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.

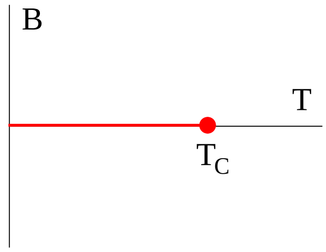

Figure 31: The phase diagram of the Ising model (again).

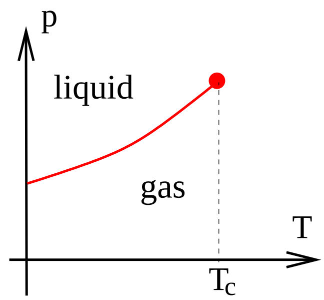

Figure 32: The phase diagram of the liquid-gas system (again).

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 allows 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

$$\phi_a(\mathbf{x}) \rightarrow R_a{}^b \phi_b(\mathbf{x})$$

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

$$F[\boldsymbol{\phi}(\mathbf{x})] = \int d^dx \ \left[\frac{\gamma}{2}\nabla\boldsymbol{\phi}\cdot\nabla\boldsymbol{\phi} + \frac{\mu^2}{2}\boldsymbol{\phi}\cdot\boldsymbol{\phi} + g(\boldsymbol{\phi}\cdot\boldsymbol{\phi})^2 + \ldots\right]$$

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, $\psi \rightarrow e^{i\alpha}\psi$. The first few terms are

$$F[\psi(\mathbf{x})] = \int d^dr \ \left[\frac{\gamma}{2}\nabla\psi^\star\cdot\nabla\psi + \frac{\mu^2}{2}|\psi|^2 + g|\psi|^4 + \ldots\right] \tag{4.1}$$

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

$$\gamma \nabla^2 \psi = \mu^2 \psi + 4g|\psi|^2 \psi + \ldots$$

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 $\alpha_2 < 0$ 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:

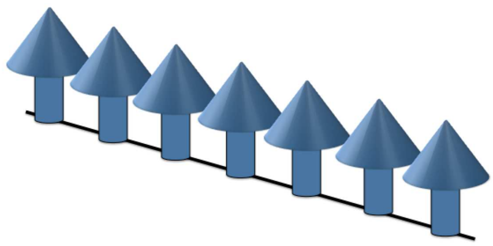

has the same energy as the configuration in which all the spins point like this:

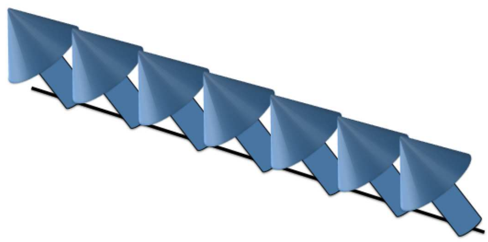

This infinitely degenerate choice of ground states gives us something new. We can consider configurations in which we stay within the space of ground states, but the direction varies in space. For such configurations, the part of the free energy $f(\boldsymbol{\phi}) = \frac{\mu^2}{2}|\boldsymbol{\phi}|^2 + g|\boldsymbol{\phi}|^4 + \ldots$ 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

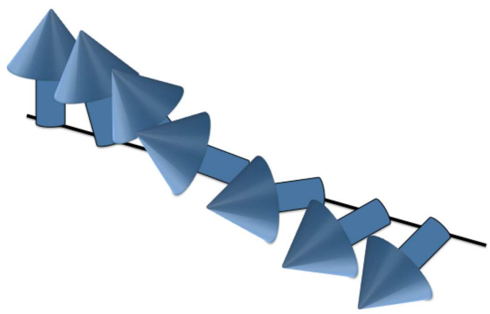

which can be made to cost an arbitrarily small amount of energy, by stretching the winding over longer and longer distance scales.

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 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 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 $\mathbf{S}^{N-1}$.

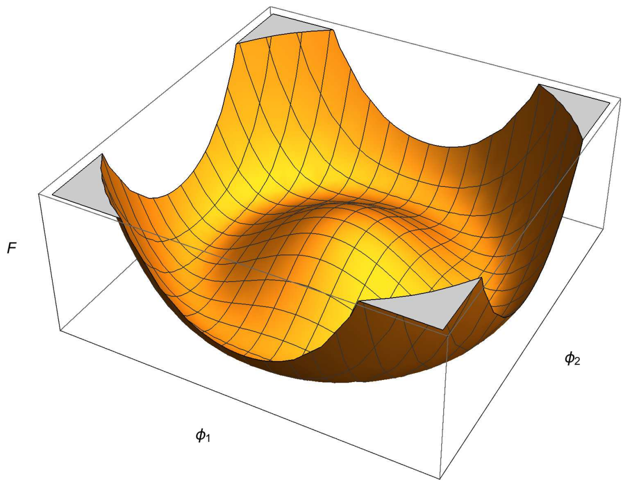

Figure 33: The so-called "Mexican hat" potential for the XY-model in the ordered phase.

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 $M = M_0 = \sqrt{-\mu^2/2g}$, 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

$$\begin{aligned} F[M,\theta,\phi] = \int d^dx \quad & \frac{\gamma}{2}(\nabla \tilde{M})^2 + |\mu^2|\tilde{M}^2 + g\tilde{M}^4 + \ldots \\ & + \frac{\gamma}{2}M_0^2\left[(\nabla\theta)^2 + \sin^2\theta\,(\nabla\phi)^2\right] + \ldots \end{aligned} \tag{4.5}$$

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

$$F[\theta] = \int d^dx\, \frac{\gamma}{2}M_0^2(\nabla\theta)^2 + \ldots$$

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

$$\langle\theta(\mathbf{x})\theta(\mathbf{y})\rangle = \frac{1}{\gamma M_0^2}\int \frac{d^dk}{(2\pi)^d}\frac{e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})}}{k^2} \tag{4.6}$$

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 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.

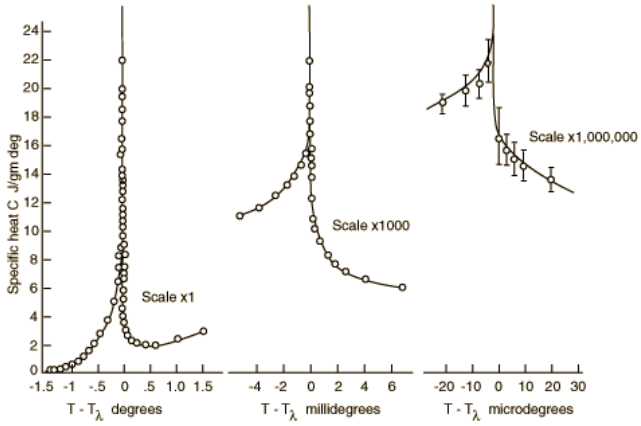

Figure 34: The heat capacity of helium at the superfluid transition. This system lies in the XY universality class. The data above is well described by the function $C \sim C_0 + A|T - T_c|^{-\alpha}$ with $\alpha \approx -0.16$ and $A < 0$.

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[9]. 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 above 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

[9]This data is taken from Buckingham and Fairbank, "The Nature of the Lambda Transition", in Progress in Low Temperature Physics III, 1961.

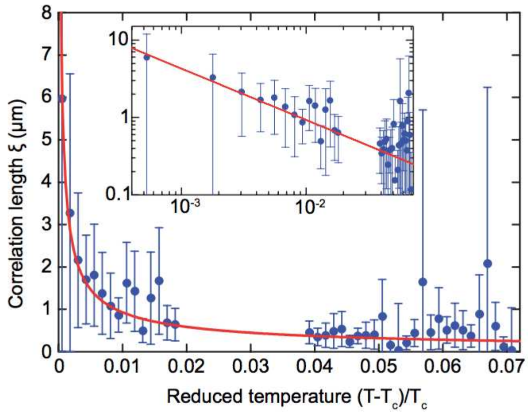

Figure 35: The correlation length for a BEC at the critical point. This system lies in the XY universality class and has $\nu = 0.67 \pm 0.13$.

to be $\nu = 0.67 \pm 0.13$, in good agreement with the theoretical prediction (albeit with fairly large error bars)[10].

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

$$\beta F[\phi] = \int d^dx \, \left[\frac{1}{2}\nabla\phi_a \cdot \nabla\phi_a + \frac{1}{2}\mu_0^2\, \phi_a\phi_a + g_0(\phi_a\phi_a)^2 + \ldots\right]$$

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.

[10]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

$$\beta F_I[\phi] = \int d^dx \; g(\phi_a\phi_a)^2$$

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

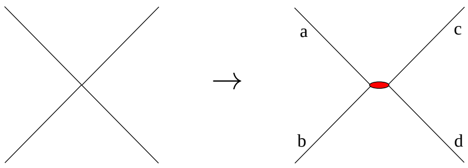

$$\sim \; g_0\,\delta_{ab}\delta_{cd}$$

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:

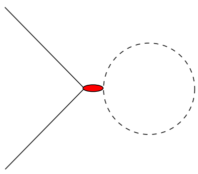

$$= 2g_0 \int \prod_{i=1}^{4} \frac{d^dk_i}{(2\pi)^d}\, \phi^-_{a,\mathbf{k}_1}\phi^-_{a,\mathbf{k}_2} \times \left\langle \phi^+_{b,\mathbf{k}_3}\phi^+_{b,\mathbf{k}_4}\right\rangle_+ \times (2\pi)^d\delta^d(\sum_i \mathbf{k}_i)$$

The other contribution has a different contraction between internal indices

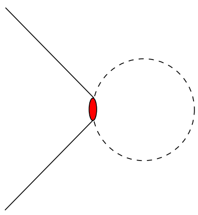

$$= 4g_0 \int \prod_{i=1}^{4} \frac{d^dk_i}{(2\pi)^d}\, \phi^-_{a,\mathbf{k}_1}\phi^-_{b,\mathbf{k}_2} \times \left\langle \phi^+_{a,\mathbf{k}_3}\phi^+_{b,\mathbf{k}_4}\right\rangle_+ \times (2\pi)^d\delta^d(\sum_i \mathbf{k}_i)$$

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 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,

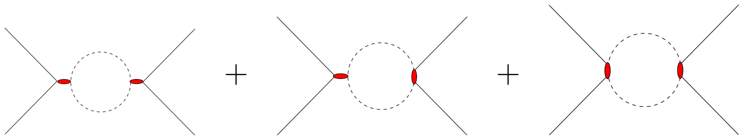

Only the first of these has a closed loop, unattached to the external legs. This comes with a factor of N, while the other two do not. A careful computation of the relevant overall factors shows that these diagrams renormalise the quartic coupling as

$$g_0 \to g_0' = g_0 - 4(N+8)g_0^2 \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d} \frac{1}{(q^2+\mu_0^2)^2}$$

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 (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

$$F[\mathbf{n}] = \int d^dx \; \frac{1}{2e^2} \nabla\mathbf{n} \cdot \nabla\mathbf{n} \tag{4.11}$$

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

$$Z = \int \mathcal{D}\mathbf{n} \; \delta\big(\mathbf{n}(\mathbf{x})^2 - 1\big) \; \exp\left(-\frac{1}{2e^2}\int d^dx \; \nabla\mathbf{n} \cdot \nabla\mathbf{n}\right) \tag{4.12}$$

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

$$\begin{aligned} F[\vec{\pi}] &= \int d^dx \; \frac{1}{2e^2} \left[\nabla\vec{\pi}\cdot\nabla\vec{\pi} + \nabla\sigma\nabla\sigma\right] \\ &= \int d^dx \; \frac{1}{2e^2} \left[\nabla\vec{\pi}\cdot\nabla\vec{\pi} + \frac{(\vec{\pi}\cdot\nabla\vec{\pi})^2}{1-\vec{\pi}^2}\right] \end{aligned} \tag{4.14}$$

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

$$(\nabla n^a)^2 = (\nabla\tilde{n}^a)^2(1-\chi^2) + (\nabla\chi_\alpha)^2 + \chi_\alpha\chi_\beta\nabla e^a_\alpha\nabla e^a_\beta + 2\nabla\tilde{n}^a\nabla(\chi_\alpha e^a_\alpha) + \mathcal{O}(\chi^3)$$

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

$$F_I[\tilde{n}^a,\chi_\alpha] = \frac{1}{2e^2}\int d^dx\ \Big[-\chi^2(\nabla\tilde{n}^a)^2 + \chi_\alpha\chi_\beta\nabla e^a_\alpha\nabla e^a_\beta + 2\nabla\tilde{n}^a\nabla(\chi_\alpha e^a_\alpha)\Big]$$

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,

$$Z = \int \mathcal{D}\tilde{n}\ \delta(\tilde{n}^2-1)\ e^{-\frac{1}{2e^2}\int d^2x(\nabla\tilde{n})^2}\ \Big\langle e^{-F_I[\tilde{n},\chi]}\Big\rangle$$

The expectation value can be Taylor expanded,

$$\Big\langle e^{-F_I[\tilde{n},\chi]}\Big\rangle = 1 - \langle F_I\rangle + \frac{1}{2}\langle F_I^2\rangle + \ldots \tag{4.17}$$

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

$$\langle F_I \rangle = \frac{1}{2e^2} \int d^dx \, \left(-\delta_{\alpha\beta}(\nabla \tilde{n}^a)^2 + \nabla e^a_\alpha \nabla e^a_\beta\right) \langle \chi_\alpha(\mathbf{x})\chi_\beta(\mathbf{x})\rangle \tag{4.18}$$

where we've used the fact that $\langle \chi(\mathbf{x})\rangle = 0$ to lose the linear term. 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

$$I_d = \int_{\Lambda/\zeta}^{\Lambda} \frac{d^dq}{(2\pi)^d} \frac{1}{q^2}$$

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 then, 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.

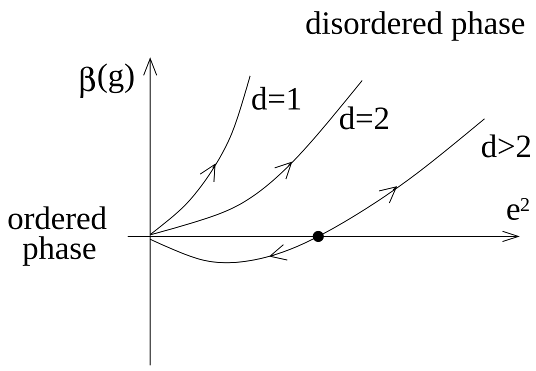

Figure 36: The beta function for e takes us away from the ordered phase at $e = 0$ when $d \leq 2$ and away from the ordered phase when $d > 2$. For $d = 2 + \epsilon$, there is an unstable, UV fixed point.

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

$$\beta(g) := \frac{de}{ds} = (N-2)\frac{e^3}{4\pi}$$

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 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're 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 $d = 2 + \epsilon$. It is natural to identify this with the Wilson-Fisher point, viewed through different eyes. 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.

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 $e^2_\star$ with the critical temperature $T_\star$. We then linearise about the fixed point by writing $e^2 = e^2_\star + \delta e^2$ 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 $\int d^dx\ \mathbf{B}\cdot\mathbf{n}(\mathbf{x})$ 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	0.61
$d = 2+\epsilon$	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

$$\begin{aligned} Z &= \int \mathcal{D}n\ \delta(\mathbf{n}^2 - 1)\ \exp\left(-\frac{1}{2e_0^2}\int d^dx\ \nabla\mathbf{n}\cdot\nabla\mathbf{n}\right) \\ &= \int \mathcal{D}\mathbf{n}\,\mathcal{D}\sigma\ \exp\left(-\frac{1}{2e_0^2}\int d^dx\ \nabla\mathbf{n}\cdot\nabla\mathbf{n} - \frac{i}{2e_0^2}\int d^dx\ \sigma\,(\mathbf{n}\cdot\mathbf{n}-1)\right) \end{aligned} \tag{4.23}$$

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 is 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 $d > 4$,

$$\frac{1}{e_0^2 N} = \int^\Lambda \frac{d^dk}{(2\pi)^d} \frac{1}{k^2 + \mu^2} \sim \begin{cases} \Lambda^{d-2} - \mu^2 \Lambda^{d-4} & d \geq 4 \\ \Lambda^{d-2} - \mu^{d-2} & 2 < d < 4 \end{cases} \tag{4.29}$$

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

$$\langle e^{-i\theta(\mathbf{x})}e^{i\theta(0)}\rangle = \frac{1}{r^\eta} \tag{4.32}$$

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 transitions 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.

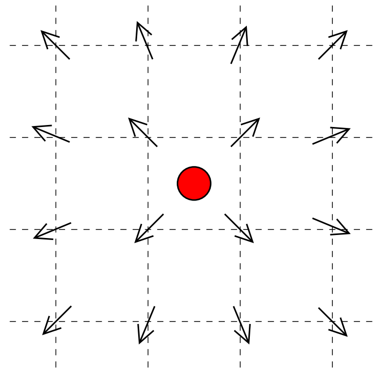

Figure 37: A vortex...

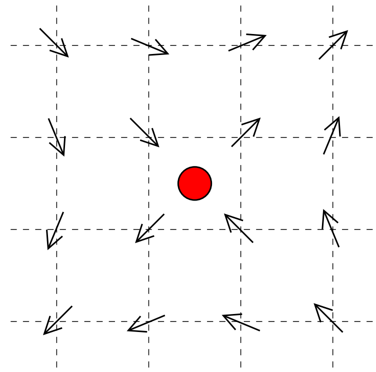

Figure 38: ... and an anti-vortex.

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

$$F_{\rm vortex} = \frac{1}{2e^2}\int d^2x\; (\nabla\theta)^2 = \frac{\pi n^2}{e^2}\log\left(\frac{L}{a}\right) + F_{\rm core} \tag{4.33}$$

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

$$p({\rm vortex}) = \left(\frac{L}{a}\right)^2 \frac{e^{-F_{\rm vortex}}}{Z} = \frac{e^{-F_{\rm core}}}{Z}\left(\frac{L}{a}\right)^{2-\pi/e^2}$$

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

$$\nabla\times\mathbf{v} = 2\pi\hat{\mathbf{z}}\sum_{i=1} n_i\,\delta^2(\mathbf{x}-\mathbf{X}_i) \tag{4.37}$$

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

$$Z_{\rm vortex} = \sum_{p=0}^{\infty} \frac{y^{2p}}{(p!)^2} \prod_{i=1}^{p} \int d^2\mathbf{X}_i^+ d^2\mathbf{X}_i^- \ \exp\left(\frac{\pi}{e^2}\sum_{i\neq j} n_i n_j \log\left(\frac{|\mathbf{X}_i - \mathbf{X}_j|}{a}\right)\right) \qquad (4.40)$$

where $y = e^{-F_{\rm core}}/a^2$ 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 on 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,

$$F = \int d^2x\ \frac{1}{2}(\nabla\phi)^2 - \lambda\cos(\beta\phi) \qquad (4.41)$$

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

$$Z_{\rm vortex} \sim \sum_{\{V_\alpha\}} \exp\left(\frac{\pi}{e^2}\sum_{\alpha\neq\beta} V_\alpha V_\beta \log\left(\frac{|\mathbf{X}_\alpha - \mathbf{X}_\beta|}{a}\right) - \sum_\alpha V_\alpha^2 F_{\rm core}\right) \tag{4.42}$$

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 $Z_{\rm vortex}$ 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

$$Z_{\rm vortex} \sim \sum_{\{V_\alpha\}} \int \mathcal{D}\phi\ \exp\left(-\int d^2x\ \frac{1}{2}(\nabla\phi)^2 + \sum_\alpha \frac{2\pi i}{e} V_\alpha \phi_\alpha - V_\alpha^2 F_{\rm core}\right)$$

where we're using a slightly unholy mix of continuous notation and discrete notation. You should think of $\phi_\alpha = \phi(X_\alpha)$ 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 that 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 leading order term is

$$\begin{aligned}\left\langle \exp\left(\lambda_0 \int d^2x \cos\left(\beta(\phi^- + \phi^+)\right)\right)\right\rangle_+ &\approx 1 + \lambda_0 \int d^2x \; \left\langle \cos\left(\beta(\phi^- + \phi^+)\right)\right\rangle_+ \\ &= 1 + \frac{\lambda_0}{2} \int d^2x \sum_{\sigma=\pm 1} e^{i\beta\sigma\phi^-} \left\langle e^{i\beta\sigma\phi^+}\right\rangle_+\end{aligned}$$

Meanwhile, we can use our handy Wick identity (3.34) to write

$$\left\langle e^{i\beta\sigma\phi^+(\mathbf{x})}\right\rangle_+ = e^{-\beta_0^2/2\, \langle\phi^+(\mathbf{x})\phi^+(\mathbf{x})\rangle_+}$$

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

$$\begin{aligned}
\langle \cos(\beta(\phi^-_{\mathbf{x}} + \phi^+_{\mathbf{x}}))\, \cos(\beta(\phi^-_{\mathbf{y}} + \phi^+_{\mathbf{y}}))\rangle_+ &= \frac{1}{4}\sum_{\sigma=\pm 1}\Big[e^{i\sigma\beta(\phi^-_{\mathbf{x}}+\phi^-_{\mathbf{y}})}\,\langle e^{i\sigma\beta(\phi^+_{\mathbf{x}}+\phi^+_{\mathbf{y}})}\rangle_+ \\
&\qquad\qquad + e^{i\sigma\beta(\phi^-_{\mathbf{x}}-\phi^-_{\mathbf{y}})}\,\langle e^{i\sigma\beta(\phi^+_{\mathbf{x}}-\phi^+_{\mathbf{y}})}\rangle_+\Big] \\
&= \frac{1}{2}\Big[\cos(\beta(\phi^-_{\mathbf{x}} + \phi^-_{\mathbf{y}}))\, e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}+\phi^+_{\mathbf{y}})^2\rangle_+} \\
&\qquad\qquad + \cos(\beta(\phi^-_{\mathbf{x}} - \phi^-_{\mathbf{y}}))\, e^{-\beta^2/2\langle(\phi^+_{\mathbf{x}}-\phi^+_{\mathbf{y}})^2\rangle_+}\Big]
\end{aligned}$$

[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 $|\mathbf{x}-\mathbf{y}| \sim \Lambda/\zeta$.

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

$$F'[\phi^-]\big|_{e^2} = \frac{1}{2}\lambda_0^2\int d^2x\Big\{A_1(\zeta)\,\cos(2\beta\phi) + A_2(\zeta)(\nabla\phi)^2\Big\} + \text{const.} \qquad (4.46)$$

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 thought of 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 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

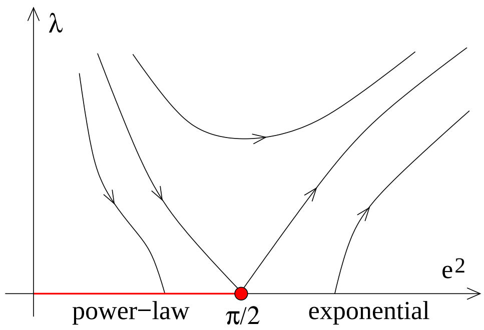

Figure 39: RG flows for the KT transition.

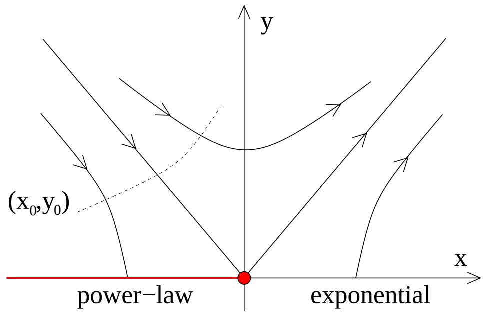

Figure 40: RG flows, zoomed in at the critical point.

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

$$\frac{dx^2}{ds} = 2xy^2 \quad \text{and} \quad \frac{dy^2}{ds} = \frac{2}{\pi}xy^2 \quad \Rightarrow \quad \frac{d}{ds}(x^2 - \pi y^2) = 0$$

In other words, close to the critical point, the flows are hyperbolae. A general flow can be written as

$$x^2 - \pi y^2 = J = x_0^2 - \pi y_0^2$$

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 $x = \pm\sqrt{\pi}y$. The line with $x = -\sqrt{\pi}y$ 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

$$\frac{dx}{ds} = y^2 = \frac{1}{\pi}(x^2 - J) = \frac{1}{\pi}(x^2 + |J|)$$

which we can solve to give

$$s = \frac{\pi}{\sqrt{|J|}}\left[\tan^{-1}\left(\frac{x}{\sqrt{|J|}}\right) - \tan^{-1}\left(\frac{x_0}{\sqrt{|J|}}\right)\right]$$

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're 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.