1

Introduction

1.1 Closed time contour

Consider a quantum many-body system governed by a time-dependent Hamiltonian $\hat{H}(t)$. Let us assume that in the distant past $t = -\infty$ the system was in a state specified by a many-body density matrix $\hat{\rho}(-\infty)$. The precise form of the latter is of no importance. It may be, e.g., the equilibrium density matrix associated with the Hamiltonian $\hat{H}(-\infty)$. We shall also assume that the time-dependence of the Hamiltonian is such that at $t = -\infty$ the particles were non-interacting. The interactions are then adiabatically switched on to reach their actual physical strength sometime prior to the observation time. *In addition*, the Hamiltonian may contain *true* time dependence through e.g. external fields or boundary conditions. Due to such true time-dependent perturbations the density matrix is driven away from equilibrium.

The density matrix evolves according to the Von Neumann equation

$$ {}_t\hat{\rho}(t) = -\mathrm{i}\big[\hat{H}(t), \hat{\rho}(t)\big], \tag{1.1}$$

where we set $\hbar = 1$. It is formally solved with the help of the unitary evolution operator as $\hat{\rho}(t) = \hat{\mathcal{U}}_{t,-\infty}\hat{\rho}(-\infty)\big[\hat{\mathcal{U}}_{t,-\infty}\big]^\dagger = \hat{\mathcal{U}}_{t,-\infty}\hat{\rho}(-\infty)\hat{\mathcal{U}}_{-\infty,t}$, where the † denotes Hermitian conjugation. The evolution operator obeys

$$ {}_t\hat{\mathcal{U}}_{t,t^*} = -\mathrm{i}\hat{H}(t)\,\hat{\mathcal{U}}_{t,t^*}; \qquad\qquad {}_{t^*}\hat{\mathcal{U}}_{t,t^*} = \mathrm{i}\,\hat{\mathcal{U}}_{t,t^*}\hat{H}(t^*).$$

Notice that the Hamiltonian operators taken at different moments of time, in general, do not commute with each other. As a result, $\hat{\mathcal{U}}_{t,t^*}$ must be understood as an infinite product of incremental evolution operators with instantaneous locally constant Hamiltonians

$$\begin{aligned}\hat{\mathcal{U}}_{t,t^*} &= \lim_{N\to\infty} \mathrm{e}^{-\mathrm{i}\hat{H}(t-\delta t)\,\delta t}\mathrm{e}^{-\mathrm{i}\hat{H}(t-2\,\delta t)\,\delta t}\ldots\mathrm{e}^{-\mathrm{i}\hat{H}(t-N\,\delta t)\,\delta t}\mathrm{e}^{-\mathrm{i}\hat{H}(t^*)\,\delta t} \\ &= \mathbb{T}\exp\left[-\mathrm{i}\int_{t^*}^{t}\hat{H}(t)\,\mathrm{d}t\right],\end{aligned} \tag{1.2}$$

where an infinitesimal time-step is $\delta t = (t - t^*)/N$ and to shorten the notations the infinite product is abbreviated as the time-ordered, or $\mathbb{T}$-exponent.

One is usually interested to know an expectation value of some observable $\hat{\mathcal{O}}$ (say density or current operator) at a time t.[1] It is defined as

$$\langle\hat{\mathcal{O}}\rangle(t) \pm \frac{\mathrm{Tr}\{\hat{\mathcal{O}}\hat{\rho}(t)\}}{\mathrm{Tr}\{\hat{\rho}(t)\}} = \frac{1}{\mathrm{Tr}\{\hat{\rho}(t)\}}\,\mathrm{Tr}\big\{\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}\hat{\rho}(-\infty)\big\}, \tag{1.3}$$

where the trace is performed over many-body Hilbert space and in the last equality we cyclically permuted the $\hat{\mathcal{U}}_{-\infty,t}$ operator under the trace sign. The expression under the last trace describes (read from right to left) evolution from $t = -\infty$, where the initial density matrix is specified, toward t, where the observable is calculated, and then back to $t = -\infty$. Therefore calculation of an observable implies evolving the initial state both *forward and backward*.

Such forward–backward evolution is avoided in the equilibrium quantum field theory with a special trick. Let us recall how it works, for example, in the zero-temperature equilibrium formalism [2]. The latter deals with the ground state expectation values of the type $\langle\mathrm{GS}|\hat{\mathcal{O}}|\mathrm{GS}\rangle$, where $|\mathrm{GS}\rangle$ is a ground state of an *interacting* many-body system. It is obtained from the known and simple ground state of the corresponding *non-interacting* system $|0\rangle$ by acting on the latter with the evolution operator $|\mathrm{GS}\rangle = \hat{\mathcal{U}}_{t,-\infty}|0\rangle$. Since we are dealing with the equilibrium situation, the only time dependence allowed for the Hamiltonian is an adiabatic switching of the interactions on and off in the distant past and distant future, respectively. The evolution operator therefore describes the evolution of a simple non-interacting ground state $|0\rangle$ toward $|\mathrm{GS}\rangle$ upon adiabatic switching of the interactions and thus $\langle\mathrm{GS}|\hat{\mathcal{O}}|\mathrm{GS}\rangle = \langle 0|\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}|0\rangle$.

Now comes the trick: one argues that

$$\hat{\mathcal{U}}_{+\infty,-\infty}|0\rangle = \mathrm{e}^{\mathrm{i}L}|0\rangle\,. \tag{1.4}$$

That is, evolution of the non-interacting ground state upon adiabatic switching of the interactions on and subsequent adiabatic switching them off brings the system back into the state $|0\rangle$, up to a phase factor $\mathrm{e}^{\mathrm{i}L}$. This statement is based on the belief that the adiabatic perturbation keeps the system in its (evolving) ground state at all times. If so, in view of normalization $\langle 0|0\rangle = 1$, the only possible change is the phase of the non-interacting ground state $\mathrm{e}^{\mathrm{i}L} = \langle 0|\hat{\mathcal{U}}_{+\infty,-\infty}|0\rangle$. Similarly $\langle 0|\hat{\mathcal{U}}_{+\infty,-\infty} = \langle 0|\mathrm{e}^{\mathrm{i}L}$. Accepting this, one proceeds as follows:

$$\begin{aligned}\langle\mathrm{GS}|\hat{\mathcal{O}}|\mathrm{GS}\rangle &= \langle 0|\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}|0\rangle = \mathrm{e}^{-\mathrm{i}L}\langle 0|\mathrm{e}^{\mathrm{i}L}\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}|0\rangle \\ &= \mathrm{e}^{-\mathrm{i}L}\langle 0|\hat{\mathcal{U}}_{+\infty,-\infty}\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}|0\rangle = \frac{\langle 0|\hat{\mathcal{U}}_{+\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}|0\rangle}{\langle 0|\hat{\mathcal{U}}_{+\infty,-\infty}|0\rangle}\,,\end{aligned} \tag{1.5}$$

[1] We work in the Schrödinger picture, where observables are t-independent operators, while the wavefunctions and the density matrix evolve.

where in the last equality we used $\hat{\mathcal{U}}_{+\infty,-\infty}\hat{\mathcal{U}}_{-\infty,t} = \hat{\mathcal{U}}_{+\infty,t}$, which is an immediate consequence of Eq. (1.2). The result of this procedure is that one needs to consider only the *forward* evolution. Indeed, the numerator in the last expression (being read from right to left) calls for evolving the non-interacting ground state $|0\rangle$ from the distant past to the observation time, where the observable operator acts, and then proceeding towards the distant future, where the overlap with the same known state $\langle 0|$ is evaluated.

The similar strategy works in the finite-temperature equilibrium formalism [11, 2, 4]. There, one treats the equilibrium density matrix $e^{-\beta\hat{H}}$, where $\beta = 1/T$ is the inverse temperature, as the evolution operator in the imaginary time τ. The latter is defined on a finite interval $0 < \tau < \beta$. The observables (or correlation functions) are also evaluated at imaginary time points $\tau_1, \tau_2, \ldots$ and the result must be analytically continued back to the real-time axis. One may argue that, since the adiabatic switching of interactions does not drive the system out of equilibrium, a statement similar to Eq. (1.4) still holds. As a result one is again left to describe only the forward evolution, albeit along the finite time interval in the imaginary direction.

Let us mention that elimination of the backward evolution comes with a price: the normalization denominator in the last expression in Eq. (1.5). It offsets the phase accumulation e^{iL} of the non-interacting ground state $|0\rangle$. In diagrammatic language it amounts to subtracting the so-called disconnected or vacuum loop diagrams. This denominator is a serious liability in the theory of disordered systems. The reason is that the accumulated phase e^{iL} sensitively depends on a specific realization of the disorder (which may be thought of as being absent at $t \to \pm\infty$ and adiabatically switched on and off in the process of evolution). Therefore the denominator absolutely must be included in any disorder averaging procedure, which complicates the treatment in a very substantial way.

The much more serious trouble with the outlined procedure is that Eq. (1.4) does *not* work in a non-equilibrium situation. If the Hamiltonian $\hat{H}(t)$ contains non-adiabatic time-dependent external fields, boundary conditions, etc., the evolution drives the system away from equilibrium. Even if all such fields are eventually switched off in the distant future, there is no guarantee that the system returns to its ground (or equilibrium) state. Therefore acting with the operator $\hat{\mathcal{U}}_{+\infty,-\infty}$ on the initial ground (or equilibrium) state results in an unknown superposition of excited states. As a result, the backward evolution, inherent to Eq. (1.3), can't be eliminated.

Nevertheless, it is still convenient to extend the evolution in Eq. (1.3) towards $t = +\infty$ and then back to t. This is achieved with the help of the trivial identity $\hat{\mathcal{U}}_{t,+\infty}\hat{\mathcal{U}}_{+\infty,t} = \hat{1}$. Inserting it into Eq. (1.3) and using $\hat{\mathcal{U}}_{-\infty,t}\hat{\mathcal{U}}_{t,+\infty} = \hat{\mathcal{U}}_{-\infty,+\infty}$, one finds

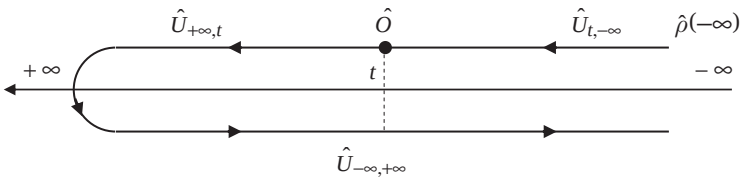

Fig. 1.1 Closed time contour $\mathcal{C}$. Evolution along such a contour is described by Eq. (1.6).

$$\langle\hat{\mathcal{O}}\rangle(t) = \frac{1}{\mathrm{Tr}\{\hat{\rho}(-\infty)\}}\,\mathrm{Tr}\{\hat{\mathcal{U}}_{-\infty,+\infty}\hat{\mathcal{U}}_{+\infty,t}\hat{\mathcal{O}}\hat{\mathcal{U}}_{t,-\infty}\hat{\rho}(-\infty)\}. \tag{1.6}$$

Here we also used that, according to the Von Neumann equation (1.1), the trace of the initial density matrix is unchanged under the unitary evolution. Equation (1.6) describes evolution along the *closed time contour* $\mathcal{C}$ depicted in Fig. 1.1. The observable $\hat{\mathcal{O}}$ is inserted at time t, somewhere along the forward branch of the contour. Notice that inserting the operator $\hat{\mathcal{U}}_{t,+\infty}\hat{\mathcal{U}}_{+\infty,t} = \hat{1}$ to the right of $\hat{\mathcal{O}}$ in Eq. (1.3), one could equally well arrange to have the observable on the backward branch of the contour. As we shall see, the most convenient choice is to take a half sum of these two equivalent representations.

Evolution along the closed time contour $\mathcal{C}$ is the central subject of this book. The fact that the field theory can be constructed with the time ordering along such a contour was first realized by Schwinger [12] and further developed in [24, 25]. About the same time Konstantinov and Perel' [13] have developed a diagrammatic technique, based on the time contour containing forward and backward branches in the real-time direction along with the imaginary time portion of length β. The formalism was significantly advanced, in particular its utility to derive the kinetic theory, in the seminal book of Kadanoff and Baym [1]. Independently Keldysh [14] (for some of the historic context see [26]) suggested a formulation which does not rely on imaginary time (and thus on the equilibrium density matrix). He also introduced a convenient choice of variables (Keldysh rotation), which made derivation of the kinetic theory particularly transparent. The time contour without the imaginary time piece, along with the Keldysh variables (which we call "classical" and "quantum") appear to be by far the most convenient choices for the functional formulation of the theory presented in this book. For this reason we occasionally refer to the construction as the Keldysh technique (this should by no way diminish the credit deserved by the other authors). Reformulation of the theory for the case of fermions, given later by Larkin and Ovchinnikov [27], became universally accepted. In fact, other theories developed about the same time, while not using the time contour explicitly, appear to be close relatives of

the Schwinger–Kadanoff–Baym–Keldysh construction. Among them are Feynman and Vernon [28], Wyld's [29] diagrammatic technique for fluid dynamics and Martin, Siggia and Rose [30] and DeDominicis' [31] calculus for classical stochastic systems.

The central object of the theory is the evolution operator along the closed contour $\hat{\mathcal{U}}_{\mathcal{C}} = \hat{\mathcal{U}}_{-\infty,+\infty}\hat{\mathcal{U}}_{+\infty,-\infty}$. If the Hamiltonian is the same on the forward and backward branches, then the forward–backward evolution of *any* state brings it back exactly to the original state. (Not even a phase factor is accumulated, indeed, any phase gained on the forward branch is exactly "unwound" on the backward branch.) As a result $\hat{\mathcal{U}}_{\mathcal{C}} = \hat{1}$ and the partition function, defined as $Z = \mathrm{Tr}\{\hat{\mathcal{U}}_{\mathcal{C}}\hat{\rho}(-\infty)\}/\mathrm{Tr}\{\hat{\rho}(-\infty)\}$, is identically equal to unity, $Z = 1$. Nevertheless, the partition function is a convenient object to develop the functional representation and the normalization identity $Z = 1$ is a useful check of its consistency. For this reason we shall use it widely in what follows.

To insert an observable somewhere along the forward (as prescribed by Eq. (1.6)) or backward branches it is convenient to modify the Hamiltonian $\hat{H}(t)$ by adding the source term $\hat{H}_V^{\pm}(t) = \hat{H}(t) \pm \hat{\mathcal{O}}V(t)$, where the plus (minus) sign refers to the forward (backward) part of the contour. Now, since the Hamiltonian is different on the two branches, the evolution operator along the contour $\hat{\mathcal{U}}_{\mathcal{C}}[V] \otimes \hat{1}$ becomes non-trivial and so does the *generating function*

$$Z[V] = \frac{\mathrm{Tr}\{\hat{\mathcal{U}}_{\mathcal{C}}[V]\,\hat{\rho}(-\infty)\}}{\mathrm{Tr}\{\hat{\rho}(-\infty)\}}. \tag{1.7}$$

The expectation value of the observable $\hat{\mathcal{O}}$, given by Eq. (1.6) (or rather by a half sum of the observable inserted along the forward and backward branches) may be found as $\langle\hat{\mathcal{O}}(t)\rangle = (\mathrm{i}/2)\,\delta Z[V]/\delta V(t)\big|_{V=0}$. This expression should be compared with the equilibrium technique [2, 4], where the observables are given by variational derivatives of the *logarithm* of the generating (or partition) function. In our case, since $Z = Z[0] = 1$, the presence of the logarithm is optional.[2] Knowledge of the generating function allows thus to find observables of interest. Therefore, after developing the functional formalism for the partition function, we extend it to include the generating function as well.

[2] It is worth mentioning that the denominators in Eqs. (1.5) and (1.7) have very different status. In the latter case $\mathrm{Tr}\{\hat{\rho}(-\infty)\}$ refers entirely to the distant past, when both interactions and disorder are switched off. It is therefore a simple constant, which may be easily evaluated. In the former case $\langle\theta|\hat{\mathcal{U}}_{+\infty,-\infty}|0\rangle$ involves evolution of the ground state upon switching on and off the interactions and disorder. It thus depends on both disorder and interactions and requires a separate calculation. The absence of a disorder-dependent denominator makes the closed time contour formalism especially suitable to deal with the averaging over the quenched disorder. The fact that observables do not require the *logarithm* is another manifestation of the absence of the non-trivial denominator.

1.2 The outline of this book

Chapter 2 is devoted to a possibly simplest many-body system of bosonic particles occupying a single quantum state. We briefly develop a second quantization representation and then proceed towards the functional formalism, based on the coherent-state functional integral along the contour $\mathcal{C}$. Here we pay close attention to a fundamental discrete time (see Eq. (1.2)) representation of the evolution operator. The model allows us to expose explicitly the discrete time structure, to verify normalization $Z = 1$ and to explain the meaning and hidden skeletons of the continuous notation. We then introduce Keldysh rotation, "classical/quantum" variables and explain the causality structure. Finally we introduce the generating function in discrete and continuous notations and explain the relation between them.

In Chapter 3 we exploit the analogy between the toy model of Chapter 2 and the harmonic oscillator to formulate the single-particle quantum mechanics as a path integral on the closed time contour. We then use it to investigate a quantum particle coupled to an equilibrium bath of harmonic oscillators. Integrating out the oscillators, degrees of freedom, we derive the real-time version of the celebrated Caldeira–Leggett model. We then use the ideas of time contour to discuss quantum mechanical tunneling in the presence of an external ac field (both with and without coupling to the bath).

In Chapter 4 we pick up the discussion of a particle coupled to a bath, introduced in Chapter 3. In particular we focus on its classical limit, where the corresponding real-time action acquires local (in time) form. We show ultimate relations of the emerging theory to the physics of classical stochastic systems. To this end we derive Langevin and Fokker–Planck equations from the time contour action and explain its connections to the Martin–Siggia–Rose–DeDominicis construction. We then focus on a few examples, which include escape from a metastable state, reaction models, fluctuation theorem and time-dependent perturbations acting on a stochastic system.

All the considerations so far were limited to systems with one (or a few) degrees of freedom, possibly coupled to an external bath. In Chapter 5 we start generalizing the formalism to true many-body systems. We introduce free bosonic fields (both complex and real) and their functional description on the closed contour. We then add interactions (collisions) between the particles and derive diagrammatic technique and the Dyson equation. Following Keldysh [14], we show that the latter contains the kinetic equation for a non-equilibrium distribution function. We demonstrate then how the collision integral emerges from a perturbative treatment of the self-energy.

Chapters 2–5 form the basis of the book. They serve to introduce most of the theoretical apparatus, notations and conventions used throughout the rest of the book. They are absolutely necessary for comprehension of what follows. We then branch into a number of applications, which are to a large extent logically independent of each other. Therefore the subsequent chapters may be read (or omitted) in an arbitrary order, without much damage to the understanding. Yet they share a number of common themes and methods which are developed in a certain sequence. The latter was meant to emphasize the connection between various fields and to reinforce common techniques at every successive encounter of them.

The first example, Chapter 6, is collisionless (i.e. where the particles interact only with a collective electromagnetic field) dynamics of a classical plasma. We derive the Vlasov equations, collective excitations – plasmons, and Landau damping. Though the system is purely classical, we treat it with the quantum formalism. Besides methodological illustration, it allows us to treat fluctuations of the collective electric potential and derive the so-called quasi-linear theory. The latter deals with the coupled kinetic equations for particles and plasmons (while Vlasov theory treats the electric field as fully deterministic, i.e. no fluctuations).

We then go to the essentially quantum system: low temperature weakly interacting Bose gas, in Chapter 7. We derive the Gross–Pitaevskii description of the condensate as a stationary point approximation of the corresponding functional integral. We then consider small fluctuations on top of the stationary field configuration and show that they bring the celebrated Bogoliubov quasiparticle spectrum. We then proceed to a description of a non-equilibrium quasiparticle cloud. To this end we derive a system of equations for the distribution function of quasiparticles coupled to the (modified) Gross–Pitaevskii equation for the condensate wavefunction. Following analogy with the collisionless plasma of Chapter 6, we show that this system contains Langevin forcing of the condensate fluctuations along with their collisionless damping. We then derive various contributions to the collision integral and use them to discuss kinetics of the condensate growth upon evaporative cooling.

Chapter 8 is devoted to the dynamics of phase transitions (mostly classical and only briefly quantum). Following Langer, we first discuss nucleation dynamics of critical droplets in first order transitions. We then switch to dynamics of continuous phase transitions, starting from equilibrium transitions and classification of their dynamic universality classes given by Hohenberg and Halperin. We then turn to essentially non-equilibrium phase transitions. The examples include absorbing state transitions in reaction-diffusion models and Kardar–Parisi–Zhang consideration of the roughening transition on growing interfaces.

The rest of the book, i.e. Chapters 9–14, is devoted to fermions. Chapter 9 presents a fermionic version of the one degree of freedom toy model. It uses a Grassmann functional integral on the closed time contour and is essentially parallel to its bosonic counterpart of Chapter 2. We then introduce fermionic fields and their interactions and derive a diagrammatic technique, the Dyson equation and finally the kinetic equation, again essentially parallel to their bosonic counterparts of Chapter 5. Finally we focus on the spin of the electron and discuss Stoner ferromagnetic transition and the spin part of the kinetic equation.

In Chapter 10 we use fermionic formalism to discuss non-equilibrium quantum transport. In particular we derive the Landauer formula for tunneling conductance, the Lesovik formula for shot-noise, the Levitov's result for the full counting statistics of transmitted charge, the Brouwer formula for adiabatic pumping of charge and, following Nazarov and Tobiska, the exact fluctuation relation and its consequences. We also deal with the spin transport, deriving the Slonczewski–Berger spin-torque term in the Landau–Lifshitz equation along with the spin-torque noise and associated Gilbert damping.

Chapters 11–14 deal with fermionic systems in the presence of a static (quenched) disorder potential. We start in Chapter 11 from the kinetic equation approach, which leads to the diffusive dynamics of density fluctuations and the concept of transport scattering time. The kinetic approximation misses quantum interference and mesoscopic fluctuation effects. To improve upon it we develop a systematic disorder averaging procedure, which takes advantage of the fundamental normalization $Z = 1$ of the closed time contour technique. It leads to the so-called matrix non-linear sigma-model, which we use to rederive the kinetic equation and diffusive density response and supplement it with quantum weak-localization corrections and the scaling theory of Anderson localization.

In Chapter 12 we focus on mesoscopic, i.e. sample-to-sample, fluctuations due to differences in disorder configurations in small metallic samples. We deal with density of states fluctuation and its particular limit, known as Wigner–Dyson statistics. We then proceed to describe universal conductance fluctuations and fluctuations of current–voltage characteristics. Finally we discuss full counting statistics of a disordered quasi-one-dimensional wire and the tunneling action.

We then include electron–electron interactions in disordered systems, which lead to non-trivial singular corrections to the density of states (the so-called zero bias anomaly) and conductivity (Altshuler–Aronov corrections). They also provide collision terms (and thus finite relaxation time) to the diffusive kinetic equation. These effects are the subject of Chapter 13.

Chapter 14 is devoted to the physics of disordered superconductors. We generalize the non-linear sigma-model to include superconducting correlations. Its stationary state condition yields the Usadel equation, which includes equations

for the spectrum of the superconductor under non-equilibrium conditions as well as the kinetic equation for the quasiparticles distribution function. Together with the self-consistency condition they provide a complete framework to study non-equilibrium superconductivity. As examples we work out the spectrum of the collective (Carlson–Goldman) mode of the superconductor, derive the time-dependent Ginzburg–Landau theory and fluctuation corrections to the conductivity above the critical temperature.

2

Bosons

The aim of this chapter is to develop a functional integral representation for the evolution operator along the closed time contour. To this end we use an example of a single quantized level populated by bosonic particles. Notations and structures introduced in this chapter are used throughout the rest of the book.

2.1 Bosonic coherent states

We start by considering a single quantum level occupied by bosonic particles. A *many-body* state with n bosons is denoted by $|n\rangle$. Such *pure number states* form a complete orthonormal basis, meaning $\langle n|n^*\rangle = \delta_{nn^*}$ and $\sum_n |n\rangle\langle n| = \hat{1}$. It is convenient to introduce bosonic annihilation and creation operators, $\hat{b}$ and $\hat{b}^\dagger$, which operate in the many-body Hilbert space of the system according to the following rules:

$$\hat{b}\,|n\rangle = \sqrt{n}\,|n-1\rangle\;; \qquad \hat{b}^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle\;. \tag{2.1}$$

By acting on an arbitrary basis state, one may check the following relations:

$$\hat{b}^\dagger\hat{b}|n\rangle = n|n\rangle\;; \qquad \hat{b}\hat{b}^\dagger|n\rangle = (n+1)|n\rangle\;; \qquad [\hat{b},\hat{b}^\dagger] = \hat{1}. \tag{2.2}$$

An extremely useful tool for our purposes is the algebra of bosonic coherent states, which we summarize briefly in this section. A coherent state, parametrized by a complex number α, is defined as a right eigenstate of the annihilation operator with the eigenvalue α:

$$\hat{b}\,|\alpha\rangle = \alpha\,|\alpha\rangle\;; \qquad \langle\alpha|\,\hat{b}^\dagger = \bar{\alpha}\langle\alpha|, \tag{2.3}$$

where the bar denotes complex conjugation. As a result, the matrix elements in the coherent state basis of any *normally ordered* operator $\hat{H}(\hat{b}^\dagger,\hat{b})$ (i.e. such that all the creation operators are to the left of all the annihilation operators) are given by

$$\langle\alpha|\hat{H}(\hat{b}^\dagger,\hat{b})|\alpha^*\rangle = H(\bar{\alpha},\alpha^*)\langle\alpha|\alpha^*\rangle. \tag{2.4}$$