8

Disorder and defects in topological phases

So far, we have largely stayed clear of the discussion of disorder. The first justification one might attempt is that for a gapped phase like the $\mathbb{Z}_2$ dimer liquid, disorder should not have too much of an effect as long as it is weak enough that states cannot move across the gap. However, this imposes a restriction on the strength of disorder that is unnecessarily strong and not satisfied by many experimental systems; for example, integer quantum Hall samples, at least for idealized non-interacting electrons, do not have a true energy gap but only a gap to spatially extended states.

More pragmatically, disorder comes in so many guises—be it vacancy sites in a spin liquid, magnetic ions coupling the two channels of a helical edge, or strong fields of ionized donors in a semiconductor heterostructure—that it would seem hard to say all that much that has a significant degree of universality to it. Nonetheless, there are a number of ways in which topology and disorder interact in a very characteristic way, and this chapter presents a selection of those in turn.

Indeed, quite generally in condensed matter physics, disorder can broadly be said to play three roles. It can be a nuisance, an asset, or a source of new physics; these roles are not mutually exclusive. The nuisance aspect is that disorder may place a veil over interesting phenomena by adding complicated artefacts to an experimental signal. Once such artefacts are systematically understood, they can become useful as experimental probes, as used e.g. in scanning tunnelling spectroscopy, where a tunnelling microscope supplies an electronic wavepacket to a sample and then probes its interference with secondary waves resulting from scattering off static impurities, thereby revealing properties of the underlying many-body system. Much of this chapter is devoted to the third aspect, with the very existence of quantum Hall plateaux, which would not be there in the absence of disorder, the most striking aspect. We use this as justification to discuss the physics of a disordered two-dimensional electron gas in more detail. This is a rewarding exercise as it allows us to visit a set of topics with interesting connections to field theories and geometric phase transitions. The role of defects is also central for topological quantum computing, and will be discussed in the chapter dedicated to this topic. At the same time, we here provide a few examples of the other two aspects, as these play such an important role for the practising condensed matter physicist.

In most of our discussions of topological phases so far, we have considered systems with no randomness, such as ideal crystals with a repeating unit cell. Real solids have imperfections of various kinds, and it is natural to regard our previous discussions as incomplete without some consideration of how topological properties can still exist with sufficiently strong disorder to modify significantly the underlying one-electron states. Here we assume for now that the disorder is quenched, i.e., constant in time, rather than thermally induced, and work at zero temperature.

It turns out that the connections between disorder and topological behavior are profound and include some remarkably general results, at least for the case of independent particles. These connections could well justify an extended treatment of their own, and we can only hope to scratch the surface here, but we can at least indicate the basic questions that arise and pictures of some key phenomena. The first challenge is to understand how topology appears when we look at an ensemble of Hamiltonians obtained from a statistical weighting of random potentials, which means that physical observables are averages over that ensemble.[1] We will then visit a number of other topics which are particularly enlightening, for instance how to glean information about the nature of a topological phase by analysing its response to the presence of disorder, be it deliberately introduced or not. This includes the physics of vortices in conventional superconductors as well as in quantum Hall states, where we will also encounter composite fermions. Another subject is the physics of defects in integer topological phases. Finally, we present a collection of topics concerning the interplay of disorder and spin liquids, including the generation of quantum dynamics from static disorder; the role of strain in the creation of a random synthetic gauge field; the capacity of a spin liquid to provide information about the level of disorder in the lattice is resides on; and finally, the response of a gapless spin liquid to dilution and distortion, which will make contact to parallel work on graphene, as well as to a different localisation problem which goes under the name of random bipartite hopping.

8.1 Introduction to disorder and localisation

The idea of a metal embodies the notion of a delocalisation: an electron can be added on one side of a sample, and extracted on the other side, by applying only an infinitesimal potential bias between the two. Without wanting to belittle the complexity of what transport in metals actually does involve in the presence of interactions and disorder, a simple cartoon picture for this is that a wavepacket can propagate through the sample as if it obeyed the Schrödinger equation of free space.

[1] Fortunately, experimentally relevant observables are often found to be self-averaging, i.e., their fluctuations become relatively small as the system size becomes large, or more precisely that the value of the observable in a single realization lies in a shrinking window around the mean value over all realizations with probability close to 1.

The alternative option is for the electron to be stuck in a finite portion of the sample. This possibility – localisation – is taken for granted nowadays, such a phenomenon having been observed in all sorts of settings involving e.g. sound, matter (cold atoms) and light waves. However, when it was first raised by Anderson in 1958, this proposal was so revolutionary that its importance was not appreciated until much later. As the Nobel laureate himself famously remarked in his lecture in 1977 (Anderson, 1977): "among those who failed to fully understand [the importance of localisation] at first was certainly its author."

The basic idea is that a wave scatters off impurities, i.e. objects which break the translational symmetry of space, so that plane waves cease to be solutions of the Schrödinger equation. The incident and scattered waves are then subject to interference. This degrades the ballistic motion of the wavepacket, and replaces it by diffusive or even localised behaviour. At its most extreme, the interference between incident and (multiply) scattered waves can be so destructive at long distances as to restrict the wave to a finite portion of space, the size of which is used to define what is called the localisation length.

Localisation is a subtle phenomenon, as attested by the fact that there have been literally thousands of publications devoted to this topic. The simplest symmetry class arises in considering the non-relativistic quantum mechanics of a particle in a random potential $V(\mathbf{r})$. We study the electronic wavefunctions that solve the nonrelativistic time-independent Schrodinger equation

$$H\psi = E\psi, \quad H = -\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}). \tag{8.1}$$

A simple guess is that at low energy E, there might be trapped states that decay at least exponentially at large distances (for example, if E is less than the typical values of V), while at high energy, there are scattering states in which on long times the electron executes a random walk, and the wavefunction extends to spatial infinity. In other words, at low energy, there should be some bound states trapped near minima of the potential, while at high energy, we expect that there should be some free states, where the electron scatters occasionally off bumps in the potential but is unbound. A mathematical distinction can be made between localised eigenstates whose magnitude falls off exponentially at spatial infinity, and extended eigenstates which fall off more slowly. Often the term critical is used for wavefunctions that fall off algebraically, and extended reserved for wavefunctions like plane waves that do not fall off at all.

In three dimensions this picture is correct: in fact there is a special energy, known as the mobility edge, that separates extended from localized states. An argument due to Mott for the existence of the mobility edge is that having extended and localized states coexist at the same energy is only possible for rare disorder potentials, as a small perturbation will mix the two types (as the energy denominator is zero) and give rise to extended states. In one and two dimensions, however, it

is now believed that all states are localized by a random potential, although the localisation length (the length scale on which the localized wavefunctions decay exponentially) becomes extremely long at high energies, especially in 2D.

How is localisation of high-energy particles by a weak potential even possible? Recall that a random walk in 1D or 2D returns to the same points over and over again, while in 3D and higher the mean number of returns to a given point is finite. As a result of this repetition, even a small bump in the random potential is amplified by being visited a large number of times: the result of constructive quantum interference is to lead to localisation by even weak disorder. In two dimensions, the number of returns of a random walk to a particular point diverges only logarithmically in the number of steps, so we might expect the localising tendency to be feeble in a weakly disordered 2D system, which is indeed the case.

More mathematically, diffusive spreading as in a random walk means that the mean squared distance after time t scales as $\langle R^2 \rangle \sim Dt$ for some diffusion constant D. We can think of the electron density at time t as concentrated in a sphere of radius proportional to $\sqrt{Dt}$. (More realistically, of course, the probability distribution of the density would be Gaussian.) Then, normalizing the overall density to 1, we have that the probability for the particle to be near the origin at time t goes as $(Dt)^{-d/2}$, where d is the spatial dimensionality, since this is the reciprocal of the sphere's volume. Now we can ask, how many times is the electron expected to have returned to the origin by time T? The expected number of returns is thus, up to a factor of order unity,

$$N = \int_{t_0}^{T} \frac{1}{(Dt)^{d/2}} dt. \tag{8.2}$$

Here we ignore any possible singularity at the origin (since we know that at short times our assumption of a Gaussian spread breaks down) and focus on the long-time behavior. The integral converges for $d > 2$. So for $d > 2$, the electron returns only a finite number of times to any particular fluctuation; if the fluctuations are weak enough, then the electron will not be localized since different fluctuations are independent for a random potential with only short-ranged correlations.

Why should the interference be constructive when the electron returns to the same point? We give two arguments. One is a heuristic picture of the relevant physics appearing in the diagrammatic approach to computing electronic properties (weak localisation). This picture gives a clue that the behavior in a magnetic field may be quite different than in zero field, which indeed is the case as the integer quantum Hall effect must somehow appear. The other approach, in Box 8.1, is a nice example of renormalization group ideas that some readers may have previously seen in statistical physics. This one-parameter scaling theory is a famous example of how a simple RG argument can be used to get a qualitative picture for a complicated system where an exact treatment is extremely difficult.

Box 8.1: One-parameter scaling approach to Anderson localisation

We can give a scaling argument that supports the result stated above, that two dimensions is the marginal dimension for localisation, and is much simpler than a serious calculation (Abrahams et al., 1979). (The four authors of this paper are widely referred to as the Gang of Four, a reference to Chinese politics of that era coined by Patrick Lee.) Be forewarned that magnetic fields or strong spin-orbit coupling give alternative behavior to what we find here; in the case of magnetic fields, for example, one needs to consider scaling as a function of two parameters (Khmelnitskii, 1984), which can be taken to be the diagonal and Hall conductivities $(\sigma_{xx}, \sigma_{xy})$, as we will touch on below; see also Fig. 8.4.

The formal version of this argument is made using the renormalization group, but the basic idea is quite simple. Let us try to understand the behavior of the function $g(L)$, which gives the conductance (not conductivity) of some material in a (hyper)cube of side L. Our goal will be, given some initial value $g(L_0)$ at a short length scale L_0, to understand what happens when we go to larger scales.

The choice of $g(L)$ as a coupling constant that can be subjected to an RG treatment may at first sight be surprising. Indeed, as a dimensionful quantity to do with electronic transport, it does not look like much of a coupling constant to start with. Its use—as a dimensionless parameter obtained by expressing it in units of the conductance quantum e^2/h which we have encountered extensively in the context of the quantum Hall effects—can nonetheless be motivated as follows (see e.g. John Chalker's article in (Comtet et al., 1999)).

Imagine what happens to the states of the small hypercube of linear size L_0 when 2^d of them are combined into a larger one of size $2L_0$. From the point of view of one level in the small hypercube, there are two interesting energy scales. One is the energy window ϵ over which it hybridises with states from adjacent hypercubes; and the other is the distance to other states in its own hypercube, the mean level spacing Δ. Their ratio ϵ/Δ therefore encodes how easily states from adjacent hypercubes can mix: if $\epsilon/\Delta \ll 1$, there is likely no state in a neighbouring hypercube available with which to hybridise, so that each state will remain localised in its own hypercube.

The connection to $g(L)$ then proceeds as follows. Let us return to the idea of particle diffusion, Eq. 8.2, where a displacement scales as $(Dt)^{d/2}$. This allows the conversion of a length scale to a timescale, and in turn an energy scale. For a diffusive conductor of size L, this is known as the Thouless energy $E_T = \hbar D/L_0^2$. We identify ϵ with this Thouless energy. At the same time, the mean level spacing is given by $\Delta = 1/(L_0^d n)$ for a hypercube with density of states n, while an Einstein relation connects the conductivity with the diffusion

Box 8.1: One-parameter scaling approach to Anderson localisation

constant as $\sigma = e^2 nD$. Putting these together makes g pop out as desired: $g(L_0) \sim \epsilon/\Delta = (h/e^2)\sigma L_0^{d-2}$.

The RG flow is conventionally parametrised in terms of the β-function defined as

$$\beta = \frac{d \log g}{d \log L}. \tag{8.3}$$

Suppose first of all that a scattering picture is correct: noninteracting electrons in the material move diffusively (rather than being localized or ballistic), and Ohm's law is satisfied. Then the conductance, once the cube is larger than the mean free path l, should go as

$$g(L) \approx \sigma L^{d-2} \tag{8.4}$$

where σ is the conductivity. Already we can see that $d = 2$ is marginal – a power 0 usually translates into a logarithmic flow. As an aside, we note that it seems possible only in $d = 2$ to have a scale-invariant conductance. From the quantum Hall effects, we are certainly familiar with the quantisation of the Hall conductance in units of e^2/h.[a]

Suppose now that instead of having diffusive electron motion, all electrons are in localized states. How then should g behave? Well, if the longest localisation length is $\xi \gg l$ (the localisation length is always longer than the mean free path), then we expect for $\xi \ll L$

$$g(L) \sim \exp(-L/\xi). \tag{8.5}$$

Here we are ignoring possible power-law factors which will be dominated for large L by the exponential. For a realization of disorder characterised by a particular strength, we expect the microscopic conductance to flow from its initial value $g(L_0)$ with increasing L until reaching one of the above two asymptotic regimes. The challenge is now to justify this picture and understand the importance of dimensionality.

The main conjecture by Abrahams et al. (1979) was that $\beta = \frac{d \log g}{d \log L}$ can be taken to be a function of g *only*. We might think that other properties such as L, the details of disorder, etc. would be important, but at least in the long-length-scale limit, it seems that β is indeed a function of g alone: this is known as one-parameter scaling. We write $\beta(g)$ henceforth to emphasize this. The idea behind one-parameter scaling is that the properties of the system on scale $2L$ are determined by the effective level of disorder at scale L, and that the dimensionless conductance g is a sufficient measure of this disorder: two types of microscopic disorder that give rise to the same conductance g

Box 8.1: One-parameter scaling approach to Anderson localisation

at a large scale L then are predicted to give the same conductance at all larger length scales. This is still an assumption that needs to be tested, but one-dimensional calculations support this picture. It is also possible to justify some of the above using perturbation theory in the disorder strength, which gives the weak localisation theory described in the main text.

Returning to our above guesses for the asymptotic form of β, we now have deep in the diffusive regime (high conductance, $g \gg 1$, $\log g > 0$)

$$\beta(g) = \frac{d \log g}{d \log L} = (d-2). \tag{8.6}$$

Deep in the localized regime (low conductance, $g \ll 1$, $\log g < 0$), we have

$$\beta(g) = \frac{d \log g}{d \log L} = (-L/\xi) \approx \log g, \tag{8.7}$$

which is negative in any dimension.

The point of one-parameter scaling is that now we can make a plot of $\beta(g)$ vs. $\log g$, as in Fig. 8.1, and argue based on continuity that $d = 1$ and $d = 2$ are very different from $d = 3$. If there were more dimensions to the plot, as occurs in a strong magnetic field, then the situation would be more complicated. In $d = 1$ and $d = 2$, the simplest continuity assumption is that $\beta(g)$ is always negative for finite g, since it is negative at $g = 0^+$ and zero or negative at $g = \infty$. (Regarding the latter, a controlled perturbation theory for $g \gg 1$ in $d = 2$ does indeed yield a negative rather than zero value.) In $d = 3$, we have a more complicated situation because $\beta(g)$ must have a zero. Above this critical point, the flow is to the diffusive regime; below this critical point, the flow is to a localized regime. This picture corresponds roughly to our intuitive idea of a mobility edge separating extended and localized states in three dimensions. So the simplest guess for how g evolves with L by interpolating between these two asymptotic regimes leads to the conclusion that in 3D there are two regimes separated by an unstable fixed point, while in 1D and 2D, the only stable fixed point is at $g = 0$.

Box 8.1: One-parameter scaling approach to Anderson localisation

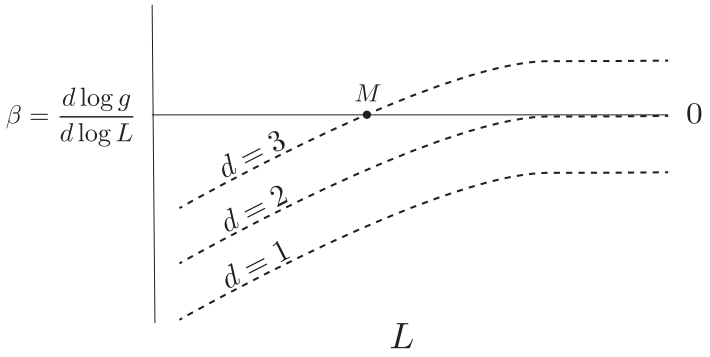

Figure 8.1 One-parameter scaling flow in various dimensions and mobility edge in $d = 3$. At low conductivity, when the system is localized, the β function describing the (logarithmic) change of conductance is negative in all dimensions. At high conductance, Ohm's law in Eq. 8.6 determines $\beta > 0$ in $d = 3$, so the conductance increases under rescaling, while in $d = 1, 2$ the conductance does not increase since $\beta \leq 0$. This leads to a critical point or mobility edge in $d = 3$ at the point where $\beta = 0$, labeled M.

[a] In $d = 1$, however, there is also a sort of conductance quantisation in units of e^2/h. The resolution to this paradox is that the 1D finite conductance results purely from the contacts; transport is quasi-ballistic in the bulk of the system.

To understand where localisation comes from and to present one famous prediction of weak-localisation theory, consider the probability that an electron initially at spatial location A in the plane propagates to another point B. We give a heuristic argument that is borne out by summing over a series of diagrams in perturbation theory in the disorder potential (Lee and Ramakrishnan, 1985). In the path-integral picture of quantum mechanics, we would expect the amplitude to be a given by a weighted sum over possible real-space paths from A to B.

The probability is the squared magnitude of the amplitude and will include cross terms between these paths. There are multiple paths to get from A to B, and some of those will have self-intersections. The relevant cross terms are those involving two paths, say 1 and 2, with the same starting and ending point, but which trace an intermediate loop in opposite directions, e.g., clockwise for path 1 but counterclockwise for path 2. In the absence of a magnetic field (i.e., because of time-reversal symmetry), these paths interfere constructively so that their cross terms are positive (Fig. 8.2):

$$\langle |\Psi|^2 \rangle = \langle |\psi_1 + \psi_2|^2 \rangle = \langle |2\psi_1|^2 \rangle = 4\langle |\psi_1|^2 \rangle > \langle |\psi_1|^2 \rangle + \langle |\psi_2|^2 \rangle. \tag{8.8}$$

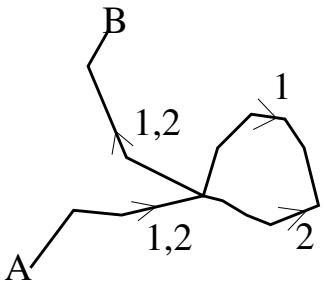

Figure 8.2 Constructive interference of time-reversed paths in weak localisation is destroyed by application of a magnetic field.

Now we know by unitarity that the total probability of the particle to get from A to any final location after some period of time should be unity. The above argument says that cross terms will act to increase the probability of paths with self-intersections, and the number of self-intersections is larger for paths where A is not too far from B. This suggests that the first quantum correction to random-walk behavior is in the localising direction.

A test of this scenario is to consider what happens when a uniform orbital magnetic field normal to the plane is added to the system. The magnetic field modifies the cross terms by generating Aharonov-Bohm phases for the two directions around the internal loop. Hence constructive interference is no longer guaranteed, with loops of different area causing various phases of the cross term. The prediction is thus that a metal showing signs of localisation should behave differently in a magnetic field than classical expectations would suggest: normally a metal's diagonal resistivity increases because the magnetic field bends the electron trajectories, but in a strongly disordered metal, applying a magnetic field famously *decreases* the resistivity by reducing the strength of the localising corrections. This negative magnetoresistance is one of the major predictions of weak localisation theory.

8.2 A semiclassical model of quantum Hall transitions

One of the central features of the quantum Hall effect is its absence in disorder-free systems: for a translationally invariant system, $\rho_{xy} = \frac{h}{\nu e^2}$ for any value of the filling factor ν not just for integer or simple rational numbers. This can be seen by a simple thought experiment involving a Lorentz transformation of a system at rest in the laboratory frame, now viewed in a frame moving with velocity $\mathbf{v} \perp \mathbf{B}$. Then, to leading order in $\beta = v/c$, the density n along with the magnetic field $B_{\hat{z}}$ remain unchanged, while an electric field $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$ appears, along with a current density $\mathbf{J} = -ne\mathbf{v}$.