1 Introduction

Solid state physics grew out of applications of quantum mechanics to the problem of electron conduction in solids. This seemingly simple problem defied solution because the presence of an ion at each lattice site seemed to present an obvious impediment to conduction. How the electrons avoid the ions was thus the basic question. Although the answer to this question is well known, it does serve to illuminate the very essence of solid state physics: there is organization in the many. Each electron adjusts its wavelength to take advantage of the periodicity of the lattice. In the absence of impurities, conduction is perfect. Hence, by understanding this simple fact that periodicity implies perfect conduction, it became clear that the experimentally observed resistivity in a metal came not from electrons running into each of the ions but rather from dirt (disorder), thermal effects mediated by dynamical motion of the ions, or electron–electron interactions. This book examines each of these effects with an eye for identifying underlying organizing principles that simplify the physics of such interactions.

1.1 Spontaneously broken symmetry

The search for organizing principles that help simplify the physics of many-body systems is at the heart of modern solid state or, more generally, condensed matter physics. One such tool is symmetry. Consider the simple case of permutation symmetry typically taught in a first class in quantum mechanics. This symmetry was introduced into quantum mechanics by W. Heisenberg in the context of the indistinguishability of identical particles. The permutation group has a finite number of elements and hence is associated with a discrete symmetry. Permutation symmetry allows us to classify fundamental particles into two groups. Bosons are even with respect to interchange of two particles and fermions odd. This symmetry can be generalized to include a non-integer phase when two particles are interchanged, as we will see in the context of the fractional quantum Hall effect.

To a large extent, the symmetries that are most relevant in condensed matter systems are typically continuous, for example rotational symmetry. Spontaneously breaking a continuous symmetry has a fundamental consequence. For example, the existence of phonons in a solid or spin waves in a magnet follows from the spontaneous breaking of a continuous symmetry. By spontaneous, we mean without the application of an external field. A periodic arrangement of ions in a crystal breaks continuous translational and rotational symmetry. Such spontaneous breaking of a continuous symmetry by the very existence of the lattice is necessarily accompanied by a massless spinless bosonic excitation. That such massless

spinless bosons, known as Nambu–Goldstone bosons (G1961; N1960), necessarily accompany the breaking of a continuous symmetry can easily be deduced from the following considerations. We consider a system with a Lagrangian

$$\mathcal{L} = T - V(\phi), \tag{1.1}$$

consisting of a kinetic energy, T, and a potential energy, $V(\phi)$, where we are allowing for ϕ to be a complex function. The claim that such a system is invariant under a symmetry operation is captured by

$$V(\phi) = V(\phi + \epsilon\delta\phi), \tag{1.2}$$

where $\epsilon\delta\phi$ is the generator of the symmetry operation. Here ϵ is an infinitesimal. We have assumed for the moment that $\delta\phi$ is independent of space. To illustrate what is meant by this identity, consider a potential of the form $V(\phi) = \epsilon_0|\phi|^2$. This potential is invariant under transformations of the form $\phi \to \phi e^{i\theta}$. Let θ be a small quantity completely independent of space. Then we can expand the exponential and retain only the first-order term. Consequently, $\phi \to \phi(1 + i\theta)$ and we identify $\epsilon\delta\phi$ as $i\theta\phi$; that is, $\epsilon = \theta$ and $\delta\phi = i\phi$. This symmetry, known as U(1), is present in models that preserve charge conservation. Expansion of $V(\phi)$ to linear order in ϵ implies that

$$\delta\phi \frac{\delta V}{\delta\phi} = 0, \tag{1.3}$$

assuming that the symmetry is intact. Now assume explicitly that the symmetry is broken such that $V \to V(\phi_0 + \chi)$, where ϕ_0 minimizes the potential and χ cannot be written as a generator of a symmetry operation as in Eq. (1.2). Since the potential has a minimum, it makes sense to expand

$$V(\phi_0 + \chi) = V(\phi_0) + \frac{1}{2}\chi^2 \left.\frac{\partial^2 V}{\partial\phi^2}\right|_{\phi=\phi_0} = V(\phi_0) + \frac{1}{2}\chi^2 m^2, \tag{1.4}$$

truncating at the restoring term at second order. The second term, which can be used to define the mass (m) in a standard harmonic expansion, is inherently positive semi-definite since we have expanded about the minimum. With this equation in hand, we differentiate Eq. (1.3),

$$\frac{\partial\delta\phi}{\partial\phi}\frac{\delta V}{\delta\phi} + \delta\phi\frac{\partial^2 V}{\partial\phi^2} = 0, \tag{1.5}$$

with respect to ϕ. The first term vanishes when evaluated at the minimum, implying that

$$\delta\phi \left.\frac{\partial^2 V}{\partial\phi^2}\right|_{\phi=\phi_0} = 0 \tag{1.6}$$

must identically vanish for any variation of ϕ in the broken symmetry state. Since $\delta\phi$ is non-zero, Eq. (1.6) is satisfied only if the second-order-derivative term vanishes or equivalently if $m^2 = 0$. That is, the mass vanishes. This is Goldstone's theorem (G1961). A zero mode exists for each generator of a continuously broken symmetry. As a result of this theorem, symmetry occupies a central place in all areas of physics, in particular particle and condensed matter physics. Typically, the massless bosons that arise in condensed matter systems represent collective excitations of the entire many-body system. In fluids, phonons are purely longitudinal and arise from spontaneous breaking of Galilean invariance. In solids, phonons are both transverse and longitudinal, though with no simple correspondence with the spontaneous breaking of Galilean, translational, and rotational symmetry. In magnets, spin waves or magnons are the collective gapless excitations that emerge from the spontaneous breaking of rotational symmetry.

We can of course relax the constraint that θ be independent of space. In so doing, we can entertain what happens under local rather than global (θ independent of space) transformations. While our analysis on the potential energy remains the same, the kinetic energy,

$$T \to \frac{1}{2}(\partial_\mu \phi^*)(\partial^\mu \phi) + \frac{1}{2}|\phi|^2(\partial_\mu \theta(x))^2, \tag{1.7}$$

does acquire a new term describing the spatial variation of the phase. If the U(1) symmetry is not broken by this transformation, then the second term must vanish. Demanding that

$$\partial_\mu \theta = 0 \tag{1.8}$$

requires that θ be spatially homogeneous for the symmetry to be preserved. As a result, a consequence of breaking the continuous U(1) symmetry is that θ must be spatially non-uniform. This is the situation in a superconductor. In fact, the current inside a superconductor arises entirely from the spatial variation of the phase, as can be seen from the quantum mechanical equation for the current,

$$j_\mu = \frac{e\hbar}{m}\mathrm{Im}\psi^\dagger \partial_\mu \psi = \frac{e^*\hbar}{m}|\Delta|^2 \partial_\mu \theta, \tag{1.9}$$

if we interpret ψ as the wavefunction for the superconducting state; that is, $\psi = \Delta e^{i\theta}$. We will see in the chapter on superconductivity precisely how this state of affairs arises. We will interpret ψ as the order parameter of a superconducting state. While the Bardeen–Cooper–Schrieffer (BCS) theory of superconductivity was certainly not formulated as an example of a broken continuous symmetry, this is the basic principle that underlies this theory. In fact, the key ingredients of superconductivity, charge $2e$ carriers and a supercurrent, all follow from breaking U(1) symmetry.

Massless bosons that emerge from broken symmetry typically generate new unexpected physics. For example, phonons mediate pairing between electrons, thereby driving the onset of superconductivity in metals such as Hg and more complicated systems, for example MgB_2. However, strict rules determine how such Nambu–Goldstone bosons can affect

any system. As shown by Adler (A1965), the interactions induced by massless bosons arising from the breaking of a continuous symmetry must be proportional to the transferred momentum. More formally, interactions mediated by the exchange of a Nambu–Goldstone boson can only obtain through derivative couplings. Consequently, the interaction vanishes for zero exchanged momentum. This principle implies that the electron–phonon interaction which mediates pairing in elemental superconductors is inherently dynamical in nature. We will verify this important principle in the context of the electron–phonon coupling through an explicit derivation. Hence, entirely from the existence of a lattice, phonons and the kinds of interactions they mediate can be easily deduced.

1.2 Tracking broken symmetry: order parameter

The idea of an order parameter is another powerful concept in condensed matter physics. Order parameters track broken symmetry. That is, they are non-zero in the broken symmetry phase and zero otherwise. Consider a ferromagnet. Locally each spin can point along any direction. This is the case at high temperature in which no symmetry is broken. In a phase transition controlled by thermal fluctuations, typically it is the high-temperature phase that has the higher symmetry. To quantify the order in a collection of spins, we sum the z-component of each of the spin operators,

$$M = \frac{1}{N}\sum_i \langle S_i^z \rangle, \tag{1.10}$$

scaled by the number of spins, N. Here S_i^z is the z-component of the spin of the atom on site i and the angle brackets indicate a thermal average over the states of the system. M is the magnetization. At high temperature before any symmetry is broken, the magnetization is identically zero. At sufficiently low temperatures, the spins order and the magnetization acquires a non-zero value. Consider iron for which the Curie or ordering temperature is 1340 K. It turns out that most parts of a block of iron below the magnetization temperature have vanishing magnetization. This state of affairs obtains because the magnetization is in general a function of space. As a result, a block of iron does not break the symmetry uniformly. In fact, the actual magnetization in bulk magnets is not acquired spontaneously but rather by some external means to align all of the individual magnetic domains. At the boundary of a domain, the magnetization changes sign, creating a domain wall. Typical domain sizes in iron are roughly 300 ions. Placing a chunk of Fe in a magnetic field will orient all of the domains in the same direction, a state of affairs that will persist long after the field is turned off. This is important since the re-oriented domain state does not constitute a minimum energy state of the system. The domains lock into place by becoming pinned to defects. One would expect then that as the magnetizing field is varied, the magnetization would not change continuously but by discontinuous jumps as domain walls de-pin from defects. This is the essence of the Barkhaussen effect, the tiny discontinuous jumps the magnetization makes in the presence of an external magnetic field and ultimately the reason why the magnetization curve in a ferromagnet exhibits hysteresis as depicted in Fig. 1.1.

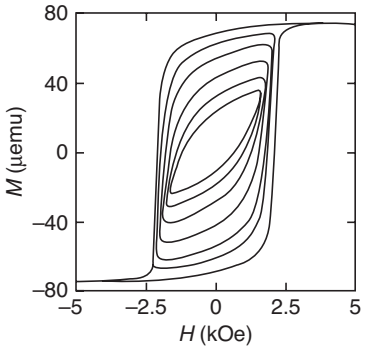

Fig. 1.1 Hysteresis curve of the magnetization as a function of the applied field for a CoPtCrB thin film. Multilayer Co/Pt is used in memory storage. Reprinted from Carpenter *et al.*, *Phys. Rev. B* **72**, 052410 (2005).

That is, ramping the field on and off is path dependent determined by which domains de-pin sequentially.

What is crucial in the magnetic system is that locally there are two degrees of freedom for each of the spins. At high temperature, both states are accessible. At sufficiently low temperature, one of the spin states is selected. Such state selectivity can be modeled with a double-well potential of the form

$$V(M) = -\frac{1}{2}\alpha M^2 + \frac{1}{4}\gamma M^4, \tag{1.11}$$

where M is the magnetization and α and γ are positive. The minima of this potential occur at

$$M_{\pm} = \pm\sqrt{\frac{\alpha}{\gamma}}. \tag{1.12}$$

Both of the minima are accessible at high temperature and no magnetization is possible. Our choice of $\alpha > 0$ ensures that deviations of M away from $M_{\pm}$ cost energy. As a result, the minimum energy of V is not zero but rather the non-zero value of $V_0 = -\alpha^2/4\gamma$. As it stands, our theory is completely symmetrical with respect to the change $M \to -M$. Surely the physics cannot change if we were to recast our theory by shifting the scalar field M by a constant such that $M \to M_+ + \phi(x)$. The new potential

$$V' = V_0 + \left(-\frac{1}{2}\alpha + \frac{3}{2}\gamma M_+^2\right)\phi^2 + \gamma M_+\phi^3 + \frac{1}{4}\gamma\phi^4 \tag{1.13}$$

no longer looks symmetrical in terms of the new field ϕ. Why? What we have done by expanding around one of the minima of V is to hide the symmetry. Essentially we have broken the symmetry by setting the magnetization to M_+. In the broken symmetry phase, up and down spins are no longer equivalent. The field M functions as the measure of the magnetic order. M is the order parameter. Unlike the old potential which was minimized by

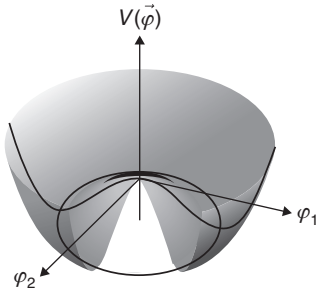

Fig. 1.2 Potential corresponding to the complex scalar field φ. The minima correspond to circles satisfying $\varphi_1^2 + \varphi_2^2 = \alpha/\gamma$.

a non-zero value of M, the minima of V' take place at $\phi(x) = 0$. This corresponds to a true vacuum. In classical models for ferromagnets, the magnetization turns on continuously,

$$M \propto |T - T_c|^{\alpha}, \tag{1.14}$$

at a non-zero temperature, T_c. It acquires its maximum value at $T = 0$. The exponent α is the critical exponent for the turn-on of the order parameter.

A slightly more complicated example of broken symmetry occurs when we modify our potential,

$$V(\varphi) = -\alpha\varphi^*\varphi + \gamma(\varphi^*\varphi)^2, \tag{1.15}$$

to allow for a complex scalar field $\varphi = \frac{1}{\sqrt{2}}(\varphi_1 + i\varphi_2)$. For $\alpha > 0$ and $\gamma > 0$, the potential is illustrated in Fig. 1.2. The corresponding Lagrangian takes the form

$$\mathcal{L} = \frac{1}{2}(\partial_\mu\varphi^*)(\partial^\mu\varphi) - \left(\alpha\varphi^*\varphi + \gamma(\varphi^*\varphi)^2\right). \tag{1.16}$$

Our Lagrangian has the global symmetry $\varphi \to \varphi e^{i\theta}$, where θ is a constant. As a result of this symmetry, the minima of V now take on a circle of values satisfying $\varphi_1^2 + \varphi_2^2 = \alpha/\gamma$, as illustrated in Fig. 1.2. That is, there are infinitely many saddle points as a result of the continuous global symmetry. As before with the single scalar field for the magnetization, we can expand about the circular minima by defining $\varphi = (\sqrt{\frac{\alpha}{\gamma}} + f(x) + ig(x))$. That is, we break the symmetry by hand. Because at the minima φ_1 and φ_2 are not independent, this transformation is not a simple translation of φ_1 and φ_2 separately. As a result, the quadratic term essentially has only one degree of freedom. We can interpret this as a vanishing of the mass for one of the scalar fields, in line with Goldstone's theorem that a massless mode must emerge upon the breaking of a continuous symmetry. Such a solution in which a complex field acquires a non-zero value is the heart of the superconducting transition. What the Bardeen–Cooper–Schrieffer solution laid plain is that the phenomenological Landau–Ginzburg treatment in terms of a complex order parameter acquiring a non-zero value in the superconducting state has a microscopic basis in the electron–phonon interaction. Such an interaction mediates pairing and the order parameter for the superconducting state is a product of an amplitude for pair formation times $e^{i\theta(r)}$, where θ is the phase of the pair field.

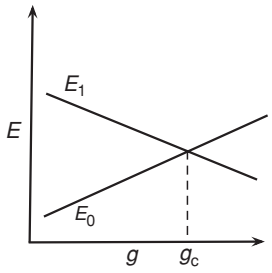

Fig. 1.3 Energy-level diagram for the ground state, E_0, and first excited state, E_1, as a function of a coupling constant, g. The crossing at g_c signals a phase transition between the ground and the first excited state.

Our statement that the low-temperature phase typically has lower symmetry is only true classically. There are many examples of symmetry breaking at $T = 0$ that have nothing to do with thermal fluctuations. Such phase transitions are governed by fluctuations of the vacuum, that is the uncertainty principle. In general, such phase transitions are governed by a transition among the quantum mechanical states of a many-body system simply by changing some system parameter. Consider a Hamiltonian $H(g)$, where g is a coupling constant. A typical energy-level diagram for this system as a function of g is depicted in Fig. 1.3. If, as a function of g, the first-excited and the ground states cross, a phase transition obtains to the first-excited state. For the transition to be continuous, we must have that $\partial E/\partial g = 0$. These types of situation are discussed explicitly in Chapter 14.

1.3 Beyond broken symmetry

Despite the utility of symmetry in classifying collective phenomena, physics is replete with examples of transitions between states of matter that share the same symmetry but are, nonetheless, distinct. An obvious example is the liquid–gas transition or the formation of the fractional quantum Hall state. However, the particular examples we focus on here, which typify the physics of strong coupling, are those in which the formation of some kind of bound state is the distinguishing feature. Consider, for example, the vulcanization or cross-linking transition in rubber. In the un-vulcanized state, rubber is a viscous liquid in which long-chain monomers move independently. Cross-linking between neighboring monomers, resulting in the formation of a highly entangled enmeshed amorphous state, defines the vulcanization transition. Although the monomers are localized in the vulcanized state, they are randomly located. Consequently, there are no Bragg peaks. Nonetheless, one can define an appropriate order parameter (GCZ1996) which reflects the fact that at $t = 0$ and $t = \infty$, the configuration of the monomer strands in the liquid changes while it is essentially static in the amorphous state. The resulting resilience and emergent static modulus of rubber both arise from the effective gluing together of the monomers. In high-energy physics, mesons or bound states of quarks are the propagating degrees of freedom at low energy in nuclei. They arise without the breaking of any continuous symmetry. In problems more

relevant for this text, a magnetic impurity in a non-magnetic host forms a bound state with all the conduction electrons below a characteristic temperature, once again without breaking any symmetry or even inducing a phase transition. The formation of such bound states is the essence of the Kondo problem which stands as one of the key triumphs of the renormalization group principle. As we will show from a systematic integration of the high-energy degrees of freedom, a bound state emerges because the coupling constant between the impurity and the host electrons diverges. Hence, at low temperatures it is not correct to think of the magnetic and conduction electron degrees of freedom separately. A new entity emerges at low energies that is not present in the starting ultraviolet (UV)-complete Lagrangian, a characteristic feature of strong-coupling systems. Although new degrees of freedom are present at strong coupling which are absent in the weakly coupled or high-temperature regime, the two states are still adiabatically connected in that by varying the system parameters, one can go smoothly from one phase to the other. Nonetheless, the phases are quite distinct. They possess different degrees of freedom, and no unified description exists of such systems in terms of a single entity. Bound-state formation is a standard paradigm in strong coupling physics and, as we will see, the Mott problem, an insulating state in a partially filled band, is no exception.

Another key example is Fermi liquid theory. The primary tenet of this theory is that the excited states of a metal stand in a one-to-one correspondence with those of a non-interacting electron gas. The interactions in a metal are of course non-zero. However, they are strongly screened and can be treated as essentially short-ranged. The Landau (L1957) assertion is that all such short-ranged repulsive interactions do not destroy the sharpness of the electron excitations in the non-interacting electron gas. That this state of affairs obtains is perhaps one of the most remarkable principles in nature. Why can the short-range interactions be ignored in a metal? The answer lies in a fundamental renormalization principle which we present in Chapter 12. The key to solving any many-body problem is to identify the propagating degrees of freedom. Identifying that the propagating degrees of freedom are single electrons with a dispersion relation given by $p^2/2m$ in an interacting electron gas is highly non-trivial. In fact, it cannot be deduced directly from the Hamiltonian. Some further fact is needed. That further fact is the existence of a Fermi surface. As we will see, the fundamental principle that makes Fermi liquid theory work is that a Fermi surface is remarkably resilient to short-range repulsive interactions. We will demonstrate explicitly that all renormalizations (P1992; SM1991; BG1990) arising from such interactions are towards the Fermi surface. As a result, such interactions can effectively be integrated out, leaving behind dressed electrons or quasi-particles, thereby justifying the key Landau tenet (L1957) that the low-energy electronic excitation spectrum of a metal is identical to that of a non-interacting Fermi gas. Consequently, breaking Fermi liquid theory in dimensions greater than or equal to two is notoriously difficult. In one spatial dimension, interactions are always relevant, as will be seen, and a new state of matter arises, termed a Luttinger liquid, in which spin and charge move but with different velocities. In higher dimensions, the problem is open and stands as the key outstanding problem in solid state physics.

As Fermi liquid theory made the BCS theory of the superconducting state possible in that it cleanly identified the propagating degrees of freedom, a similar identification of the propagating degrees of freedom in the normal state of the copper-oxide high-temperature

superconductors is necessary to know what pairs up to form the superconducting condensate. This problem is particularly difficult as the parent materials are all antiferromagnetic Mott insulators. Some of the agreed-upon physics of this remarkable problem and a forward-leaning perspective are discussed in the final chapter of this book.

References

[A1965] S. L. Adler, *Phys. Rev.* B **137**, 1022 (1965); *Phys. Rev.* **139**, B 1638 (1965).
[BG1990] G. Benefatto and G. Gallavotti, *J. Stat. Phys.* **59**, 541 (1990).
[GCZ1996] P. M. Goldbart, H. E. Castillo, A. Zippelius, *Adv. Phys.* **45**, 393 (1996).
[G1961] J. Goldstone, *Nuovo Cimento* **19**, 154 (1961).
[N1960] Y. Nambu, *Phys. Rev.* **117**, 648 (1960).
[L1957] L. D. Landau, *Sov. Phys. JETP* **3**, 920 (1957); **5**, 101 (1957); **8** 70 (1959).
[P1992] J. Polchinski, arXiv:hep-th/9210046.
[SM1991] R. Shankar, *Physica* **A177**, 530 (1991).

2 Non-interacting electron gas

At the close of the previous chapter, we noted that one can understand the elementary properties of metals in terms of non-interacting electrons and phonons. For example, the low-temperature specific heat of a metal is the sum of a term linear in the temperature, T, from the electrons and a term proportional to T^3 from the phonons. This result follows from a non-interacting particle picture. The electrical conductivity limited by non-magnetic impurity scattering is also well described by a non-interacting electron gas. In addition, from a knowledge of single-electron band theory, one can discern qualitatively the differences between metals, insulators, and semiconductors. The remarkable success of the non-interacting model is paradoxical because electrons and ions are strongly interacting both with themselves and with one another. Along with its successes, the non-interacting picture has colossal shortcomings, most notably its inability to describe old problems such as cohesive energies, superconductivity, magnetism, and newer phenomena such as doped Mott insulators, the Kondo problem, and the fractional quantum Hall effects. We first review the physics of the non-interacting electron gas. It is only after we develop methodology for dealing with electron interactions than we can lay plain the reasons why the non-interacting model works so well.

Electrons in metals are quantum mechanical particles with spin $\hbar/2$ obeying Fermi–Dirac statistics. The Hamiltonian of a single electron is $\widehat{p}^2/2m$ where $\widehat{\mathbf{p}}$ is the electron momentum (operator) and m the electron mass. Its eigenstates are plane waves of the form $e^{i\mathbf{p}\cdot r/\hbar}/\sqrt{V}$ times a spinor which specifies the electron spin projection on a convenient axis (usually $\hat{z}$), $\hbar\sigma/2$ where $\sigma = \pm 1$; here V is the system volume. The Hamiltonian (operator) for N such non-interacting electrons,

$$\widehat{H} = \sum_{i=1}^{N} \frac{\widehat{p}_i^2}{2m}, \tag{2.1}$$

is simply the sum of the kinetic energies of the individual particles. In this case, the eigenstates are products of the occupied single-particle plane-wave states. Each plane-wave state can be occupied at most by one electron of a given spin. We label these eigenstates by the distribution function $f_{\mathbf{p}\sigma}$, which is 1 if the single-particle momentum-spin state is occupied and 0 otherwise. In the ground state, the lowest $N/2$ single-particle states are doubly occupied with electrons of opposite spin. Consequently, in the ground state (temperature $T = 0$), the distribution function is

$$f_{\mathbf{p}\sigma} = \Theta(\mu_0 - p^2/2m), \tag{2.2}$$

where $\Theta(x)$ is the Heaviside function, $\Theta(x > 0) = 1$, and 0 otherwise. Here μ_0 is the zero-temperature electron chemical potential, which in this case is simply the Fermi energy,