

1

Overview of Condensed Matter Physics

Matter that we encounter in our everyday life comes in various forms: air is gaseous, water (between 0°C and 100°C under ambient pressure) is a liquid, while ice and various kinds of metals and minerals are crystalline solids. We also encounter other familiar forms of matter from our daily experience, including glasses and **liquid crystals**, which do not fall into the categories of gas, liquid, or solid/crystal. More exotic forms of matter exist under extreme conditions, like very low (all the way to almost absolute zero) or very high temperatures, extremely high pressures, very far from equilibrium, etc. Roughly speaking, “condensed matter physics” studies physical properties of matter in “condensed” forms (where the density is high enough that interaction among the constituent particles is crucial to these properties). In the rest of this chapter we will attempt to give a more precise (but far from unique) definition of condensed matter physics, and discuss how to classify, theoretically describe, and experimentally probe various forms or phases of condensed matter. In this book we will deal exclusively with condensed matter systems that are made of atoms or molecules and, in particular, the electrons that come with them (though on occasion we will study collective excitations in which photons play an important role). On the other hand, the methodology developed here and many specific results apply to more exotic condensed matter systems, like neutron stars and quark–gluon plasmas that are best described in terms of quarks, gluons, or nucleons.

1.1 Definition of Condensed Matter and Goals of Condensed Matter Physics

Matter which surrounds us is made of huge numbers (of order 10^{23}) of atoms or molecules, which have a characteristic size of 10^{-10} m or 1 Å. In the gaseous form, the typical interparticle distance is much larger than this characteristic size, thus the particles interact weakly (except when occasional interparticle collisions occur) and retain their integrity; in particular, electrons are attached to individual atoms or molecules. As a result, the physical properties of such gaseous matter are usually dictated by the properties of *individual* atoms or molecules, and we do *not* refer to the gaseous matter as condensed matter in most cases.¹

¹ There are exceptions to this. For example, at very low temperatures, collective behavior that involves large numbers of atoms or molecules (such as Bose–Einstein condensation) may occur. Interactions may be weak in some sense, but not on the scale of the extremely low temperature. In such cases, the physical properties are dominated by the collective behavior, and we call such gaseous matter condensed matter as well, in the spirit of following discussions.

For matter in liquid/fluid or solid forms, on the other hand, the constituent atoms are in sufficiently close proximity that the distance between them is comparable to the size of individual atoms. As a result, these atoms interact with each other strongly. In many cases, some of the electrons (mostly the outer-shell ones) which were attached to individual atoms may be able to move throughout the system. (In a metal the electrons can move more or less freely. In an insulator they can move only by trading places.²) We take these interactions among the atoms (and their electrons) as the defining characteristic of a condensed matter system.

This characteristic of condensed matter, namely important interactions among the atoms (and possibly the loss of integrity of individual atoms due to detachment of outer-shell electrons), leads to a fundamental difference from gaseous matter, in that many properties of fluids and solids differ qualitatively from the properties of an aggregate of isolated atoms. When the atoms are in close proximity, low-energy states of the system have strongly correlated atomic positions. Owing to the motion of electrons throughout the system, low-energy states may then also have strongly correlated electronic positions. Low-energy excitations of the system usually involve subtle changes in the atomic or electronic degrees of freedom and have an energy scale much smaller than the binding energy scale for isolated atoms. Many physical properties of a system at a temperature T depend on those excitations which have an energy less than the thermal energy $k_B T$. The Boltzmann constant $k_B \approx 8.167 \times 10^{-5} \text{ eV K}^{-1}$ so that the thermal energy at room temperature is $\sim 2.35 \times 10^{-2} \text{ eV}$. The binding energy of an isolated hydrogen atom, the Rydberg $\sim 13.6 \text{ eV}$, represents the vastly larger chemical energy scale. For this reason, many physical properties of condensed matter systems, from the absolute zero of temperature to temperatures many times higher than room temperature, reflect the possibilities for small rearrangement of correlations among the huge numbers of atomic degrees of freedom. Thus, the low-energy/temperature properties of condensed matter systems are *emergent*; that is, these are collective properties that make the “whole greater than the sum of the parts.” These collective properties (examples of which include the rigidity of solids, superfluidity, superconductivity, and the various quantum Hall effects) emerge through such subtle correlated motions involving very large number of particles, and are associated with various quantum or classical phase transitions in the thermodynamic limit. Understanding these subtle correlations and how the physical properties of condensed matter systems depend on them is the business of condensed matter physics.

An isolated atom consists of a small positively charged nucleus composed of protons and neutrons and a surrounding “cloud” of negatively charged electrons. The size of the nucleus is given by the length scales of nuclear physics (that is, the length scale of the strong force) $\sim 10^{-15} \text{ m}$ to 10^{-14} m . Furthermore, the energy scale needed to excite the nucleus out of its ground state is roughly six orders of magnitude larger than typical chemical energy scales. Hence, it is almost always the case in condensed matter physics that the nucleus can be taken to be an inert point particle with only translational degrees of freedom. (The orientation of the spin of the nucleus is occasionally important.) The familiar picture of an atom has electrons occupying a set of atomic bound states. Electrons in the most strongly bound states are most likely to be close to the nucleus, while the most weakly bound states are most likely to be found far from the nucleus. The size of the electron cloud is given by the length scales of atomic physics $\sim 10^{-10} \text{ m}$ to 10^{-9} m . Of course, this picture of the atom is only approximate, since it is based on the quantum mechanics of a single electron moving in an average potential of all the other electrons. In fact the positions of electrons in an atom are correlated. It is now possible to solve the many-electron quantum mechanics problem for an atom with extremely high accuracy. (That is a story we will not allow ourselves to be distracted by here.) In a condensed matter system it is often the case that electrons in states which evolve from some of the more weakly

² Since all electrons of the same spin are indistinguishable we have to be a little careful in what we mean by “moving by trading places.”




bound atomic orbitals move relatively freely through the entire system. The presence of these *itinerant* electrons is very important. For example, they allow electric currents to flow through the system with relative ease. Correlations among itinerant electrons are very intricate and, unlike the atomic case, we frequently do not yet understand them adequately. This is one of the frontiers of condensed matter physics.

It is usually appropriate to think of a condensed matter system as consisting of ion cores and (possibly) itinerant valence electrons. The ion cores are composed of the nucleus and those atomic orbitals which are still tightly bound to an individual atom in the condensed matter system. Itinerant electrons are not present in every system and, as we will understand better later, the distinction between ion-core electrons and itinerant electrons can become fuzzy.

We are interested in understanding how the macroscopic physical properties of condensed matter systems depend on their microscopic underpinning. We are interested in thermodynamic properties like the specific heat and the magnetic susceptibility and in structural properties, i.e. how the ion cores are distributed within the system. We will learn how a system carries electric and heat currents, and about many other interesting properties. All of these macroscopic properties reflect the quantum and statistical mechanics of the ion cores and the itinerant electrons.

Condensed matter physics is no longer content with understanding on a fundamental level the properties of those condensed matter systems which nature readily provides. Present frontiers are often found in attempts to understand the properties of artificial materials which frequently have unusual and sometimes unexplained properties. Often artificial materials are designed to have desired physical properties. Occasionally, new knowledge makes it possible to fabricate materials whose properties are extremely useful. Probably the most spectacular example of this, at present, is the engineered electronic material used in semiconductor devices (e.g. the microwave amplifiers in every cell phone). The hope of condensed matter physicists is that there will be many more equally spectacular examples in the future.

Condensed matter physics is such a vast, diverse, and rapidly growing field that it is impossible to do justice to all of it in a single volume. The field is roughly divided into two halves: “ $\hbar = 0$ ” (“soft” condensed matter) and “ $\hbar = 1$ ” (electronic/quantum). This book focuses mostly on the electronic/quantum subfield and attempts to include some of the numerous important scientific advances made since the time of the publication of the classic text by Ashcroft and Mermin [1]. Readers interested in the “ $\hbar = 0$ ” side of things should consult the text by Chaikin and Lubensky [2].

We have attempted to keep the level of discussion accessible to beginning graduate students in physics. More advanced sections that can be skipped on first reading are marked with Knuth’s “dangerous bend” T_EX symbols  and  .

1.2 Classification (or Phases) of Condensed Matter Systems

Every system is different in some way.³ It is, however, neither possible nor necessary to study every condensed matter system in nature, or every theoretical model one can write down. We would like, instead, to group different systems with *qualitatively* similar properties together and study their common and robust (often called “universal”) properties. In order to do this, we need to classify condensed matter systems into various classes or phases. This type of analysis shows us, for example, that (rather remarkably) the thermodynamic critical point of a liquid/vapor system has the same

³ Following Tolstoy’s famous quote from *Anna Karenina*, “Happy families are all alike; every unhappy family is unhappy in its own way,” this suggests that every system is an unhappy family of atoms.

universal properties as the magnetic critical point of the three-dimensional (3D) Ising model of classical magnetism.

Owing to the vast domain of condensed matter physics, there is not a single complete scheme to classify all condensed matter systems. In the following we discuss several different but overlapping schemes; which one is more appropriate depends on the specific class of systems being studied, and the specific properties of interest.

1.2.1 Atomic Spatial Structures

Since our condensed matter systems are made of atoms, we can classify them in terms of the spatial structures (or patterns of the positions) of the atoms, or more precisely, their ion cores or atomic nuclei. The two most familiar forms of condensed matter, namely solid and liquid, have their most significant difference precisely in this respect: in a solid the atoms form periodic arrays and have long-range positional order, while in a liquid the atomic positions do not have any such global pattern or long-range order. This difference in spatial structure also leads to the biggest difference in their physical properties: a solid has rigidity but cannot flow, because once the position of a single atom is fixed, so are those of all other atoms. A liquid, on the other hand, can flow easily and has no rigidity because there is no fixed pattern for the atomic positions. One immediate consequence is that a liquid can take any shape dictated by the container in which it is placed, while the shape of a piece of solid does not change as easily. We have learned in thermal physics that solids and liquids are different phases of matter, with a thermodynamic phase boundary separating them. Frequently, these phase boundaries are first-order ones, meaning that there is a finite discontinuity in the internal energy and the entropy in crossing the boundary.

1.2.2 Electronic Structures or Properties

Many physical properties of condensed matter systems, especially those important for applications, are dictated by the behavior of the electrons within them. We often classify solids (and sometimes even liquids) into insulators, semiconductors, metals, and superconductors. Such classifications, of course, are based on the ability of the system to transport electric charge. Electric current in condensed matter systems is almost always carried by electrons (nuclei, while charged objects as well, are too heavy to make a significant contribution, though there do exist materials which are fast-ion conductors of interest in the construction of batteries); in general, electric transport properties are dominated by electronic structure.

Many condensed matter systems are magnetic, and we can classify them in terms of their magnetic properties as paramagnets, diamagnets, ferromagnets, antiferromagnets, etc. Such magnetic properties are also dominated by electrons, through the magnetic moments from their spins and orbital motions. Very often the electric transport and magnetic properties are closely related. For example, non-magnetic metals tend to be **paramagnets**, while non-magnetic insulators tend to be **diamagnets**; **ferromagnets** tend to be metals while **antiferromagnets** are often insulators of a specific type, known as **Mott insulators**. Superconductors, on the other hand, are also “perfect diamagnets” in a very precise sense that we will discuss later.⁴ We thus refer to electric transport and magnetic properties collectively as “electronic properties.”

⁴ A paramagnet has no magnetic polarization unless an external field is applied, in which case the magnetization is proportional to (and reinforces) the applied field. A diamagnet is similar, but its response opposes an applied field. Superconductors can act as perfect diamagnets in some regimes, completely resisting the penetration of any applied field into the sample by generating an exactly opposing field.

1.2.3 Symmetries

Many condensed matter phases can be characterized by their symmetry properties. A symmetry is a transformation which leaves the Hamiltonian H invariant. Typical examples include translation (in a system with no external potential), inversion about a special point ($\vec{r} \rightarrow -\vec{r}$ for a certain choice of origin), reflection about a particular plane, and time-reversal. A specific phase, either thermal (i.e. classical, at finite temperature T) or quantum ($T = 0$), may break some of the symmetries of the Hamiltonian. That is, the phase may have less symmetry than the Hamiltonian that describes the system. For example, the space translation and rotation symmetries are (spontaneously) broken by a crystalline solid, while they are respected by a liquid. A liquid (at least when averaged over a period of time) is completely uniform and translation-invariant. A solid, however, is not. In a perfect crystal, the atoms sit in particular places and once you know the position of a few of them, you can predict the positions of all the rest (modulo small thermal and quantum fluctuations), no matter how far away. If we translate the crystal by some arbitrary amount, the atoms end up in different places but (absent an external potential) the energy is left invariant. The initial position among all these degenerate possibilities is an accident of the history. This is what we mean when we say that the symmetry breaking is “spontaneous.”

Ferromagnets and antiferromagnets both break spin rotation symmetry and time-reversal symmetry, while both of these symmetries are respected in paramagnets and diamagnets. An antiferromagnet has two sublattices of opposite spin and thus further breaks lattice **translation symmetry**, which is *unbroken* in a ferromagnet.

In a phase that breaks one or more symmetries, the pattern of symmetry breaking can be characterized by the so-called “**order parameters**.” An order parameter is a measurable physical quantity, which transforms non-trivially under a symmetry transformation that leaves the Hamiltonian invariant. To understand the meaning of this definition, consider a ferromagnet. Here the order parameter is the magnetization. Magnetization changes sign under time reversal (i.e. “transforms non-trivially under a symmetry transformation that leaves the Hamiltonian invariant”). Hence, if the order parameter is non-zero, the system has less symmetry than its Hamiltonian. Furthermore, we can deduce from the symmetry of the Hamiltonian that states with opposite values of the order parameter are degenerate in energy.

In a crystal the order parameter is the (Fourier transform of the) atomic density. The Fourier transform of the density will have sharp peaks at wave vectors commensurate with the lattice spacing. Under translation of the crystal by some arbitrary small amount, the phases of the (complex) Fourier amplitudes will change, but not their magnitudes. These are thus order parameters associated with broken translation symmetry. The order parameters are zero in the corresponding symmetry-unbroken phases, namely liquid phases (which are translation-symmetric). Similarly, paramagnet/diamagnet phases have zero magnetic order parameters since they are non-magnetic in the absence of an applied field, i.e. when the Hamiltonian is time-reversal symmetric.

Superfluids (which are in some ways closely related to superconductors yet different from them in a fundamental way because they are charge-neutral) have a complex order parameter with continuously variable phase. They break the continuous symmetry corresponding to particle number conservation.⁵ This is a subtler and less familiar form of symmetry breaking, but its mathematical structure (including that of its order parameter) turns out to be identical to that of a ferromagnet whose magnetization vector is restricted to a plane (sometimes referred to as an “ XY ” or “easy-plane” magnet).

⁵ In physics there is a one-to-one correspondence between conservation laws and *continuous* symmetries; in the context of field theory, this is known as Noether’s theorem.

Symmetry is a unifying concept in physics, and is of paramount importance. It is therefore not surprising that it appears both in the atomic crystal structure and in the electronic structure classification schemes.

1.2.4 Beyond Symmetries

It was the great Soviet physicist Lev Landau who first advocated using symmetry as a classification scheme for phases. The Landau scheme has been tremendously successful, both for classifying phases and for developing phenomenological models of symmetry breaking in thermodynamic phase transitions. Owing to the success and the great influence of Landau and his followers, it was felt for many years that we could classify *all* condensed matter phases using symmetry. In recent years, however, especially since the discovery of the fractional quantum Hall effect in 1982, examples of *distinctive* phases with the *same* symmetry have been accumulating. Such phases obviously do *not* fit into the Landau symmetry scheme. In the following we discuss an elementary example, which the reader should be familiar with from her/his undergraduate-level solid state or modern physics courses.

The simplest types of metal and insulator are the so-called band metal and band insulator, formed by filling electron bands of a (perfect) crystal with non-interacting electrons (an idealization). Metals and insulators have the same symmetry properties, but are obviously different phases. So what is the (qualitative) difference between them? The difference, of course, lies in the fact that all bands are either completely filled or empty in a band insulator, while there is at least one partially filled band in a metal, resulting in one or more Fermi surface⁶ sheet(s). This is a *topological* difference, as the number of Fermi surface sheets is a topological invariant (quantum number) that does *not* change (at least not continuously) when the shape or *geometry* of a Fermi surface is varied due to the band structure. Furthermore, it can be shown that when the number of Fermi surface sheets changes, the system must undergo a quantum phase transition (known as a Lifshitz transition) at which the ground-state energy or its derivatives become singular. However unlike the solid–liquid transition, there is no change in the symmetry properties of the system in a Lifshitz transition.

With the discovery of the fractional quantum Hall effect, physicists started to realize that many phases with excitation gaps separating their ground states from excited states have non-trivial topological properties. They are thus termed *topological phases* and are said to possess *topological order* [15]. All quantum Hall phases, as well as the recently discovered topological insulators, are examples of topological phases, which will be discussed in this book. Perhaps the most familiar but also under-appreciated example is an ordinary superconductor. In contrast to common wisdom and unlike superfluids (which are uncharged), there is no spontaneous symmetry breaking in a superconductor [16]. The topological nature of the superconducting phase is reflected in the fact that its ground-state degeneracy is dependent on the topology of the space in which it is placed, but *nothing else*. As we will learn, robust and topology-dependent ground-state degeneracy is one of the most commonly used methods to probe and characterize topological phases.

1.3 Theoretical Descriptions of Condensed Matter Phases

In atomic physics one typically attempts to give a full microscopic description of the atom or molecule by solving Schrödinger's equation to obtain the (few-electron) wave functions. Except for the simplest

⁶ Because electrons are fermions, they obey the Pauli exclusion principle and each must occupy a different quantum state. As we will learn later, the ground state for non-interacting electrons has all states occupied up to the so-called Fermi energy, above which no states are occupied. The set of states having energy equal to the Fermi energy defines a surface with one or more sheets in momentum space.

case of a hydrogen-like atom with one electron, analytic exact solution is not possible, and numerical solutions become out of reach very quickly with increasing electron number. In an interacting many-electron condensed matter system, microscopic descriptions based on electron wave functions are still widely used. Such descriptions are of course approximate (except for very few highly idealized model systems), often based on some mean-field type of approximations or variational principles. Such descriptions are highly effective when the approximate or variational wave function is simple yet captures the most important correlations of the phase that the system is in. The most famous examples of this are the Bardeen–Cooper–Schrieffer (BCS) wave function for superconductors and the Laughlin wave function for fractional quantum Hall liquids.

Very often, even approximate microscopic descriptions are beyond reach. Fortunately, to understand physical properties that are probed experimentally and important for applications, we often need only understand how a condensed matter system responds to an external perturbation at low frequency or energy (compared with microscopic or atomic energy scales, typically eV), and/or at long wavelength (again compared with the atomic scale, 1 Å). Most microscopic degrees of freedom do *not* make significant contributions to the response in this limit. We thus need only focus on the low-energy, long-wavelength degrees of freedom that dominate such responses. Also fortunate is the fact that, very often, the physics simplifies significantly in the low-energy limit, rendering an accurate description in terms of these (often heavily “renormalized” or “effective”) degrees of freedom possible. Such simplification can often be understood theoretically in terms of a **renormalization group** analysis.⁷ A theoretical description of this type goes under the name of a “low-energy effective theory,” and we will encounter several examples. Condensed matter systems in the same phase share the same low-energy effective theory (but possibly with different parameters), while different phases can be characterized by different low-energy effective theories. These concepts will become clearer as we study particular examples in later chapters.

Band metals again serve as a simple example of the discussion above. In the idealization of non-interacting electrons, the ground state is simply the single Slater determinant in which all electron levels below the Fermi energy are occupied and all above are empty, and excitations are created by moving one or more electrons from below the Fermi energy (or inside the Fermi surface) to above (or outside). It is already clear in this simple case that low-energy excitations involve electron states near the Fermi surface(s) *only*. Remarkably, this remains true even when electron–electron interactions are turned on (as long as no phase transition is caused by the interaction), and in some sense electrons near the Fermi surface continue to behave as if they do *not* interact with each other. It is thus possible to describe the low-energy/temperature properties of a metal in terms of the electrons that live near the Fermi surface (often called Landau quasiparticles), and such a low-energy effective theory is known as Landau’s Fermi-liquid theory (yes, the same Landau, but symmetry or its breaking plays no role here!).

In symmetry-breaking phases, the low-energy degrees of freedom are often those giving rise to the order parameter. This is particularly true when a continuous symmetry is broken spontaneously, as guaranteed by the **Goldstone theorem**, which states that there is a *gapless* collective mode (known as a “**Goldstone mode**”) associated with each **spontaneously broken continuous symmetry**. The gapless mode is nothing but the long-wavelength fluctuation of the order parameter. As a result of this, the corresponding low-energy effective theory often takes the form of a field theory (either quantum or classical, depending on whether we are at zero or finite temperature), in which the field is nothing but the local order parameter. Such low-energy effective field theories go under the name of Ginzburg–Landau theories (yes, the same Landau again).

⁷ This is an analysis that relates the microscopic theory describing the system at short length and time scales to an effective theory (possibly having different effective masses and interactions for the particles) that describes the physics at longer length and time scales.

Ginzburg–Landau theories can also be used to describe topological phases. Actually, the birth place of (the original) Ginzburg–Landau theory is superconductivity, which is a topological phase. With proper generalizations they can also be used to describe the simplest types of quantum Hall phases. These types of Ginzburg–Landau theory are also written in terms of a local order parameter field. One might wonder how one can have an order parameter without a broken symmetry. The difference here is that in these topological phases the order parameters couple to gauge fields and are therefore *not* (by themselves) gauge-invariant; as a result, the presence of an order parameter does *not* signal symmetry breaking. Our understanding of topological phases is far from complete, and it is not clear what the low-energy effective theories are for some of the known examples of topological phases. On the other hand, many well-understood topological phases (in particular quantum Hall phases) can be well described by topological quantum field theories. This discussion is doubtless mysterious to the reader at the beginning, but the issues will be clarified as we proceed.

1.4 Experimental Probes of Condensed Matter Systems

Most experimental probes of condensed matter systems are based on linear response. As explained in Appendix A, one weakly perturbs the system, and measures how the system responds to the perturbation. In the linear response regime, the response is proportional to the perturbation, and one thus measures the ratio between the two, known as the response function (at the frequency and wavelength of the perturbation). For example, electric current is the response to a (perturbing) electric field or voltage drop, and the ratio between them is the conductance that one measures in a transport experiment, while magnetization is the response to an external magnetic field, and the ratio between them is the magnetic susceptibility measured in a magnetic measurement. In many cases, the frequency of the probing perturbation is low, and the wavelength is long compared with the characteristic microscopic scales of the system, and that is why we need only focus on the low-energy and long-wavelength properties of the system.

Different probes are used to probe different properties of the system. Atomic spatial structures are most often probed using X-ray scattering, while neutron scattering can probe both atomic lattice vibrations and magnetic excitations. Thermodynamic measurements (like specific heat) probe contributions from all degrees of freedom. Electronic contributions, on the other hand, dominate electric responses like conductivity. We will discuss these and many other experimental probes in later chapters.