
PART I

On a fundamental level, all forms of quantum matter can be formulated in terms of a many-body Hamiltonian for a macroscopically large number of constituent particles. However, in contrast with many other areas of physics, the structure of this operator conveys as much information about the properties of the system as, say, the knowledge of the basic chemical constituents tells us about the behavior of a living organism. Rather, it has been a long-standing tenet that the degrees of freedom relevant to the low-energy properties of a system usually are not the microscopic ones. It is a hallmark of many “deep” problems of modern condensed matter physics that the passage between the microscopic and the effective degrees of freedom involves complex and, at times, even controversial mappings. To understand why, it is helpful to review the process of theory building in this field of physics.

adiabatic
continuity

The development of early condensed matter physics often hinged on the “unreasonable” success of *non-interacting* theories. The impotency of interactions observed in a wide range of physical systems reflects a principle known as **adiabatic continuity**: the quantum numbers characterizing an (interacting) many-body system are determined by fundamental symmetries – translational, rotational, particle exchange, etc. As long as these symmetries are maintained, the system’s elementary excitations, or *quasi-particles*, can usually be traced back “adiabatically” to the bare particles of the non-interacting limit. This principle, embodied in Landau’s Fermi-liquid theory, has provided a platform for the investigation of a wide range of systems, from conventional metals to ³helium fluids and cold atomic Fermi gases.

However, being contingent on symmetry, it must be abandoned at phase transitions, where interactions effect a rearrangement of the many-body ground state into a state of different, or “broken” symmetry. Symmetry-broken phases generically show excitations different from those of the parent non-interacting phase. They either require classification in terms of new species of quasi-particles, or they may be *collective modes* engaging the cooperative motion of many bare particles. For example, when atoms condense from a liquid into a solid phase, translational symmetry

is broken and the elementary excitations (phonons) involve the motion of many individual bare particles.

In this way, each phase of matter is associated with a unique “non-interacting” reference state with its own characteristic quasi-particle excitations – a product of only the relevant symmetries. Within each individual phase, a continuity principle keeps the effects of interactions at bay. This hierarchical picture delivers two profound implications. First, within the quasi-particle framework, the underlying “bare” or elementary particles remain invisible. (To quote from P. W. Anderson’s famous article *More is different*, *Science* **177**, 393 (1972), “The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe.”) Second, while one may conceive an almost unlimited spectrum of interactions, there are comparatively few non-interacting or free theories, constrained by the set of fundamental symmetries. These arguments go a long way in explaining the principle of “universality” observed in condensed matter.

How can these concepts be embedded into a concrete theoretical framework? At first sight, problems with macroscopically many particles seem overwhelmingly daunting. However, our discussion above indicates that representations of manageable complexity may be obtained by focusing on symmetries and restricted sets of excitations. Quantum field theory provides the keys to making this reduction concrete. Starting from an efficient microscopic formulation of the many-body problem, it allows one to systematically develop effective theories for collective degrees of freedom. Such representations afford a classification of interacting systems into a small number of universality classes defined by their fundamental symmetries. This form of complexity reduction has become a potent source of unification in modern theoretical physics. Indeed, several subfields of theoretical physics (such as conformal field theory, random matrix theory, etc.) now define themselves not so much through any specific application as by a certain conceptual or methodological framework.

As mentioned in the preface, the first part of this text is a primer aimed at elevating graduate students to a level where they can engage in independent research. While the discussion of conceptual aspects stands in the foreground, we have endeavored to keep the text firmly rooted in experimental application. As well as routine exercises, it includes extended problems meant to train research-oriented thinking. Some of these answered problems are deliberately designed to challenge. (We all know from experience that the intuitive understanding of formal structures can be acquired only by persistent, and at times even frustrating training.)

1 From Particles to Fields


SYNOPSIS To introduce some basic concepts of field theory, we begin by considering two simple model systems – a one-dimensional “caricature” of a solid and a freely propagating electromagnetic wave. As well as exemplifying the transition from discrete to continuous degrees of freedom, these examples introduce the basic formalism of classical and quantum field theory as well as the notions of elementary excitations, collective modes, symmetries and universality – concepts that will pervade the rest of the text.

One of the appealing facets of condensed matter physics is that phenomenology of remarkable complexity can emerge from a Hamiltonian of comparative simplicity. Indeed, microscopic “condensed matter Hamiltonians” of high generality can be constructed straightforwardly. For example, a prototypical metal or insulator may be described by the **many-particle Hamiltonian**, $H = H_e + H_i + H_{ei}$, where

many-
particle
Hamiltonian

$$\begin{aligned}
 H_e &= \sum_i \frac{\mathbf{p}_i^2}{2m} + \sum_{i<j} V_{ee}(\mathbf{r}_i - \mathbf{r}_j), \\
 H_i &= \sum_I \frac{\mathbf{P}_I^2}{2M} + \sum_{I<J} V_{ii}(\mathbf{R}_I - \mathbf{R}_J), \\
 H_{ei} &= \sum_{iI} V_{ei}(\mathbf{R}_I - \mathbf{r}_i).
 \end{aligned} \tag{1.1}$$

Here, $\mathbf{r}_i(\mathbf{R}_I)$ denotes the coordinates of valence electrons (ion cores), while H_e, H_i , and H_{ei} describe the dynamics of electrons, ions and the interaction of electrons and ions, respectively (see the figure). Of course, the Hamiltonian (1.1) can be made more realistic by, for example, remembering that electrons and ions carry spin, adding potential disorder, or introducing host lattices with multi-atomic unit cells. However, for developing our present line of thought, the prototype Hamiltonian H will suffice.



The fact that a seemingly innocuous Hamiltonian like Eq. (1.1) is capable of generating a plethora of phenomenology can be read in reverse: normally, one will not be able to make progress theoretically by approaching the problem in an

- ab initio approach** “**ab initio**” manner, i.e., by an approach that treats all microscopic constituents as equally relevant degrees of freedom. How, then, can successful analytical approaches be developed? The answer lies in several basic principles inherent in generic condensed matter systems.
- reduction principle**
1. **Structural reducibility:** Not all components of the Hamiltonian (1.1) need to be treated simultaneously. For example, when our interest is in the vibrational motion of the ion lattice, the dynamics of the electron system can often be neglected or, at least, treated in an “effective” manner. Similarly, the dynamics of the electron system can often be considered independent of the ions, etc.
- universality principle**
2. In the majority of condensed matter applications, one is interested not so much in the full profile of a given system but, rather, in its energetically low-lying dynamics. This is motivated partly by practical aspects (in daily life, iron is normally encountered at room temperature and not at its melting point), and partly by the tendency of large systems to behave in a “universal” manner at low temperatures. Here, **universality** implies that systems differing in microscopic detail (i.e., with different types of interaction potentials, ion species, etc.) exhibit common collective behavior at low energy or long length scales. As a physicist, one will normally seek for unifying principles in collective phenomena rather than to describe the peculiarities of individual elements or compounds. However, universality is equally important in the *practice* of condensed matter theory. In particular, it implies that, at low temperatures, system-specific details of microscopic interaction potentials are often of secondary importance, i.e., one may employ *simple* model Hamiltonians.
- statistical principles**
3. For most systems of interest, the number of degrees of freedom $N = \mathcal{O}(10^{23})$ is formidably large. However, contrary to first impressions, the magnitude of this figure is rather an advantage: in addressing condensed matter problems, the **principles of statistics** imply that statistical errors tend to be negligibly small.¹
- symmetries**
4. Finally, condensed matter systems typically possess intrinsic **symmetries**. For example, the Hamiltonian (1.1) is invariant under the simultaneous translation and/or rotation of all coordinates, which expresses the global Galilean invariance of the system (these are continuous symmetries). Invariance under spin rotation (continuous) or time reversal (discrete) are other examples of common symmetries. The general importance of symmetries cannot be overemphasized: symmetries support the conservation laws that simplify any problem. Yet, in

¹ The importance of this point is illustrated by the empirical observation that the most challenging systems in physical sciences are of *medium*, and not large, scale, e.g., metallic clusters, medium-sized nuclei or large atoms consisting of some $\mathcal{O}(10^1\text{--}10^2)$ fundamental constituents. Such systems are beyond the reach of few-body quantum mechanics while not yet accessible to reliable statistical modeling. The only viable path to approaching systems of this type is often through numerical simulation or the use of phenomenology.

condensed matter, symmetries are even more important. A conserved observable is generally tied to an energetically low-lying excitation. In the universal, low-temperature, regime in which we will typically be interested, it is precisely the dynamics of these excitations that govern the gross behavior of the system. Generally, the identification of fundamental symmetries is the first step in the sequence “symmetry \mapsto conservation law \mapsto low-lying excitations” and one that we will encounter time and again.


To understand how these basic principles can be used to formulate and explore effective low-energy field theories of solid state systems, we will begin by focusing on the **harmonic chain**, a collection of atoms bound by a harmonic potential. In doing so, we will observe that the universal characteristics encapsulated by the low-energy dynamics² of large systems relate naturally to concepts of **field theory**.

1.1 Classical Harmonic Chain: Phonons

[Classical Harmonic Chain: Phonons]

Returning to the prototype Hamiltonian (1.1), let us focus on the dynamical properties of the positively charged *core ions* that constitute the host lattice of a crystal. For the moment, we will neglect the fact that atoms are quantum objects and treat the ions as *classical* entities. To further simplify the problem, let us consider a one-dimensional atomic chain rather than a generic d -dimensional solid. In this case, the positions of the ions can be specified by a sequence of coordinates with average lattice spacing a . Relying on the structural reducibility principle 1, we will first argue that, to understand the behavior of the ions, the dynamics of the conduction electrons are of secondary importance, i.e., we will set $H_e = H_{ei} = 0$.

At zero temperature, the system freezes into a regularly spaced array of ion cores at coordinates $R_I = \bar{R}_I \equiv Ia$. Any deviation from a perfectly regular configuration incurs a potential energy cost. For low enough temperatures (principle 2), this energy will be approximately quadratic in the small deviation of the ion from its equilibrium position. The “reduced” low-energy Hamiltonian of the system then reads



$$H = \sum_{I=1}^N \left[\frac{P_I^2}{2m} + \frac{k_s}{2} (R_{I+1} - R_I - a)^2 \right], \quad (1.2)$$

² In this text, we will focus on the *dynamical* behavior of large systems, as opposed to their *static* structural properties. In particular, we will not address questions related to the formation of definite crystallographic structures in solid state systems.

where the coefficient k_s determines the steepness of the lattice potential. Note that H can be interpreted as the Hamiltonian of N point-like particles of mass m connected by elastic springs with spring constant k_s (see the figure).

1.1.1 Lagrangian formulation and equations of motion

classical
harmonic
chain

What are the elementary low-energy excitations of the **classical harmonic chain**? To answer this question we might, in principle, attempt to solve Hamilton's equations of motion. Indeed, since H is quadratic in all coordinates, such a program is feasible. However, few of the problems encountered in solid state physics enjoy this property. Further, it seems unlikely that the low-energy dynamics of a macroscopically large chain – which we know from our experience will be governed by *large-scale* wave-like excitations – is adequately described in terms of an “atomistic” language; the relevant degrees of freedom will be of a different type. We should, rather, draw on the basic principles 1–4 set out above. Notably, so far, we have paid attention to neither the intrinsic symmetry of the problem nor the fact that the number of ions, N , is large.

Joseph-Louis Lagrange
1736–1813

was a French mathematician and astronomer (though born in Turin) who excelled in all fields of analysis, number theory and celestial mechanics. In 1788, he published *Mécanique Analytique*, which summarized the field of mechanics since the time of Newton, and is notable for its use of the theory of differential equations. In this text, he transformed mechanics into a branch of mathematical analysis.



To reduce a microscopic model to an effective low-energy theory, often the Hamiltonian is not a very convenient starting point. Usually, it is more efficient to start from the **classical action**, S . In the present case, $S = \int dt L(R, \dot{R})$, where $(R, \dot{R}) \equiv \{R_I, \dot{R}_I\}$ represent the coordinates and their time derivatives. The corresponding **classical Lagrangian** L related to the Hamiltonian (1.2) is given by

classical
action
classical
Lagrangian

$$L = T - U = \sum_{I=1}^N \left[\frac{m}{2} \dot{R}_I^2 - \frac{k_s}{2} (R_{I+1} - R_I - a)^2 \right], \quad (1.3)$$

where T and U denote, respectively, the kinetic and potential energy.

Since we are interested in the properties of the large- N system, we can expect boundary effects to be negligible. In this case, we may impose periodic boundary conditions, making the identification $R_{N+1} = R_1$. Further, anticipating that the effect of lattice vibrations on the solid is weak (i.e., long-range atomic order is maintained), we may assume that the deviation of ions from their equilibrium position is small (i.e., $|R_I(t) - \bar{R}_I| \ll a$). For $R_I(t) = \bar{R}_I + \phi_I(t)$, with $\phi_{N+1} = \phi_1$, the Lagrangian (1.3) assumes the simplified form


$$L = \sum_{I=1}^N \left[\frac{m}{2} \dot{\phi}_I^2 - \frac{k_s}{2} (\phi_{I+1} - \phi_I)^2 \right].$$

To make further progress, we will now make use of the fact that we are not concerned with behavior on “atomic” scales. For such purposes, our model would,

continuum
limit

in any case, be much too primitive! Rather, we are interested in experimentally observable behavior that becomes manifest on macroscopic length scales (principle 2). For example, one might wish to study the specific heat of the solid in the limit of infinitely many atoms (or at least a macroscopically large number, $\mathcal{O}(10^{23})$). Under these conditions, microscopic models can usually be simplified substantially (principle 3). In particular, it is often permissible to subject a discrete lattice model to a **continuum limit**, i.e., to neglect the discreteness of the microscopic entities and to describe the system in terms of effective continuum degrees of freedom.

In the present case, taking a continuum limit amounts to describing the lattice fluctuations ϕ_I in terms of smooth functions of a continuous variable x . (See the figure, where the [horizontal] displacement of the point particles is plotted along the vertical axis.)



Clearly such a description makes sense only if the relative fluctuations on atomic scales are weak (for otherwise the smoothness condition would be violated). However, if this condition is met – as it will be for sufficiently large values of the stiffness constant k_s – the continuum description is much more powerful than the discrete encoding in terms of the “vector” $\{\phi_I\}$. The steps that we will need to take to go from the Lagrangian to concrete physical predictions will then be much easier to formulate.

Introducing continuum degrees of freedom $\phi(x)$, and applying a first-order Taylor expansion,³ let us define

$$\phi_I \rightarrow a^{1/2} \phi(x) \Big|_{x=Ia}, \quad (\phi_{I+1} - \phi_I) \rightarrow a^{3/2} \partial_x \phi(x) \Big|_{x=Ia}, \quad \sum_{I=1}^N \rightarrow \frac{1}{a} \int_0^L dx,$$

where $L = Na$. Note that, as defined, the functions $\phi(x, t)$ have dimensionality $[\text{length}]^{1/2}$. Expressed in terms of the new degrees of freedom, the continuum limit of the Lagrangian then reads

$$L[\phi] = \int_0^L dx \mathcal{L}(\partial_x \phi, \dot{\phi}), \quad \mathcal{L}(\partial_x \phi, \dot{\phi}) = \frac{m}{2} \dot{\phi}^2 - \frac{k_s a^2}{2} (\partial_x \phi)^2, \quad (1.4)$$

Lagrangian
density

where the **Lagrangian density** \mathcal{L} has dimensionality $[\text{energy}]/[\text{length}]$. Similarly, the classical action assumes the continuum form

$$S[\phi] = \int dt L[\phi] = \int dt \int_0^L dx \mathcal{L}(\partial_x \phi, \dot{\phi}) \quad (1.5)$$

classical
field

We have thus succeeded in abandoning the N point-particle description in favor of one involving *continuous* degrees of freedom, a **(classical) field**. The dynamics of the latter are specified by the **functionals** L and S , which represent the continuum generalizations of the discrete classical Lagrangian and action, respectively.


³ Indeed, for reasons that will become clear, higher-order contributions to the Taylor expansion are immaterial in the long-range continuum limit.

field **INFO** The continuum variable ϕ is our first encounter with a **field**. Before proceeding with our example, let us pause to make some preliminary remarks on the general definition of these objects. This will help to place the subsequent discussion of the atomic chain in a broader context. Formally, a field is a smooth mapping

$$\begin{aligned} \phi : M &\rightarrow T \\ z &\mapsto \phi(z) \end{aligned}$$

field manifold

from a certain manifold M ,⁴ often called the “**base manifold**,” into a “**target**” or “**field manifold**” T (see the figure).⁵ In our present example, $M = [0, L] \times [0, t] \subset \mathbb{R}^2$ is the product of intervals in space and time, and $T = \mathbb{R}$. In fact, the factorization $M \subset \mathcal{R} \times \mathcal{T}$ into a space-like manifold \mathcal{R} multiplied by a one-dimensional time-like manifold \mathcal{T} is inherent in most applications of condensed matter physics.⁶



However, the individual factors \mathcal{R} and \mathcal{T} may, of course, be more complex than in our prototypical problem above. As for the target manifold, not much can be said in general; depending on the application, the realizations of T range from real or complex numbers over vector spaces and groups to the “fanciest objects” of mathematical physics.

functionals

In applied field theory, fields appear not as final objects, but rather as input to **functionals**. Mathematically, a functional $S : \phi \mapsto S[\phi] \in \mathbb{R}$ is a mapping that takes a field as its argument and maps it into the real numbers. The functional profile $S[\phi]$ essentially determines the character of a field theory. Notice that the argument of a functional is commonly indicated in square brackets $[\]$.

While these formulations may appear unnecessarily abstract, remembering the mathematical backbone of the theory often helps to avoid confusion. At any rate, it takes some time and practice to get used to the concept of fields and functionals. Conceptual difficulties in handling these objects can be overcome by remembering that any field in condensed matter physics arises as the limit of a *discrete* mapping. In the present example, the field $\phi(x)$ is obtained as a continuum limit of the discrete vector $\{\phi_I\} \in \mathbb{R}^N$; the functional $L[\phi]$ is the continuum limit of the function $L : \mathbb{R}^N \rightarrow \mathbb{R}$, etc. While, in practice, fields are usually easier to handle than their discrete analogs, it is sometimes helpful to think about problems of field theory in a discrete language. Within the discrete picture, the mathematical apparatus of field theory reduces to finite-dimensional calculus.

equations of motion

Although the Lagrangian (1.4) contains the full information about the model, we have not yet learned much about its actual behavior. To extract concrete physical information from the Lagrangian, we need to derive **equations of motion**. At first sight, it may not be entirely clear what is meant by the term “equations of motion” in the context of an infinite-dimensional model: the equations of motion relevant for

⁴ If you are unfamiliar with the notion of manifolds (for a crash course, see appendix section A.1), think of M and T as subsets of some vector space. For the moment, this limitation won’t do any harm.

⁵ In some (rare) cases it becomes necessary to define fields in a more general sense (e.g., as sections of mathematical objects known as fiber bundles). However, in practically all condensed matter applications, the more restrictive definition above will suffice.

⁶ By contrast, the condition of Lorentz invariance implies the absence of such factorizations in relativistic field theory. In classical statistical field theories, i.e., theories probing the thermodynamic behavior of large systems, M is just space-like.

the present problem are obtained as the generalization of the conventional Lagrange equations of N point-particle classical mechanics to a model with infinitely many degrees of freedom. To derive these equations, we need to generalize Hamilton’s extremal principle (i.e., the route from an action to the associated equations of motion) to infinite dimensions. As a warm-up, let us briefly recapitulate how the extremal principle works for a system with one degree of freedom.

Suppose the dynamics of a classical *point* particle with coordinate $x(t)$ is described by the classical Lagrangian $L(x, \dot{x})$ and action $S[x] = \int dt L(x, \dot{x})$. **Hamilton’s extremal principle** states that the configurations $x(t)$ that are *realized* are those that extremize the action, $\delta S[x] = 0$, i.e., for any smooth curve $t \mapsto y(t)$,

Hamilton’s
extremal
principle

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (S[x + \epsilon y] - S[x]) = 0. \tag{1.6}$$


(For a more rigorous discussion, see section 1.2 below.) To first order in ϵ , the action has to remain invariant. Applying this condition, one finds that it is fulfilled if and only if x satisfies **Lagrange’s equation of motion**

Lagrange’s

$$\frac{d}{dt} (\partial_{\dot{x}} L) - \partial_x L = 0 \tag{1.7}$$

EXERCISE Recapitulate the derivation of (1.7) from the classical action.

In Eq. (1.5) we are dealing with a system of infinitely many degrees of freedom, $\phi(x, t)$. Yet Hamilton’s principle is general and we may see what happens if Eq. (1.5) is subjected to an extremal principle analogous to Eq. (1.6). In this case, we require the action (1.5) to be invariant under variations $\phi(x, t) \rightarrow \phi(x, t) + \epsilon \eta(x, t)$, to first order in ϵ . Note that field variations must respect boundary conditions, if present. For example, if $\phi|_{\text{boundary}} = \text{const.}$, then $\eta|_{\text{boundary}} = 0$ (see the figure). When applied to the specific Lagrangian (1.4), substituting the “varied” field leads to



$$S[\phi + \epsilon \eta] = S[\phi] + \epsilon \int dt \int_0^L dx (m \dot{\phi} \dot{\eta} - k_s a^2 \partial_x \phi \partial_x \eta) + \mathcal{O}(\epsilon^2).$$

Integrating by parts and requiring the contribution linear in ϵ to vanish, one obtains

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (S[\phi + \epsilon \eta] - S[\phi]) = - \int dt \int_0^L dx (m \ddot{\phi} - k_s a^2 \partial_x^2 \phi) \eta \stackrel{!}{=} 0.^7$$

(Notice that the boundary terms vanish identically.) Now, since η was defined to be an arbitrary smooth function, the integral above can vanish only if the factor in

⁷ Here and throughout $a \stackrel{!}{=} b$ means “we require a to be equal to b .”

classical
wave
equation

parentheses is globally vanishing. Thus the equation of motion takes the form of a **classical wave equation**

$$\boxed{(m\partial_t^2 - k_s a^2 \partial_x^2) \phi = 0} \quad (1.8)$$

The solutions of (1.8) have the form $\phi_+(x-vt) + \phi_-(x+vt)$, where $v = a\sqrt{k_s/m}$, and ϕ_{\pm} are arbitrary smooth functions of the argument.



sound
waves

From this we can deduce that the low-energy **elementary excitations** of our model are lattice vibrations propagating as **classical sound waves** to the left or right at a constant velocity v (see figure).⁸ The trivial behavior of the model is, of course, a direct consequence of its simplistic definition – no dissipation, dispersion, or other nontrivial ingredients. Adding these refinements leads to the general classical theory of lattice vibrations (see, e.g., Ashcroft and Mermin⁹). Finally, notice that the elementary excitations of the chain have little in common with its “microscopic” constituents (the atoms). Rather they are **collective excitations**, i.e., elementary excitations comprising a macroscopically large number of microscopic degrees of freedom.

collective
excitations

INFO The “relevant” **excitations of a condensed matter system** can, but need not, be of collective type. For example, the interacting electron gas (a system to be discussed in detail below) supports microscopic excitations – charged quasi-particles standing in 1:1 correspondence with the electrons of the original microscopic system – while the collective excitations are plasmon modes of large wavelength that involve many electrons. Typically, the nature of the fundamental excitations cannot be straightforwardly inferred from the microscopic definition of a model. Indeed, the mere *identification* of the relevant excitations often represents the most important step in the solution of a condensed matter problem.

1.1.2 Hamiltonian formulation

An important characteristic of any excitation is its *energy*. How much energy is stored in the sound waves of the harmonic chain? To address this question, we need to switch back to a Hamiltonian formulation. Once again, this is achieved by generalizing standard manipulations from point-particle mechanics to the continuum. Remembering that, in the Lagrangian formulation of

Sir William Rowan Hamilton 1805–1865

was an Irish mathematician credited with the discovery of quaternions, the first non-commutative algebra to be studied. He also made seminal contributions to the study of geometric optics and classical mechanics.



⁸ Strictly speaking, the modeling of our system enforces a periodicity constraint $\phi_{\pm}(x+L) = \phi_{\pm}(x)$. However, in the limit of large system sizes, this aspect becomes inessential.

⁹ N. W. Ashcroft and N. D. Mermin, *Solid State Physics* (Holt–Saunders International, 1983).