = \int d^d x \ \left[\frac{1}{2} \nabla\phi \cdot \nabla\phi + \frac{1}{2}\mu^2\phi^2 + g\phi^4 + \ldots \right] \tag{3.1}$$

In what follows, we will look at what happens as we approach the critical point from above $T \to T_c$. All the important temperature dependence in (3.1) is sitting in the quadratic term, with

$$\mu^2 \sim T - T_c \tag{3.2}$$

In contrast to the previous section, we will allow μ^2 to take either sign: $\mu^2 > 0$ in the disordered phase, and $\mu^2 < 0$ in the ordered phase where $\langle\phi\rangle \neq 0$.

There is one important change in convention from our earlier discussion: we have rescaled the coefficient of the gradient term to be $1/2$; we will see the relevance of this shortly. All other terms have arbitrary coefficients.

So far we've focussed on just a few couplings, as shown in the free energy (3.1). Here we're going to expand our horizons. We'll consider all possible terms in the free energy, subject to a couple of restrictions. We'll insist that the free energy is analytic around $\phi = 0$, so has a nice Taylor expansion, and we will insist on the $\mathbf{Z}_2$ symmetry $\phi \to -\phi$, so that only even powers of ϕ arise. This means, for example, that we will include the term ϕ^6 and $(\nabla^2\phi)^2$ and ϕ^{14} and $\phi^{137}(\nabla\phi\cdot\nabla\phi)\nabla^2\phi$ and so so. Each of these terms comes with its own coupling constant. However, we don't include terms like ϕ^{17} because this violates the $\mathbf{Z}_2$ symmetry, nor $1/\phi^2$ because this is not analytic at $\phi = 0$.

Next, consider the infinite dimensional space, parameterised by the infinite number of coupling constants. We will call this the *theory space* (although I should warn you that this isn't standard terminology). You should have in your mind something like this:

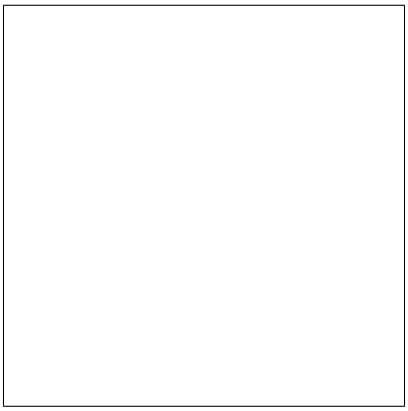

But possibly bigger.

As we've seen, our interest is in computing the partition function

$$Z = \int \mathcal{D}\phi \; e^{-F[\phi]} \tag{3.3}$$

Note that I've written the exponent as e^{-F} rather than $e^{-\beta F}$. This is because the overall power of β does nothing to affect the physics; all the relevant temperature dependence is in the coefficient (3.2) while, for the quantities of interest near the critical point, this overall factor can be set to $\beta \approx 1/T_c$. You can think that we've simply rescaled this into the field ϕ.

There is one more ingredient that we need to make sense of the path integral (3.3). This is the UV cut-off Λ. Recall that, implicit in our construction of the theory is the requirement that the Fourier modes $\phi_{\mathbf{k}}$ vanish for suitably high momenta

$$\phi_{\mathbf{k}} = 0 \quad \text{for } k > \Lambda$$

This arises because, ultimately, our spins sit on some underlying lattice which, in turn, was coarse-grained into boxes of size a. The UV cut-off is given by $\Lambda \sim 1/a$.

Until now, the UV cut-off has taken something of a back seat in our story, although it was needed to render some of the path integral calculations in the previous section finite. Now it's time for Λ to move centre stage. As we will explain, we can use the cut-off to define a *flow* in the space of theories.

Suppose that we only care about physics on long distance scales, L. Then we have no real interest in the Fourier modes $\phi_{\mathbf{k}}$ with $k \gg 1/L$. This suggests that we can write down a different theory, that has a lower cut-off,

$$\Lambda' = \frac{\Lambda}{\zeta}$$

for some $\zeta > 1$. As long as $\Lambda' \gg 1/L$, the scale of interest, our theory can tell us everything that we need to know. Moreover, we know, at least in principle, how to construct such a theory. We write our Fourier modes as

$$\phi_{\mathbf{k}} = \phi^-_{\mathbf{k}} + \phi^+_{\mathbf{k}}$$

where $\phi^-_{\mathbf{k}}$ describe the long-wavelength fluctuations

$$\phi^-_{\mathbf{k}} = \begin{cases} \phi_{\mathbf{k}} & k < \Lambda' \\ 0 & k > \Lambda' \end{cases}$$

and $\phi^+_{\mathbf{k}}$ describe the short-wavelength fluctuations that we don't care about

$$\phi^+_{\mathbf{k}} = \begin{cases} \phi_{\mathbf{k}} & \Lambda' < k < \Lambda \\ 0 & \text{otherwise} \end{cases}$$

There are several other names for these variables that are used interchangeably. The modes $\phi^-_{\mathbf{k}}$ and $\phi^+_{\mathbf{k}}$ are also referred to as low- and high-energy modes or, importing language from quantum mechanics, slow and fast modes, respectively. In a rather quaint nod to the electromagnetic spectrum, the short-distance, microscopic physics that we care little about is often called the *ultra-violet*; the long-distance physics that we would like to understand is the *infra-red*.

Similarly, we decompose the free energy, written in Fourier space, as

$$F[\phi_{\mathbf{k}}] = F_0[\phi^-_{\mathbf{k}}] + F_0[\phi^+_{\mathbf{k}}] + F_I[\phi^-_{\mathbf{k}}, \phi^+_{\mathbf{k}}]$$

Here $F_I[\phi^-_{\mathbf{k}}, \phi^+_{\mathbf{k}}]$ involves the terms which mix the short and long-wavelength modes. The partition function (3.3) can then be written as

$$Z = \int \prod_{k<\Lambda} d\phi_{\mathbf{k}}\ e^{-F} = \int \prod_{k<\Lambda'} d\phi^-_{\mathbf{k}}\ e^{-F_0[\phi^-_{\mathbf{k}}]} \int \prod_{\Lambda'<k<\Lambda} d\phi^+_{\mathbf{k}}\ e^{-F_0[\phi^+_{\mathbf{k}}]} e^{-F_I[\phi^-_{\mathbf{k}}, \phi^+_{\mathbf{k}}]}$$

We write this as

$$Z = \int \mathcal{D}\phi^-\ e^{-F'[\phi^-]}$$

where $F'[\phi]$ is known as the *Wilsonian effective free energy*. (In fairness, this term is rarely used: you're more likely to hear "Wilsonian effective action" to describe the analogous object in a path integral describing a quantum field theory.) We're left with a free energy which describes the long-wavelength modes, but takes into account the effects of the short wavelength modes. It is defined by

$$e^{-F'[\phi^-]} = e^{-F_0[\phi^-_{\mathbf{k}}]} \int \prod_{\Lambda'<k<\Lambda} d\phi^+_{\mathbf{k}}\ e^{-F_0[\phi^+_{\mathbf{k}}]} e^{-F_I[\phi^-_{\mathbf{k}}, \phi^+_{\mathbf{k}}]} \tag{3.4}$$

In subsequent sections, we'll put some effort into calculating this object. However, at the end of the day, the new free energy $F'[\phi^-]$ must take the same functional form as the original free energy (3.1), simply because we started from the most general form possible. The only effect of integrating out the high-momentum modes is to shift the infinite number of coupling constants, so we now have

$$F'[\phi] = \int d^d x\ \left[\frac{1}{2}\gamma' \nabla\phi \cdot \nabla\phi + \frac{1}{2}\mu'^2 \phi^2 + g'\phi^4 + \ldots\right] \tag{3.5}$$

We would like to compare the new free energy (3.5) with the original (3.1). However, we're not quite there yet because the two theories are different types of objects – like apples and oranges – and shouldn't be directly compared. This is because the theory is defined by both the free energy *and* the UV cut-off and, by construction, our two theories have different cut-offs. This means that the original theory $F[\phi]$ can describe things that the new theory $F'[\phi^-]$ cannot, namely momentum modes above the cut-off Λ'.

It is straightforward to remedy this. We can place the two theories on a level playing field by rescaling the momenta in the new theory. We define

$$\mathbf{k}' = \zeta \mathbf{k}$$

Now k' takes values up to Λ, as did k in the original theory. The counterpart of this scaling in real space is

$$\mathbf{x}' = \frac{\mathbf{x}}{\zeta}$$

This means that all lengths scales are getting smaller. You can think of this step as zooming out, to observe the system on larger and larger length scales. As you do so, all features become smaller.

There is one final step that we should take. The new theory $F'[\phi]$ will typically have some coefficient $\gamma' \neq 1$ in front of the leading, quadratic gradient term. To compare with the original free energy (3.1), we should rescale our field. We define

$$\phi'_{\mathbf{k}'} = \sqrt{\gamma'}\,\phi^-_{\mathbf{k}}$$

Which, in position space, reads

$$\phi'(\mathbf{x}') = \sqrt{\gamma'}\,\phi^-(\mathbf{x}) \tag{3.6}$$

Now, finally, our free energy takes the form

$$F_\zeta[\phi'] = \int d^dx\ \left[\frac{1}{2}\nabla\phi'\cdot\nabla\phi' + \frac{1}{2}\mu^2(\zeta)\phi'^2 + g(\zeta)\phi'^4 + \ldots\right] \tag{3.7}$$

We see that this procedure induces a continuous map from $\zeta \in [1,\infty)$ onto the space of coupling constants. Our original coupling constants in (3.1) are those evaluated at $\zeta = 1$. As we increase ζ, we trace out curves in our theory space, that look something like the picture shown in Figure 22.

We say that the coupling constants *flow*, where the direction of the flow is telling us what the couplings look like on longer and longer length scales. The equations which describe these flows – which we will derive shortly – are known, for historic reasons, as *beta functions*.

These, then, are the three steps of what is known as the *renormalisation group* (RG):

- Integrate out high momentum modes, $\Lambda/\zeta < k < \Lambda$.
- Rescale the momenta $\mathbf{k}' = \zeta\mathbf{k}$.
- Rescale the fields so that the gradient term remains canonically normalised.

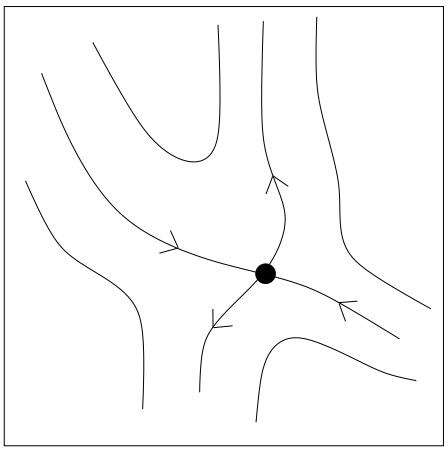

Figure 22: Flows in theory space; the arrows are in the direction of increasing ζ.

You may wonder why we didn't just include a coupling constant $\gamma(\zeta)$ for the gradient term, and watch that change too. The reason is that we can always scale this away by redefining ϕ. But ϕ is just a dummy variable which is integrated over the path integral, so this rescaling can't change the physics. To remove this ambiguity, we should pin down the value of one of the coupling constants, and the gradient term $(\nabla\phi)^2$ is the most convenient choice. If we ever find ourselves in a situation where $\gamma(\zeta) = 0$ for some ζ then we would have to re-evaluate this choice. (We'll actually come across an example where it's sensible to make a different choice in Section 4.3.)

The "renormalisation group" is not a great name. It has a hint of a group structure, because a scaling by ζ_1 followed by a scaling by ζ_2 gives the same result as a scaling by $\zeta_1\zeta_2$. However, unlike for groups, there is no inverse: we can only integrate out fields, we can't put them back in. A more accurate name would be the "renormalisation semi-group".

The Renormalisation Group in Real Space

The procedure we've described above is the renormalisation group in momentum space: to get an increasingly coarse-grained description of the physics, we integrate out successive momentum shells. This version of the renormalisation group is most useful when dealing with continuous fields and will be the approach we will focus on in this course.

There is a somewhat different, although ultimately equivalent, phrasing of the renormalisation group which works directly in real space. This approach works best when dealing directly with lattice systems, like the Ising model. As we explained rather

briefly in Section 1.3, one constructs a magnetisation field $m(\mathbf{x})$ by coarse-graining over boxes of size a, each of which contains many lattice sites. One can ask how the free energy changes as we increase a, a procedure known as *blocking*. Ultimately this leads to the same picture that we built above.

3.1.1 Universality Explained

Even before we do any calculations, there are general lessons to be extracted from the framework above. Let's suppose we start from some point in theory space. This can be arbitrarily complicated, reflecting the fact that it contains information about all the microscopic, short-distance degrees of freedom.

Of course, we care little about most of these details so, in an attempt to simplify our lives, we perform a renormalisation group transformation, integrating out short distance degrees of freedom to generate a new theory which describes the long wavelength physics. And then we do this again. And then we do this again. Where do we end up?

There are essentially two possibilities: we could flow off to infinity in theory space, or we could converge towards a *fixed point*. These are points which are invariant under a renormalisation group transformation. (One could also envisage further possibilities, such as converging towards a limit cycle. It turns out that these can be ruled out in many theories of interest.)

Our interest here lies in the fixed points. The second step in the renormalisation group procedure ensures that fixed points describe theories that have no characteristic scale. If the original theory had a correlation length scale ξ, then the renormalised theory has a length scale $\xi' = \xi/\zeta$. (We will derive this statement explicitly below when we stop talking and start calculating.) Fixed points must therefore have either $\xi = 0$ or $\xi = \infty$.

In the disordered phase, with $T > T_c$, enacting an RG flow reduces the correlation length. Pictorially, we have

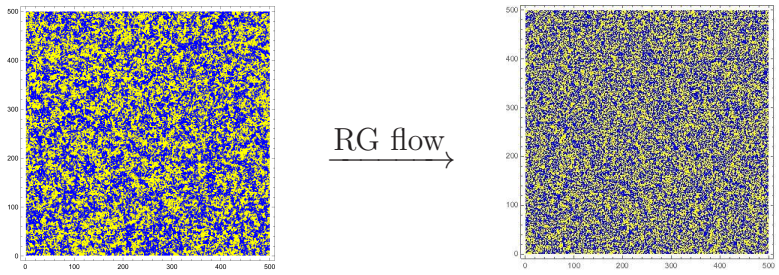

In this case, shrinking the correlation length is equivalent to increasing the temperature. The end point of the RG flow, at $\xi = 0$, is the infinite temperature limit of the theory.

This is rather like flowing off to infinity in theory space. As we will see, it is not uncommon to end up here after an RG flow. But it is boring.

Similarly, in the ordered phase the RG flow again reduces the correlation length,

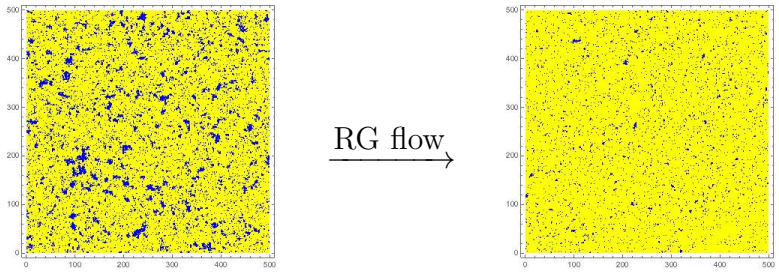

Now the end point at $\xi = 0$ corresponds to the zero temperature limit; again, it is a typical end point of RG flow but is dull.

Theories with $\xi = \infty$ are more interesting. As we saw above, this situation occurs at a critical point where the theory contains fluctuations on all length scales. Now, if we do an RG flow, the theory remains invariant. In terms of our visual configurations,

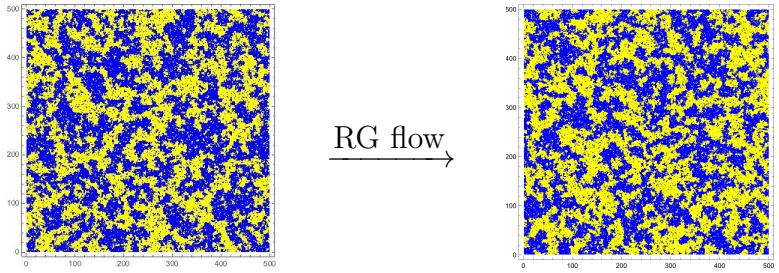

Note that the configuration itself doesn't stay the same. (It is, after all, merely a representative configuration in the ensemble.) However, as the fluctuations on small distance scales shrink away due to RG, they are replaced by fluctuations coming in from larger distance scales. The result is a theory which is scale invariant. For this reason, the term "critical point" is often used as a synonym for "fixed point" of the RG flow.

This picture is all we need to understand the remarkable phenomenon of universality: it arises because many points in theory space flow to the same fixed point. Thus, many different microscopic theories have the same long distance behaviour.

Relevant, Irrelevant or Marginal

It is useful to characterise the properties of fixed points by thinking about the theories in their immediate neighbourhood. Obviously, there are an infinite number of ways we

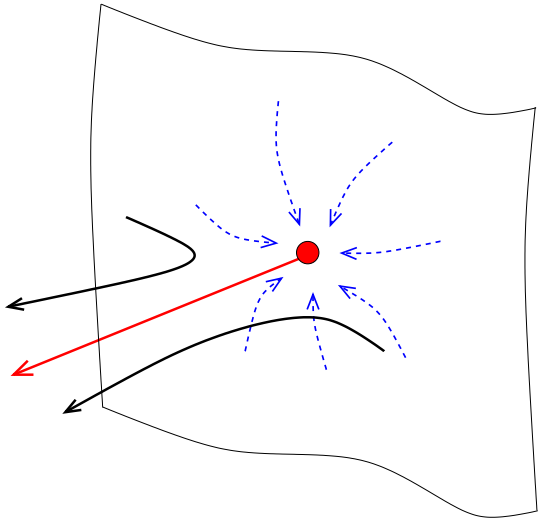

Figure 23: The critical surface. Irrelevant deformations from the critical point are shown as dotted blue lines; the relevant deformation is red.

can move away from the fixed point. If we move in some of these directions, the RG flow will take us back towards the fixed point. These deformations are called *irrelevant* because if we add any such terms to the free energy we will end up describing the same long-distance physics.

In contrast, there will be some directions in which the RG flow will sweep us away from the fixed point. These deformations are called *relevant* because if we add any such terms to the free energy, the long-distance physics will be something rather different. Examples of relevant and irrelevant deformations are shown in Figure 23. Much of the power of universality comes from the realisation that the vast majority of directions are irrelevant. For a given fixed point, there are typically only a handful of relevant deformations, and an infinite number of irrelevant ones. This means that our fixed points have a large basin of attraction, huge slices of the infinite dimensional theory space all converging to the same fixed point. The basin of attraction for a particular fixed point is called the *critical surface*.

Finally, it's possible that our fixed point is not a point at all, but a line or a higher dimensional surface living within theory space. In this case, if we deform the theory in the direction of the line, we will not flow anywhere, but simply end up on another fixed point. Such deformations are called *marginal*; they are rare, but not unheard of.

Why High Energy Physics is Hard

Universality is a wonderful thing if you want to understand the low-energy, long-wavelength physics. It tells you that you can throw away many of the microscopic details because they are irrelevant for the things that you care about.

In contrast, if you want to understand the high-energy, short distance physics then universality is the devil. It tells you that you have very little hope of extracting any information about microscopic degrees of freedom if you only have access to information at long distances. This is because many different microscopic theories will all give the same answer.

As we saw in Section 2.3, quantum field theory is governed by the same mathematical structure as statistical field theory, and the comments above also apply. Suppose, for example, that you find yourself living in a technologically adolescent civilisation that can perform experiments at distance scales of 10^{-16} cm or so, but no smaller. Yet, what you really care about is physics at, say, 10^{-32} cm where you suspect that something interesting is going on. The renormalisation group says that you shouldn't pin your hopes on learning anything from experiment.

The renormalisation group isn't alone in hiding high-energy physics from us. In gravity, cosmic censorship ensures that any high curvature regions are hidden behind horizons of black holes while, in the early universe, inflation washes away any trace of what took place before. Anyone would think there's some kind of conspiracy going on....

3.2 Scaling

The idea that second order phase transitions coincide with fixed points of the renormalisation group is a powerful one. In particular, it provides an organising principle behind the flurry of critical exponents that we met in Section 1.

As we explained above, at a fixed point of the renormalisation group any scale must be washed away. This is already enough to ensure that correlation functions must take the form of a power-law,

$$\langle \phi(\mathbf{x})\phi(0)\rangle \sim \frac{1}{r^{d-2+\eta}} \tag{3.8}$$

Any other function would require a scale on dimensional grounds. The only freedom that we have is in the choice of exponent which we have chosen to parameterise as η. One of the tasks of the RG procedure is to compute η, and we will see how this works in Section 3.5.

However, even here there's something of a mystery because usually we can figure out the way things scale by doing some simple dimensional analysis. (If you would like to refresh your memory, some examples of dimensional analysis can be found in Chapter 3 of the lectures on Dynamics and Relativity.) What does that tell us in the present case?

We will measure dimension in units of inverse length. So, for example, $[x] = -1$ while $[\partial/\partial x] = +1$. The quantity $F[\phi]$ must be dimensionless because it sits in the exponent of the partition function as e^{-F}. The first term is

$$F[\phi] = \int d^d x \; \frac{1}{2} \nabla\phi \cdot \nabla\phi + \ldots$$

From this we learn that

$$[\phi] = \frac{d-2}{2} \tag{3.9}$$

Which, in turn, tells us exactly what the exponent of the correlation function must be: $\eta = 0$.

This is sobering. Dimensional analysis is one of the most basic tools that we have, and yet it seems to be failing at critical points where experiment is showing that $\eta \neq 0$. What's going on?

A better way to think about dimensional analysis is to think in terms of scaling. Suppose that we rescale all length as $\mathbf{x} \to \mathbf{x}' = \mathbf{x}/\zeta$. How should other quantities scale so that all formulae remain invariant? Stated this way, it's clear that there's a close connection between dimensional analysis and RG. The correlation function (3.8) is telling us that we should rescale $\phi(\mathbf{x}) \to \phi'(\mathbf{x}') = \zeta^{\Delta_\phi}\phi(\mathbf{x})$, where

$$\Delta_\phi = \frac{d-2+\eta}{2} \tag{3.10}$$

This is called the *scaling dimension*. It differs from the naive "engineering dimension" $[\phi]$ by the extra term $\eta/2$ which is referred to as the *anomalous dimension*.

We still haven't explained why the scaling dimension differs from engineering dimension. The culprit turns out to be the third step of the RG procedure (3.6) where the field ϕ gets rescaled. In real space, this is viewed as coarse-graining ϕ over blocks of larger and larger size a. As we do so, it dresses ϕ with this UV cut-off scale $\Lambda \sim 1/a$, often in a complicated and non-intuitive way. This means that the correlation function (3.8) is actually

$$\langle \phi(\mathbf{x})\phi(0)\rangle \sim \frac{a^\eta}{r^{d-2+\eta}}$$

which is in full agreement with naive dimensional analysis. We can work with usual engineering dimensions if we keep track of this microscopic distance scale a. But it is much more useful to absorb this into ϕ and think of a coarse-grained observable, with dimension Δ_ϕ, that is appropriate for measuring long distance correlations.

3.2.1 Critical Exponents Revisited

The critical exponents that we met in Section 1.2.3 are all a consequence of scale invariance, and dimensional analysis based on the scaling dimension. Let's see how this arises.

We know that as we move away from the critical point by turning on $\mu^2 \sim T - T_c$, we introduce a new length scale into the problem. This is the correlation length, given by

$$\xi \sim t^{-\nu} \quad \text{with} \quad t = \frac{|T - T_c|}{T_c} \tag{3.11}$$

Here t is called the *reduced temperature*, while ν is another critical exponent that we will ultimately have to calculate. Since ξ is a length scale, it transforms simply as $\xi \to \xi/\zeta$. In other words, it has scaling dimension $\Delta_\xi = -1$. The meaning of the critical exponent ν is that the reduced temperature scales as $t \to \zeta^{\Delta_t} t$, with

$$\Delta_t = \frac{1}{\nu} \tag{3.12}$$

In what follows, our only assumption is that the correlation length ξ is the only length scale that plays any role.

We start with the thermodynamic free energy, $F_{\text{thermo}}(t)$, evaluated at $B = 0$. This takes the form

$$F_{\text{thermo}}(t) = \int d^d x \; f(t)$$

Because F_{thermo} is scale invariant at the fixed point, $f(t)$ must have scaling dimension d, which immediately tells us that

$$f(t) \sim t^{d\nu}$$

There is an intuitive way to understand this. At T close to T_c, the spins are correlated over distance scales ξ, and can be viewed as moving as one coherent block. The free energy F_{thermo} is extensive, and so naturally scales as $F \sim (L/\xi)^d \sim t^{d\nu}$.

From the thermodynamic free energy, we can compute the singular contribution to the heat capacity near $t = 0$. It is

$$c \sim \frac{\partial^2 f}{\partial t^2} \sim t^{d\nu-2} \sim t^{-\alpha}$$

where the second relationship is there to remind us that we already had a name for the critical exponent related to heat capacity. We learn that

$$\alpha = 2 - d\nu \tag{3.13}$$

This is called the *Josephson relation* or, alternatively, the *hyperscaling* relation.

The next critical exponent on the list is β. Recall that this relates the magnetisation in the ordered phase – which we used to call m and have now called ϕ – to the temperature as

$$\phi \sim t^\beta$$

But the scaling dimensions of this equation only work if we have

$$\Delta_\phi = \beta\Delta_t \quad \Rightarrow \quad \beta = \nu\Delta_\phi = \frac{(d-2+\eta)\nu}{2} \tag{3.14}$$

The next two critical exponents require us to move away from the critical point by turning on a magnetic field B. This is achieved through the addition of a linear term $\int d^dx\ B\phi$ in the free energy. (We didn't include such a linear term in our previous discussion of RG, but it can be added without changing the essence of the story.) The scaling dimensions of this term must add to zero, giving

$$\Delta_\phi + \Delta_B = d \quad \Rightarrow \quad \Delta_B = \frac{d+2-\eta}{2}$$

Now we can look at the various relationships. The behaviour of the susceptibility near the critical point is

$$\chi = \left.\frac{\partial\phi}{\partial B}\right|_T \sim t^{-\gamma}$$

Once again, the scaling dimensions are enough to fix γ to be

$$\Delta_\phi - \Delta_B = -\frac{\gamma}{\nu} \quad \Rightarrow \quad \gamma = \nu(2-\eta) \tag{3.15}$$

which is sometimes called *Fisher's identity*. Once again, there is an intuitive way to understand this. The meaning of ξ is that the spins are no longer correlated at distances

$r \gg \xi$. This can be seen, for example, in our original formula (2.29). Using our earlier expression (2.28) for the susceptibility, we have

$$\chi \sim \int_0^\xi d^d x \; \frac{1}{r^{d-2+\eta}} \sim \xi^{2-\eta} \sim t^{-\nu(2-\eta)}$$

which again gives $\gamma = \nu(2-\eta)$.

The final critical exponent relates the magnetisation ϕ to the magnetic field B when we sit at the critical temperature $t = 0$. It should come as little surprise by now to learn that this is again fixed by scaling analysis

$$\phi \sim B^{1/\delta} \quad \Rightarrow \quad \delta = \frac{\Delta_B}{\Delta_\phi} = \frac{d+2-\eta}{d-2+\eta} \tag{3.16}$$

We end up with four equations, relating α (3.13), β (3.14), γ (3.15) and δ (3.16) to the critical exponents η and ν. For convenience, let's recall what values we claimed these exponents take:

	α	β	γ	δ	η	ν
MF	$\frac{4-d}{2}$	$\frac{1}{2}$	1	3	0	$\frac{1}{2}$
$d=2$	0	$\frac{1}{8}$	$\frac{7}{4}$	15	$\frac{1}{4}$	1
$d=3$	0.1101	0.3264	1.2371	4.7898	0.0363	0.6300

where we've used the result (2.14), including quadratic fluctuations, for the mean field value of α. We see that the relations are satisfied exactly for $d = 2$ and to within the accuracy stated for $d = 3$. However, there's a wrinkle because they only agree with the mean field values when $d = 4$!

This latter point is an annoying subtlety and will be explained in Section 3.3.2. Our main task is to understand why the mean field values don't agree with experiment when $d < 4$.

3.2.2 The Relevance of Scaling

The kind of dimensional analysis above also determines whether a given interaction is relevant, irrelevant or marginal.

Consider an interaction term $\mathcal{O}(\mathbf{x})$ in the free energy,

$$F[\phi] \sim \int d^d x \; g_{\mathcal{O}} \, \mathcal{O}(\mathbf{x}) \tag{3.17}$$

Here $\mathcal{O}$ can be ϕ^n or $\phi^m(\nabla\phi)^2$ or any of the other infinite possibilities. In a spillover from quantum field theory, the different interaction terms $\mathcal{O}(\mathbf{x})$ are referred to as *operators*.

We're interested in operators which, in the vicinity of a given point, transform simply under RG. Specifically, suppose that, under the rescaling $\mathbf{x} \to \mathbf{x}' = \mathbf{x}/\zeta$, the operator has a well defined scaling dimension, transforming as

$$\mathcal{O}(\mathbf{x}) \to \mathcal{O}'(\mathbf{x}') = \zeta^{\Delta_\mathcal{O}} \mathcal{O}(\mathbf{x}) \tag{3.18}$$

You can think of such operators as eigenstates of the RG process. From the free energy (3.17), the scaling dimension of the coupling is

$$\Delta_{g_\mathcal{O}} = d - \Delta_\mathcal{O}$$

Under an RG flow, these couplings scale as $g_\mathcal{O} \to \zeta^{d-\Delta_\mathcal{O}} g_\mathcal{O}$. We can see immediately that $g_\mathcal{O}$ either diverges or vanishes as we push forwards with the RG. Invoking our previous classification, $\mathcal{O}$ is:

- Relevant if $\Delta_\mathcal{O} < d$
- Irrelevant if $\Delta_\mathcal{O} > d$
- Marginal if $\Delta_\mathcal{O} = d$

The tricky part of the story is that it's not always easy to identify the operators $\mathcal{O}$ which have the nice scaling property (3.18). As we'll see in the examples below, these are typically complicated linear combinations of the operators ϕ^n and $\phi^n(\nabla\phi)^2$ and so on.

3.3 The Gaussian Fixed Point

It's now time to start calculating. We will start by sitting at a special point in theory space and enacting the renormalisation group. At this special point, only two quadratic terms are turned on:

$$F_0[\phi] = \int d^d x \ \left[\frac{1}{2}\nabla\phi\cdot\nabla\phi + \frac{1}{2}\mu_0^2\,\phi^2\right] = \int^\Lambda \frac{d^d k}{(2\pi)^d}\,\frac{1}{2}(k^2+\mu_0^2)\phi_{\mathbf{k}}\phi_{-\mathbf{k}} \tag{3.19}$$

where we've added a subscript to the coefficient μ_0^2 in anticipation the fact that this quantity will subsequently change under RG flow.

Because the free energy is quadratic in ϕ, it has the property that there is no mixing between the short and long wavelength modes, and so factorises as

$$F_0[\phi] = F_0[\phi^-] + F_0[\phi^+]$$

Integrating over the short wavelength modes is now easy, and results in an overall constant in the partition function

$$e^{-F'[\phi^-]} = \left[\int \mathcal{D}\phi^+ e^{-F_0[\phi^+]}\right] e^{-F_0[\phi^-]} = \mathcal{N} e^{-F_0[\phi^-]}$$

This constant $\mathcal{N}$ doesn't change any physics; it just drops out when we differentiate $\log Z$ to compute correlation functions. However, we're not yet done with our RG; we still need to do the rescaling

$$\mathbf{k}' = \zeta \mathbf{k} \quad \text{and} \quad \phi'_{\mathbf{k}'} = \zeta^{-w} \phi^-_{\mathbf{k}} \tag{3.20}$$

where w is constant that we will determine. Written in terms of the rescaled momenta, we have

$$\begin{aligned} F_0[\phi^-] &= \int^{\Lambda/\zeta} \frac{d^d k}{(2\pi)^d} \frac{1}{2}(k^2 + \mu_0^2)\phi^-_{\mathbf{k}}\phi^-_{-\mathbf{k}} \\ &= \int^{\Lambda} \frac{d^d k'}{(2\pi)^d} \frac{1}{2\zeta^d}\left(\frac{k'^2}{\zeta^2} + \mu_0^2\right)\zeta^{2w}\phi'_{\mathbf{k}'}\phi'_{-\mathbf{k}'} \end{aligned}$$

We can put this back in the form we started with if we take

$$w = \frac{d+2}{2} \tag{3.21}$$

leaving us with

$$F'_0[\phi'] = \int^{\Lambda} \frac{d^d k}{(2\pi)^d} \frac{1}{2}(k^2 + \mu^2(\zeta))\phi'_{\mathbf{k}}\phi'_{-\mathbf{k}}$$

The only price that we've paid for this is that the coefficient of the quadratic term has become

$$\mu^2(\zeta) = \zeta^2 \mu_0^2 \tag{3.22}$$

This illustrates how the length scales in the problem transform under RG. Recall that the correlation length (2.22) is $\xi^2 \sim 1/\mu^2$. We see that, under an RG procedure,

$$\xi \to \frac{\xi}{\zeta}$$

The fixed points obey $d\mu^2/d\zeta = 0$. As we anticipated previously, there are two of them. The first is $\mu^2 = \infty$. This corresponds to a state which has infinite temperature. It is not where our interest lies. The other fixed point is at $\mu^2 = 0$. This is known as the *Gaussian fixed point*.

3.3.1 In the Vicinity of the Fixed Point

As we mentioned previously, we would like to classify fixed points by thinking about what happens when you sit near them. Do you flow into the fixed point, or get pushed away?

We already have the answer to this question in one direction in coupling space. If we add the term $\mu^2\phi^2$, the scaling (3.22) tells us that μ^2 gets bigger as we flow towards the infra-red. This is an example of a relevant coupling: turning it on pushes us away from the fixed point.

Here is another example: it is simple to repeat the steps above including the term $\alpha_0(\nabla^2\phi)^2$ in the free energy. Upon RG, this coupling flows as $\alpha(\zeta) = \zeta^{-2}\alpha_0$. It is an example of an irrelevant coupling, one which becomes less important as we flow towards the infra-red.

More interesting are the slew of possible couplings of the form

$$F[\phi] = \int d^dx \ \left[\frac{1}{2}\nabla\phi\cdot\nabla\phi + \frac{1}{2}\mu_0^2\phi^2 + \sum_{n=4}^{\infty} g_{0,n}\phi^n\right] \tag{3.23}$$

where, to keep the $\mathbf{Z}_2$ symmetry, we restrict the sum to n even. Here things are a little more subtle because, once we turn these couplings on, the first step of the RG procedure is no longer so simple. Integrating out the short distance modes will shift each of these couplings,

$$g_{0,n} \to g'_n = g_{0,n} + \delta g_n$$

We will learn how to calculate the δg_n in section 3.4. But, for now, let's ignore this effect and concentrate on the second and third parts of the RG procedure, in which we rescale lengths and fields as in (3.20). In this approximation, the operators ϕ^n enjoy the nice scaling property (3.18),

$$\mathbf{x}' = \mathbf{x}/\zeta \quad \text{and} \quad \phi'(\mathbf{x}') = \zeta^{\Delta_\phi}\phi(\mathbf{x})$$

The free energy is then rescaled by

$$F[\phi'] = \int d^dx' \ \zeta^d \left[\frac{1}{2}\zeta^{-2-2\Delta_\phi}\nabla'\phi'\cdot\nabla'\phi' + \frac{1}{2}\mu_0^2\zeta^{-2\Delta_\phi}\phi'^2 + \sum_{n=4}^{\infty} g_{0,n}\zeta^{-n\Delta_\phi}\phi'^n\right]$$

To restore the coefficient of the gradient term, we pick the scaling dimension

$$\Delta_\phi = \frac{d-2}{2}$$

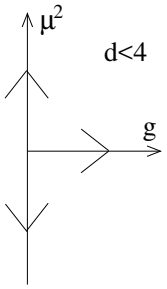

Figure 24: RG flows when $d < 4$

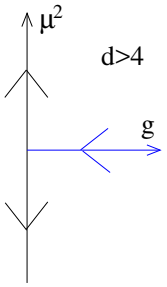

Figure 25: RG flows when $d > 4$

For once, the scaling dimension coincides with the engineering dimension (3.9): $\Delta_\phi = [\phi]$. This is because we're looking at a particularly simple fixed point. Note that this is related to our earlier result (3.21) by $\Delta_\phi = d - w$, with the extra factor of d coming from the $\int d^d k$ in the definition of the Fourier transform.

Our free energy now takes the same form as before,

$$F[\phi'] = \int d^d x' \left[\frac{1}{2} \nabla' \phi' \cdot \nabla' \phi' + \frac{1}{2} \mu^2(\zeta) \phi'^2 + \sum_{n=4}^{\infty} g_n(\zeta) \phi'^n \right]$$

where

$$g_n(\zeta) = \zeta^{d - n\Delta_\phi} g_{0,n} = \zeta^{(1-n/2)d+n} g_{0,n} \tag{3.24}$$

We see that the way these coupling scale depends on the dimension d. For example, the coefficient for ϕ^4 scales as

$$g_4(\zeta) = \zeta^{4-d} g_{0,4}$$

We learn that ϕ^4 is irrelevant for $d > 4$ and is relevant for $d < 4$. According to the analysis above, when $d = 4$, we have $g_4(\zeta) = g_{0,4}$ and the coupling is marginal. In this case, however, we need to work a little harder because the leading contribution to the scaling will come from the corrections δg_4 that we neglected. We'll look at this in the next section.

Restricting to the plane of couplings parameterised by μ^2 and g_4, we see that (if we neglect the interactions) the RG flow near the origin is very different when $d > 4$ and $d < 4$. These are shown in the two figures. In the former case, we need to tune only $\mu^2 \sim T - T_c$ if we want to hit the fixed point; the other couplings will take care of themselves. In contrast, when $d < 4$ both of these couplings are relevant. This means that we would need to tune both to zero if we want to hit the Gaussian fixed point.

We can tally this with our discussion in Section 3.2. The fact that the scaling dimension Δ_ϕ coincides with the naive engineering dimension $[\phi]$ immediately tells us that $\eta = 0$. Meanwhile, the scaling of $\mu^2 \sim t$ is given by $\Delta_t = [g_2] = 2$, which tells us that $\nu = 1/2$. From this we can use (3.13) – (3.16) to extract the remaining critical exponents. These agree with mean field for $d = 4$, but not for $d < 4$. (We will address the situation in $d > 4$ in Section 3.3.2.)

It is no coincidence that this behaviour switches at $d = 4$, which we previously identified as the upper critical dimension. In an experiment, one can always change μ^2 by varying the temperature. However, one may not have control over the ϕ^4 couplings which typically correspond to some complicated microscopic property of the system. If ϕ^4 is irrelevant, we don't care: the system will drive itself to the Gaussian fixed point. In contrast, if ϕ^4 is relevant the system will drive itself elsewhere. This is why we don't measure mean field values for the critical exponents: these are the critical exponents of the Gaussian fixed point.

The coupling for ϕ^6 scales as

$$g_6(\zeta) = \zeta^{6-2d} g_{0,6}$$

This is irrelevant in $d > 3$, relevant in $d < 3$ and, naively, marginal in $d = 3$.

Note that in dimension $d = 2$ all of the couplings $g_n\phi^n$ are relevant.

So far, this all looks rather trivial. However, things become much more interesting at other fixed points. For example, around most fixed points $\Delta_{\phi^n} \neq n\Delta_\phi$. Indeed, around most fixed points neither ϕ nor ϕ^n will have well-defined scaling dimension; instead, those operators to which one can assign a scaling dimension consist of some complicated linear combination of the ϕ^n. We will start to understand this better in Section 3.4.

The Meaning of Mean Field

The meaning of the phrase "mean field theory" has evolved as these lectures have progressed. We started in Section 1.1.2 by introducing mean field as a somewhat dodgy approximation to the partition function. Subsequently, we used the expression "mean field theory" to mean writing down a free energy $F[\phi]$ and focussing on the saddle point equations. This saddle point is a good approximation to the partition function only when the couplings are small; this is true only in the vicinity of the Gaussian fixed point. For this reason, using mean field theory is usually synonymous with working at the Gaussian fixed point, and ignoring the effect of operators like ϕ^4 on fluctuations.

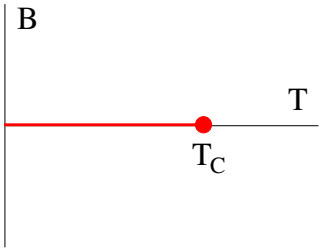

Figure 26: The phase diagram of the Ising model.

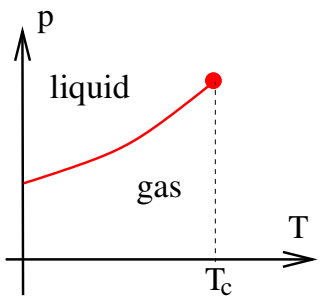

Figure 27: The phase diagram of the liquid-gas system.

(Of course, mean field still retains the ϕ^4 term in the ordered phase, where it is needed to stabilise the potential.)

Interactions that Break $\mathbf{Z}_2$ Symmetry

Until now, we have have restricted ourselves to interactions ϕ^n with n even, to zealously safeguard the $\mathbf{Z}_2$ symmetry $\phi \to -\phi$. One particularly nice aspect of RG is that if we restrict ourselves to a class of free energies that obey a certain symmetry, then we will remain in that class under RG. We'll see examples of this in Section 3.4.

However, suppose that we sit outside of this class and turn on interactions ϕ^n with n odd. The leading order effect is the magnetic field $B\phi$ that we included in our original Ising model. This is always a relevant interaction. This means that if we want to hit the critical point, we must tune this to zero.

It may be more natural to tune $B = 0$ in some systems than others. For example, a magnet in the Ising class automatically has $B = 0$ unless you choose to submit it to a background magnetic field. This means that it's easy to hit the critical point: just heat up a magnet and it will exhibit a second order phase transition.

In contrast, in the liquid gas system, setting "$B = 0$" is less natural. Unlike in the Ising model, there is no $\mathbf{Z}_2$ symmetry manifest in the microscopic physics of gases. Instead, it is an emergent symmetry which relates the density of liquid and gas states at the phase transition. Correspondingly, if we simply take a liquid and heat it up then we're most likely to encounter a first order transition, or no transition at all. If we want to hit the critical point, we must now tune the two relevant operators: temperature μ^2 and pressure, which corresponds to the linear term with coefficient B.

In both situations above we really need to tune two relevant couplings to zero to hit the critical point. Of these, one is even under $\mathbf{Z}_2$ and one is odd under $\mathbf{Z}_2$. Doing this

will allow us to hit a fixed point with two relevant deformations, one even one odd. This is the Gaussian fixed point in $d > 4$ and is something else (to be described below) in $d < 4$.

What about higher order interactions ϕ^n with n odd. If we have to tune ϕ, do we not also need to tune ϕ^3? It turns out that the ϕ^3 interaction is redundant. If you have a free energy with no $\mathbf{Z}_2$ symmetry, and all powers of ϕ^n, then you can always redefine your field as $\phi \to \phi + c$ for some constant c. This freedom allows you to eliminate the ϕ^3 term. Note that if your free energy enjoys the $\mathbf{Z}_2$ symmetry $\phi \to -\phi$ then it prohibits you from making this shift.

3.3.2 Dangerously Irrelevant

We've learned that the ϕ^4 interaction is irrelevant for $d > 4$, and so one can hit the Gaussian fixed point by tuning just one parameter: $\mu^2 = 0$.

However, there's one tricky issue that we haven't yet explained: the mean field exponents agree with the scaling analysis of Section 3.2 only when $d = 4$. Comparing the two results, we have

	α	β	γ	δ	η	ν
MF	$\frac{4-d}{2}$	$\frac{1}{2}$	1	3	0	$\frac{1}{2}$
Scaling	$\frac{4-d}{2}$	$\frac{d-2}{4}$	1	$\frac{d+2}{d-2}$	0	$\frac{1}{2}$

where we've used the result (2.14), including quadratic fluctuations, for the mean field value of α. This agrees with the scaling analysis. However, for $d > 4$, the exponents β and δ differ. It turns out that the results from Landau mean field are correct, and those from the scaling analysis are wrong. Why?

To understand this, let's recall our scaling argument from Section 3.2. We set $B = 0$ and focus on the critical exponent β. The magnetisation scales with the temperature $t = |T - T_c|/T_c$ as

$$m \sim t^\beta$$

Here m is identified with the scalar field ϕ. Scaling analysis gives $\Delta_\phi = \beta\Delta_t$. But both mean field and scaling analysis agree that $\Delta_t = 1/\nu = 2$ and $\Delta_\phi = (d-2)/2$, and this gives $\beta = (d-2)/4$, rather than the mean field result $\beta = 1/2$.

However, we were a little quick in the scaling analysis because we neglected the quartic coupling g_4. Mean field really told us (1.31),

$$m_0 \sim \left(\frac{t}{g_4}\right)^{1/2}$$

But both t and g_4 scale under RG flow. The scaling dimension of g_4 is $\Delta_{g_4} = 4 - d$ and now the mean field result, with $\beta = 1/2$ is fully compatible with scaling.

There's a more general lesson to take from this. It is tempting, when doing RG, to think that we can just neglect the irrelevant operators because their coefficients flow to zero as we approach the infra-red. However, sometimes we will be interested in quantities – such as the magnetisation above – which have the irrelevant coupling constants sitting in the denominator. In this case, one cannot just blindly ignore these irrelevant couplings as they affect the scaling analysis. When this happens, the irrelevant coupling is referred to as *dangerously irrelevant*.

3.3.3 An Aside: The Emergence of Rotational Symmetry

This is a good point to revisit an issue that we previously swept under the rug. We started our discussion with a lattice model, but very quickly moved to the continuum, field theory. Along the way we stated, without proof, that we expect the long distance physics to enjoy rotational invariance and we restricted our attention to field theories with this property. Why are we allowed to do this?

To make the discussion concrete, consider a square lattice in $d = 2$ dimensions. This has a discrete $\mathbf{Z}_4$ rotational symmetry, together with a $\mathbf{Z}_2$ reflection symmetry. These combine together into the dihedral group D_8. (More precisely, and more annoyingly, they combine into what group theorists call D_8 and what many other mathematicians and physicists call D_4.)

Our field theory description will respect the D_8 symmetry of the underlying lattice model, together with the $\mathbf{Z}_2$ symmetry $\phi \to -\phi$ which ensures that fields come in pairs. But this would appear to be much less powerful than the full $O(2)$ continuous rotation and reflection symmetry. Have we cheated?

Let's see what kind of terms we might expect. First, there are some simple terms that are prohibited by the dihedral symmetry. For example, a lone term $(\partial_1\phi)^2$ would break the $x_1 \to x_2$ discrete rotational symmetry and so would not appear in the free energy. Similarly, a term $\phi\partial_1\phi$ breaks the $x_1 \to -x_1$ symmetry. (On top of this, it is also a total derivative and so doesn't contribute to the free energy.) The lowest

dimension term that includes derivatives and is compatible with the discrete symmetry is

$$\mathcal{O}_2 \sim (\partial_1\phi)^2 + (\partial_2\phi)^2$$

But this term happens to be invariant under the full, continuous $O(2)$ rotational symmetry. We should keep going. The first term that preserves D_8, but not $SO(2)$, is

$$\mathcal{O}_4 \sim \phi\partial_1^4\phi + \phi\partial_2^4\phi$$

There is no reason not to add such terms to the free energy and, in general, we expect that these will be present in any field theoretic description that accurately describes the microscopic physics. However, this operator has dimension $\Delta_{\mathcal{O}_4} = d + 2$ and so is irrelevant. This means that it gets washed away by the renormalisation group, and the long wavelength physics is invariant under the full $O(2)$ symmetry. We say that the continuous rotational symmetry is *emergent*. Sometimes it is referred to as an *accidental symmetry*. A similar argument holds for higher dimensions.

3.4 RG with Interactions

The previous section left two questions hanging. What happens to the renormalisation of the coupling g_4 in $d = 4$ dimensions? And where does the flow of g_4 take us in $d < 4$ dimensions? In this section we will answer the first of these. In Section 3.5 we will see that our analysis also contains the answer to the second.

We now repeat our RG procedure, but with a different starting point in theory space,

$$F[\phi] = \int d^dx \ \left[\frac{1}{2}\nabla\phi\cdot\nabla\phi + \frac{1}{2}\mu_0^2\phi^2 + g_0\phi^4\right]$$

The renormalisation group procedure tells us to split the Fourier modes of the field into long and short wavelengths,

$$\phi_{\mathbf{k}} = \phi^-_{\mathbf{k}} + \phi^+_{\mathbf{k}} \tag{3.25}$$

and write the free energy as

$$F[\phi] = F_0[\phi^-] + F_0[\phi^+] + F_I[\phi^-, \phi^+]$$

where we take $F_0[\phi]$ to coincide with the quadratic terms (3.19), and the interaction terms are

$$F_I[\phi] = \int d^dx \ g_0\phi^4$$

Note that we've chosen to include, for example, $(\phi^-)^4$ in the interaction terms rather than F_0. This is a matter of convention.

The effective free energy for $\phi^-_{\mathbf{k}}$, defined in (3.4), is given by

$$e^{-F'[\phi^-]} = e^{-F_0[\phi^-_{\mathbf{k}}]} \int \mathcal{D}\phi^+ \; e^{-F_0[\phi^+_{\mathbf{k}}]} e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]}$$

There is a nice interpretation of this functional integral $\int \mathcal{D}\phi^+_{\mathbf{k}}$. We can think of it as computing the expectation value of $e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]}$, treating $\phi^+_{\mathbf{k}}$ as the random variable with Gaussian distribution $e^{-F_0[\phi^+_{\mathbf{k}}]}$. In other words, we can write this as

$$e^{-F'[\phi^-]} = e^{-F_0[\phi^-_{\mathbf{k}}]} \left\langle e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]} \right\rangle_+$$

where the subscript on $\langle\cdot\rangle_+$ is there to remind us that we are averaging over the $\phi^+_{\mathbf{k}}$ modes only. We take the definition of the path integral to be suitably normalised so that $\langle 1\rangle_+ = 1$. Taking the log of both sides,

$$F'[\phi^-] = F_0[\phi^-_{\mathbf{k}}] - \log \left\langle e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]} \right\rangle_+ \tag{3.26}$$

Our task is to compute this.

We can't do this functional integral exactly. Instead, we resort to perturbation theory. We assume that g_0 is suitably small, and expand. (The dimensionless small parameter is $g_0\mu_0^{d-4}$.) We first Taylor expand the exponential,

$$\log \left\langle e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]} \right\rangle_+ = \log \left\langle 1 - F_I + \frac{1}{2}F_I^2 + \ldots \right\rangle_+$$

and we then Taylor expand $\log(1+x)$, to get

$$\log \left\langle e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]} \right\rangle_+ = -\langle F_I\rangle_+ + \frac{1}{2}\left[\langle F_I^2\rangle_+ - \langle F_I\rangle_+^2\right] + \ldots \tag{3.27}$$

where, in general, the n^{th} term is $(-1)^n \times n^{\text{th}}$ cumulant of F_I. This also follows from the same kind of manipulations that we did at the beginning of Section 2.2. We will deal with each of terms above in turn.

3.4.1 Order g_0

At leading order in g_0, we need to compute $\left\langle F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]\right\rangle_+$. The first order of business is to expand the interaction terms (3.26) in Fourier modes. We have

$$F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}] = g_0 \int \prod_{i=1}^4 \frac{d^dk_i}{(2\pi)^d} \times (\text{terms with } \phi) \times (2\pi)^d\delta^d(\sum_i \mathbf{k}_i)$$

There are five different "terms with ϕ", most of which do not give anything interesting. These five terms are:

i) $\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\phi^-_{\mathbf{k}_3}\phi^-_{\mathbf{k}_4}$: This term doesn't include any $\phi^+_{\mathbf{k}}$, the average is trivial. It carries over to give the term $g_0\int d^dx\ \phi^{-4}$ in the effective free energy.

ii) $4\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\phi^-_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}$: This term has just a single $\phi^+_{\mathbf{k}}$ and so vanishes when averaged over the Gaussian ensemble.

iii) $6\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\phi^+_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}$: This term is interesting. We will look at it more closely below. For now, note that the factor of 6 comes from the different choices of momentum labels.

iv) $4\phi^-_{\mathbf{k}_1}\phi^+_{\mathbf{k}_2}\phi^+_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}$: This term is cubic in $\phi^+_{\mathbf{k}}$ and, like any term with an odd number of insertions, vanishes when averaged over the Gaussian ensemble.

v) $\phi^+_{\mathbf{k}_1}\phi^+_{\mathbf{k}_2}\phi^+_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}$: This term doesn't include any $\phi^-_{\mathbf{k}}$, it simply gives a constant to the free energy. It will not be important here.

We learn that we need to compute just a single term,

$$\langle F_I\rangle_+ = 6g_0\int\prod_{i=1}^4\frac{d^dk_i}{(2\pi)^d}\,\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\times\left\langle\phi^+_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}\right\rangle_+\times(2\pi)^d\delta^d(\sum_i\mathbf{k}_i) \tag{3.28}$$

But this is the same kind of correlation function that we computed in Section 2.2: it is given by

$$\left\langle\phi^+_{\mathbf{k}}\phi^+_{\mathbf{k}'}\right\rangle_+ = (2\pi)^d\delta^d(\mathbf{k}+\mathbf{k}')\,G_0(k)\quad\text{with}\quad G_0(k)=\frac{1}{k^2+\mu_0^2} \tag{3.29}$$

After playing around with the delta-functions, and relabelling momentum variables, we're left with our first correction to the free energy,

$$\langle F_I\rangle_+ = 6g_0\int_0^{\Lambda/\zeta}\frac{d^dk}{(2\pi)^d}\,\phi^-_{\mathbf{k}}\phi^-_{-\mathbf{k}}\int_{\Lambda/\zeta}^{\Lambda}\frac{d^dq}{(2\pi)^d}\frac{1}{q^2+\mu_0^2}$$

where the limits on the $\int dq$ integral reflect the fact that we've only integrated out the short wavelength modes, whose momenta lie within this band. We see that, at order g_0, we get a correction only to the quadratic term whose coefficient becomes

$$\mu_0^2\to\mu'^2=\mu_0^2+12g_0\int_{\Lambda/\zeta}^{\Lambda}\frac{d^dq}{(2\pi)^d}\frac{1}{q^2+\mu_0^2} \tag{3.30}$$

Finally, we should enact the rescaling steps of the renormalisation group. This takes the same form as before (3.20),

$$\mathbf{k}'=\zeta\mathbf{k}\quad\text{and}\quad\phi'_{\mathbf{k}}=\zeta^{-w}\phi^-_{\mathbf{k}/\zeta}\quad\text{with}\quad w=\frac{d+2}{2}$$

This gives the same scaling of the parameters that we saw in Section 3.3. We have

$$\mu^2(\zeta) = \zeta^2 \left(\mu_0^2 + 12 g_0 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^d q}{(2\pi)^d} \frac{1}{q^2 + \mu_0^2} \right) \quad \text{and} \quad g(\zeta) = \zeta^{4-d} g_0 \tag{3.31}$$

The upshot of this calculation is that turning on a ϕ^4 coupling will give rise to a quadratic ϕ^2 coupling under RG flow. This is typical of these kinds of calculations: couplings of one type will induce others.

The coefficient of the ϕ^2 term is particularly important for our story, since the critical point is defined to be the place where this vanishes. We see that it's not so easy to make this happen. You can't simply set $\mu_0^2 = 0$ at some high scale and expect to hit criticality. Indeed, the result (3.31) tells us that, at long wavelengths, the "natural" value is $\mu^2 \sim g_0 \Lambda^{d-2}$, which is typically large. If you want to hit the critical point, you must "fine tune" the original coefficient μ_0^2 to cancel the new terms that are generated by RG flow.

You might think that this calculation answers the question of what happens to the theory in $d = 4$ when we turn on $g\phi^4$. It certainly tells us that turning on this coupling will induce the relevant coupling $\mu^2\phi^2$ and so take us away from the Gaussian fixed point. However, a closer look at (3.31) reveals that it's possible to turn on a combination of g_0 and μ_0^2, so that $\mu^2(\zeta)$ remains zero. This combination is a marginal coupling. We learn that, at this order, there remains one relevant and one marginal deformation.

3.4.2 Order g_0^2

The corrections to the ϕ^4 terms first arise at order g_0^2. Here we have the contribution

$$F'[\phi^-] \sim -\frac{1}{2} \left(\langle F_I^2 \rangle - \langle F_I \rangle^2 \right) \tag{3.32}$$

Expanding out $\langle F_I^2 \rangle$, we find 256 different terms. We will see how to organise them shortly, but for now we make a few comments before focussing on the term of interest.

Some of the terms in $\langle F_I^2 \rangle$ will result in corrections that cannot be written as a local free energy, but are instead of the form $\left(\int d^d x f(\phi^-) \right)^2$ for some $f(\phi)$. These terms will be cancelled by the $\langle F_I \rangle^2$ terms. This is a general phenomena which you can learn more about in the lectures on Quantum Field Theory. In terms of Feynman diagrams, which we will introduce below, these kind of terms correspond to disconnected diagrams.

The terms that we care about in $\langle F_I^2 \rangle$ are those which can be written as local corrections to the free energy. Of immediate interest for us are a subset of terms in $F_I^2 \sim \int d^dx\, d^dy\; \phi^4(\mathbf{x})\, \phi^4(\mathbf{y})$, given by

$$\frac{1}{2}\langle F_I^2\rangle_+ \sim \frac{1}{2}\binom{4}{2}^2 g_0^2 \int_0^{\Lambda/\zeta} \prod_{i=1}^4 \left[\frac{d^dk_i}{(2\pi)^d}\; \phi^-_{\mathbf{k}_i}\right] \int_{\Lambda/\zeta}^{\Lambda} \prod_{j=1}^4 \left[\frac{d^dq_j}{(2\pi)^d}\right]\; \langle \phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_2}\phi^+_{\mathbf{q}_3}\phi^+_{\mathbf{q}_4}\rangle_+ \qquad (3.33)$$
$$\times\; (2\pi)^{2d}\delta^d(\mathbf{k}_1+\mathbf{k}_2+\mathbf{q}_1+\mathbf{q}_2)\delta^d(\mathbf{k}_3+\mathbf{k}_4+\mathbf{q}_3+\mathbf{q}_4)$$

Let's explain what's going on here. Each $\phi(\mathbf{x})$ is decomposed into Fourier modes ϕ^- and ϕ^+. The same is true for each $\phi(\mathbf{y})$. In the above term, we have chosen two ϕ^- out of the $\phi^4(\mathbf{x})$ and two ϕ^- out of the $\phi^4(\mathbf{y})$; the remaining terms are ϕ^+. Each combinatoric factor $\binom{4}{2} = 6$ out front reflects the choice of picking two ϕ^- from ϕ^4. Meanwhile, the two delta functions come from doing the $\int d^dx$ and $\int d^dy$ integrals respectively. Matching the momenta in the arguments of the delta functions to the $\phi^\pm$ tells us that we've picked two ϕ^- from the $\phi^4(\mathbf{x})$ and two ϕ^- from $\phi^4(\mathbf{y})$ (as opposed to, say, all four from $\phi^4(\mathbf{y})$).

To proceed, we need to compute the four-point function $\langle \phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_2}\phi^+_{\mathbf{q}_3}\phi^+_{\mathbf{q}_4}\rangle_+$. To do this we need a result known as *Wick's theorem*.

Wick's Theorem

As we proceed in our perturbative expansion, the integrals start to blossom. From the form of the expansion (3.27), we can see that the integrand will involve expectation values of the form $\langle \phi^+_{\mathbf{k}_1} \ldots \phi^+_{\mathbf{k}_l}\rangle_+$. There is a simple way to compute expectation values of this type in Gaussian ensembles. This follows from:

Lemma: Consider n variables ϕ_a drawn from a Gaussian ensemble. This means that, for any function $f(\phi)$, the expectation value is

$$\langle f(\phi)\rangle = \frac{1}{\mathcal{N}}\int_{-\infty}^{\infty} d^n\phi\; f(\phi)\, e^{-\frac{1}{2}\phi\cdot G^{-1}\phi}$$

for some symmetric, invertible, positive-definite $n \times n$ matrix G. The normalisation factor is $\mathcal{N} = \det{}^{1/2}(2\pi G)$ and ensures that $\langle 1\rangle = 1$. The following identity then holds:

$$\left\langle e^{B_a\phi_a}\right\rangle = e^{\frac{1}{2}B_a\langle\phi_a\phi_b\rangle B_b} \qquad (3.34)$$

for any constant B_a.

Proof: This is straightforward to show since we can just evaluate both sides

$$\begin{aligned}\left\langle e^{B_a\phi_a}\right\rangle &= \frac{1}{\mathcal{N}}\int_{-\infty}^{\infty} d^n\phi\; e^{-\frac{1}{2}\phi\cdot G^{-1}\phi + B\cdot\phi} \\ &= \frac{1}{\mathcal{N}}\int_{-\infty}^{\infty} d^n\phi\; e^{-\frac{1}{2}(\phi-GB)\cdot G^{-1}(\phi-GB)} e^{\frac{1}{2}B\cdot GB} = e^{\frac{1}{2}B\cdot GB} = e^{\frac{1}{2}B_a\langle\phi_a\phi_b\rangle B_b}\end{aligned}$$

where, in the last step, we used the fact that $\langle\phi_a\phi_b\rangle = G_{ab}$. □

The Taylor expansion of the identity (3.34) gives us the expressions that we want. The left-hand-side is

$$\left\langle e^{B_a\phi_a}\right\rangle = 1 + B_a\langle\phi_a\rangle + \frac{1}{2}B_aB_b\langle\phi_a\phi_b\rangle + \frac{1}{3!}B_aB_bB_c\langle\phi_a\phi_b\phi_c\rangle + \frac{1}{4!}B_aB_bB_cB_d\langle\phi_a\phi_b\phi_c\phi_d\rangle + \ldots$$

Meanwhile, the right-hand-side is

$$e^{\frac{1}{2}B_a\langle\phi_a\phi_b\rangle B_b} = 1 + \frac{1}{2}B_aB_b\langle\phi_a\phi_b\rangle + \frac{1}{8}B_aB_bB_cB_d\langle\phi_a\phi_b\rangle\langle\phi_c\phi_d\rangle + \ldots$$

Now we just match powers of B_a on both sides. We immediately learn that

$$\langle\phi_{a_1}\ldots\phi_{a_l}\rangle = 0 \quad \text{for } l \text{ odd}$$

Our real interest is in l even. Here we have to be a little careful because multiplying by the string of B's automatically symmetrises the products of $\langle\phi_a\phi_b\rangle$ over the $a = 1, \ldots, n$ indices. So, for example, comparing the B^4 terms gives

$$\langle\phi_a\phi_b\phi_c\phi_d\rangle = \langle\phi_a\phi_b\rangle\langle\phi_c\phi_d\rangle + \langle\phi_a\phi_c\rangle\langle\phi_b\phi_d\rangle + \langle\phi_a\phi_d\rangle\langle\phi_b\phi_c\rangle \tag{3.35}$$

It's not hard to convince yourself that

$$\langle\phi_{a_1}\ldots\phi_{a_{2l}}\rangle = \langle\phi_{a_1}\phi_{a_2}\rangle\ldots\langle\phi_{a_{2l-1}}\phi_{a_{2l}}\rangle + \text{ symmetrised}$$

This leaves us with a sum over all pairwise contractions. This result is known as *Wick's theorem*.

Back to the Free Energy

We can now apply Wick's theorem to our free energy (3.33),

$$\langle\phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_2}\phi^+_{\mathbf{q}_3}\phi^+_{\mathbf{q}_4}\rangle_+ = \langle\phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_2}\rangle_+\langle\phi^+_{\mathbf{q}_3}\phi^+_{\mathbf{q}_4}\rangle_+ + \langle\phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_3}\rangle_+\langle\phi^+_{\mathbf{q}_2}\phi^+_{\mathbf{q}_4}\rangle_+ + \langle\phi^+_{\mathbf{q}_1}\phi^+_{\mathbf{q}_4}\rangle_+\langle\phi^+_{\mathbf{q}_2}\phi^+_{\mathbf{q}_3}\rangle_+$$

Recall that each of these propagators comes with a delta function,

$$\langle\phi^+_{\mathbf{q}}\phi^+_{\mathbf{q}'}\rangle_+ = (2\pi)^d\delta^d(\mathbf{q}+\mathbf{q}')\,G_0(q)$$

The trick is to see how these new delta functions combine with the original delta functions in (3.33). There are two different structures that emerge. The first term, gives (ignoring factors of 2π for now)

$$\int \prod_{j=1}^{4} d^d q_j \ \langle \phi^+_{\mathbf{q}_1} \phi^+_{\mathbf{q}_2} \rangle_+ \langle \phi^+_{\mathbf{q}_3} \phi^+_{\mathbf{q}_4} \rangle_+ \delta^d(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{q}_1 + \mathbf{q}_2)\, \delta^d(\mathbf{k}_3 + \mathbf{k}_4 + \mathbf{q}_3 + \mathbf{q}_4)$$
$$\sim \int d^d q_2 d^d q_4 \ G_0(q_2) G_0(q_4)\, \delta^d(\mathbf{k}_1 + \mathbf{k}_2)\, \delta^d(\mathbf{k}_3 + \mathbf{k}_4) \tag{3.36}$$

We're still left with two delta functions over the $\mathbf{k}$ variables. This means that when we go back to real space, this term does not become a local integral. Instead, if you follow it through, it becomes a double integral of the form $\left(\int d^d x\ \phi^-(\mathbf{x})^2\right)^2$. As we explained after (3.32), these terms are ultimately cancelled by corresponding terms in $\langle F_I \rangle^2$. They are not the terms of interest.

Instead, we care about the second and third terms in the Wick expansion of $\langle \phi^+_{\mathbf{q}_1} \phi^+_{\mathbf{q}_2} \phi^+_{\mathbf{q}_3} \phi^+_{\mathbf{q}_4} \rangle_+$. Each of them gives a contribution of the form

$$\int \prod_{j=1}^{4} d^d q_j \ \langle \phi^+_{\mathbf{q}_1} \phi^+_{\mathbf{q}_3} \rangle_+ \langle \phi^+_{\mathbf{q}_2} \phi^+_{\mathbf{q}_4} \rangle_+ \delta^d(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{q}_1 + \mathbf{q}_2)\, \delta^d(\mathbf{k}_3 + \mathbf{k}_4 + \mathbf{q}_3 + \mathbf{q}_4)$$
$$\sim \int d^d q_1 d^d q_2 \ G_0(q_1)\, G_0(q_2)\, \delta^d(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{q}_1 + \mathbf{q}_2)\, \delta^d(\mathbf{k}_3 + \mathbf{k}_4 - \mathbf{q}_1 - \mathbf{q}_2)$$
$$\sim \int d^d q \ G_0(q)\, G_0(|\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{q}|)\, \delta^d(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 + \mathbf{k}_4) \tag{3.37}$$

where, in going to the last line, we have done the integral $\int d^d q_2$ and relabelled $q_1 = q$. Now we have just a single delta function over $\mathbf{k}_i$ and, correspondingly, when we go back to real space this will give a local contribution to the free energy. Indeed, the terms (3.33) now become

$$\frac{1}{2}\langle F_I^2 \rangle_+ \sim \binom{4}{2}^2 g_0^2 \int_0^{\Lambda/\zeta} \prod_{i=1}^{4} \left[\frac{d^d k_i}{(2\pi)^d}\ \phi^-_{\mathbf{k}_i} \right] f(\mathbf{k}_1 + \mathbf{k}_2)\ (2\pi)^d \delta^d(\sum_i \mathbf{k}_i) \tag{3.38}$$

where the factor of $\frac{1}{2}$ in (3.33) has disappeared because we get two contributions from the Wick expansion, each of which gives the same contribution (3.37). The remaining integral over $\int d^d q$ is hidden in the function $f(\mathbf{k})$, which is given by

$$f(\mathbf{k}_1 + \mathbf{k}_2) = \int_{\Lambda/\zeta}^{\Lambda} \frac{d^d q}{(2\pi)^d}\ \frac{1}{q^2 + \mu_0^2} \frac{1}{(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{q})^2 + \mu_0^2} \tag{3.39}$$

This is not as complicated as it looks. We can write it as

$$f(\mathbf{k}_1 + \mathbf{k}_2) = \int_{\Lambda/\zeta}^{\Lambda} \frac{d^d q}{(2\pi)^d} \frac{1}{(q^2 + \mu_0^2)^2} \left(1 + \mathcal{O}(\mathbf{k}_1, \mathbf{k}_2)\right)$$

All the terms that depend on the external momenta $\mathbf{k}_1$ and $\mathbf{k}_2$ will generate terms in the free energy of the form $k^2(\phi^-)^4 \sim (\phi^-)^2\nabla^2(\phi^-)^2$. These are irrelevant terms that will not be interesting for us other than to note that once we let loose the dogs of RG, we will no longer sit comfortably within some finite dimensional subspace of the coupling constants. Integrating out degrees of freedom generates all possible terms consistent with symmetries; flowing to the IR allows us to focus on the handful of relevant ones.

The contribution (3.38) to the free energy is what we want. Translating back to real space, we learn that the quartic term gets corrections at this order. We have

$$g_0 \to g_0' = g_0 - 36 g_0^2 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^d q}{(2\pi)^d} \frac{1}{(q^2 + \mu_0^2)^2} \tag{3.40}$$

The minus sign in (3.40) is important. It can be traced to the minus sign in (3.32), and it determines the fate of the would-be marginal coupling $g\phi^4$ in $d = 4$ dimensions. Recall that, in $d = 4$, there is no contribution to the running of $g(\zeta)$ from the second and third steps of RG. Here, the leading contribution comes from the first step and, as we see above, this causes $g(\zeta)$ to get smaller as ζ increases. This means that the theory in $d = 4$ is similar in spirit to those in $d > 4$, with ϕ^4 an irrelevant coupling.

However, in $d = 4$, the RG flow for g happens much more slowly than other couplings. For this reason it is sometimes called *marginally irrelevant*, to highlight the fact that it only failed to be marginal when the perturbative corrections were taken into account. This is a general phenomenon: most couplings which naively appear marginal will end up becoming either marginally relevant or marginally irrelevant due to such corrections. In the vast majority of cases, the coupling turns out to be marginally irrelevant. However, there are a number of very important examples – the Kondo effect and non-Abelian gauge theories prominent among them – where a marginal coupling turns relevant. We'll see such an example in Section 4.3.

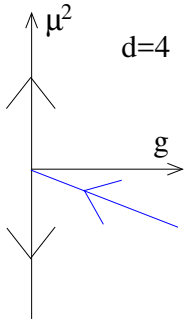

Figure 28:

Finally, just because g is marginally irrelevant in $d = 4$ does not mean that you can turn it on and expect to flow back to the Gaussian fixed point. As depicted in the diagram, the coupling mixes with μ^2. If you want to flow back to the Gaussian fixed point, you need to turn on a particular combination of μ_0^2 and g_0.

3.4.3 Feynman Diagrams

The calculation above needed some care. As we go to higher order terms in the expansion, the number of possibilities starts to blossom. Fortunately, there is a simple graphical formalism to keep track of what's going on.

Suppose that we're interested in a term in the expansion (3.27) of the form $g_0^p(\phi^-)^n(\phi^+)^l$. This can be represented by one or more *Feynman diagrams*. Here is the game:

- Each $\phi^-_{\mathbf{k}}$ is represented by an external, solid line.
- Each $\phi^+_{\mathbf{k}}$ is represented by a dotted line.
- The dotted lines are connected with each other to form internal loops. They are paired up in all possible ways, reflecting the pair contraction of Wick's theorem. No dotted lines can be left hanging which means that the diagrams only make sense for l even.
- Each factor of g_0 is represented as a vertex at which four lines meet.
- Each line has an attached momentum $\mathbf{k}$ which is conserved as we move around the diagram.

Each of these diagrams is really shorthand for an integral. The dictionary is as follows:

- Each internal line corresponds to an insertion of the propagator $\left\langle \phi^+_{\mathbf{k}}\phi^+_{\mathbf{k}'}\right\rangle_+$ defined in (3.29)
- For each internal loop, there is an integral $\int d^dq/(2\pi)^d$.
- Each vertex comes with a power of $g_0(2\pi)^d\delta^d(\sum_i \mathbf{k}_i)$ where the delta function imposes momentum conservation, with the convention that all momenta are incoming.
- There are a bunch of numerical coefficients known as symmetry factors.

This means that a term in the effective action of the form $g_0^p(\phi^-)^n$ will correspond to a diagram with n external lines and p vertices. We can see what these diagrams look like for some of the terms that we've met so far. At order g_0, the rules don't allow us to draw diagrams with an odd number of $\phi^-_{\mathbf{k}}$. The term $\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\phi^-_{\mathbf{k}_3}\phi^-_{\mathbf{k}_4}$ in the expansion gives the trivial contribution to ϕ^4. In diagrams, it is

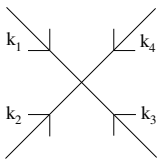

$$= g\int d^dx\ (\phi^-)^4 \tag{3.41}$$

Similarly, the $(\phi^+)^4$ terms are $\phi^+_{\mathbf{k}_1}\phi^+_{\mathbf{k}_2}\phi^+_{\mathbf{k}_3}\phi^+_{\mathbf{k}_4}$ $\sim$  , but because these are internal lines we should join them up to get a diagram that looks like  .

The interesting term at order g_0 is $\phi^-_{\mathbf{k}_1}\phi^-_{\mathbf{k}_2}\phi^+_{\mathbf{q}}\phi^+_{-\mathbf{q}}$, which includes both $\phi^-_{\mathbf{k}}$ and $\phi^+_{\mathbf{k}}$. This was evaluated in (3.28). It is represented by the diagram

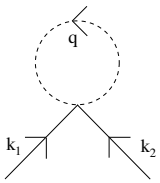

$$= 6g_0 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d}\frac{1}{q^2+\mu_0^2}\int d^dx\ (\phi^-)^2$$

where the integral arises because the momentum $\mathbf{q}$ of the ϕ^+ excitation running in the loop is not determined by momentum conservation of the external legs. Such corrections are said to arise from *loop* diagrams, as opposed to the *tree* diagrams of (3.41).

Finally, at order g_0^2, the correction (3.38) to the ϕ^4 coupling comes from the

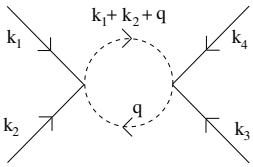

$$= 36g_0^2 \int_0^{\Lambda/\zeta} \prod_{i=1}^{4}\left[\frac{d^dk_i}{(2\pi)^d}\ \phi^-_{\mathbf{k}_i}\right] f(\mathbf{k}_1+\mathbf{k}_2)\ (2\pi)^d\delta^d(\sum_i \mathbf{k}_i)$$

where the function $f(\mathbf{k}_1+\mathbf{k}_2)$ was given in (3.39) and contains two propagators associated to the ϕ^+ fields running in the loop. Evaluating this leads to the correction (3.40), as well as a slew of higher derivative couplings.

Any graph that we can draw will appear somewhere in the expansion of $\log\langle e^{-F_I[\phi^-_{\mathbf{k}},\phi^+_{\mathbf{k}}]}\rangle_+$ given in (3.27). As we noted above, this is a cumulant expansion which has the rather nice graphical interpretation that only *connected* graphs appear in the expansion. For example, at order g_0^2, the disconnected graph that appears in the expansion of $\langle F_I^2\rangle_+$ that looks like 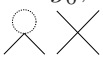 will be cancelled by the same disconnected graph appearing in $-\langle F_I\rangle_+^2$.

In the language of quantum field theory, it's tempting to view the lines in the Feynman diagrams as the worldlines of particles. There is no such interpretation in the present case: they're simply useful.

More Diagrams

We can now also look at other diagrams and see what role they play. For example, you might be worried about the diagram 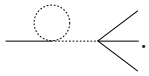 . This is strictly vanishing, because

the incoming momentum for lone ϕ^- leg is forced to be equal to the intermediate momentum for the ϕ^+ propagator. Yet the momentum for ϕ^- and ϕ^+ can never be equal.

There are, however, two further diagrams that we neglected which do have an interesting role to play. Both of these are two loop diagrams. They will not affect what we're going to do later but, nonetheless, are worth highlighting. The first diagram is

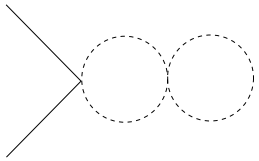

$$= g_0^2 \int \frac{d^dk}{(2\pi)^d}\ \frac{1}{2}C(\Lambda)\phi^-_{\mathbf{k}}\phi^-_{-\mathbf{k}}$$

for some $C(\Lambda)$ whose exact form will not be important. This gives rise to a shift in the quadratic term, so that (3.30) is replaced by

$$\mu_0^2 \to \mu'^2 = \mu_0^2 + 12g_0 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d}\frac{1}{q^2+\mu_0^2} + g_0^2 C(\Lambda)$$

The second diagram is

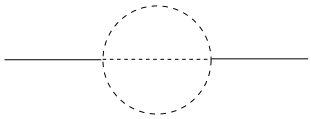

$$= g_0^2 \int \frac{d^dk}{(2\pi)^d}\ \frac{1}{2}A(k,\Lambda)\phi^-_{\mathbf{k}}\phi^-_{-\mathbf{k}} \tag{3.42}$$

We've called the result of this diagram $A(k,\Lambda)$; again we won't need its detailed form. Importantly, it is now a function of the external momentum k. This means that it gives rise to two effects that are (in a technical sense) relevant. The first is that there is yet another renormalisation of μ^2, this time depending on $A(0,\Lambda)$. The second is novel: upon Taylor expanding $A(k) = A(0) + \frac{1}{2}k^2A''(0) + \ldots$, we get a contribution to the the gradient term, which is now

$$F'[\phi] = \int d^dx\ \frac{1}{2}\gamma'\nabla\phi\cdot\nabla\phi + \ldots \qquad \text{with } \gamma' = 1 - 2g_0^2A''(0,\Lambda)$$

This, in turn, means that we need a new rescaling of the field. To order g_0^2, we can write this as

$$\mathbf{k}' = \zeta\mathbf{k} \quad \text{and} \quad \phi'_{\mathbf{k}} = \frac{\zeta^{-(d+2)/2}}{1-g_0^2A''(0,\Lambda)}\,\phi^-_{\mathbf{k}/\zeta}$$

This last, additional step is known as *field renormalisation*. (Actually, that's not completely true. It should be known as "field renormalisation", but instead is known as "wavefunction renormalisation". This a terrible name, one that betrays the long and deep confusion that permeated the origins of this subject. Even in the context of quantum field theory, this rescaling has nothing to do with wavefunctions. It is a rescaling of fields!)

Although we won't compute this field renormalisation exactly, it is nonetheless important for this is what gives rise to the anomalous dimension of ϕ, and this was underlying the whole scaling analysis of Section 3.2.

3.4.4 Beta Functions

It is useful to write down equations which describe the flow of the coupling constants. These are first order differential equations which, for historic reasons, are known as *beta functions*. It turns out to be convenient to parameterise the change in cut-off as

$$\Lambda' = \frac{\Lambda}{\zeta} = \Lambda e^{-s}$$

The renormalisation group transformation described above tells us that each coupling changes with scale, $g_n = g_n(s)$. The beta function is defined as

$$\frac{dg_n}{ds} = \beta_n(g_n)$$

Note that s increases as we flow towards the IR. This means that a positive beta function tells us that g_n gets stronger in the IR, while a negative beta function means that g_n gets weaker in the IR. (As an aside: this is the opposite to how beta functions are sometimes defined in quantum field theory, where one parameterises the flow in terms of energy rather than length.)

Before we jump straight in, it's useful to take a step backwards and build up the beta functions. Let's go back to our original scaling analysis around the Gaussian fixed point (3.24), where the running of the couplings is given by $g_n(s) = e^{(d-nd/2+n)s} g_{0,n}$. The beta functions are

$$\frac{dg_n}{ds} = \left(d - \frac{1}{2}nd + n\right) g_n \tag{3.43}$$

Notice that, at this leading order, there's no mixing between different couplings: turning on one coupling g_n does not induce another to flow. As we saw above, this state of affairs no longer holds when we include interactions.

We now focus on the two most important couplings, μ_0^2 and g_0. At order g_0, the RG equations are given by (3.31); the additional correction at order g_0^2, given in (3.40), means that these get replaced by

$$\mu^2(\zeta) = \zeta^2\left(\mu_0^2 + a g_0\right) \quad \text{and} \quad g(\zeta) = \zeta^{4-d}(g_0 - b g_0^2) \tag{3.44}$$

where

$$a = 12\int_{\Lambda/\zeta}^{\Lambda}\frac{d^dq}{(2\pi)^d}\frac{1}{q^2+\mu_0^2} \quad \text{and} \quad b = 36\int_{\Lambda/\zeta}^{\Lambda}\frac{d^dq}{(2\pi)^d}\frac{1}{(q^2+\mu_0^2)^2}$$

Note that we have kept our original scaling dimensions in (3.44); the corrections in scaling due to the diagram (3.42) will be subleading and not needed in what follows.

When we differentiate $\mu^2(\zeta)$ and $g(\zeta)$ to derive the beta functions, we will get two terms: the first is (3.43) and comes from the scaling; the second comes from the corrections, given by the integrals a and b. Differentiating these integrals is particularly easy. For small s, we write

$$\int_{\Lambda e^{-s}}^{\Lambda} dq\ f(q) \approx [\Lambda - \Lambda e^{-s}]f(\Lambda) \approx \Lambda f(\Lambda)s \quad \Rightarrow \quad \frac{d}{ds}\int_{\Lambda e^{-s}}^{\Lambda} dq\ f(q) = \Lambda f(\Lambda)$$

Let's restrict to $d = 4$ dimensions. The beta function equations are

$$\frac{d\mu^2}{ds} = 2\mu^2 + \frac{3g}{2\pi^2}\frac{\Lambda^4}{\Lambda^2+\mu^2} \quad \text{and} \quad \frac{dg}{ds} = -\frac{9}{2\pi^2}\frac{\Lambda^4}{(\Lambda^2+\mu^2)^2}g^2 \tag{3.45}$$

These don't (yet) contain any new physics, but it's worth reiterating what information we can extract from these equations.

First, the beta function for μ^2 has two terms; the first term comes from the second and third steps in RG (scaling), while the second comes from the first step in RG (integrating out). Meanwhile, the beta function for g has only a single term. There is no term linear in g because it was marginal under scaling, but it does receive a contribution when we integrate out the high momentum modes at order g^2. This contribution is negative, which tells us the coupling is marginally irrelevant. (A repeat of the warning above: this is the opposite convention to quantum field theory where one flows in decreasing energy, rather than increasing length, which means that a marginally irrelevant interaction is usually said to have a positive beta function.)

3.4.5 Again, the Analogy with Quantum Field Theory

The calculations above are very similar to the kind of loop integrals that you do in quantum field theory in $d = 3 + 1$ dimensions. There are, however, some philosophical differences between the approaches.

In statistical mechanics, the field $\phi(\mathbf{x})$ is, by construction, a coarse grained object: at the microscopic level, it dissolves into constituent parts, whether spins or atoms or something else. This has the practical advantage that we have no expectation that the statistical field theory will describe physics on arbitrarily short distance scales. In contrast, when we talk about quantum field theory in the context of high energy physics, it is tempting to think of the fields as "fundamental", a basic building block of our Universe. We may then wish for the theory to make sense down to arbitrarily small distance scales.

This ambition leads to a subtly different viewpoint on renormalisation. In quantum field theory one must introduce a cut-off, as we have above, to render integrals finite. However, this cut-off is very often viewed as an artefact, one which we would ultimately like to get rid of and make sense of the theory as $\Lambda \to \infty$. The trouble is that the renormalised quantities – things that we've called μ^2 and g – typically depend on this cut-off. We saw this, for example, in (3.44). Often this makes it tricky to take the limit $\Lambda \to \infty$.

To avoid this problem, one makes the so-called bare couplings – things we've called μ_0^2 and g_0 – depend on Λ. This is not such a dumb thing to do; after all, these quantities were defined at the cut-off scale Λ. The original game of renormalisation was to find a way to pick $\mu_0^2(\Lambda)$ and $g_0(\Lambda)$ such that all physical quantities remain finite as $\Lambda \to \infty$. It is by no means obvious that this is possible. Theories which can be rendered finite in this way are said to be *renormalisable*.

The high-energy approach to renormalisation predates the statistical physics approach and is now considered rather old-fashioned. The idea that a theory needs to make sense up to arbitrarily high energy scales smacks of hubris. The right way to view renormalisation – whether in statistical mechanics or in high energy physics – is through the renormalisation group procedure that has been our main focus in this chapter, in which one integrates out short wavelength modes to leave an effective long-distance theory.

Nonetheless, the high-energy approach to renormalisation has its advantages. Once one goes beyond the calculations described above, things are substantially easier with a high-energy viewpoint. You will learn more about these issues in the lectures on *Advanced Quantum Field Theory*.

3.5 The Epsilon Expansion

We have learned that ϕ^4 interaction is irrelevant for $d \geq 4$ and relevant for $d < 4$,

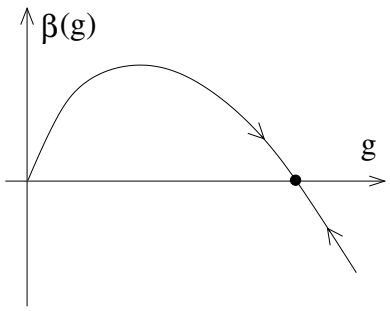

Figure 29: The beta function for g when $d < 4$. The arrows show the flow of g as we move towards the infra-red. The fixed point lies within the regime of perturbation theory when $d = 4 - \epsilon$.

sweeping us away from the Gaussian fixed point. But we seem to be no closer to figuring out where we end up. All we know is that we're not in Kansas anymore.

The difficulty is that we're limited in what we can calculate. We can't do the path integral exactly in the presence of ϕ^4 interactions, and are forced to work perturbatively in the coupling g. Yet, as we have seen, in dimension $d < 4$ the RG flow increases g, taking us to a regime where perturbation theory is no longer valid.

Nonetheless, the calculations that we did above do contain information about where we might expect to end up. But to see it, we have to do something rather dramatic. We will consider our theories not in $d = 4$ dimensions, but in

$$d = 4 - \epsilon$$

dimensions where ϵ is a small number, much less than 1. Clearly, this is an odd thing to do. You could view it as an act of wild creativity or one of utter desperation. Probably it is a little bit of both. But, as we shall see, it will give us the insight we need to understand critical phenomena.

First, we should ask whether it makes sense to work in a non-integer dimension. The lattice models that we started with surely need to be defined in dimension $d \in \mathbf{Z}^+$. Similarly, it was important for us that the free energy is local, meaning that it is written as an integral over space, and this too requires $d \in \mathbf{Z}^+$. However, by the time we get to the beta function equations, it makes mathematical, if not physical, sense to relax this and work in arbitrary $d \in \mathbf{R}$. We can read off these beta functions from the RG

equations (3.44): they are

$$\frac{d\mu^2}{ds} = 2\mu^2 + \frac{12\Omega_{d-1}}{(2\pi)^d}\frac{\Lambda^4}{\Lambda^2+\mu^2}\tilde{g} + \ldots$$
$$\frac{d\tilde{g}}{ds} = \epsilon\tilde{g} - \frac{36\Omega_{d-1}}{(2\pi)^d}\frac{\Lambda^4}{(\Lambda^2+\mu^2)^2}\tilde{g}^2 + \ldots$$

where Ω_{d-1} is the area of the unit sphere $\mathbf{S}^{d-1}$. We've introduced the dimensionless coupling $\tilde{g} = \Lambda^{-\epsilon}g$. Note that the beta function for $\tilde{g}$ now includes a term linear in $\tilde{g}$ arising from the scaling.

It can be checked that the two-loop diagrams we neglected contribute only at order ϵ^2. This means that it's consistent to truncate to the beta functions above. We could use the general formula for the area of a sphere, $\Omega_{d-1} = 2\pi^{d/2}/\Gamma(d/2)$, but this will give corrections of order ϵ^2, so instead we simply use $\Omega_3 = 2\pi^2$. Similar comments apply to $(2\pi)^d$. We're left with

$$\frac{d\mu^2}{ds} \approx 2\mu^2 + \frac{3}{2\pi^2}\frac{\Lambda^4}{\Lambda^2+\mu^2}\tilde{g}$$
$$\frac{d\tilde{g}}{ds} \approx \epsilon\tilde{g} - \frac{9}{2\pi^2}\frac{\Lambda^4}{(\Lambda^2+\mu^2)^2}\tilde{g}^2$$

The novelty of these beta functions is that they have two fixed points. There is the Gaussian fixed point $\mu^2 = \tilde{g} = 0$ that we discussed before. And there is a new fixed point,

$$\mu_\star^2 = -\frac{3}{4\pi^2}\frac{\Lambda^4}{\Lambda^2+\mu_\star^2}\tilde{g}_\star \quad \text{and} \quad \tilde{g}_\star = \frac{2\pi^2}{9}\frac{(\Lambda^2+\mu_\star^2)^2}{\Lambda^4}\epsilon$$

Since we're working to leading order in ϵ, the solution is

$$\mu_\star^2 = -\frac{3}{4\pi^2}\Lambda^2\tilde{g}_\star = -\frac{1}{6}\Lambda^2\epsilon \quad \text{and} \quad \tilde{g}_\star = \frac{2\pi^2}{9}\epsilon$$

This is the *Wilson-Fisher fixed point*. Importantly, when ϵ is small then the fixed point $\tilde{g}_\star$ is also small, so our calculation is self-consistent (although, since we are in a fractional dimension, arguably unphysical!).

3.5.1 The Wilson-Fisher Fixed Point

To understand the flows in the vicinity of the new fixed point, we write $\mu^2 = \mu_\star^2 + \delta\mu^2$ and $\tilde{g} = \tilde{g}_\star + \delta\tilde{g}$. Linearising the beta functions, we find

$$\frac{d}{ds}\begin{pmatrix}\delta\mu^2\\ \delta\tilde{g}\end{pmatrix} = \begin{pmatrix}2-\epsilon/3 & \frac{3}{2\pi^2}\Lambda^2\left(1+\frac{\epsilon}{6}\right)\\ 0 & -\epsilon\end{pmatrix}\begin{pmatrix}\delta\mu^2\\ \delta\tilde{g}\end{pmatrix}$$

where, as with all our other calculations, the entries in the matrix hold only up to $\mathcal{O}(\epsilon^2)$.

The eigenvalues of a triangular matrix coincide with the diagonal entries. We see that this fixed point has one positive and one negative eigenvalue,

$$\Delta_t = 2 - \frac{\epsilon}{3} + \mathcal{O}(\epsilon^2) \quad \text{and} \quad \Delta_g = -\epsilon + \mathcal{O}(\epsilon^2)$$

In other words, the Wilson-Fisher fixed point has one relevant direction and one irrelevant. The flows are shown in the figure.

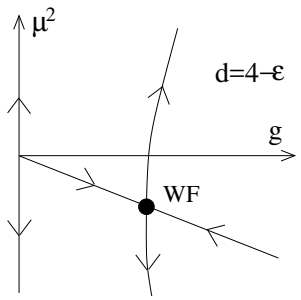

Figure 30:

We see that the epsilon expansion provides us with a global picture of the RG flows in $d < 4$ dimensions. One can check that all other couplings are also irrelevant at the Wilson-Fisher fixed point. Crucially, the fixed point sits at small g where our perturbative analysis is valid.

Now suppose that we increase ϵ. The Wilson-Fisher fixed point moves to higher g, and our perturbative approach breaks down. Nonetheless, it is not unreasonable to suppose that the qualitative picture of the flows remains the same. Indeed, this is thought to happen. Because the Wilson-Fisher fixed point has just a single relevant operator, it means that we will generically end up there if we we're willing to tune just a single parameter, namely $T \to T_c$.

It is now a simple matter to compute the critical exponents in the epsilon expansion. Recall from Section 3.2 that these are related to the scaling dimensions of various terms. The relevant direction away from the Wilson-Fisher fixed point is temperature, $t = |T - T_c|/T_c$. Its dimension is determined by the way it scales as we approach the critical point, $t \to \zeta^{\Delta_t} t = e^{s\Delta_t} t$. But this is precisely the eigenvalue Δ_t that we just computed.

The critical exponent ν, defined by $\xi \sim t^{-\nu}$, is then given by (3.12)

$$\nu = \frac{1}{\Delta_t} = \frac{1}{2} + \frac{\epsilon}{12} + \mathcal{O}(\epsilon^2)$$

We can then use the hyperscaling relation $\alpha = 2 - d\nu$, given in (3.13), to compute the critical exponent for the heat capacity

$$c \sim t^{-\alpha} \quad \text{with} \quad \alpha = \frac{\epsilon}{6} + \mathcal{O}(\epsilon^2)$$

To compute the other critical exponents, we need to evaluate the anomalous dimension η. As we mentioned briefly above, this is related to the diagram  and turns out to be $\eta = \epsilon^2/6$, which is higher order in the expansion. We then have, from (3.10),

$$\Delta_\phi \approx \frac{d-2}{2} = 1 - \frac{\epsilon}{2}$$

Equations (3.14), (3.15) and (3.16) then give

$$\beta \approx \frac{1}{2} - \frac{\epsilon}{6} \quad , \quad \gamma \approx 1 + \frac{\epsilon}{6} \quad , \quad \delta \approx 3 + \epsilon$$

where all expressions hold up to corrections of order $\mathcal{O}(\epsilon^2)$.

Of course, our real interest lies in the system at $d = 3$, or $\epsilon = 1$. It would be in poor taste to simply plug in $\epsilon = 1$. But I know that you're curious. Here's what we get:

	α	β	γ	δ	η	ν
MF	0	$\frac{1}{2}$	1	3	0	$\frac{1}{2}$
$\epsilon = 1$	0.17	0.33	1.17	4	0	0.58
$d = 3$	0.1101	0.3264	1.2371	4.7898	0.0363	0.6300

Our answers are embarrassingly close to the correct values given the dishonest method we used to get there. One can, however, make this approach more respectable. The ϵ expansion has been carried out to order $\mathcal{O}(\epsilon^5)$. It is not a convergent series. Nonetheless, sophisticated resummation techniques can be used to make sense of it, and the resulting expressions give a fairly accurate account of the critical exponents.

The real power of the epsilon expansion, however, is more qualitative than quantitative; it usually – but not always – gives a reliable view of the structure of RG flows.

3.5.2 What Happens in $d = 2$?

We have not yet discussed much about what happens in $d = 2$ dimensions. Here the story is somewhat richer. The first hint of this can be seen in a simple analysis of the Gaussian fixed point, which shows that

$$\Delta_\phi = [\phi] = 0$$

This means that the Gaussian fixed point has an infinite number of relevant deformations since ϕ^n, for each n, is relevant.

It turns out that, in contrast to $d = 3$, there are actually an infinite number of fixed points in $d = 2$. Roughly speaking, the n^{th} fixed point can be reached from the Gaussian fixed point by turning on

$$F[\phi] = \int d^2x \; \frac{1}{2}(\nabla\phi)^2 + g_{2(n+1)}\phi^{2(n+1)}$$

Of course, as we've seen above, the RG flow is not quite so simple. When we turn on the coupling $g_{2(n+1)}\phi^{2(n+1)}$ we will generate all other terms, including ϕ^2 and ϕ^4 and so on. To reach the n^{th} fixed point, we should tune all such terms to zero as we flow towards the infra-red.

One can show that the n^{th} fixed point has n relevant operators: schematically, these can be thought of as $\phi^2, \phi^4, \ldots, \phi^{2n}$ although, as above, there will be mixing between these. By turning on the least relevant operator, one can flow from the n^{th} fixed point to the $(n-1)^{\text{th}}$ fixed point.

The results above are not derived using the $\epsilon = 4-d$ expansion which, unsurprisingly, is not much use in $d = 2$. Instead, they rely on something new which we will briefly describe in Section 3.6.

3.5.3 A History of Renormalisation

> "After about a month of work [at General Atomic Corp] I was ordered to write up my results, as a result of which I swore to myself that I would choose a subject for research where it would take at least five years before I had anything worth writing about. Elementary particle theory seemed to offer the best prospects of meeting this criterion."
>
> *Kenneth Wilson*

Renormalisation first entered physics in the context of quantum field theory, with the need to make sense of the UV divergences that arise in quantum electrodynamics. The theory, developed by Schwinger, Feynman, Tomonaga, Dyson and others, amounts to finding a consistent way to subtract one infinity from another, leaving behind a finite answer. This method yields excellent agreement with experiment but is, in the words of Feynman, a "dippy process", in which the infinities are not so much understood as swept under a very large rug.

The first hint of something deeper – and the first hint of a (semi)-group action – was seen in the work of Gell-Mann and Low in 1954. They realised that one could define an effective charge of the electron, $e(\mu)$ where μ denotes the energy scale at which the experiment takes place. This interpolates between the physical charge, as $\mu \to 0$, and the so called bare charge at high energies.

Meanwhile, throughout the 50s and 60s, a rather different community of physicists were struggling to understand second order phase transitions. It had long been known that Landau theory fails at critical points, but it was far from clear how to make progress, and the results that we've described in this course took several decades to uncover. In Kings College London, a group led by Cyril Domb stressed the importance of focussing on critical exponents; at Cornell University, Benjamin Widom showed that the relationships between critical exponents could be derived by invoking a scale invariance, albeit with little understanding of where such scale invariance came from; and at the University of Illinois, Leo Kadanoff introduced the idea of "blocking" in lattice models, a real-space renormalisation group in which one worked with successively coarser lattice models.

While many people contributed to these developments, the big picture, linking ideas from particle physics, statistical physics and condensed matter physics, is due mostly to...

Kenneth Wilson: 1936-2013

Ken Wilson received his PhD in 1961, working with Murray Gell-Mann on an assortment of topics in particle physics that fed his interest in the renormalisation group. He went on to spend much of the 1960s confused, scrabbling to understand the physics of scale, first in quantum field theory and later in the context of critical phenomena. He wrote very few papers in this time, but his reputation was strong enough to land him postdocs at Harvard and CERN and later even tenure at Cornell. (Some career advice for students: the strategy of being a genius and not writing anything rarely leads to such success.)

The floodgates opened in 1971 when Wilson set out his grand vision of the renormalisation group and, with his colleague Michael Fisher, suggested the epsilon expansion as a perturbative method to compute critical exponents in a paper charmingly titled "Critical Exponents in 3.99 Dimensions". Wilson used these methods to solve the "Kondo problem" in which an isolated spin, sitting in a bath of mobile electrons, exhibits asymptotic freedom, and he was among the first to understand the importance of numerical approaches to solve statistical and quantum field theories, pioneering the

subject of lattice gauge theory. In 1982 he was awarded the Nobel prize for his contributions to critical phenomena.

3.6 Looking Forwards: Conformal Symmetry

There are many questions that we have not yet answered? How do we know the critical exponents in $d = 2$ exactly? How do we know that there are an infinite number of fixed points in $d = 2$? Why are the critical exponents in $d = 2$ rational numbers while, in $d = 3$ they have no known closed form? How are we able to compute the $d = 3$ critical exponents to 5 significant figures?

The answers to all these questions can be found in the emergence of a rich mathematical structure at the critical point. As we've seen throughout these lectures, the basic story of RG ensures that physics at the critical point is invariant under scale transformations

$$\mathbf{x} \to \lambda \mathbf{x} \tag{3.46}$$

More surprising is the fact that the physics is invariant under a larger class of symmetries, known as conformal transformations. These transformations consist of any map

$$\mathbf{x} \to \tilde{\mathbf{x}}(\mathbf{x}) \quad \text{such that} \quad \frac{\partial \tilde{x}^i}{\partial x^k}\frac{\partial \tilde{x}^j}{\partial x^l}\,\delta_{ij} = \phi(\mathbf{x})\delta_{kl} \tag{3.47}$$

for some function $\phi(\mathbf{x})$. Such conformal transformations have the property that they preserve angles between lines.

The equation (3.47) has obvious solutions, such as translations and rotations, for which $\phi(\mathbf{x}) = 1$. Furthermore, it is simple to see that the scaling (3.46) falls into the class of conformal transformations, with $\phi(\mathbf{x}) = \lambda^2$. However, it turns out that there is one further, less intuitive transformation that obeys this condition. This is known as the *special conformal transformation* and is given by

$$x'^i = \frac{x^i - (\mathbf{x}\cdot\mathbf{x})a^i}{1 - 2(\mathbf{x}\cdot\mathbf{a}) + (\mathbf{a}\cdot\mathbf{a})(\mathbf{x}\cdot\mathbf{x})} \tag{3.48}$$

parameterised by an arbitrary vector $\mathbf{a}$.

The first question that we should ask is: why are theories at the fixed points invariant under the larger group of conformal transformations, rather than just scale transformations? The answer to this, which goes somewhat beyond this course, involves a deeper understanding of the nature of the RG flows and hinges, crucially, on being able to

construct a quantity which decreases monotonically along the flow. This quantity is, unhelpfully, called c in $d = 2$ dimensions, f in $d = 3$ dimensions and a in $d = 4$ dimensions, and the fact that it decreases monotonically is referred to as the c-theorem, f-theorem and a-theorem respectively. Using this machinery, it is then possible to prove that scale invariance implies conformal invariance. (The proof is more clear-cut in $d = 2$; it relies on some extra assumptions in higher d, and there is a general feeling that there is more to understand here.)

The existence of an extra symmetry (3.48) brings a newfound power to the study of fixed points. The d translational symmetries, $\frac{1}{2}d(d-1)$ rotational symmetries, d special conformal symmetries and single scale transformation combine to form the conformal group, which can be shown to be isomorphic to $SO(d+1,1)$. All fields and, in particular, all correlation functions must fall into representations of this group, a fact which restricts their form. In recent years, our understanding of these representations has allowed new precision in the computation of critical exponents in $d = 3$ dimensions. This programme goes by the name of the *conformal bootstrap*.

Conformal Symmetry in $d = 2$

In $d = 2$ dimensions, conformal symmetry turns out to be particularly powerful. The group of finite conformal transformations follows the pattern in higher dimensions, and is $SO(3,1) \cong SL(2,\mathbf{C})$. However, something rather special happens if you look at infinitesimal transformations where one finds that many many more are allowed. In fact, there are an infinite number. This means that there is a powerful, infinite dimensional algebra, known as the *Virasoro algebra*, underlying conformal theories in $d = 2$ dimensions. This places much stronger constraints on these fixed points, ultimately rendering many of them solvable without resorting to perturbation theory. This is the reason why the critical exponents are rational numbers which can be computed exactly. This is also what allows us to understand the structure of the infinite number of multi-critical fixed points described in Section 3.5.2.

Conformal field theory in $d = 2$ dimensions is a vast subject which arises in many different areas of physics. Although originally developed to understand critical phenomena, it also plays an important role in the lectures on the Quantum Hall Effect and the lectures on String Theory, where you can find an introduction to the basics of the subject.

4. Continuous Symmetries

So far we have focussed almost exclusively on the Ising model. Now it is time to diversify. First, however, there is one more lesson to wring from Landau's approach to phase transitions...

4.1 The Importance of Symmetry

Phases of matter are characterised by symmetry. More precisely, phases of matter are characterised by two symmetry groups. The first, which we will call G, is the symmetry enjoyed by the free energy of the system. The second, which we call H, is the symmetry of the ground state.

This structure can be seen in the Ising model. When $B = 0$, the free energy has a $G = \mathbf{Z}_2$ symmetry. In the high temperature, disordered phase this symmetry is unbroken; here $H = \mathbf{Z}_2$ also. In contrast, in the low temperature ordered phase, the symmetry is spontaneously broken as the system must choose one of two ground states; here $H = \varnothing$. The two different phases – ordered and disordered – are characterised by different choices for H.

In the ordered phase we have two different ground states, whose phase diagram is reproduced on the next page. Whenever a discrete symmetry group like $\mathbf{Z}_2$ is spontaneously broken, it results in multiple ground states. One can move from one ground state to another by acting with the broken generators of G.

In contrast, when $B \neq 0$ the free energy does not have a $\mathbf{Z}_2$ symmetry, so $G = \varnothing$. According to Landau's criterion, this means that there is only a single phase. Indeed, by going to temperatures $T > T_c$, it is possible to move from any point in the phase diagram to any other point without passing through a phase transition, so there is no preferred way to carve the phase diagram into different regions. However, this also means that, by varying B at low temperatures $T < T_c$, we can have a first order phase transition between two states which actually lie in the same phase. This can also be understood on symmetry grounds because the first order transition does not occur at a generic point of the phase diagram, but instead only when G is enhanced to $\mathbf{Z}_2$.

The discussion carries over identically to any system which lies in the Ising universality class, including the liquid-gas system. This leaves us with the slightly disconcerting idea that a liquid and gas actually describe the same phase of matter. As with the Ising model, by taking a path through high pressures and temperatures one can always convert one smoothly into the other, which means that any attempt to label points in the phase diagram as "liquid" or "gas" will necessarily involve a degree of arbitrariness.

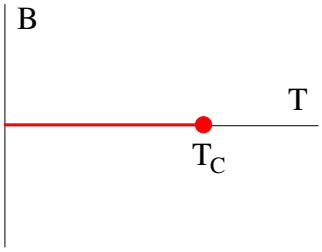

Figure 31: The phase diagram of the Ising model (again).

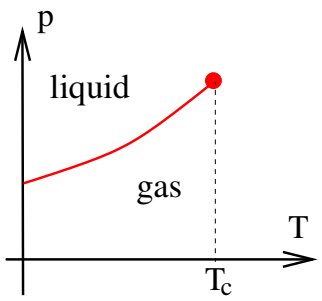

Figure 32: The phase diagram of the liquid-gas system (again).

It is really only possible to unambiguously distinguish a liquid from a gas when we sit on the line of first order phase transitions. Here there is an emergent $G = \mathbf{Z}_2$ symmetry, which is spontaneously broken to $H = \varnothing$, and the two states of matter – liquid and gas – are two different ground states of the system. In everyday life, we sit much closer to the line of first order transitions than to the critical point, so feel comfortable extending this definition of "liquid" and "gas" into other regimes of the phase diagram, as shown in the figure.

Beyond Ising

The idea of symmetry, and of broken symmetry, turns out to be useful in characterising nearly all phases of matter. In each case, one should first determine an order parameter and a symmetry group G under which it transforms. Sometimes the choice of order parameter is obvious; sometimes it is more subtle. One then writes down the most general Landau-Ginzburg free energy, subject to the requirement that it is invariant under G. The different phases of matter within this class are characterised by the group H preserved by the ground state.

There are a number of reasons why it is useful to characterise states of matter in terms of their (broken) symmetry. The original idea of Landau was that, as we've seen with the Ising model, symmetry provides a powerful mechanism to understand when a phase transition will take place. In particular, there must be a phase transition whenever H changes.

However, it turns out that this is not the only thing symmetry is good for. As we will see below, knowledge of G and H is often sufficient to determine many of the low energy properties of a system, both through a result known as *Goldstone's theorem* (that we will describe in Section 4.2) and through various topological considerations (some of which we will see in Section 4.4).

Finally, and particularly pertinent for this course is the role that symmetry plays in the renormalisation group and specifically in universality. One can ask: when do two systems lie in the same universality class? Although the full answer to this question is not yet understood, a fairly good guess is: when they share the same symmetry G.

There are many different systems and choices of G that we could look at. A particularly interesting class occurs when we take $G = \mathbf{R}^d \times SO(d)$, the group of spatial translations and rotations. The pattern of symmetry breaking provides, for example, a clean distinction between a liquid/gas and a solid, with the latter breaking G down to its crystal group. In this framework, there is not one solid phase of matter, but many, with each different crystal structure preserving a different H and hence representing a different phase of matter. The different breaking patterns of spatial rotations also allow us to define novel phases of matter, such as liquid crystals. Viewed in this way, even soap, which can undergo a discontinuous change to become slippery, constitutes a new phase of matter. We will not discuss this form of symmetry breaking in this course, but you can learn more about it in next term's course on *Soft Matter*.

Here, instead, we will be interested in phases of matter that are characterised by "internal" symmetry groups G that are continuous, as opposed to the discrete symmetry of the Ising model. This includes materials like magnets, where the spin is a vector that is free to rotate. It also includes more exotic quantum materials such as Bose-Einstein condensates, superfluids and superconductors. We will see that systems with continuous symmetry groups G exhibit a somewhat richer physics than we've seen in the Ising model.

Beyond the Landau Classification

The idea that phases of matter can be classified by (broken) symmetries has proven crucial in placing some order on the world around us. However, it is not the last word. Over the past twenty years, it has become increasingly clear that certain highly entangled quantum systems defy a simple characterisation by symmetry. The first, and most prominent, examples are the quantum Hall states. To understand these, one needs a new ingredient: topology. We will not touch upon this here, but you can read more in the lecture notes on the Quantum Hall Effect.

4.2 $O(N)$ Models

Phases of matter that are characterised by continuous, as opposed to discrete, symmetries offer a rich array of new physics. The simplest such models contain N real scalar fields, which we arrange in a vector

$$\boldsymbol{\phi}(\mathbf{x}) = (\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \ldots, \phi_N(\mathbf{x}))$$

We will ask that the free energy is invariant under the $O(N)$ symmetry

$$\phi_a(\mathbf{x}) \to R_a{}^b \phi_b(\mathbf{x})$$

where $R \in O(N)$ so that $R^T R = 1$. Now, when constructing the free energy we write down only the terms invariant under $O(N)$. The first few are

$$F[\boldsymbol{\phi}(\mathbf{x})] = \int d^d x \ \left[\frac{\gamma}{2} \nabla\boldsymbol{\phi} \cdot \nabla\boldsymbol{\phi} + \frac{\mu^2}{2} \boldsymbol{\phi} \cdot \boldsymbol{\phi} + g(\boldsymbol{\phi} \cdot \boldsymbol{\phi})^2 + \ldots \right]$$

where rotational invariance requires $\nabla\boldsymbol{\phi} \cdot \nabla\boldsymbol{\phi} = \partial_i \phi_a \partial_i \phi_a$. These kind of theories are known, not unreasonably, as $O(N)$ models. They are of interest for all N, but $N = 2$ and $N = 3$ play particularly prominent roles.

$O(2)$: The XY-Model

When $N = 2$, it is often convenient to pair the two real scalar fields into a single complex field

$$\psi(\mathbf{x}) = \phi_1(\mathbf{x}) + i\phi_2(\mathbf{x})$$

The free energy now consists of all terms which are invariant under $U(1)$ phase rotations, $\psi \to e^{i\alpha}\psi$. The first few terms are

$$F[\psi(\mathbf{x})] = \int d^d x \ \left[\frac{\gamma}{2} \nabla\psi^\star \cdot \nabla\psi + \frac{\mu^2}{2} |\psi|^2 + g|\psi|^4 + \ldots \right] \tag{4.1}$$

This is also known as the XY-model or, sometimes, the rotor model.

There are at least two physical systems which sit in this universality class. The first are magnets where, in contrast to the Ising model, each atom has a continuous spin $\mathbf{s}$ which can rotate in a plane. (This is usually taken to be the $X - Y$ plane, which is where the name comes from.) The microscopic Hamiltonian is the generalisation of the Ising model (1.1) to

$$E = -J \sum_{\langle ij \rangle} \mathbf{s}_i \cdot \mathbf{s}_j \tag{4.2}$$

where $|\mathbf{s}_i| = 1$. This is also written as

$$E = -J \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j)$$

where θ_i is the angle that the spin s_i makes with some, fixed reference direction. Coarse-graining this microscopic model gives rise to the free energy (4.1). One could also add a magnetic field term $\sum_i \mathbf{B} \cdot \mathbf{s}_i$, where $\mathbf{B}$ is also a two-component vector. Such a term would break the $O(2)$ symmetry, and introduce terms in (4.1) that are odd in ψ.

The second physical system described by (4.1) is rather different in nature: it is a Bose-Einstein condensate, or its strongly coupled counterpart, a superfluid. Here, the origin of the order parameter ψ is rather more subtle, and is related to off-diagonal long-range order in the one-particle density matrix. In this case, the saddle point of the free-energy leads to the equation of motion

$$\gamma \nabla^2 \psi = \mu^2 \psi + 4g|\psi|^2 \psi + \dots$$

which is known as the Gross-Pitaevskii equation.

It is sometimes, rather lazily, said that $\psi(\mathbf{x})$ can be thought of as the macroscopic wavefunction of the system, and the Gross-Pitaevskii equation is then referred to as a non-linear Schrödinger equation. This is misleading for the simple reason that quantum mechanics is always linear.

$O(3)$: The Heisenberg Model

The case $N = 3$ also describes magnets. The microscopic energy again takes the form (4.2), but now where each $\mathbf{s}_i$ is free to rotate in three dimensions. This is referred to as the $O(3)$ model or, alternatively, as the *Heisenberg model*.

4.2.1 Goldstone Bosons

The real novelty of continuous symmetries arises in the ordered phase, where $\mu^2 < 0$ and, correspondingly, $\langle \boldsymbol{\phi} \rangle \neq 0$ in the ground state. For the Ising model, we had two possible choices: $m = \pm m_0$. The system had to pick one, and in doing so spontaneously broke the $\mathbf{Z}_2$ symmetry. With a continuous symmetry, we have an infinite number of choices. The minimum of the free energy constrains only the magnitude of $\boldsymbol{\phi}$ which is given by

$$\langle |\boldsymbol{\phi}| \rangle = M_0 = \sqrt{-\frac{\mu^2}{4g}}$$

However, minimising the free energy does not determine the direction of $\boldsymbol{\phi}$. We are left with a space of ground states which is the sphere $\mathbf{S}^{N-1}$. Each point on the sphere parameterises the direction of $\boldsymbol{\phi}$ and has the same energy. The configuration in which all the spins point like this:

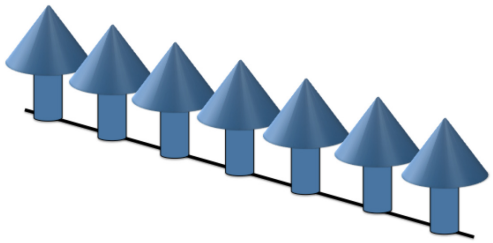

has the same energy as the configuration in which all the spins point like this:

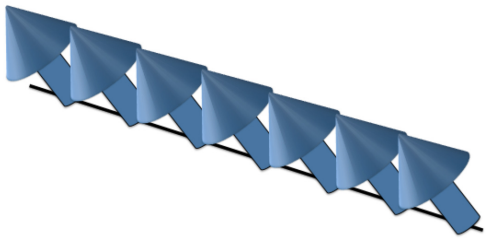

This infinitely degenerate choice of ground states gives us something new. We can consider configurations in which we stay within the space of ground states, but the direction varies in space. For such configurations, the part of the free energy $f(\boldsymbol{\phi}) = \frac{\mu^2}{2}|\boldsymbol{\phi}|^2 + g|\boldsymbol{\phi}|^4 + \ldots$ remains minimised, but we pick up contributions from the gradient terms $|\nabla\boldsymbol{\phi}|^2$. However, we can always lower this free energy by making the variation take place over longer and longer distances. The upshot is that there are excitations of the system that look like this

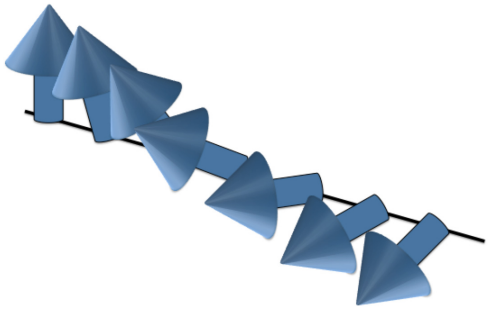

which can be made to cost an arbitrarily small amount of energy, by stretching the winding over longer and longer distance scales.

These kind of excitations, which arise from the spontaneous breaking of continuous symmetries, are known as *Goldstone bosons*, or sometimes *Nambu-Goldstone bosons*. In the particular context of magnets, they are called *spin waves*.

There is a dizzying array of names for these kind of excitations, reflecting their ubiquity and importance. In general, an excitation whose energy cost vanishes as the wavelength goes to infinity is referred to as a *soft mode* or, alternatively, is said to be *gapless*. These are to be contrasted with *gapped* excitations whose energy remains finite in this limit. In the context of quantum field theory, "gapless" = "massless", and "gapped" = "massive", with the energy gap coming from $E = mc^2$.

Gapless excitations often dominate the low-temperature behaviour of a system, where they are the only excitations that are not Boltzmann suppressed. In many systems, these gapless modes arise from the breaking of some symmetry. A particularly important example, that we will not discuss in these lectures, are phonons in a solid. These can be thought of as Goldstone bosons for broken translational symmetry.

The Symmetry Behind Goldstone Bosons

The intuitive idea described above can be placed on more rigorous footing in the form of *Goldstone's theorem*. This states that, in any system the spontaneous breaking of a continuous symmetry gives rise to a gapless excitation, the eponymous Goldstone boson. This can be stated in the language of group theory.

For our $O(N)$ model, the $G = O(N)$ symmetry is broken by a choice of $\langle \boldsymbol{\phi} \rangle$ to $H = O(N-1)$. (To see this, note that if $\boldsymbol{\phi} = (M_0, 0, \ldots, 0)$ then there is a surviving $O(N-1)$ symmetry which acts on the string of zeros.) The space of ground states has a group theoretic interpretation as the coset space

$$\mathbf{S}^{N-1} = \frac{O(N)}{O(N-1)}$$

This idea generalises. If a continuous symmetry G is spontaneously broken to H, then the manifold of ground states is given by G/H. We get a Goldstone boson for each broken symmetry generator, so the total number is

$$\#\text{ Goldstone Bosons } = \dim G - \dim H$$

For the $O(N)$ model, $G = O(N)$ and $H = O(N-1)$ so the number of Goldstone modes is then $\frac{1}{2}N(N-1) - \frac{1}{2}(N-1)(N-2) = N-1$, which is indeed the dimension of the sphere $\mathbf{S}^{N-1}$.

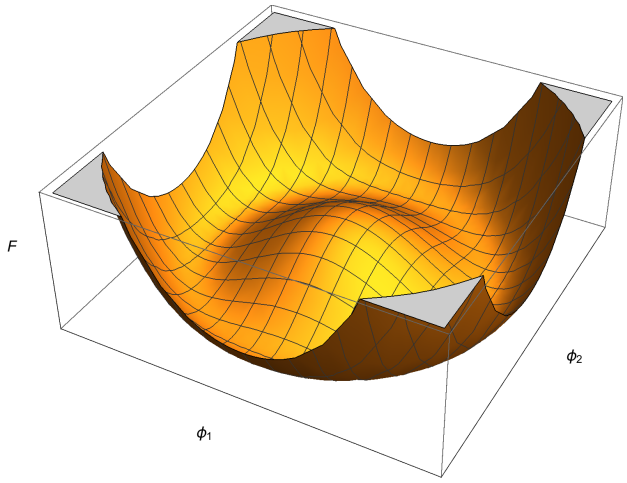

Figure 33: The so-called "Mexican hat" potential for the XY-model in the ordered phase.

An Example: the XY-Model

It's simple to write some equations to go with the pictures above. Let's start with the XY-model. In the ordered phase, we get a so-called Mexican hat potential, as shown in the figure. We can see that there is a circle, $\mathbf{S}^1$ of minima. It's useful to decompose the field as $\psi(\mathbf{x}) = M(\mathbf{x})e^{i\theta(\mathbf{x})}$. In the ground state $M = M_0 = \sqrt{-\mu^2/4g}$, while θ is arbitrary. If we write

$$M(\mathbf{x}) = M_0 + \tilde{M}(\mathbf{x}) \tag{4.3}$$

then the free energy has the expansion

$$\begin{aligned} F[M,\theta] = \int d^dx \quad & \frac{\gamma}{2}(\nabla\tilde{M})^2 + |\mu^2|\tilde{M}^2 + g\tilde{M}^4 + \ldots \\ & + \frac{\gamma}{2}M_0^2(\nabla\theta)^2 + \gamma M_0\tilde{M}(\nabla\theta)^2 + \ldots \end{aligned} \tag{4.4}$$

Here, the Goldstone boson is $\theta(\mathbf{x})$. There can be no terms of the form θ^2 or θ^4 arising in the free energy. Instead, it has only derivative interactions.

Another Example: The Heisenberg Model

For the $O(3)$ model, we decompose the field in spherical polar coordinates,

$$\boldsymbol{\phi} = M(\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$$

with $\theta \in [0, \pi)$ and $\phi \in [0, 2\pi]$. Once again, in the ordered phase we have $M = M_0 \neq 0$, with θ and ϕ arbitrary. Expanding M as (4.3), the free energy now takes the form

$$F[M, \theta, \phi] = \int d^d x \quad \frac{\gamma}{2}(\nabla \tilde{M})^2 + |\mu^2|\tilde{M}^2 + g\tilde{M}^4 + \ldots \\ + \frac{\gamma}{2} M_0^2 \left[(\nabla\theta)^2 + \sin^2\theta\,(\nabla\phi)^2\right] + \ldots \tag{4.5}$$

Here θ and ϕ are the two Goldstone modes and, correspondingly, have only derivative interactions. Note, however, that this time the Goldstone modes interact with each other, as seen in the $\sin^2\theta\,(\nabla\phi)^2$ term.

The kinetic terms for the Goldstone bosons above take the form of the metric on the two-sphere $\mathbf{S}^2$, i.e. $ds^2 = d\theta^2 + \sin^2\theta d\phi^2$. This is no coincidence: the Goldstone bosons describe fluctuations around the minima of the free energy $F[\boldsymbol{\phi}]$. In the present case, this set of minima is $\mathbf{S}^2$, and this geometry gets imprinted on the dynamics of the Goldstone modes. We will explore this more in Section 4.3.

Correlation Functions

We saw in Section 2.2 that the quadratic term in the free energy is related (inversely) to the correlation length ξ. For Goldstone bosons this quadratic term necessarily vanishes and so they have infinite correlation length.

This manifests itself in the correlation function, which decays as a power-law rather than exponential. This is simplest to see in the XY-model. (We will discuss $O(N)$ models with $N \geq 3$ in more detail in Section 4.3.) The free energy (4.4) is

$$F[\theta] = \int d^d x\ \frac{\gamma}{2} M_0^2 (\nabla\theta)^2 + \ldots$$

where the higher order terms are all derivatives and will not affect the discussion below. To compute the correlation function $\langle\theta(\mathbf{x})\theta(\mathbf{y})\rangle$, we can simply import the result (2.20). (There are some subtleties in doing the path integral because $\theta(\mathbf{x})$ is periodic, now valued in $[0, 2\pi)$ rather than $\mathbf{R}$. These subtleties turn out not to be important here but we will revisit them in Section 4.4.) The long distance behaviour is

$$\langle\theta(\mathbf{x})\theta(\mathbf{y})\rangle = \frac{1}{\gamma M_0^2} \int \frac{d^d k}{(2\pi)^d} \frac{e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})}}{k^2} \tag{4.6}$$

This is similar to the behaviour of the correlation function at the critical point. Indeed, a critical point can be thought of as having gapless excitations. But there are differences.

First, the power-law decay at the critical point requires some fine-tuning of a parameter; we must pick the temperature to be exactly $T = T_c$. In contrast, spontaneous symmetry breaking is more robust, and we get power-law decay for all $T < T_c$. In other words, we have a *phase* with long range correlations, rather than just a point in the phase diagram. (For $T > T_c$, where there is no symmetry breaking, all modes still decay exponentially as in the Ising model.)

The second difference is that Goldstone bosons are much simpler to understand than the gapless modes at a critical point. As we have seen, at critical points the power-law decay of correlation functions suffers a correction due to integrating out short distance modes, resulting in the critical exponent $\eta \neq 0$. There are no such subtleties for Goldstone modes since all the dynamics is constrained by symmetry, and the correlation function (4.6). There are two caveats to this statement, both of which we will elaborate upon below. The first is that the simplicity only holds when $T < T_c$; when we sit at the critical point $T = T_c$, things become interesting once again. The second caveat is that we have to be above the lower critical dimension for the Goldstone bosons to exist.

4.2.2 The $d = 4 - \epsilon$ Expansion

At the critical temperature, $T = T_c$, the $O(N)$ models exhibit critical behaviour. The mean field approach to the $O(N)$ model gives the same answer as we saw for the $N = 1$ Ising field theory in previous sections. By now, you will not be surprised to learn that these mean field exponents are not always correct. However, the system now does not flow to the Ising critical point. Instead, they lie in a different universality class.

First, in $d = 2$ there are no critical points with $G = O(N)$ symmetry. We'll see why in Section 4.2.3 and explore the physics more in Sections 4.3 and 4.4.

In $d = 3$, the theories flow to a different critical point for each N. The critical exponents are known to be:

	η	ν
MF	0	$\frac{1}{2}$
Ising	0.0363	0.6300
$N = 2$	0.0385	0.6719
$N = 3$	0.0386	0.702

where the other critical exponents, α, β, γ and δ all follow from the scaling relations that we saw in Section 3.2.

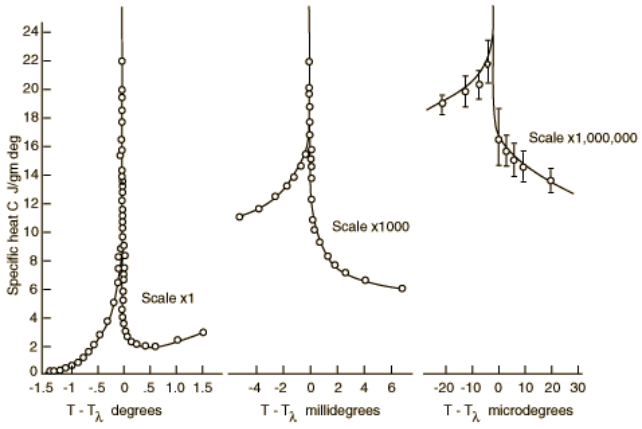

Figure 34: The heat capacity of helium at the superfluid transition. This system lies in the XY universality class. The data above is well described by the function $C \sim C_0 + A|T - T_c|^{-\alpha}$ with $\alpha \approx -0.16$ and $A < 0$.

While the values of η and ν do not look very different from the Ising exponents, there is an important difference in the critical exponent for the heat capacity $c \sim |T - T_c|^{-\alpha}$, which is given by $\alpha = 2 - 3\nu$. For both the $O(2)$ and $O(3)$ transition, α is negative. For example, $\alpha \approx -0.16$ for the $O(2)$ transition. This means that the heat capacity exhibits a cusp, rather than a true divergence.

For example, the superfluid transition of helium lies in the XY universality class. The heat capacity has long been known to exhibit cusp-like behaviour as shown in Figure 34[8]. This characteristic shape means that the second order superfluid transition is sometimes referred to as the "lambda transition". It turns out that the accuracy in these experiments is limited by the effect of the Earth's gravitational field. In the 1990s, these measurements were made on a space shuttle flight, in broad (but not perfect) agreement with theoretical prediction of $c \sim A_\pm - Bt^{-\alpha}$ for the critical exponent $\alpha \approx -0.16$ and suitable coefficients $A_\pm$ and B.

The transition to Bose-Einstein condensate also sits in the XY universality class. This is a particularly clean system which allows precision experiments. For example, the data in Figure 35 shows the behaviour of the correlation length as a gas of ultracold rubidium-87 atoms passes through the critical point. The critical exponent is found

[8]This data is taken from Buckingham and Fairbank, "The Nature of the Lambda Transition", in Progress in Low Temperature Physics III, 1961.

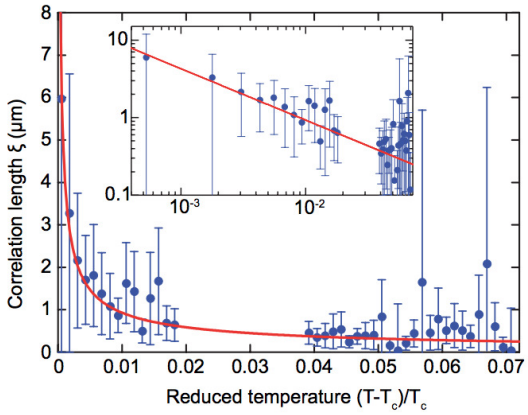

Figure 35: The correlation length for a BEC at the critical point. This system lies in the XY universality class and has $\nu = 0.67 \pm 0.13$.

to be $\nu = 0.67 \pm 0.13$, in good agreement with the theoretical prediction (albeit with fairly large error bars)[9].

It is not too difficult to repeat the RG calculations that we did in Sections 3.3 - 3.5 for the $O(N)$ model. As before, we rescale fields so that our order parameter – which we now call $\phi_a(\mathbf{x})$, with $a = 1, \ldots, N$ – has free energy

$$\beta F[\phi] = \int d^d x \ \left[\frac{1}{2} \nabla \phi_a \cdot \nabla \phi_a + \frac{1}{2} \mu_0^2 \, \phi_a \phi_a + g_0 (\phi_a \phi_a)^2 + \ldots \right]$$

The study of the Gaussian fixed point, at $\mu_0^2 = g_0 = 0$, goes through much as before. Indeed, a simple dimensional analysis argument tells us that

$$[\phi_a] = \Delta_\phi = \frac{d-2}{2}$$

and, so

$$[\mu_0^2] = 2 \quad \text{and} \quad [g_0] = 4 - d$$

so that μ_0^2 is always a relevant deformation, while g_0 is relevant in $d < 4$ dimensions. So far, nothing depends on N.

[9]This data is taken from the paper "Critical behaviour of a trapped interacting Bose gas" by T. Donner et. al., arXiv:0704.1439.

The differences arise in perturbation theory. The part of the free energy which mixes long and short wavelength modes is

$$\beta F_I[\phi] = \int d^d x \; g(\phi_a \phi_a)^2$$

The presence of the internal indices, $a = 1, \ldots, N$, means that the interaction has more structure than previously. To reflect this, we need to change our rules for drawing diagrams. First, each line should now be accompanied by an internal index a. Second, it is useful to split the interaction vertex as

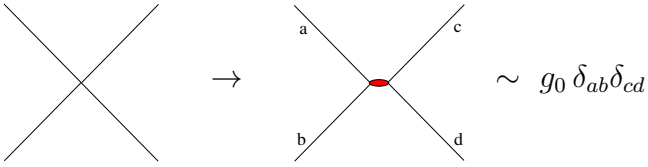

$$\sim \; g_0 \, \delta_{ab} \delta_{cd}$$

where the red ellipse splits the four legs into two pairs, each of which is a singlet under the $O(N)$ symmetry, as shown in the delta function structure. (You may have to squint in some of the following pictures to see which pairs of legs are contracted.)

Order g_0

We can now run through our previous calculation to see how things change when we have N fields. At order g_0, we previously found just a single diagram which renormalised μ^2. Now the index structure means that splits into two different contributions. The first is:

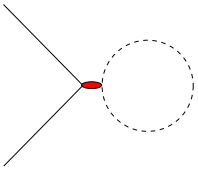

$$= 2g_0 \int \prod_{i=1}^{4} \frac{d^d k_i}{(2\pi)^d} \, \phi^-_{a,\mathbf{k}_1} \phi^-_{a,\mathbf{k}_2} \times \left\langle \phi^+_{b,\mathbf{k}_3} \phi^+_{b,\mathbf{k}_4} \right\rangle_+ \times (2\pi)^d \delta^d(\sum_i \mathbf{k}_i)$$

The other contribution has a different contraction between internal indices

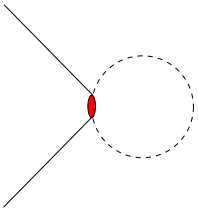

$$= 4g_0 \int \prod_{i=1}^{4} \frac{d^d k_i}{(2\pi)^d} \, \phi^-_{a,\mathbf{k}_1} \phi^-_{b,\mathbf{k}_2} \times \left\langle \phi^+_{a,\mathbf{k}_3} \phi^+_{b,\mathbf{k}_4} \right\rangle_+ \times (2\pi)^d \delta^d(\sum_i \mathbf{k}_i)$$

Note that the overall coefficients are $2 + 4 = 6$, which agrees with our earlier counting (3.28). Each of these gives the same result we saw for a single scalar field, but with an important difference: the first diagram has an extra factor of N, arising from the fact that any of the N species can run in the loop. This is a general result: any closed dotted loop gives an extra factor of N.

The rest of the calculation proceeds as in Section 3.4.1. We find that, at this order, we have a renormalisation of the quadratic term given by

$$\mu_0^2 \to \mu'^2 = \mu_0^2 + 4(N+2)g_0 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d} \frac{1}{q^2+\mu_0^2}$$

This agrees with our earlier result (3.30) when $N = 1$.

Order g_0^2

A similar thing happens at the next order, with the single diagram at $N = 1$ splitting into three different diagrams,

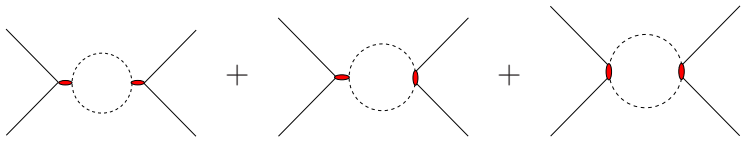

Only the first of these has a closed loop, unattached to the external legs. This comes with a factor of N, while the other two do not. A careful computation of the relevant overall factors shows that these diagrams renormalise the quartic coupling as

$$g_0 \to g_0' = g_0 - 4(N+8)g_0^2 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d} \frac{1}{(q^2+\mu_0^2)^2}$$

Again, this reproduces our earlier result (3.40) when $N = 1$.

The Epsilon Expansion

We learn that the general structure of RG flow around the Gaussian fixed point is essentially the same as in the Ising field theory discussed in Section 3; only the coefficients of the beta functions differ. The same structure is also seen in the epsilon expansion. Working in dimension

$$d = 4 - \epsilon$$

the beta function equations become, at leading order in ϵ and g,

$$\begin{aligned} \frac{d\mu^2}{ds} &= 2\mu^2 + \frac{N+2}{2\pi^2} \frac{\Lambda^4}{\Lambda^2+\mu^2} \tilde{g} \\ \frac{d\tilde{g}}{ds} &= \epsilon\tilde{g} - \frac{N+8}{2\pi^2} \frac{\Lambda^4}{(\Lambda^2+\mu^2)^2} \tilde{g}^2 \end{aligned}$$

where, as in our earlier discussion, we've introduced the dimensionless coupling $\tilde{g} = \Lambda^{-\epsilon} g$. The analog of the Wilson-Fisher fixed point sits at

$$\mu_\star^2 = -\frac{1}{2}\frac{N+2}{N+8}\Lambda^2\epsilon \quad \text{and} \quad \tilde{g}_\star = \frac{2\pi^2}{N+8}\epsilon$$

Around this fixed point, the linearised beta functions take the form

$$\frac{d}{ds}\begin{pmatrix}\delta\mu^2\\ \delta\tilde{g}\end{pmatrix} = \begin{pmatrix}2-\frac{N+2}{N+8}\epsilon & C\\ 0 & -\epsilon\end{pmatrix}\begin{pmatrix}\delta\mu^2\\ \delta\tilde{g}\end{pmatrix}$$

where the off-diagonal entry is $C = \frac{N+2}{2\pi^2}\Lambda^2 + \frac{(N+2)^2}{4\pi^2(N+8)}\Lambda^2\epsilon$. This does not affect the eigenvalues which are given by the diagonal entries,

$$\Delta_t = 2 - \frac{N+2}{N+8}\epsilon + \mathcal{O}(\epsilon^2) \quad \text{and} \quad \Delta_g = -\epsilon + \mathcal{O}(\epsilon^2)$$

The interacting fixed point has one relevant and one irrelevant direction, just like for the Ising model. To leading order in ϵ, the critical exponents are

$$\alpha = \frac{4-N}{2(N+8)}\epsilon\ , \quad \beta = \frac{1}{2} - \frac{3}{2(N+8)}\epsilon\ , \quad \gamma = 1 + \frac{N+2}{2(N+8)}\epsilon$$

and

$$\nu^{-1} = 2 - \frac{N+2}{N+8}\epsilon \tag{4.7}$$

Meanwhile $\delta = 3 + \epsilon$ is independent of N, and the anomalous dimension turns out to be $\eta = (N+2)\epsilon^2/2(N+8)$

4.2.3 There Are No Goldstone Bosons in $d = 2$

We learned in Section 1 that field theories have a lower critical dimension, in which the ordered phase ceases to exist. For theories characterised by any discrete symmetry, such as the $\mathbf{Z}_2$ of the Ising field theory, this lower critical dimension is $d_l = 1$. As we explained in Section 1.3.3, the lack of an ordered phase in $d = 1$ dimensions can be traced to the existence of domain walls.

The story is different when we have continuous symmetries. Now there are no domain walls because the space of ground states is continuously connected. However, there is a more prominent phenomena which means that the lower critical dimension is raised to $d_l = 2$.

An Example: the XY model

The lack of ordered phase in $d = 2$ dimensions arises due to the presence of the would-be Goldstone modes. This is simplest to explain in the XY model. Let's sit in the broken phase and focus only on the Goldstone bosons. We must pick a ground state

for θ: this is the essence of spontaneous symmetry breaking. Let's choose $\langle\theta(\mathbf{x})\rangle = 0$. We now look at fluctuations around this ground state,

$$\langle[\theta(\mathbf{x}) - \theta(0)]^2\rangle = 2\langle\theta(\mathbf{x})^2\rangle - 2\langle\theta(\mathbf{x})\theta(0)\rangle \tag{4.8}$$

From the correlation function (4.6), the long distance behaviour is

$$\langle\theta(\mathbf{x})\theta(0)\rangle = \frac{1}{\gamma M_0^2}\int^\Lambda \frac{d^dk}{(2\pi)^d}\frac{e^{-i\mathbf{k}\cdot\mathbf{x}}}{k^2} \sim \begin{cases} \Lambda^{d-2} - r^{2-d} & d > 2 \\ \log(\Lambda r) & d = 2 \\ r - \Lambda^{-1} & d = 1 \end{cases} \tag{4.9}$$

with $r = |\mathbf{x}|$.

We see that there is a qualitative difference between $d > 2$ and $d \leq 2$. For $d > 2$, the two point correlator $\langle\theta(\mathbf{x})\theta(0)\rangle$ decays to a constant as $r \to \infty$. This constant is cancelled by the other term in (4.8), which means that the phase returns to its original value $\langle\theta\rangle = 0$.

In contrast, for $d \leq 2$, the fluctuation of the phase grows indefinitely as we go to larger distances. You may have thought that you had placed the system in a fixed ground state, but the thermal fluctuations of the Goldstone mode mean that it doesn't stay there. The interpretation is that there is no ordered phase with $\langle\boldsymbol{\phi}\rangle \neq 0$ in $d = 2$ dimensions or below.

This is a general result, known as the *Mermin-Wagner theorem*. A continuous symmetry cannot be spontaneously broken in $d = 2$ dimensions or below: there are no Goldstone bosons in $d = 2$ dimensions.

This leaves us with a delicate question. The existence of the gapless Goldstone modes was predicated on the idea of spontaneous symmetry breaking. But for $d \leq 2$ dimensions, no such symmetry breaking happens. What, then, is the resulting physics?

In $d = 1$, the physics is straightforward: there are no gapless modes. As before, this can also be understood in the language of quantum mechanics, where the spectrum of a particle moving on $\mathbf{S}^{N-1}$ is discrete and gapped. For $d = 2$, the physics is more interesting. It turns out that the answer is somewhat different for $O(N)$ models with $N \geq 3$ and for the XY-model with $N = 2$. We will discuss the fate of the Goldstone modes for each of these in Sections 4.3 and 4.4 respectively.

4.3 Sigma Models

We have still to understand the fate of the Goldstone bosons in $d = 2$ dimensions. In this section, we will tell their story. As a spin-off, we will see that we also a get a new handle on the critical point in $d = 3$ dimensions.

We place ourselves firmly in the ordered phase, with $T < T_c$. Mean field considerations tell us that $\langle|\boldsymbol{\phi}|\rangle \neq 0$, leaving us with a space of possible ground states which is identified with the sphere $\mathbf{S}^{N-1}$. As we saw in Section 4.2.1, fluctuations in the directions parallel to the $\mathbf{S}^{N-1}$ have only power-law decay; these are the Goldstone modes. In contrast the "longitudinal" fluctuation, in which $\delta\boldsymbol{\phi} \sim \langle\boldsymbol{\phi}\rangle$, acts very much like in the Ising model and has exponential decay with a correlation length $\xi \neq 0$. This suggests that the long distance dynamics is dominated purely by the Goldstone modes.

Here we will study the theory of these Goldstone modes. First, rather than working with $\boldsymbol{\phi} \cdot \boldsymbol{\phi} = M_0^2$, we rescale the field $\boldsymbol{\phi}(\mathbf{x})$ to a new field, $\mathbf{n}(\mathbf{x})$ which has unit length,

$$\mathbf{n} \cdot \mathbf{n} = 1 \tag{4.10}$$

For now, we will keep the dimension d arbitrary. The free energy is given by

$$F[\mathbf{n}] = \int d^d x \; \frac{1}{2e^2} \nabla\mathbf{n} \cdot \nabla\mathbf{n} \tag{4.11}$$

where the coefficient $e^2 = 1/\gamma M_0^2$ is the price that we pay for rescaling $\mathbf{n}$ to be a unit vector. The free energy (4.11) looks like that of a free theory; all the interactions come from the constraint (4.10) which say that the fields $\mathbf{n}(\mathbf{x})$ must lie on the unit sphere $\mathbf{S}^{N-1}$.

The theory defined by (4.11), together with the constraint (4.10), lies in a class of theories referred to as *non-linear sigma models*. These are theories in which the fields can be viewed as coordinates on some manifold $\mathcal{M}$. In the present context, this manifold is $\mathcal{M} = \mathbf{S}^{N-1}$.

We would like to understand the path integral for the sigma model. Schematically, this can be written as

$$Z = \int \mathcal{D}\mathbf{n} \; \delta\big(\mathbf{n}(\mathbf{x})^2 - 1\big) \; \exp\left(-\frac{1}{2e^2}\int d^d x \; \nabla\mathbf{n} \cdot \nabla\mathbf{n}\right) \tag{4.12}$$

Here we've imposed the constraint through a delta function. Note that the only coefficient in the game is e^2; this will play the role of our coupling constant. Recall that, long ago, before we grew up and set $\beta = 1$, we used to write the partition function as $e^{-\beta F}$. Comparing to this form suggests that e^2 can be viewed as temperature, with large e^2 corresponding to high temperature. This interpretation will be useful later.

We can do some simple dimensional analysis. The field $\mathbf{n}(\mathbf{x})$ must be dimensionless since it obeys the constraint (4.10). So,

$$[e^2] = 2 - d \tag{4.13}$$

In particular, e^2 is dimensionless in $d = 2$. Here the theory is weakly coupled when $e^2 \ll 1$, in the sense that field configurations $\mathbf{n}(\mathbf{x})$ with wild spatial variations are suppressed in the path integral. Before our rescaling, this corresponds to the case where $\boldsymbol{\phi}$ parameterise a large $\mathbf{S}^{N-1}$ sphere. In contrast, when $e^2 \gg 1$ these configurations are unsuppressed and the theory is strongly coupled; in this case the $\boldsymbol{\phi}$ parameterise a small sphere.

It is possible to write the sigma model in a more explicit form. We can decompose the vector $\mathbf{n}$ as

$$\mathbf{n}(\mathbf{x}) = (\vec{\pi}(\mathbf{x}), \sigma(\mathbf{x}))$$

where $\vec{\pi}(\mathbf{x})$ is an $(N-1)$-dimensional vector and $\sigma(\mathbf{x})$ is given by

$$\sigma(\mathbf{x})^2 = 1 - \vec{\pi}(\mathbf{x}) \cdot \vec{\pi}(\mathbf{x})$$

which ensures that the fields sit on the ground state manifold $\mathbf{n} \cdot \mathbf{n} = 1$. The free energy is then given by

$$\begin{aligned} F[\vec{\pi}] &= \int d^d x \, \frac{1}{2e^2} \left[\nabla\vec{\pi} \cdot \nabla\vec{\pi} + \nabla\sigma \cdot \nabla\sigma \right] \\ &= \int d^d x \, \frac{1}{2e^2} \left[\nabla\vec{\pi} \cdot \nabla\vec{\pi} + \frac{(\vec{\pi} \cdot \nabla\vec{\pi})^2}{1 - \vec{\pi}^2} \right] \end{aligned} \tag{4.14}$$

This form of the sigma model does not have any constraint; it is an interacting theory of the Goldstone modes $\vec{\pi}(\mathbf{x})$. However, we have paid a price: only an $O(N-1)$ symmetry is now manifest in the free energy, rather than the full $O(N)$ symmetry. This is because we have had to make a choice of which of the redundant $\mathbf{n}$ variables to eliminate in order to solve the constraint (4.10). Related to this, our free energy (4.14) is now only valid as long as $\sigma(\mathbf{x}) \neq 0$ anywhere, in which case the second term would diverge. This is because the $\vec{\pi}$ fields are coordinates on the space $\mathbf{S}^{N-1}$ and it is impossible to introduce coordinates which are well behaved over the entire manifold.

As an aside: the name "*sigma model*" is, obviously, completely uninformative. It has its roots in particle physics where a theory of this type describes the interactions of pions. Strangely, the eponymous "sigma" meson is the one particle not described by the sigma-model; it is analogous to the longitudinal mode $\sigma(\mathbf{x})$ which is determined in terms of the $\vec{\pi}$ fields in our description above.

4.3.1 The Background Field Method

We would like to perform a renormalisation group analysis on the sigma model. There are a number of ways to proceed. First, we could Taylor expand the second term in the free energy for small $\vec{\pi}$. This would result in an infinite tower of interactions. We could then restrict attention to the first few, and do the kind of Wilsonian RG treatment we've seen before. This method works, but it butchers the underlying geometry and, in doing so, disguises what's really going on.

Instead, there is a better approach, first introduced in this context by Polyakov, called the *background field method*. First, suppose that $\mathbf{n}(\mathbf{x})$ takes some profile which varies slowly in space,

$$n^a(\mathbf{x}) = \tilde{n}^a(\mathbf{x})$$

This profile must obey $\tilde{\mathbf{n}} \cdot \tilde{\mathbf{n}} = 1$. This will play the role of our long wavelength modes.

On top of this background, we want to introduce short wavelength modes which change rapidly in space. To parameterise these modes, we first introduce *frame fields*. These are a basis of $N-1$ unit vectors $e^a_\alpha(\mathbf{x})$, with $a = 1, \ldots, N$ and $\alpha = 1, \ldots N-1$, which are orthogonal to $\tilde{n}^a(\mathbf{x})$,

$$\tilde{n}^a(\mathbf{x}) e^a_\alpha(\mathbf{x}) = 0 \quad \forall\ \alpha \qquad \text{and} \qquad e^a_\alpha(\mathbf{x}) e^a_\beta(\mathbf{x}) = \delta_{\alpha\beta} \tag{4.15}$$

The frame fields are, like $\tilde{n}^a$, slowly varying in space. There is an ambiguity in the definition of these frame fields; we can always rotate them by a local $O(N-1)$ transformation and we will still have a good set of frame fields.

The short wavelength modes sit on top of our original field $\tilde{\mathbf{n}}(\mathbf{x})$ and fluctuate in the direction of the frame fields. We call these $\chi_\alpha(\mathbf{x})$, and write the full configuration as

$$n^a(\mathbf{x}) = \tilde{n}^a(\mathbf{x}) \left(1 - \chi(\mathbf{x})^2\right)^{1/2} + \sum_{\alpha=1}^{N-1} \chi_\alpha(\mathbf{x}) e^a_\alpha(\mathbf{x}) \tag{4.16}$$

By construction, this configuration still satisfies the constraint (4.10). This is morally equivalent to our previous Fourier space decomposition $\phi = \phi_- + \phi_+$, but now in real space. We will integrate out the short wavelength modes χ and determine their effect on the long wavelength mode $\tilde{\mathbf{n}}$.

Integrating out Short Wavelengths

We have a short calculation ahead of us. Our plan is to expand the free energy to quadratic order in the short wavelength fields $\chi_\alpha(\mathbf{x})$ and then integrate them out, in exactly the same way that we integrated out the Fourier modes ϕ_+ previously. We will then interpret this in terms of an effective free energy for the long wavelength fields $\tilde{n}^a$ and, in particular, in terms of a renormalisation of the coupling e^2.

First, we have

$$\begin{aligned}\nabla n^a &= \nabla\tilde{n}^a\,(1-\chi^2)^{1/2} + \tilde{n}^a\nabla(1-\chi^2)^{1/2} + \nabla(\chi_\alpha e^a_\alpha)\\ &= \nabla\tilde{n}^a(1-\frac{1}{2}\chi^2) - \tilde{n}^a\chi_\alpha\nabla\chi_\alpha + \nabla(\chi_\alpha e^a_\alpha) + \mathcal{O}(\chi^2)\end{aligned}$$

The gradient term then becomes

$$\begin{aligned}(\nabla n^a)^2 = (\nabla\tilde{n}^a)^2(1-\chi^2) + (\nabla\chi_\alpha)^2 + \chi_\alpha\chi_\beta\nabla e^a_\alpha\nabla e^a_\beta + 2(\nabla\chi_\alpha)\chi_\beta e^a_\alpha\nabla e^a_\beta\\ + 2\nabla\tilde{n}^a\nabla(\chi_\alpha e^a_\alpha) + \mathcal{O}(\chi^3)\end{aligned}$$

where we have used the identities (4.15). One of the cross-terms vanishes by dint of the fact that $\tilde{n}^a\tilde{n}^a = 1$ so that $\tilde{n}^a\nabla\tilde{n}^a = 0$.

Our partition function is now

$$Z = \int \mathcal{D}\tilde{n}\;\delta(\tilde{n}^2-1)\;e^{-\frac{1}{2e^2}\int d^dx(\nabla\tilde{n})^2}\int\mathcal{D}\chi\;e^{-\frac{1}{2e^2}\int d^dx(\nabla\chi)^2}e^{-F_I[\tilde{n},\chi]}$$

where the interaction between $\tilde{n}^a$ and χ_α are captured in

$$\begin{aligned}F_I[\tilde{n}^a,\chi_\alpha] = \frac{1}{2e^2}\int d^dx\;\Big[\;&-\chi^2(\nabla\tilde{n}^a)^2 + \chi_\alpha\chi_\beta\nabla e^a_\alpha\nabla e^a_\beta\\ &+ 2(\nabla\chi_\alpha)\chi_\beta e^a_\alpha\nabla e^a_\beta + 2\nabla\tilde{n}^a\nabla(\chi_\alpha e^a_\alpha)\Big]\end{aligned}$$

As previously, we interpret the functional integral over χ as computing the expectation value of e^{-F_I} using the probability distribution $\exp(-\frac{1}{2e^2}\int d^dx\;(\nabla\chi)^2)$. In other words,

$$Z = \int\mathcal{D}\tilde{n}\;\delta(\tilde{n}^2-1)\;e^{-\frac{1}{2e^2}\int d^2x(\nabla\tilde{n})^2}\left\langle e^{-F_I[\tilde{n},\chi]}\right\rangle$$

The expectation value can be Taylor expanded,

$$\left\langle e^{-F_I[\tilde{n},\chi]}\right\rangle = 1 - \langle F_I\rangle + \frac{1}{2}\langle F_I^2\rangle + \ldots \tag{4.17}$$

As usual, the renormalisation group will generate many terms when we integrate out χ. Our interest is in how the leading order kinetic terms $(\nabla\tilde{n})^2$ are affected; all other

terms will turn out to be irrelevant. At leading order, it is sufficient to focus on $\langle F_I \rangle$. (Given the term linear in χ, one might wonder if such a term can be generated from $\langle F_I^2 \rangle$; a closer inspection shows that this in fact gives rise only to terms like $(\nabla \tilde{n})^4$.) We have

$$\langle F_I \rangle = \frac{1}{2e^2} \int d^d x \ \left(-\delta_{\alpha\beta} (\nabla \tilde{n}^a)^2 + \nabla e^a_\alpha \nabla e^a_\beta \right) \ \langle \chi_\alpha(\mathbf{x}) \chi_\beta(\mathbf{x}) \rangle \tag{4.18}$$

where we've used the fact that $\langle \chi(\mathbf{x}) \rangle = 0$ to lose the linear term and, on rotational grounds, $\langle (\nabla \chi_\alpha) \chi_\beta \rangle = 0$ to lose another term. The correlator that we want takes the same form that we calculated in previous sections. (See, for example, (2.21).)

$$\langle \chi_\alpha(\mathbf{x}) \chi_\beta(\mathbf{x}) \rangle = e^2 \delta_{\alpha\beta} I_d$$

where

$$I_d = \int_{\Lambda/\zeta}^{\Lambda} \frac{d^d q}{(2\pi)^d} \ \frac{1}{q^2}$$

Here we've introduced the limits on the integral to reflect the fact that, as in our earlier RG analysis, the short wavelength modes – which are here $\chi_\alpha(\mathbf{x})$ – have support only in $\Lambda/\zeta < k < \Lambda$. The integral is simple to perform; we have

$$I_d = \frac{\Omega_{d-1}}{(2\pi)^d} \Lambda^{d-2} \times \begin{cases} \zeta - 1 & d = 1 \\ \log \zeta & d = 2 \\ 1 - \zeta^{2-d} & d \geq 3 \end{cases}$$

where Ω_{d-1} is the area of the unit sphere $\mathbf{S}^{d-1}$. Substituting this into (4.18), it is clear that the first term corrects the kinetic term in the sigma model. But what of the second term? Using the fact that the correlator is proportional to $\delta_{\alpha\beta}$, it takes the form $\nabla e^a_\alpha \nabla e^a_\alpha$. We can massage this into the form we need using some identities. Between them, the fields $(\tilde{n}^a, e^a_\alpha)$ provide an orthonormal basis of $\mathbf{R}^N$. Inverting this, we have

$$\tilde{n}^a \tilde{n}^b + e^a_\alpha e^b_\alpha = \delta^{ab} \tag{4.19}$$

Using this, we can write

$$\nabla e^a_\alpha \nabla e^a_\alpha = \nabla e^a_\alpha \nabla e^b_\alpha (\tilde{n}^a \tilde{n}^b + e^a_\beta e^b_\beta)$$

But since $\tilde{n}^a e^a_\alpha = 0$, we have $\tilde{n}^a \nabla e^a_\alpha = -(\nabla \tilde{n}^a) e^a_\alpha$ so

$$\begin{aligned} \nabla e^a_\alpha \nabla e^a_\alpha &= e^a_\alpha e^b_\alpha (\nabla \tilde{n}^a)(\nabla \tilde{n}^b) + (e^a_\beta \nabla e^a_\alpha)(e^b_\beta \nabla e^b_\alpha) \\ &= \nabla \tilde{n}^a \nabla \tilde{n}^a + (e^a_\beta \nabla e^a_\alpha)(e^b_\beta \nabla e^b_\alpha) \end{aligned}$$

where, in going to the second line, we've used (4.19) again, together with the fact that $\tilde{n}^a \nabla \tilde{n}^a = 0$. The first term is just what we want; the second term is a new term that we can add to the sigma model and will be generated by RG. It is related to the geometric concept of torsion; it turns out to be irrelevant and we will not discuss it further here.

Both terms in (4.18) now give a contribution to $(\nabla \tilde{n}^a)^2$; the first is $-\delta_{\alpha\beta}\delta^{\alpha\beta} = -(N-1)$; the second is simply $+1$. The upshot of this is that the $\langle F_I \rangle$ includes the term

$$\langle F_I \rangle = (2-N) I_d \int d^d x \; \frac{1}{2} (\nabla \tilde{n}^a)^2$$

We now exponentiate this so that, to the order we're working at, (4.17) becomes $\langle e^{-F_I} \rangle = e^{-\langle F_I \rangle}$. This gives us an effective free energy for the long wavelength field $\tilde{\mathbf{n}}$,

$$F[\tilde{\mathbf{n}}] = \int d^d x \; \frac{1}{2e'^2} \nabla \tilde{\mathbf{n}} \cdot \nabla \tilde{\mathbf{n}}$$

with

$$\frac{1}{e'^2} = \frac{1}{e_0^2} + (2-N) I_d$$

Usually there are two further steps in the RG programme. First, we need to rescale our momentum cut-off back up to Λ, and in doing so rescale all length by $1/\zeta$. This proceeds as before. The second step, advertised in Section 3, is to rescale the fields so that the kinetic term is canonically normalised. This step is not for us, since the normalisation of the kinetic term is the only coupling we have. Instead, the fields are normalised correctly by imposing the constraint (4.10). The upshot is that we have the running coupling constant

$$\frac{1}{e^2(\zeta)} = \zeta^{d-2} \left(\frac{1}{e_0^2} + (2-N) I_d \right) \tag{4.20}$$

The first term comes from the naive dimensional analysis $[e^2] = 2 - d$ that we saw in (4.13); the second term is the one-loop correction from integrating out the high momentum modes.

Note that this second term vanishes when $N = 2$. This reflects the fact that the Goldstone boson in the XY-model is free; this can be seen in (4.4) where there are no interaction terms. Interesting things happen in the XY-model but we will postpone discussion to Section 4.4. In contrast, for $N \geq 3$ the Goldstone bosons are interacting, as can be seen for the $O(3)$ model in (4.5), and this drives the running of the coupling constant.

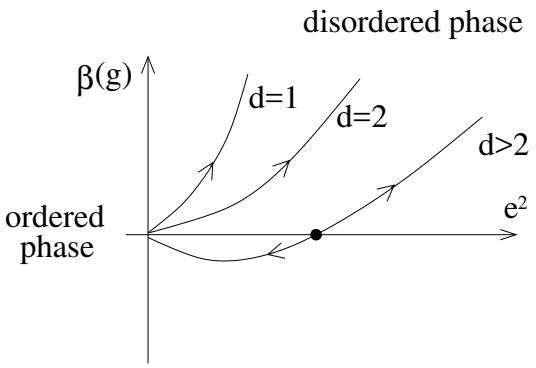

Figure 36: The beta function for e takes us away from the ordered phase at $e = 0$ when $d \leq 2$ and towards the ordered phase when $d > 2$. For $d = 2 + \epsilon$, there is an unstable, UV fixed point.

4.3.2 Asymptotic Freedom and the $d = 2 + \epsilon$ Expansion

Let's look more closely at the running coupling (4.20). To start, consider what happens in $d = 2$ dimensions. As previously, we write

$$\Lambda' = \frac{\Lambda}{\zeta} = \Lambda e^{-s}$$

and compute the beta function

$$\beta(e) := \frac{de}{ds} = (N-2)\frac{e^3}{4\pi}$$

The beta function vanishes for $N = 2$. This is because, as we mentioned above, the Goldstone mode in the XY model is non-interacting. However, for $N \geq 3$ the beta function is positive. It means that e^2 is an example of a *marginally relevant* coupling: as we flow to the infra-red, e^2 gets larger and the theory becomes strongly coupled. Correspondingly, the theory is weakly coupled in the ultra-violet. This property is known as *asymptotic freedom*, which refers to the fact that the theory is free at asymptotically high energies.

Asymptotically free theories are rather rare in physics. Perhaps the best known example is QCD or, in general, Yang-Mills theory with some small number of matter fields.

The sign of the beta function is telling us that, in $d = 2$ (and indeed in $d = 1$), the weakly coupled ordered phase that we started with is unstable. This is a manifestation of the Mermin-Wagner theorem that we mentioned in Section 4.2.3; there

are no Goldstone bosons in $d \leq 2$. Unfortunately, to really understand the infra-red physics in these low dimensions we will have to figure out how to deal with the strongly interacting theory. We will introduce a particularly useful approach in Section 4.3.3.

In higher dimensions, say $d = 3$, the beta function is negative. This means that the sigma-model flows towards weak coupling in the infra-red, telling us that the ordered phase is stable.

However, there is now something new we can do. We can look at what happens in dimension

$$d = 2 + \epsilon$$

Here the beta function takes the form

$$\frac{de}{ds} = -\frac{\epsilon}{2}e + (N-2)\,\frac{e^3}{4\pi}\Lambda^\epsilon$$

This has a fixed point that lies within the remit of perturbation theory, namely

$$e_\star^2 = \frac{2\pi\epsilon}{N-2}\Lambda^{-\epsilon} \tag{4.21}$$

However, in contrast to the story of Section 4.2, this is now a UV fixed-point, rather than an IR fixed point. How should we interpret this?

To understand this, let's recap the story so far. The $O(N)$ model, with unconstrained fields $\boldsymbol{\phi}$, has a Wilson-Fisher fixed point in dimension $d = 3$. This has one relevant deformation which is, roughly speaking, $\boldsymbol{\phi}^2$. If we turn on this relevant deformation with negative sign, we flow to the ordered phase which is described by our sigma model.

What we've seen above is this story in reverse. Starting from the ordered phase, described by the sigma model, we have managed to claw our way back up the RG flow to find a UV fixed point, at least in dimensions $d = 2+\epsilon$. One possibility is that this is identified with the Wilson-Fisher point, viewed through different eyes[10]. This provides us with a different handle on the $O(N)$ Wilson-Fisher fixed points for $N \geq 3$; we can either approach them from above using the $d = 4 - \epsilon$ expansion, or from below using the $d = 2 + \epsilon$ expansion.

[10]This was believed to be true for many years, but a detailed analysis suggests that the $2 + \epsilon$ expansion might be accessing a different fixed point. For example, the $2 + \epsilon$ expansion for $N = 3$ is conjectured to give a fixed point in which hedgehog defects are suppressed. For more details, see de Cesare and Rychkov arXiv:2505.21611.

To extract the critical exponent ν, we need to understand how e^2 is related to the temperature. From our definition of the path integral (4.12), we see that $1/e^2$ sits in the exponent where β would sit in a usual partition function. This motivates us to identify e^2 with temperature T and the fixed point $e_\star^2$ with the critical temperature $T_\star$. We then linearise about the fixed point by writing $e^2 = e_\star^2 + \delta e^2$ to find

$$\frac{d(\delta e^2)}{ds} = +\epsilon\, \delta e^2$$

This gives $\Delta_t = \epsilon$ and, correspondingly, the critical exponent

$$\nu = \frac{1}{\epsilon} \tag{4.22}$$

independent of N. To compute the critical exponent η, one could add the interaction $\int d^d x\ \mathbf{B} \cdot \mathbf{n}(\mathbf{x})$ or, alternatively, extract the anomalous dimension of $\mathbf{n}$ from the calculation above. One finds that

$$\eta = \frac{\epsilon}{N-2}$$

The $d = 2 + \epsilon$ expansion does not give great results if we just go ahead and plug in $\epsilon = 1$. But then, there is little reason that it should! For example, for $N = 3$, we can compare the best known results with mean field, and with the $d = 4 - \epsilon$ and $d = 2 + \epsilon$ expansions, where we work to first order and plug in $\epsilon = 1$. We have

	η	ν
MF	0	$\frac{1}{2}$
$d = 4 - \epsilon$	0	0.65
$d = 2 + \epsilon$	1	1
Actual	0.0386	0.702

Nonetheless, there is some utility in having two expansion parameters, coming from different ends. By going to higher powers in ϵ, one can try to use sophisticated matching techniques to join together the two expansions and get a better handle on the values of the critical exponents for $d = 3$.

4.3.3 Large N

So far, we have seen that the dynamics of interacting Goldstone modes (i.e. $N \geq 3$) becomes strongly coupled in $d = 2$ dimensions. But we have yet to figure out what actually happens.

Questions like this are typically hard. As $e^2 \to \infty$, it naively appears that all field configurations contribute equally to the path integral, no matter how wildly they vary and how far they are from the saddle point. We have very few techniques to deal with such situations. Often we have to turn to some hidden and surprising symmetry, or to some unusual limit where the theory is soluble.

In the present case, it turns out that such a limit exists: it is $N \to \infty$. To proceed, we first rewrite the delta-function in the path integral (4.12) as

$$\begin{aligned} Z &= \int \mathcal{D}\mathbf{n}\ \delta(\mathbf{n}^2 - 1)\ \exp\left(-\frac{1}{2e_0^2}\int d^dx\ \nabla\mathbf{n}\cdot\nabla\mathbf{n}\right) \\ &= \int \mathcal{D}\mathbf{n}\,\mathcal{D}\sigma\ \exp\left(-\frac{1}{2e_0^2}\int d^dx\ \nabla\mathbf{n}\cdot\nabla\mathbf{n} - \frac{i}{2e_0^2}\int d^dx\ \sigma\,(\mathbf{n}\cdot\mathbf{n}-1)\right) \end{aligned} \tag{4.23}$$

Here the field $\sigma(\mathbf{x})$ plays the role of a Lagrange multiplier; integrating it out gives us back the delta-function, imposing the field constraint $\mathbf{n}^2 = 1$.

Now, however, we're left with a free energy which is quadratic in the $\mathbf{n}$. Instead of integrating out σ, we can instead integrate out $\mathbf{n}$. This gives us

$$Z = \int \mathcal{D}\sigma\ \det{}^{-N/2}\left(-\nabla^2 + i\sigma(\mathbf{x})\right)\ \exp\left(\frac{i}{2e_0^2}\int d^dx\ \sigma\right)$$

Here the determinant of the differential operator should be viewed, in the usual way, as the product of all its eigenvalues, with a truncation associated to the UV cut-off Λ, reflecting the fact that the eigenfunctions can't oscillate at high frequencies. This determinant will, in general, be a complicated function of σ, and it does not look as if we are any closer to evaluating the path integral. We can, however, use the standard "$\log\det = \mathrm{tr}\log$" identity to write the partition function as

$$Z = \int \mathcal{D}\sigma\ \exp\left(-\frac{N}{2}\mathrm{tr}\log\left(-\nabla^2 + i\sigma\right) + \frac{i}{2e_0^2}\int d^dx\ \sigma\right) \tag{4.24}$$

The factor of N in front of the first term is what gives us hope because, in the limit $N \to \infty$, this term is then crying out to be evaluated by saddle point. However, we're still left with the second term. We can only apply the saddle point technique to this too if we scale the coupling e_0^2 with N in a particular way. Specifically, we send $N \to \infty$, keeping $e_0^2 N$ fixed.

The path integral (4.24) is then dominated by the minimum. We use the identity $\delta\,\mathrm{tr}\log X = \mathrm{tr}\,X^{-1}\delta X$, to find that the saddle point is

$$\frac{N}{2}G(\mathbf{x},\mathbf{x}) = \frac{1}{2e_0^2} \tag{4.25}$$

where $G(\mathbf{x}, \mathbf{x}')$ is the Green's function for the operator $(-\nabla^2 + i\sigma(\mathbf{x}))$. This equation looks somewhat foreboding, but is rather simple in Fourier space. First, we look for constant solutions, of the form

$$\sigma(\mathbf{x}) = -i\mu^2$$

Note the factor of i; our saddle point sits on the complex plane, but is nonetheless still applicable. The saddle point (4.25) then becomes simpler in Fourier space: we have

$$\int^{\Lambda} \frac{d^d k}{(2\pi)^d} \frac{1}{k^2 + \mu^2} = \frac{1}{e_0^2 N} \tag{4.26}$$

where we've explicitly included the UV cut-off Λ in the integral. This equation should now be viewed as an equation for μ^2.

Large N in $d = 2$

It is perhaps no surprise by now to learn that solutions to (4.26) depend on the dimension d. Our main concern was with the fate of the Goldstone bosons in $d = 2$. Here the integral gives us

$$\frac{1}{4\pi} \log\left(\frac{\Lambda^2 + \mu^2}{\mu^2}\right) = \frac{1}{e_0^2 N} \tag{4.27}$$

If we start with a weakly coupled theory in the UV, $e_0^2 N \ll 1$, then we can self-consistently assume that $\mu \ll \Lambda$ to find the solution

$$\mu \approx \Lambda e^{-2\pi/e_0^2 N} \tag{4.28}$$

This simple formula is interesting for several reasons. First, let's understand the physical interpretation of setting $\sigma = -i\mu^2 \neq 0$. Referring back to (4.23), we see that it induces an effective quadratic term for $\mathbf{n}^2$ in the free energy. This kind of term was supposed to be prohibited by Goldstone's theorem, but here we see that it is generated – at least in the large N limit – by thermal fluctuations in $d = 2$ dimensions. This means that the Goldstone bosons in $d = 2$ are no longer gapless. Correspondingly, if we compute their correlator using (4.23), we will see that it decays exponentially, with a finite correlation length given by $\xi \sim 1/\mu$.

The second interesting fact about (4.28) is that the dynamically generated scale μ is exponentially smaller than the UV cut-off Λ. Indeed, the function $e^{-1/x}$ has the lovely property that its Taylor expansion around $x = 0$ vanishes at every order in x. This means that the gap μ will not show up in any order in perturbation theory in e_0. We say that it is a non-perturbative effect.

Although the calculation we presented above is valid for $N \gg 1$, it turns out that the conclusions hold for all $N \geq 3$; that is, for any theory of interacting Goldstone bosons in $d = 2$ dimensions. This means that there is no phase transition for $O(N)$ models with $N \geq 3$ in $d = 2$ dimensions. As we lower the temperature, mean field theory suggests that we enter an ordered phase with gapless excitations, but this is misleading: instead, thermal fluctuations destroy both the order and gapless modes.

The discussion above carries over directly to quantum field theory, where non-linear sigma models in $d = 1+1$ dimensions are also of great interest. Here the interpretation of the calculation is that the Goldstone modes – which appear to be massless in the classical action – get a mass due to quantum effects. If one didn't think carefully about the meaning of quantum field theory this appears miraculous because the sigma-model in $d = 1+1$ dimensions has only a dimensionless coupling e_0^2. Yet somehow, the theory generates a mass out of this dimensionless coupling, a phenomenon that is known as *dimensional transmutation*. The reason that this is mathematically possible is because a quantum field theory, like its statistical counterpart, is not defined by the classical action (or free energy) alone. It also requires a UV cut-off Λ. And, as we see in (4.28), it is this UV cut-off which provides the dimensional scale for the mass.

Finally, I should mention that if you can do a calculation like the one above for Yang-Mills in $d = 3 + 1$ dimensions (or, indeed, in $d = 4$ dimensions) then fame and fortune awaits. The massless gauge bosons that appear in the classical action are strongly believed to get a mass through quantum effects, yet this remains to be proven. This is the famous "Yang-Mills mass gap" problem. The $O(N)$ sigma model in $d = 1 + 1$ dimensions provides a useful analogy for how this might happen.

Large N in $d > 2$

We can also ask if our large N analysis can shed any light on the Wilson-Fisher fixed point in $d = 3$. (Or, if you're willing for dimensions to wander, in $2 < d < 4$.) Here we find something interesting. The saddle point equation (4.26) has different behaviour for $2 < d < 4$ and for $d > 4$,

$$\frac{1}{e_0^2 N} = \int^{\Lambda} \frac{d^d k}{(2\pi)^d} \frac{1}{k^2 + \mu^2} \sim \begin{cases} \Lambda^{d-2} - \mu^2 \Lambda^{d-4} & d \geq 4 \\ \Lambda^{d-2} - \mu^{d-2} & 2 < d < 4 \end{cases} \tag{4.29}$$

where we haven't been careful about the coefficients in front of either term, except to stress that the second term comes with a negative sign relative to the first. (We also analysed the behaviour of this integral in (2.13), but there only kept the leading term.) This equation now has rather different behaviour than the corresponding equation

(4.27) in $d = 2$. In particular, when the theory is weakly coupled at the cut-off scale, in the sense that

$$e_0^2 N \lesssim \Lambda^{2-d}$$

there are no solutions to (4.29) for μ^2. In this case, one finds that the saddle point of the free energy actually arises when $\mathbf{n}$ gets an expectation value. In other words, it's reconfirming our expectation that the low-energy physics is that of Goldstone bosons.

In contrast, as the theory becomes more strongly coupled at the cut-off scale, there is a critical value

$$e_\star^2 N \sim \Lambda^{2-d}$$

at which solutions to (4.29) for μ start to appear.

As in the previous section, we identify the deviation from $e_\star^2$ with the temperature,

$$T - T_c \sim e^2 - e_\star^2$$

We can then ask how the correlation length $\xi \sim 1/\mu$ diverges as we approach this critical coupling from above. Here the story is different for $2 < d < 4$ and $d > 4$, because of the different behaviour of the subleading term in (4.29). For $2 < d < 4$, we have

$$T - T_c \sim \xi^{2-d}$$

which gives the critical exponent

$$\nu = \frac{1}{d-2}$$

Rather wonderfully, this agrees precisely (for all $2 < d < 4$) with the result of our $d = 2 + \epsilon$ expansion (4.22), and with the large N limit of our result from the $d = 4 - \epsilon$ expansion (4.7). Indeed, this result is exact in the $N \to \infty$ limit and can be used as the starting point for a $1/N$ expansion

Meanwhile, when $d > 4$ we can read off the behaviour from (4.29); we have

$$T - T_c \sim \xi^{-2} \quad \Rightarrow \quad \nu = \frac{1}{2}$$

This, of course, is the mean field value that we expect.

4.4 The Kosterlitz-Thouless Transition

The Mermin-Wagner theorem means that any system with a continuous symmetry has no ordered phase in $d = 2$ dimensions. As we saw in the previous sections, for the $O(N)$ model with $N \geq 3$, the would-be Goldstone modes are interacting and become gapped as a result of the thermal fluctuations. This means that these models do not exhibit a phase transition as the temperature is lowered.

However, the results of the previous section do not hold for the XY model with $N = 2$. In this case, the sigma-model coupling does not run, and the system remains gapless at low temperatures. As we will now see, the resulting physics is rather more subtle and interesting.

The first surprise is that the $d = 2$ XY model does exhibit a phase transition as the temperature is lowered. However, it is somewhat different from the kind of phase transitions that we have met so far. In particular, as we saw in Section 4.2.3, thermal fluctuations mean that there can be no spontaneous breaking of continuous symmetry in $d = 2$ and, correspondingly, there is no local order parameter that distinguishes the two phases. Instead, that task falls on the correlation function.

In the high temperature phase, we work with the complex field ψ. The free energy has a quadratic term $\mu^2|\psi|^2$ and, as we've now seen many times (starting in (2.29)) the correlation function decays exponentially

$$\langle \psi^\dagger(\mathbf{x})\psi(0)\rangle = \frac{e^{-r/\xi}}{\sqrt{r}} \tag{4.30}$$

with $\xi \sim 1/\mu^2$. In contrast, in the low temperature phase we have $\mu^2 < 0$ and, as we described in Section 4.2, we can write $\psi = Me^{i\theta}$, with the long distance physics dominated by θ. To leading order, we can write the free energy as

$$F[\theta] = \frac{1}{2e^2}\int d^2x\ (\nabla\theta)^2 \tag{4.31}$$

Very low temperatures correspond to $e^2 \ll 1$.

The correlation function for this Goldstone mode exhibits a log divergence (4.9),

$$\langle \theta(\mathbf{x})\theta(0)\rangle = -\frac{e^2}{2\pi}\log(\Lambda r)$$

To compare to (4.30), we should look at

$$\langle e^{-i\theta(\mathbf{x})}e^{i\theta(0)}\rangle = \langle e^{-i(\theta(\mathbf{x})-\theta(0))}\rangle = e^{-\langle(\theta(\mathbf{x})-\theta(0))^2\rangle/2}$$

where the final equality follows because we are dealing with a Gaussian theory (4.31) and so can employ Wick's identity (3.34). We learn that, in the low temperature phase, the correlation function for the XY model takes power-law form

$$\langle e^{-i\theta(\mathbf{x})} e^{i\theta(0)} \rangle \sim \frac{1}{r^\eta} \tag{4.32}$$

where the anomalous dimension η is given by

$$\eta = \frac{e^2}{2\pi}$$

Note that this power-law does not occur just at a critical point, but for a range of temperatures. As we increase the coupling e^2, which is equivalent to increasing the temperature, the anomalous dimension increases. We see that the coupling e^2 in the XY model (which can be traced to the $\mu^2 < 0$ coupling in the original theory) is something rather rare: it is an example of a genuinely marginal coupling.

The correlation function exhibits two different behaviours in the high temperature (4.30) and low temperature phases (4.32). This suggests that there may be a phase transition between them. The fact that the order parameter for this phase transition is non-local – it involves the position of fields at two distinct points rather than one – is our first hint that this phase transition has a slightly different smell from others. As we will now see, this is not the only thing that sets it apart.

4.4.1 Vortices

The mechanism for the phase transition can be found within the sigma model approach (4.31), but involves something a little novel. The novelty arises from the fact that, in contrast to the Ising field $\phi(\mathbf{x})$ that we worked with in Section 3, the field $\theta(\mathbf{x})$ is periodic. There can be field configurations, localised around a point $\mathbf{x} = \mathbf{X}$, in which $\theta(\mathbf{x})$ winds some number of times,

$$\oint \nabla\theta \cdot d\mathbf{x} = 2\pi n \qquad \text{with } n \in \mathbf{Z}$$

Crucially, the winding number n must be an integer so that θ comes back to itself up to a 2π shift. A configuration with $n = 1$ is referred to as a *vortex*; when $n = -1$, it is an *anti-vortex*. These are examples of *topological defects*. The configurations of lattice spins that correspond to a vortex and anti-vortex are shown in the figures.

At the location of the vortex, $\mathbf{x} = \mathbf{X}$ the field $\theta(\mathbf{x})$ is not well defined. One way to proceed is to revert to the original XY model and allow the magnitude of ψ to vary close to the core. However, for our purposes it will suffice to do something simpler: we just admit ignorance on short distance scales, and say that the vortex has some core size which we denote as a. This will later play the role of the UV cut-off in our system.

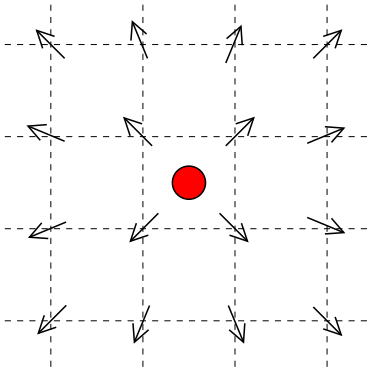

Figure 37: A vortex...

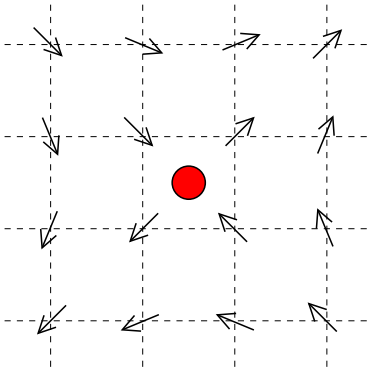

Figure 38: ... and an anti-vortex.

We'll start by giving a rough and ready derivation of the effect of vortices. A configuration with winding n has $\nabla\theta = \frac{n}{r^2}(y, -x)$, and so free energy

$$F_{\text{vortex}} = \frac{1}{2e^2}\int d^2x\ (\nabla\theta)^2 = \frac{\pi n^2}{e^2}\log\left(\frac{L}{a}\right) + F_{\text{core}} \tag{4.33}$$

where, in addition to the UV cut-off a, we also need to place the system in a finite size L to avoid a long-distance divergence in the energy. We've also included a contribution from the vortex core region $r < a$ which depends on the microscopic details. Note that the free energy of multi-vortices, with $|n| > 1$, scales as n^2 and so is energetically disfavoured. For this reason we focus on configurations with $n = \pm 1$.

The logic now is very similar to the story of domain walls in dimension $d = 1$ that we met in Section 1.3.3. The probability of a vortex configuration arising in the system is enhanced by the fact that it can sit anywhere; this gives an extra factor of $(L/a)^2$. We then have

$$p(\text{vortex}) = \left(\frac{L}{a}\right)^2 \frac{e^{-F_{\text{vortex}}}}{Z} = \frac{e^{-F_{\text{core}}}}{Z}\left(\frac{L}{a}\right)^{2-\pi/e^2}$$

We see that, when e^2 surpasses a critical value,

$$e^2 > e^2_{KT} = \frac{\pi}{2} \tag{4.34}$$

then there is no suppression of vortices; their entropy, coming from the fact that they can sit anywhere on the plane, wins out over their energetic cost. As in the previous section, e^2 can be viewed as the temperature of the system, and e^2_{KT} translates into a temperature scale T_{KT}, above which vortices proliferate. This, it turns out, is responsible for the change in the behaviour of the correlation function, with the vortices randomising the phase θ, destroying the delicate power-law fall off (4.32).

This phase transition, driven by proliferation of vortices, is known as the *Kosterlitz-Thouless transition*, and is important both for superfluid films, and for the melting of two-dimensional lattices, where the defects play the role of vortices. It is also known as the BKT transition, as the Russian theorist Berezinskii was the first to appreciate that such a transition is possible, although he didn't fill in all the details. It is sometimes referred to as a *topological phase transition*, because it is driven by the proliferation of topological defects.

Michael Kosterlitz and David Thouless are both Brits, educated in Cambridge, who subsequently moved to the US. In fact, Thouless was the first Director of Studies of physics in Churchill College. They shared the 2016 Nobel prize in physics for their work on this transition.

A Coulomb Gas of Vortices

The quick discussion above shows that vortices proliferate when e^2 gets too large. But we can do better. The first step is to appreciate that what really emerges as we increase the coupling is a gas of vortices and anti-vortices. The Kosterlitz-Thouless transition is better thought of as an unbinding of vortex-anti-vortex pairs.

To see this in more detail, we will first look at the interactions between vortices. To this end, it's useful to think in terms of the vector field

$$\mathbf{v} = \nabla\theta \tag{4.35}$$

In the context of superfluids, this is the velocity field. The equation of motion for θ is

$$\nabla^2\theta = 0 \quad \Rightarrow \quad \nabla\cdot\mathbf{v} = 0 \tag{4.36}$$

A smooth vector field defined by (4.35) would obey $\nabla\times\mathbf{v} = 0$. However, in the presence of vortices, the θ field admits singularities and, correspondingly, the velocity field obeys

$$\nabla\times\mathbf{v} = 2\pi\hat{\mathbf{z}}\sum_i n_i\,\delta^2(\mathbf{x}-\mathbf{X}_i) \tag{4.37}$$

where $\hat{\mathbf{z}}$ is the unit vector out of the plane, and $n_i = \pm 1$ determines the charge of the vortex at position $\mathbf{x} = \mathbf{X}$.

We can perform a change of variables to transform (4.36) and (4.37) into more familiar equations. We define

$$E_i = \epsilon_{ij}v_j \quad \Rightarrow \quad \begin{pmatrix} E_1 \\ E_2 \end{pmatrix} = \begin{pmatrix} v_2 \\ -v_1 \end{pmatrix}$$

and the equations of motion then become

$$\nabla \times \mathbf{E} = 0 \quad \text{and} \quad \nabla \cdot \mathbf{E} = 2\pi \sum_i n_i \delta^2(\mathbf{x} - \mathbf{X}_i) \tag{4.38}$$

These are the Maxwell equations for the auxiliary electric field $\mathbf{E}$, with the vortices acting as "electric charge". This means that we can import our machinery from our course on Electromagnetism; the only difference is that our electric field lives in $d = 2$ spatial dimensions. For example, to determine the interaction between two vortices, we need to solve the Gauss' law equation in (4.38). We do this by writing $\mathbf{E} = -\nabla\chi$, where

$$\chi(\mathbf{x}) = -\sum_i n_i \log\left(\frac{|\mathbf{x} - \mathbf{X}_i|}{a}\right) \tag{4.39}$$

The free energy (4.31) can be expressed in terms of the electric field as

$$F = \int d^2x \ \frac{1}{2e^2}\mathbf{E}\cdot\mathbf{E} = \int d^2x \ \frac{1}{2e^2}(\nabla\chi)^2$$

This looks very similar to our starting point (4.31), except that the relationship between the original field $\theta(\mathbf{x})$ and the new field $\chi(\mathbf{x})$ is given by $\partial_i\theta = \epsilon_{ij}\partial_j\chi$, which is not straightforward to solve. However, it is now straightforward to compute the free energy. First integrating by parts, we have

$$F = \int d^2x \ -\frac{1}{2e^2}\chi\nabla\cdot\mathbf{E} = \frac{\pi}{e^2}\sum_{i\neq j} n_i n_j \log\left(\frac{|\mathbf{X}_i - \mathbf{X}_j|}{a}\right) + \sum_i n_i^2 F_{\text{core}}$$

where, to get the second equality, we've substituted in the expressions (4.38) and (4.39) and, for the cases $i = j$, replaced our expression with the energy of the core of the vortex. We learn that the interaction between vortices grows logarithmically. This is the Coulomb force in $d = 2$ dimensions; it is repulsive for vortex pairs, and attractive for a vortex-anti-vortex pair.

We can now use this expression to write down an expression for the partition function of the XY sigma model. There are two contributions. To isolate these, we decompose the velocity field as

$$\mathbf{v} = \mathbf{v}_{sw} + \mathbf{v}_{\text{vortex}}$$

The first of these obeys

$$\nabla \times \mathbf{v}_{sw} = 0$$

This is circulation-free flow in the absence of vortices. It describes the contribution from the fluctuations of θ: we call these "spin waves". The second contribution comes from vortices and obeys $\nabla \cdot \mathbf{v}_{\rm vortex} = 0$. The free energy, and hence the partition function, then factorise into two

$$Z = Z_{\rm sw}\, Z_{\rm vortex}$$

The spin wave piece is harmless; it shows no sign of a phase transition. Meanwhile, the vortex piece contains contributions from all number of vortices and anti-vortices. We restrict attention to configurations that have equal number of vortices and anti-vortices, as these don't suffer the IR divergence (4.33) in their free energy. We're left with

$$Z_{\rm vortex} = 1 + \sum_{p=1}^{\infty} \frac{y^{2p}}{(p!)^2} \prod_{i=1}^{p} \int d^2\mathbf{X}_i^+ d^2\mathbf{X}_i^- \ \exp\left(\frac{\pi}{e^2} \sum_{i\neq j} n_i n_j \log\left(\frac{|\mathbf{X}_i - \mathbf{X}_j|}{a} \right) \right) \quad (4.40)$$

where the initial 1 comes from the configuration with no vortices, and the $y = e^{-F_{\rm core}}/a^2$ can be thought of as the fugacity of vortices. Here $\mathbf{X}_i^+$ denote the positions of p vortices, and $\mathbf{X}_i^-$ the positions of p anti-vortices. Meanwhile, the argument of the logarithm involves the sum over the separations $\mathbf{X}_i - \mathbf{X}_j$, $i \neq j$ of all $2p$ (anti)-vortices, regardless of their charge. Finally, the integral should be taken over all $|\mathbf{X}_i - \mathbf{X}_j| > a$ so that the cores of vortices do not overlap.

$Z_{\rm vortex}$ is the partition function of a neutral Coulomb gas in the grand canonical ensemble, with the $\pm$ charges interacting through the 2d Coulomb force.

We would like to understand the phase structure of $Z_{\rm vortex}$ as the coupling e^2 is varied. There are different ways to go about this. One possibility is to implement the RG directly on $Z_{\rm vortex}$. This proceeds by integrating out the vortices that are separated by some short distance scale $\tilde{a}$, effectively increasing the UV cut-off scale a. Here we will take an alternative approach. We will first map the Coulomb gas to a seemingly very different problem, one which will be more amenable to the traditional RG methods that we've been using in this course.

4.4.2 From Coulomb Gas to Sine-Gordon

The Coulomb gas (4.40) lies in the same universality class as the so-called Sine-Gordon model. This is a theory of a real scalar field $\phi(\mathbf{x})$, with free energy,

$$F = \int d^2x \ \frac{1}{2}(\nabla\phi)^2 - \lambda\cos(\beta\phi) \quad (4.41)$$

The name is a physicist's version of a joke: it is a play on "Klein-Gordon" theory[11].

[11] Sidney Coleman has a famous paper on this model which starts with the sentence "The Sine-Gordon equation is the sophomoric but unfortunately standard name for...".

We start by giving a quick derivation of the equivalence between the Sine-Gordon model and the Coulomb gas. We will be fairly heuristic. It turns out that this mapping is somewhat simpler if we revert back to a spatial lattice, rather than working in the continuum.

To this end, we introduce a lattice with spacing a with lattice sites $\mathbf{X}_\alpha$. On each lattice site, we include a variable V_α which can take values $V_\alpha = -1, 0, +1$. The interpretation is that if $V_\alpha = +1$, there is a vortex at this site; if $V_\alpha = -1$ there is an anti-vortex; and if $V_\alpha = 0$ the site is empty. We allow V_α to only take these three values to reflect the fact that two vortices feel a large repulsion, which means that they effectively have a hard core, while a vortex and an anti-vortex annihilate to nothing if they come too close.

The grand canonical partition function (4.40) can then be rewritten as

$$Z_{\text{vortex}} \sim \sum_{\{V_\alpha\}} \exp\left(\frac{\pi}{e^2} \sum_{\alpha \neq \beta} V_\alpha V_\beta \log\left(\frac{|\mathbf{X}_\alpha - \mathbf{X}_\beta|}{a} \right) - \sum_\alpha V_\alpha^2 F_{\text{core}} \right) \tag{4.42}$$

We restrict the sum $\{V_\alpha\}$ to configurations that are neutral, so $\sum_\alpha V_\alpha = 0$. This mimics the sum over all numbers and positions of vortex-anti-vortex pairs.

To proceed, we will use the fact that the log that appears in Z_{vortex} is the Green's function for the 2d Laplacian ∇^2. In general, we have

$$\int \mathcal{D}\phi \ \exp\left(-\int d^2x \ \frac{1}{2}(\nabla\phi)^2 + f(\mathbf{x})\phi(\mathbf{x}) \right) \sim \exp\left(-\frac{1}{4\pi} \int d^2x d^2y \ f(\mathbf{x}) \log|\mathbf{x} - \mathbf{y}| f(\mathbf{y}) \right)$$

where we've dropped a factor of the determinant $\det(-\nabla^2)^{-1/2}$ which gives an unimportant overall contribution to the partition function. Using this, the partition function (4.42) can be rewritten yet again as

$$Z_{\text{vortex}} \sim \sum_{\{V_\alpha\}} \int \mathcal{D}\phi \ \exp\left(-\int d^2x \ \frac{1}{2}(\nabla\phi)^2 + \sum_\alpha \frac{2\pi i}{e} V_\alpha \phi_\alpha - V_\alpha^2 F_{\text{core}} \right)$$

where we're using a slightly unholy mix of continuous notation and discrete notation. You should think of $\phi_\alpha = \phi(\mathbf{X}_\alpha)$ as the value of $\phi(\mathbf{x})$ at the lattice site, and write your preferred discretised version of the kinetic term. Now we can do the sum over the V_α;

we have

$$\begin{aligned} Z_{\text{vortex}} &= \int \mathcal{D}\phi \ \exp\left(-\frac{1}{2}\int d^2x \ (\nabla\phi)^2\right) \ \prod_\alpha \sum_{V_\alpha=-1,0,+1} e^{\frac{2\pi i}{e}V_\alpha\phi_\alpha - V_\alpha^2 F_{\text{core}}} \\ &= \int \mathcal{D}\phi \ \exp\left(-\frac{1}{2}\int d^2x \ (\nabla\phi)^2\right) \ \prod_\alpha \left[1 + 2e^{-F_{\text{core}}}\cos(2\pi\phi_\alpha/e)\right] \\ &\approx \int \mathcal{D}\phi \exp\left(-\frac{1}{2}\int d^2x \ (\nabla\phi)^2 + \frac{2}{a^2}e^{-F_{\text{core}}}\int d^2x \ \cos\left(\frac{2\pi\phi}{e}\right)\right) \end{aligned}$$

This is the Sine-Gordon model, as promised. Although our derivation used an underlying lattice, the final result is expressed as a continuum field theory, and this is the form we will use moving forward. As always, however, the memory of the lattice will remain in the UV cut-off scale a. The dictionary between the couplings in (4.41) and those of the original XY-model are

$$\lambda = \frac{2e^{-F_{\text{core}}}}{a^2} \quad \text{and} \quad \beta = \frac{2\pi}{e}$$

We will now see how these couplings fare under the renormalisation group.

4.4.3 RG Flows in Sine-Gordon

We apply our standard RG programme to the Sine-Gordon model,

$$F = \int d^2x \ \frac{1}{2}(\nabla\phi)^2 - \lambda_0\cos(\beta_0\phi)$$

where we've added the subscript 0 to reflect the fact that this free energy is defined at the cut-off scale Λ.

What follows next is familiar. We work in Fourier space and decompose the field ϕ into low and high momentum modes,

$$\phi_{\mathbf{k}} = \phi^-_{\mathbf{k}} + \phi^+_{\mathbf{k}}$$

where $\phi^+_{\mathbf{k}}$ includes all modes in the momentum shell $\Lambda/\zeta < k < \Lambda$. We also define $\phi^-(\mathbf{x})$ and $\phi^+(\mathbf{x})$ in real space as the inverse Fourier transform of $\phi^-_{\mathbf{k}}$ and $\phi^+_{\mathbf{k}}$ respectively.

We then integrate out the high momentum modes to leave ourselves with an effective free energy,

$$F'[\phi^-] = F_0[\phi^-] - \log\left\langle e^{-F_I[\phi^- + \phi^+]}\right\rangle_+ \tag{4.43}$$

where

$$F_0[\phi] = \int d^2x\ \frac{1}{2}(\nabla\phi)^2 \quad \text{and} \quad F_I[\phi] = -\lambda_0 \int d^2x\ \cos(\beta_0\phi)$$

and the expectation value reflects the fact that we're integrating out the fast momentum modes, weighted with

$$\langle e^{-F_I[\phi^- + \phi^+]} \rangle_+ = \int \mathcal{D}\phi^+\ e^{-F_0[\phi^+]}\, e^{-F_I[\phi^- + \phi^+]}$$

Our goal is to compute this effective free energy.

First Order in λ_0

We will assume that λ is suitably small so that the leading order term is

$$\begin{aligned}\left\langle \exp\left(\lambda_0 \int d^2x \cos\big(\beta(\phi^- + \phi^+)\big)\right)\right\rangle_+ &\approx 1 + \lambda_0 \int d^2x\ \left\langle \cos\big(\beta(\phi^- + \phi^+)\big)\right\rangle_+ \\ &= 1 + \frac{\lambda_0}{2}\int d^2x\ \sum_{\sigma=\pm 1} e^{i\beta\sigma\phi^-}\left\langle e^{i\beta\sigma\phi^+}\right\rangle_+\end{aligned}$$

Meanwhile, we can use our handy Wick identity (3.34) to write

$$\left\langle e^{i\beta\sigma\phi^+(\mathbf{x})}\right\rangle_+ = e^{-\beta^2/2\,\langle\phi^+(\mathbf{x})\phi^+(\mathbf{x})\rangle_+}$$

The propagator for the fast mode, evaluated at the same point, is

$$\langle\phi^+(0)\phi^+(0)\rangle_+ = \int_{\Lambda/\zeta}^{\Lambda} \frac{d^2k}{(2\pi)^2}\ \frac{1}{k^2} = \frac{1}{2\pi}\log\zeta \tag{4.44}$$

The upshot of this calculation is that the interaction term is renormalised after integrating out the high momentum modes, and becomes

$$F'[\phi^-] = \int d^2x\ \frac{1}{2}(\nabla\phi^-)^2 - \lambda_0\zeta^{-\beta^2/4\pi}\ \cos(\beta\phi^-)$$

and the coupling λ_0 becomes

$$\lambda' = \lambda_0\zeta^{-\beta^2/4\pi}$$

The next step of the RG is the rescaling, $\mathbf{x} \to \mathbf{x}' = \mathbf{x}/\zeta$. It's simple to check that the rescaling of the field is trivial. With this, our free energy becomes

$$F'[\phi] = \int d^2x\ \frac{1}{2}(\nabla\phi)^2 - \lambda(\zeta)\ \cos(\beta\phi)$$

with

$$\lambda(\zeta) = \lambda_0\, \zeta^{2-\beta^2/4\pi} \tag{4.45}$$

Already we can see the essence of the Kosterlitz-Thouless phase transition in this equation. When β is suitably large,

$$\beta^2 > 8\pi \quad \Rightarrow \quad e^2 < \frac{\pi}{2}$$

then the effect of RG is to reduce λ. This means that the coupling $\lambda\cos(\beta\phi)$ is an irrelevant operator and, as we flow towards the infra-red, $\lambda \to 0$. In this case, the free energy for ϕ is given just by the gradient terms, and the correlation function will exhibit a power-law fall off.

In contrast, when β is small,

$$\beta^2 < 8\pi \quad \Rightarrow \quad e^2 > \frac{\pi}{2}$$

the operator $\lambda\cos(\beta\phi)$ becomes relevant, growing as we go towards the infra-red[12]. Now the minimum of the potential is at $\phi = 0 \bmod 2\pi/\beta$. Expanding the cos potential about this minimum gives a quadratic term for ϕ and correlation functions will now be exponentially suppressed, with a finite correlation length. It is perhaps surprising that our sophisticated RG analysis gives exactly the same value for the critical coupling $g_2 = \pi/2$ as our previous, hand-waving discussion of vortex proliferation (4.34).

Using our earlier result (4.32) for the anomalous dimension, we see that at the phase transition, $e^2 = \pi/2$, the system exhibits a universal anomalous dimension,

$$\langle e^{-i\theta(\mathbf{x})}\, e^{i\theta(0)}\rangle \sim \frac{1}{r^{1/4}}$$

Second Order in λ_0

At order λ_0^2, we find ourselves with the double cosine

$$\begin{aligned}\langle\cos(\beta(\phi^-_{\mathbf{x}} + \phi^+_{\mathbf{x}}))\, \cos(\beta(\phi^-_{\mathbf{y}} + \phi^+_{\mathbf{y}}))\rangle_+ &= \frac{1}{4}\sum_{\sigma=\pm1}\Big[e^{i\sigma\beta(\phi^-_{\mathbf{x}}+\phi^-_{\mathbf{y}})}\,\langle e^{i\sigma\beta(\phi^+_{\mathbf{x}}+\phi^+_{\mathbf{y}})}\rangle_+ \\ &\qquad + e^{i\sigma\beta(\phi^-_{\mathbf{x}}-\phi^-_{\mathbf{y}})}\,\langle e^{i\sigma\beta(\phi^+_{\mathbf{x}}-\phi^+_{\mathbf{y}})}\rangle_+\Big] \\ &= \frac{1}{2}\Big[\cos(\beta(\phi^-_{\mathbf{x}} + \phi^-_{\mathbf{y}}))\, e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}+\phi^+_{\mathbf{y}})^2\rangle_+} \\ &\qquad + \cos(\beta(\phi^-_{\mathbf{x}} - \phi^-_{\mathbf{y}}))\, e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}-\phi^+_{\mathbf{y}})^2\rangle_+}\Big]\end{aligned}$$

[12]This same result is derived in a very different way, using conformal field theory, in the lectures on String Theory. It can be found in Claim 2 of Section 4.3.3 where, in the notation of that course, the operator e^{ikX} is shown to have dimension $\Delta = \alpha' k^2/2$. A quick check of the conventions for the propagator shows that we should set $\alpha' = 1/2\pi$ so that $\Delta < 2$ and the operator is relevant if $k^2 > 8\pi$.

where, as a space-saving measure, I've put the spatial position as a subscript, $\phi(\mathbf{x}) = \phi_{\mathbf{x}}$. Taking the logarithm in (4.43) means that we subtract the disconnected diagrams, $\langle\cos(\beta(\phi^-_{\mathbf{x}} + \phi^+_{\mathbf{x}}))\rangle^2_+$. The upshot is that, at order λ_0^2, the effective free energy includes the piece

$$F'[\phi^-]\big|_{e^2} = \frac{\lambda_0^2}{4}\int d^2x d^2y \ \Big\{ \cos(\beta(\phi^-_{\mathbf{x}} + \phi^-_{\mathbf{y}}))\ \Big[e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}+\phi^+_{\mathbf{y}})^2\rangle_+} - e^{-\beta^2/2\langle\phi^+_{\mathbf{x}}\phi^+_{\mathbf{x}}\rangle_+}e^{-\beta^2/2\langle\phi^+_{\mathbf{y}}\phi^+_{\mathbf{y}}\rangle_+}\Big] \\ +\ \cos(\beta(\phi^-_{\mathbf{x}} - \phi^-_{\mathbf{y}}))\ \Big[e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}-\phi^+_{\mathbf{y}})^2\rangle_+} - e^{-\beta^2/2\langle\phi^+_{\mathbf{x}}\phi^+_{\mathbf{x}}\rangle_+}e^{-\beta^2/2\langle\phi^+_{\mathbf{y}}\phi^+_{\mathbf{y}}\rangle_+}\Big]\ \Big\}$$

The expectation values that sit in the exponents are given by

$$G(\mathbf{x}-\mathbf{y};\zeta) = \langle\phi^+(\mathbf{x})\phi^+(\mathbf{y})\rangle = \int_{\Lambda/\zeta}^{\Lambda}\frac{d^2k}{(2\pi)^2}\ \frac{e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})}}{k^2}$$

Previously we needed only $G(0;\zeta) = \frac{1}{2\pi}\log\zeta$; now we see that the correlator at spatially separated points also arises. We have

$$F'[\phi^-]\big|_{e^2} = \frac{1}{4}\lambda_0^2\zeta^{-\beta^2/2\pi}\int d^2x d^2y \ \Big\{ \cos(\beta(\phi^-_{\mathbf{x}} + \phi^-_{\mathbf{y}}))\ \Big[e^{-\beta^2 G(\mathbf{x}-\mathbf{y})} - 1\Big] \\ +\ \cos(\beta(\phi^-_{\mathbf{x}} - \phi^-_{\mathbf{y}}))\ \Big[e^{+\beta^2 G(\mathbf{x}-\mathbf{y})} - 1\Big]\ \Big\}$$

At first sight, we seem to have a non-local free energy involving a double integral. To massage it into something more familiar, we need to realise that the function $G(\mathbf{r})$ receives contributions from a small sliver of Fourier modes, and so decays quickly at distances $r > \zeta/\Lambda$. This means that the functions $[e^{\pm\beta G(\mathbf{x}-\mathbf{y})} - 1]$ are non-zero only in a small window $|\mathbf{x}-\mathbf{y}| \sim \zeta/\Lambda$.

We write $\mathbf{y} = \mathbf{x} + \mathbf{v}$ and Taylor expand the cos factors in the integral for small $\mathbf{v}$. For the first, we have simply

$$\cos(\beta(\phi^-_{\mathbf{x}} + \phi^-_{\mathbf{y}})) \approx \cos(2\beta\phi^-_{\mathbf{x}})$$

The second is more interesting for us; we have

$$\cos(\beta(\phi^-_{\mathbf{x}} - \phi^-_{\mathbf{y}})) \approx 1 - \frac{\beta^2 v^2}{2}(\nabla\phi^-)^2$$

Our free energy, at order λ_0^2, then becomes

$$F'[\phi^-]\big|_{e^2} = \frac{1}{2}\lambda_0^2\int d^2x\Big\{A_1(\zeta)\ \cos(2\beta\phi) + A_2(\zeta)(\nabla\phi)^2\ \Big\} + \text{const.} \tag{4.46}$$

where all the messy details have been absorbed into two functions

$$A_1(\zeta) = \frac{1}{2}\zeta^{-\beta^2/2\pi}\int d^2v\ \left[e^{-\beta^2 G(v)} - 1\right]$$
$$A_2(\zeta) = -\frac{1}{4}\zeta^{-\beta^2/2\pi}\beta^2\int d^2v\ v^2\left[e^{+\beta^2 G(v)} - 1\right] \tag{4.47}$$

We see that the RG flow has generated two terms for us in (4.46). The first, $\cos(2\beta\phi)$, is something new: it can be viewed as the effect of two vortices, which we didn't include in our original Sine-Gordon model but is generated upon integrating out high momentum modes. We will not need this here. The second term is something familiar: it is renormalisation of our kinetic term. The final steps of the RG procedure tell us to rescale space, $\mathbf{x} \to \mathbf{x}' = \mathbf{x}/\zeta$, but also rescale the field ϕ so that the kinetic term remains canonically normalised, as in (3.6). We have

$$\phi'(\mathbf{x}') = \sqrt{1 + \lambda_0^2 A_2(\zeta)}\,\phi(\mathbf{x})$$

This rescaling gets spat out inside the potential, where it has the effect of renormalising our other coupling, β. We have

$$\beta(\zeta) = \beta_0\ \left(1 + \lambda_0^2 A_2(\zeta)\right)^{-1/2} \approx \beta_0(1 - \frac{1}{2}\lambda_0^2 A_2(\zeta)) \tag{4.48}$$

Recall that, in terms of our original XY model, $\beta^2 = 4\pi^2/e^2$. Looking back to our XY sigma model (4.31), we see that it renormalises the $1/e^2$ coefficient of the kinetic term. This is sometimes called the spin wave "stiffness", since it measures how difficult it is to twist the spins. The intuition behind the result (4.48) is that a gas of vortex-anti-vortex pairs screens the spins, reducing their stiffness.

Beta Functions

Our task now is to understand the global properties of the resulting RG flow. We write down the beta function in terms of $\zeta = e^s$. From (4.45) we have

$$\frac{d\lambda}{ds} = \left(2 - \frac{\beta^2}{4\pi}\right)\lambda = \left(2 - \frac{\pi}{e^2}\right)\lambda$$

Meanwhile, from (4.48), we get

$$\frac{d\beta}{ds} = -C(\beta)\beta^3\lambda^2$$

where $C(\beta) > 0$ is a positive function that we could extract from the formula (4.47); it's exact value will not be important for us. Written in terms of $e^2 = 4\pi^2/\beta^2$, this

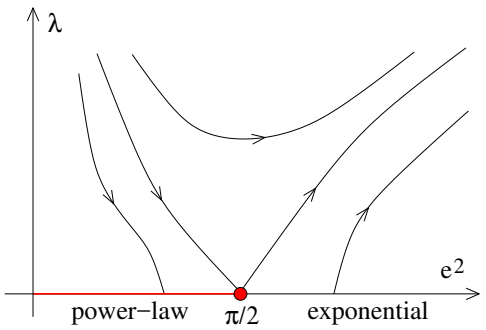

Figure 39: RG flows for the KT transition.

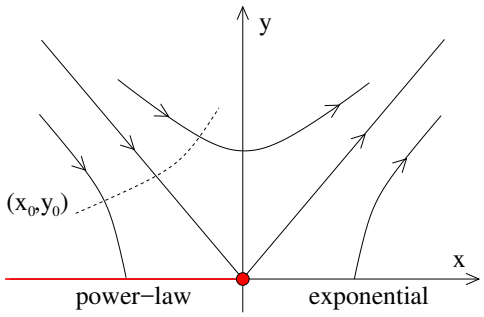

Figure 40: RG flows, zoomed in at the critical point.

latter RG equation becomes

$$\frac{de^2}{ds} = 8\pi^2 C\lambda^2$$

The global structure of the RG flow is shown in the figure. To get a sense for this, first note that if λ is small, and $e^2 \ll \pi/2$, then λ will be rapidly driven to zero; this is the low-temperature phase in which correlation functions drop off with a power-law. Meanwhile if $e^2 \gg \pi/2$ then λ will be pushed large. This is the high temperature phase, with a non-vanishing correlation length. In the high temperature phase, we see that $e^2(s) \to \infty$ as $s \to \infty$. Meanwhile, in the low temperature phase $e^2(s)$ is finite as $s \to \infty$. At the transition, there is a jump

$$\Delta \left.\frac{1}{e^2}\right|_{s\to\infty} = \frac{2}{\pi}$$

To see the larger picture, it's best to zoom in to the phase transition itself. We define

$$x = e^2 - \frac{\pi}{2} \quad \text{and} \quad y = \sqrt{8\pi^2 C}\,\lambda$$

For x and y small, the beta functions become

$$\frac{dx}{ds} = y^2 \quad \text{and} \quad \frac{dy}{ds} \approx \frac{4}{\pi}xy$$

From this we can compute

$$\frac{dx^2}{ds} = 2xy^2 \quad \text{and} \quad \frac{dy^2}{ds} = \frac{8}{\pi}xy^2 \quad \Rightarrow \quad \frac{d}{ds}\left(x^2 - \frac{\pi}{4}y^2\right) = 0$$

In other words, close to the critical point, the flows are hyperbolae. A general flow can be written as

$$x^2 - \frac{\pi}{4}y^2 = J = x_0^2 - \pi y_0^2$$

where (x_0, y_0) are the initial "bare" values of the couplings at the cut-off scale. There are two regions with $J > 0$; these correspond to the low and high temperature regimes that we discussed above. The separatrix at $J = 0$ is given by $x = \pm\sqrt{\pi}y/2$. The line with $x = -\sqrt{\pi}y/2$ flows directly to the critical point. The line with $x = +\sqrt{\pi}y$ flows away from the critical point.

Suppose that we start on the left of the figure, with $x < 0$. Then the initial data, shown as a dotted line in the figure, is (x_0, y_0). As we vary this data, we pass through the phase transition. In this sense, it is natural to think of this initial data as a function of the temperature $(x_0(T), y_0(T))$, with

$$J(T) \sim T_c - T$$

ensuring that we hit the critical point when $J = 0$.

Finally, we can ask about the correlation length ξ. To compute this, we can look at flows with $J < 0$ which don't quite hit the $y = 0$ axis. We have

$$\frac{dx}{ds} = y^2 = \frac{4}{\pi}(x^2 - J) = \frac{4}{\pi}(x^2 + |J|)$$

which we can solve to give

$$s = \frac{\pi}{4\sqrt{|J|}}\left[\tan^{-1}\left(\frac{x}{\sqrt{|J|}}\right) - \tan^{-1}\left(\frac{x_0}{\sqrt{|J|}}\right)\right]$$

This has the slightly odd property that $x \to \infty$ as s remains finite. This is an artefact of our approximation but, nonetheless, can be used to our advantage. By the time $x \approx 1$, we also have $y \approx 1$ and the theory is in the gapped phase. We can stop the RG flow here and use this as a proxy for our correlation length which, as we approach the phase transition from above, scales as

$$\xi \sim ae^s \sim \exp\left(\frac{1}{\sqrt{|J|}}\right) \sim \exp\left(\frac{1}{\sqrt{T - T_c}}\right)$$

We are used to a fairly soft divergence in the correlation length as we approach the critical temperature. For the Kosterlitz-Thouless transition, the change is much more dramatic. This also has an effect on the thermodynamic free energy which, as an extensive quantity, scales as $F \sim (L/\xi)^2$. As we approach the phase transition from above, we have

$$F_{\text{thermo}} \sim \frac{1}{\xi^2} \sim \exp\left(-\frac{1}{\sqrt{T - T_c}}\right)$$

This is a very weak singularity. There is no discontinuity in the heat capacity. Moreover, there is no discontinuity in any derivative of the free energy. In terms of Ehrenfest's original classification, the Kosterlitz-Thouless transition is rather strange: it is a phase transition of infinite order.