

# Chapter 1

## Introduction:

# Why are quantum many-body simulations difficult and what can be done about it?

Here it is considered how stochastic methods can allow first-principles quantum simulations despite the seemingly intractably large Hilbert space required for such a task. The main approaches: path integral Monte Carlo, quantum trajectories, and phase-space distributions are briefly compared (without going into mathematical details of the methods themselves). This thesis develops and applies the stochastic gauge technique to make improvements to the phase-space distribution methods, and the choice to concentrate on these methods is motivated by their greater versatility. The emphasis throughout this thesis on methods applicable to open systems (leading to mostly mode-based formulations) is also discussed.

## 1.1 Hilbert space complexity

In classical mechanics, an exhaustive simulation of an interacting  $N$ -particle system is (in principle) computable<sup>1</sup> on a classical computer. For the trivial case of a non-interacting three-dimensional (3D) gas in a potential, one would start from initial conditions in  $6N$  real variables<sup>2</sup>, and then generically simulate by solving  $6N$  differential equations. For particles with ranged binary interactions<sup>3</sup> there are  $6N$  real differential equations, each with  $\mathcal{O}(N)$  terms, so simulation times scale as  $N^2$  — grudgingly computable.

Let us now look at the quantum case. Consider a generic mesoscopic system, divided up (arbitrarily) into  $N$  subsystems. Most typically these subsystems would be  $N$  particles,  $N$  spatial or momentum modes, or  $N$  orbitals. If we allow each mode or orbital to be occupied by up to  $P_{\max}$  particles ( $P_{\max} = 1$  for fermions, and more for bosons), then the full state vector of a pure system in such a model has  $(P_{\max} + 1)^N$  complex number components. In the other popular formulation with a set number of particles  $N$ , one would discretize  $\mathcal{D}$ -dimensional space into  $M_{\text{latt}}$  lattice points in each dimension, leading to  $M_{\text{latt}}^{\mathcal{D}}$  lattice points in all. The state vector of such a system would then have  $M_{\text{latt}}^{N\mathcal{D}}$  complex number components. Both of these cases are *exponential* in the system size  $N$ , and the amount of memory needed to hold the state (let alone the time needed to evolve the corresponding number of coupled differential equations) becomes intractable very rapidly. This is the Hilbert space complexity problem, and is generic to all mesoscopic quantum calculations.

For example, if one had 8000 Gigabytes of memory just to hold the state vector<sup>4</sup>, one could have 40 fermion modes, or 12 boson modes with each one occupied by at most 9 atoms. In the set particle number formulation, one could get away with 4 particles in a (very sparse!)  $10 \times 10 \times 10$  three-dimensional lattice, or 12 particles in such a one-dimensional lattice. This indicates where the estimate of about “five

---

<sup>1</sup>“Computable” here having the usual meaning that simulation time scales linearly, or (grudgingly) polynomially with  $N$ . Naturally if  $N$  is truly macroscopic such a system still cannot be simulated in practice, but in such large classical systems the primary hindrance to accurate predictions is usually deterministic chaos rather than simulation speed.

<sup>2</sup>Position and momentum in each degree of freedom.

<sup>3</sup>As in the majority of first-principles classical models.

<sup>4</sup>Assuming 4 byte floating point data.

particles” quoted in the Thesis Rationale on page 1 came from.

For open systems, one needs even more data space to hold the state, as the density matrix  $\hat{\rho}$  will have approximately half the square of the number of components in a state vector — exponential in system size again, but with twice the exponent. The general conclusion is that for mesoscopic (and even many microscopic) systems a brute force approach is doomed to failure, except for perhaps a few special cases that are separable<sup>5</sup>.

## 1.2 Bulk properties are the key

The exponential growth of Hilbert space mentioned above appears prohibitive, but let us first stand back from the quantum problem for a moment. In classical mechanics there are enormous amounts of variables to keep track of once a system has macroscopic numbers of particles — Avogadro’s number is very big! The reason that statistical mechanics is so successful at dealing with this issue is that for a large system, one is only really interested in *bulk* properties (temperature, pressure, two-particle correlations, . . .). The exact position or momentum of particle number 9501292851748 is not of interest to anyone. Even more importantly, knowledge of such a “raw” observable calculated from a model has no physical predictive power, because in such a large system it is not feasible to measure the initial conditions exhaustively. Only a statistical description (based on bulk properties) can be input as initial conditions into a model and because of this, only bulk statistical properties of the model will correspond accurately to the behavior of the physical system under study.

In quantum systems, we have (in a general sense) an analogous problem, but it is much more acute because it turns up when we have  $\mathcal{O}(5)$  subsystems, not  $\mathcal{O}(10^5)$ . This is simply because in quantum mechanics there is much more to keep track of. The full classical description contains the positions and momenta of all the particles (e.g. system variables record things like “the position and momentum of particle 753024”), whereas a full quantum description contains the amplitude and phase of

---

<sup>5</sup>e.g. an ideal gas

each one of all the possible configurations between the entire set of particles<sup>6</sup>. (e.g. a quantum mechanics variable would be expressing something like “the probability and phase (relative to a general reference) of the particular configuration of particles for which: number 1 is at  $x_1$ , ..., number 753024 is at  $x_{753024}$ , ..., and number  $N$  is at  $x_N$ ”). In classical mechanics the entire description is just one of these possible correlations.

Conceptually, what lets one vanquish the enormous numbers of variables is similar in both cases: We’re only interested in *bulk properties*, be they themselves classical or quantum in nature<sup>7</sup>.

### 1.3 Sampling system configurations

There are two kinds of conceptually distinct kinds of correlations to sample (rather than characterize exhaustively) to estimate the position of the quantum system in Hilbert space.

Firstly, in models that include interactions with an external environment, the system is described fully by a density matrix rather than a pure state vector. Here, the mixture of pure states that combine to the reduced density matrix description of the system must be sampled. While this kind of sampling reduces the number of variables substantially (We go from some huge number to its square root), it is not enough to make a qualitative difference. The pure state vectors still reside in a space whose dimension scales exponentially with system size  $N$ .

An example of an approach that carries out only this first kind of sampling is the quantum trajectories method for open systems. (Some reviews of this approach are e.g. the articles of Zoller and Gardiner[21], and Plenio and Knight[22], or Wiseman and Milburn[23] for continuous variable measurements.). This method can be successful in simulations of several particles, especially when microscopic details of their entanglement are important — for example, simulations with up to 24 qubits have been reported[24].

---

<sup>6</sup>Or, more generally, “subsystems”, particularly when particles are not conserved.

<sup>7</sup>An experimentally accessible example of a macroscopic property of a “quantum” nature is the condensate phase at a given position in a Bose-Einstein Condensate.

What is really needed to attack at least mesoscopic systems, is to change the scaling with number of subsystems from exponential down to at least polynomial. This is achieved by going from the pure system states (which are in general superpositions of various separable system configurations) to a *sample of the separable system configurations*. (examples of such separable system configurations are: the set of occupations of all modes, or the set of positions of all the particles.) The essence of the two sampling steps can be schematically expressed as

$$\begin{aligned} \hat{\rho} &\longrightarrow \sum_{j'=1}^{\mathcal{S}'} |\psi_{j'}\rangle \langle \tilde{\psi}_{j'}|, \\ |\psi_{j'}\rangle &\longrightarrow \sum_{j=1}^{\mathcal{S}} \otimes_{k=1}^N |C_k^{(j)}\rangle_k, \end{aligned} \tag{1.1}$$

where the sums are over  $\mathcal{S}$  and  $\mathcal{S}'$  samples, and the tensor products over  $N$  subsystems, having local subsystem parameters  $C_k$  ( $C$  stands for separable Configuration). The pure states  $|\psi_{j'}\rangle$  and  $|\tilde{\psi}_{j'}\rangle$  are not necessarily separable product states, but the samples  $\otimes_k |C_k^{(j)}\rangle_k$  are. In the limit of many samples ( $\mathcal{S}, \mathcal{S}' \rightarrow \infty$ ), the correspondence becomes exact.

This second sampling is really the crucial step to allow any kind of first-principles quantum simulation of realistic mesoscopic models<sup>8</sup>, as the number of variables to keep track of now finally scales proportionally to  $N$ .

The unavoidable price paid is loss of accuracy. Fortunately most of this price gets transferred onto the detailed non-bulk properties, which one is not interested in anyway. This occurs almost automatically, because there are just so many more of these detailed properties. In many cases bulk properties of quite large systems can be calculated with useful precision, where no simulation at all was possible with brute force methods. The calculation time required for a given uncertainty  $\Delta$  now scales (by the central limit theorem) roughly proportionally to  $N/\Delta^2$ , whereas for a brute force non-stochastic method it would be<sup>9</sup> proportional to  $e^N(\log \Delta)^2$ . Clearly, one won't be getting more than two- or three-digit accuracy with the stochastic methods in most cases, but this is often sufficient for a good comparison between theory and experiment.

---

<sup>8</sup>Obviously an enormous amount of physical prediction can be made with mean field, or other semiclassical methods, but by first-principles models I refer to those where no semiclassical approximations are made.

<sup>9</sup>The time required for multiplication operations on floating point numbers with  $d_{10} \approx -\log_{10} \Delta$  digits scales as  $d_{10}^2$

## 1.4 Path integral Monte Carlo and static thermal calculations

The most widely used approach for many-mode/body calculations is path integral Monte Carlo. The basis for these methods were Feynman's imaginary time path integrals[25]. This method, however, is only successful for static calculations of thermal equilibrium states (reasons explained below).

In the static thermal case, the un-normalized density matrix at temperature  $T$  is  $\hat{\rho}_u(T) = \exp\left[-\hat{H}/k_B T\right]$ , and path integrals make use of the simple (and exact) property that  $\hat{\rho}_u(T) = \hat{\rho}_u(\mathcal{M}T)^{\mathcal{M}}$ . If  $\mathcal{M}$  is sufficiently large that  $\mathcal{M}k_B T \gg \langle \hat{H} \rangle$ , then one can use a high-temperature approximation<sup>10</sup> to evaluate  $\hat{\rho}_u(\mathcal{M}T)$ . The density matrix  $\hat{\rho}_u(T)$  can be written as a convolution of the high temperature density matrix elements (in complete bases of separable system configurations  $C$  labeled by  $C_{(0)}, \dots, C_{(\mathcal{M})}$ ):

$$\hat{\rho}_u(T) = \int \dots \int dC_{(0)} dC_{(1)} \dots dC_{(\mathcal{M})} \quad (1.2)$$

$$\langle C_{(0)} | \hat{\rho}_u(\mathcal{M}T) | C_{(1)} \rangle \langle C_{(1)} | \hat{\rho}_u(\mathcal{M}T) | C_{(2)} \rangle \dots \langle C_{(\mathcal{M}-1)} | \hat{\rho}_u(\mathcal{M}T) | C_{(\mathcal{M})} \rangle.$$

To obtain properties diagonal in the configuration basis chosen, one takes samples while setting  $C_{(0)} = C_{(\mathcal{M})}$ , and ends up sampling a ‘‘ring polymer’’ of configurations, the number of variables in a full sample scaling as  $\mathcal{M}N$ . The distribution of the samples is given by the product of the expectation values:

$$P(C_{(0)}, C_{(1)}, \dots, C_{(\mathcal{M})}) = \otimes_{j=1}^{\mathcal{M}} \langle C_{(j-1)} | \hat{\rho}_u(\mathcal{M}T) | C_{(k)} \rangle, \quad (1.3)$$

which is real positive, or largely so. Off-diagonal properties require separate calculations with slightly different sample properties. An extensive discussion of the details of the path integral approach can be found e.g. in Ceperley[27, 28] and references therein.

While, formally, this approach might appear to be also applicable for dynamical evolution via  $\hat{\rho}(t) = \hat{\rho}(0) \exp\left[-i\hat{H}t/\hbar\right]$ , in practice this does not work. The reason

---

<sup>10</sup>Typically one uses the Feynman-Kacs or Trotter formula[26]  $\exp\left[-(\hat{H}_V + \hat{H}_K)/k_B T\right] = \lim_{\mathcal{M} \rightarrow \infty} \left(\exp\left[-\hat{H}_V/\mathcal{M}k_B T\right] \exp\left[-\hat{H}_K/\mathcal{M}k_B T\right]\right)^{\mathcal{M}}$  to evaluate the effect of the potential  $\hat{H}_V$  and kinetic  $\hat{H}_K$  terms of the Hamiltonian separately.

is that now  $P(C_{(j)})$  is primarily a product of complex phase factors, rather than real positive values. One obtains a very wide range of configurations, all with similar weighting, but mutually canceling phase factors. In the limit of infinite trajectories, almost all of these configurations would destructively interfere leaving just the observed dynamics, but for finite numbers of samples in a complex system there is no automatic way to tell which configurations are the important ones. The end result is that no sensible observable averages emerge from the noise even for very short times.

## 1.5 phase-space evolution methods

The density matrix of a system can often be written in a form

$$\hat{\rho} = \int P(C) \hat{\Lambda}(C) dC, \quad (1.4)$$

where  $C$  is a separable system configuration<sup>11</sup>, containing some number of variables linear in  $N$  (the number of subsystems).  $P(C)$  is a positive function, and  $\hat{\Lambda}(C)$ , the *kernel*, is a projector or similar tensor product operator parameterized by  $C$ .

In a generic case where pure quantum states of each subsystem (labeled  $k$ ) can be fully described by the set of parameters  $C_k$ , a viable kernel would be  $\hat{\Lambda}(C) = e^{i\theta} \otimes_k |C_k\rangle_k \langle \tilde{C}_k|_k$ . Here  $|C_k\rangle_k$  is the pure state of subsystem  $k$  given by parameters  $C_k$ , and  $\theta$  is a phase. The full configuration then would be  $C = \{\theta, C_1, \dots, C_N, \tilde{C}_1, \dots, \tilde{C}_N\}$ .

The basic idea is that  $P(C)$  can be interpreted as a probability distribution of the configuration  $C$ . Then, the evolution equation for  $\hat{\rho}$  (with time or with temperature) produces a corresponding evolution equation for  $P$ , and finally evolution equations for the configuration samples  $C$ . Depending on the system and the kernel  $\hat{\Lambda}$  chosen, this procedure can often be carried out, and evolution equations for the configuration samples of the system derived. In particular, *dynamical evolution can be simulated* while the number of variables remains linear or polynomial in  $N$ .

In fact phase-space distribution methods are much more versatile than path integral approaches on a number of counts:

---

<sup>11</sup>The variables of  $C$  can be discrete or continuous. In the former case the integral over  $C$  is replaced by a sum.

1. Dynamics can be simulated.
2. Results are given for a range of times (or temperatures) by one simulation
3. All observables can be calculated in one simulation

In both phase-space distribution and path integral approaches, the amount of computer effort scales linearly with system size  $N$ . More details of how this is achieved are given in Chapter 3. See also a tabular comparison of the various quantum many-body/mode simulation approaches in Tables 1.1 and 1.2.

Unfortunately, a major stumbling block for phase-space distribution methods is that the equations for the configuration variables in  $C$  often have stability problems, which can lead to large noise and/or statistical bias. This is one of the main reasons that these methods have not been as widely investigated nor used as path integral Monte Carlo, despite their other advantages.

This thesis concerns itself with methods of overcoming this stability problem using the stochastic gauge technique, and some demonstration of subsequent applications. It is also intended to show that when stability issues are taken care of, these phase-space distribution methods can be competitive with path integrals for static thermal ensemble calculations as well.

## 1.6 Open systems and the mode formulation

Actual physical systems do not exist in isolation, and there is always some degree of coupling to the external environment. This coupling becomes ever more pronounced as one moves from microscopic to mesoscopic systems, and is considered responsible for the reduction of quantum phenomena to classical behavior as the macroscopic limit is reached.

Traditionally in most many-body studies the system has been taken to be composed of a set number of  $N$  particles. A different formalism is required if one is to investigate open systems that can exchange particles with a reservoir having chemical potential  $\mu$ , or for dynamical calculations if particles are lost or injected into the system. A natural choice is to consider lattice models where rather than dividing

Table 1.1: A comparison of different methods to exactly simulate quantum dynamics from first principles. As applicable to many-body/mode simulations with  $N$  subsystems (e.g. particles, modes).

Variety:	EXPLICIT DENSITY MATRIX	QUANTUM TRAJECTORIES	PATH INTEGRAL	PHASE-SPACE DISTRIBUTIONS
Stochastic	No	Yes	Yes	Yes
Sample size	$\propto e^{2N}$	$\propto e^N$	$\propto \mathcal{M}N$	$\propto N$
Phase factor catastrophe <sup>a</sup>	No	No	Yes	No <sup>b</sup>
Unstable equations or sampling bias	No	No	No	Often <sup>c</sup>
Observables from one calculation	All	All	Limited subset	All
Simulation scope	Range of $t$	Range of $t$	Single $t$	Range of $t$

<sup>a</sup>Where samples of similar weight but different phase factors interfere, masking any physically meaningful results.

<sup>b</sup>At long times some gauged schemes may develop interfering phase factors. This occurs only for some gauges in some systems.

<sup>c</sup>However, in a wide range of cases the instability can be removed using the stochastic gauge technique, as shown in this thesis.

---

the system into  $N$  subsystems consisting of a particle, we divide the system up into  $N$  spatial (or momentum) lattice points that can be occupied by a variable number of particles. The kernel  $\hat{\Lambda}$  then becomes a separable tensor product of lattice-point states, described by some local parameters. Their total number for a system configuration sample is again proportional to  $N$ . Because we want the methods developed in this thesis to be applicable to systems that allow particle exchange with an environment, this is the kind of model that we will mostly concentrate on here. The general methods developed are, however, by no means restricted to mode-based models.

A feature to be aware of in such a model is that in general the state will be in a superposition or mixture of pure states with different numbers of total particles. This is not expected to hinder comparison with experiment, however. A subsystem

Table 1.2: A comparison of different methods to exactly calculate static quantum properties of equilibrium thermodynamic ensembles from first principles. As applicable to many-body/mode simulations with  $N$  subsystems (e.g. particles, modes).

Variety:	EXPLICIT DENSITY MATRIX	QUANTUM TRAJECTORIES	PATH INTEGRAL	PHASE-SPACE DISTRIBUTIONS
Stochastic	No	Yes	Yes	Yes
Sample size	$\propto e^{2N}$	$\propto e^N$	$\propto MN$	$\propto N$
Phase factor catastrophe <sup>b</sup>	No	No	For fermions <sup>a</sup> $k_B T \lesssim \Delta E_{FB}$	Can occur for low $T$
Unstable equations or sampling bias	No	No	No	Often <sup>c</sup>
Observables from one calculation	All	All	Limited subset	All
Simulation scope	Range of $T$	Range of $T$	Single $T$	Range of $T$

<sup>a</sup>For fermion calculations the “fermion sign problem” appears after  $k_B T \lesssim \Delta E_{FB}$ , where  $\Delta E_{FB}$  is the free energy difference between the fermion system, and an analogous system of bosons. For a review see Schmidt and Kalos[29]

<sup>b</sup>Where samples of similar weight but different phase factors interfere, masking any physically meaningful results.

<sup>c</sup>However, in a wide range of cases the instability can be removed using the stochastic gauge technique, as shown in this thesis.

such as e.g. a trapped gas is perfectly entitled to be in a superposition of number states, and e.g. evaporative cooling of Bose atoms will produce BECs of different sizes in each individual experimental run.

For open systems with particle transfer to the environment, the boundary where the system ends and the environment starts will be at best hazy, and particles will take some time to travel from one to the other. There will always be a gray area, and in it some particles, where the labeling of particles as belonging to the system or environment is dubious. Superpositions and/or mixtures of different particle number states are actually trying to reflect this labeling difficulty which is encountered in experiments.

Also, in many situations the effect of such superpositions will be small: A rea-

sonably accurate description of a realistic mesoscopic system in most cases requires a fairly large number of lattice points to resolve spatial and momentum variations in the system. In such a case, for statistical reasons, the variation in particle number of the whole system will be small in comparison with the total number of particles.

Lastly, it should also be pointed out that if desired, one can always simply turn off the external couplings, and obtain results for closed number-conserving Hamiltonian models, which are just special cases with external interaction strength zero. (Although a simulation method that is “hard-wired” to a conservative system is likely to be somewhat more efficient than an open systems method.)