



Universiteit
Leiden

The Netherlands

The power of one qubit in quantum simulation algorithms

Polla, S.

Citation

Polla, S. (2024, February 22). *The power of one qubit in quantum simulation algorithms*. *Casimir PhD Series*. Retrieved from <https://hdl.handle.net/1887/3719849>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3719849>

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 1

Introduction

1.1 Preface

Quantum theory, the mathematical framework that governs the behavior of elementary physical systems, stands as a cornerstone of modern science. While its first principles allow modeling complex systems like molecules and materials, managing the complexity and predicting emergent behaviors demands additional tools [1]. Analytical approximations played a key role in early successes of nuclear, atomic, and solid-state physics; together with computational simulation techniques they now drive the fields of quantum chemistry and material science. Since the 1980s *quantum computers* have been proposed as an additional tool, promising access to simulations impossible for classical methods [2]. Recent years saw the realization of the first programmable *quantum devices* that can outperform classical computers in benchmark tasks [3–7], but lack practical applications yet. The next quantum computing milestone is demonstrating the ability to solve genuinely useful problems beyond the capabilities of other tools. Achieving this requires hardware enhancements, but also research on pertinent target problems, and the development of quantum algorithms that exploit the hardware to solve the problem. The study of complex quantum systems offer a variety of challenging problems, supported by heuristic insight from physics: a valuable resource in the effective development of

quantum algorithms.

This thesis collects the proposals of quantum algorithms developed in this context, tackling some known hurdles in simulating complex systems with near-term quantum devices. The studied techniques focus on questions like the preparation of ground states of natural systems, the extraction of information about observables from quantum states, and the resilience of quantum simulations to noise. A red line connecting the proposed algorithms is the focus on the support which a single qubit – the smallest unit of quantum information – can provide in implementing a variety of tasks in quantum simulation.

The remainder of this chapter is structured as follows: Section 1.2 provides an overview of key concepts in quantum information processing, discussing the advantages and limitations of quantum devices. Section 1.3 introduces quantum simulation by examining the computational problems naturally defined by a Hamiltonian description of a quantum system. Section 1.4 offers a high-level description of some significant quantum simulation algorithms relevant to this thesis. Section 1.5 delves deeper into a specific quantum simulation target, namely, the study of molecular systems. Finally, in Section 1.6, we present an overview of the chapters that constitute the main body of this work.

1.2 Processing quantum information

A *state* is a specific configuration that a physical system can assume at some point in time. In classical physics, states are described by variables taking defined values, such as the position \mathbf{x} and velocity $\dot{\mathbf{x}}$ of a cannonball in Newtonian mechanics, or the ON/OFF state of a transistor in digital electronics. Conversely, the state of a quantum system can consist of a *superposition* of multiple, distinguishable states. In Dirac notation [8]

$$|\psi\rangle = \alpha_A |A\rangle + \alpha_B |B\rangle + \dots, \quad |\alpha_A|^2 + |\alpha_B|^2 + \dots = 1, \quad (1.1)$$

where $\{A, B, \dots\}$ are a set of states that can be distinguished deterministically by a sufficiently precise measurement, each representing a different set of values of the classical variables that characterize the system. The state of a quantum switch will then have the form $\alpha_{\text{ON}} |\text{ON}\rangle + \alpha_{\text{OFF}} |\text{OFF}\rangle$, while the description of a quantum particle will need one value of $\alpha_{\mathbf{x}}$ for each possible position \mathbf{x} .

Consider a system which can assume N distinguishable states. A deterministic description of a classical state of such a system just needs to

identify one out of the N possible states; this information can be encoded in $\log_2(N)$ bits of memory. A full description of a quantum state for the same system will require $N - 1$ complex amplitudes, each stored to precision ϵ_C in $m = 2 \log_2(\epsilon_C^{-1})$ bits of classical memory. For a composite system, made up of n elements, the number of total distinguishable states scales exponentially, i.e. $N = \mathcal{O}(2^n)$. A classical state of such system can be described in $\mathcal{O}(n)$ bits, but a full representation of the quantum amplitudes requires $\mathcal{O}(m2^n)$ bits of classical memory. This exponential scaling quickly makes unfeasible to fully represent quantum states of complex systems.

In contrast to a classical computer, a quantum device can natively store and process quantum states of the form Eq. (1.1). Typical quantum devices are constructed by assembling a number of *two-level systems* – elements with two distinguishable states labeled $|0\rangle$ and $|1\rangle$. Combining n two-level systems allows to store a quantum superposition of 2^n distinguishable states, or n qubits (a quantum analogue for bits) of quantum information. Each two-level system stores a qubit of quantum information; each additional qubit doubles the dimension of the Hilbert space of states available to the quantum device.

1.2.1 Classical input-output

While quantum theory introduces one significant extension to the concept of information by allowing for superposition storage, this comes with a notable limitation. Contrary to a classical state, a single copy of an arbitrary quantum state cannot be fully characterized by observation. The conversion of quantum information to classical information implies a loss, typically associated with the randomness of a measurement's outcome. Holevo's theorem [9] quantifies this, stating the amount of classical information extracted from a n -qubit quantum state is bounded by n bits. As a corollary, it is impossible to copy an arbitrary state of a quantum device (*no-cloning*), unlike for classical information which can always be observed and transcribed.

A quantum algorithm needs to ensure that, at the end of the computation, the quantum state of the device encodes the relevant information in a way that makes its extraction easy. The whole n qubit Hilbert space can be used to perform computation, but the output should be encoded in $\mathcal{O}(n)$ bits of classical information. The final measurement performed on the state needs to be designed such that it can extract the necessary information from the state as efficiently as possible. In Chapter 4, we explore this optimization of measurements for the expectation value of observables,

under a further restriction on the amount of classical information that can be extracted from the device.

Similarly, the input to any quantum computation – the list of instructions – should be given in terms of classical information. A common framework to describe this list of instructions is the *quantum circuit model*. There, operations on the quantum memory are described as *gates* involving a small number of qubits at a time. Gates are typically unitary operations, although in some cases it is convenient to consider quantum channels to model noisy gates or non-unitary gates such as *reset* (a key element in the algorithm presented in Chapter 2). The computation is then realized by applying a sequence of gates (a circuit) to a fiducial initial state.

Another model of computation relevant to this thesis is that employed by *analog quantum simulators*. These quantum devices can natively implement the time evolution generated by a Hamiltonian $H(\theta)$, function of parameters θ that can eventually be changed during the computation. The Hamiltonian acts on the whole system at once, but it can be described in terms of classical information as explained in Section 1.3. In contrast to the circuit model, here the computation is described by stating the Hamiltonian throughout the computation time. Quantum simulators allowing for general enough Hamiltonian are universal [10]. In fact, gate-based quantum computers are practically realized as specialized simulators, where each gate is generated by a pre-calibrated time-dependent Hamiltonian. The term *analog quantum simulators* tends to be reserved for devices that can implement a limited set of Hamiltonians, modeling analogously the evolution of some quantum system of interest.

1.2.2 Noise, error correction and mitigation

All physical processes, including computation, suffer from a certain amount of noise due to unpredictable perturbations. Classical computers can deal with noise by storing information redundantly, both at the hardware level (e.g. using macroscopic bistable systems) and by backing up data. Redundancy-based error correction cannot be trivially applied in quantum computation, as quantum information cannot be copied (no-cloning) and compared throughout the computation. This leads to an accumulation of noise over the computation, leading to degradation of quantum information and, ultimately, a success probability that decays exponentially with the space-time volume of the computation. [11, 12]

Quantum error correction codes aim to protect quantum information by encoding it in a subspace of *logical states*, living in a much larger space of *physical states* of the device. The logical subspace is engineered to ensure

that physical noise processes scattering one logical state to the other are exceedingly rare. For example, in *topological codes* [13], information is encoded non-locally in terms of correlations between a large number of two-level systems (a.k.a. *physical qubits*). A perturbation can only convert one logical state into the other by acting on many two-level systems in a globally-correlated fashion. The probability of such process occurring due to natural noise is exponentially small in the number of physical qubits. Non-logical states can be led back to logical states by following the geodesic defined by error probability, effectively correcting the most-probable errors.

Quantum computation can be made completely *fault-tolerant* by ensuring information is stored in a large enough logical subspace throughout the computation (including initialization, gate operations and measurements). The implementation of a large scale fault tolerant quantum computer is the only known way towards many key applications of quantum computers, including breaking cryptosystems using Shor's algorithm [14] and performing challenging chemistry simulations of industrial importance [15]. However, such a computer requires an underlying physical device of very large scale (millions of qubits) and with sufficiently small physical error rates, far out of reach of today's technologies. Large amounts of resources are being invested towards the goal of fault-tolerant quantum computing, and the last two years saw the first small scale proof-of-concept demonstration of successful error correction [16] and fault tolerant operation of a quantum device [17]. While defining the timeline to full-scale fault tolerant quantum computing is not yet possible, even the most optimistic estimates accepted by the community measure in the order of decades.

Another approach, more adapted to today's noisy intermediate-scale quantum (NISQ, [18]) devices, consists in keeping computations short to prevent accumulated noise to reach a disruptive level. Short computations are obviously limited in capabilities, but have been demonstrated to outperform classical devices on benchmark tasks [3–7]. The NISQ paradigm can be summarized as *prepare, sample, repeat*, and it focuses on performing many short-time quantum computations (called *circuit runs* or *shots*), used as subroutines by a classical algorithm. This results in a hybrid quantum-classical computation, where the usefulness of the quantum subroutine is typically heuristic. An archetypal category of NISQ algorithms are variational quantum algorithms, such as the variational quantum eigensolver [19] described in Section 1.4.4 and considered in many chapters of this thesis.

1.3 Targets of quantum simulation

The term *quantum simulation*, in its most general sense, refers to the use of a quantum device to study a physical phenomenon through a representative model. Phenomena happening within an isolated quantum system can be modeled through the Hamiltonian picture. The Hamiltonian H is a Hermitian linear operator on the Hilbert space of states, which defines its *dynamics* (i.e. how its state changes in time) through the *time-dependent Schrödinger equation*

$$-i \frac{\partial}{\partial t} |\psi(t)\rangle = H |\psi(t)\rangle. \quad (1.2)$$

If the system is influenced by a classical environment, the Hamiltonian could depend on external parameters, possibly changing in time. In the case of more complex interactions with the environment, describing the dynamics of the system will require a more complex equation (such as the Liouville equation, or a Schrödinger equation for a larger model including the environment). Nevertheless, weak interactions with the environment can often be studied through linear response theory, with Hamiltonian eigenstates and dynamics playing a key role.

A system's Hamiltonian naturally defines a set of natural targets for quantum simulation: the synthesis of its dynamics (Section 1.3.1), the preparation of thermal states and eigenstates (Section 1.3.2), and the measurement of the system's energy (Section 1.3.3). In this section, we formalize these primitive target as problems with classical input-output relations (where one of the inputs is always the Hamiltonian). We define the size of each problem as the number qubits n defining the model's Hilbert space. We require all inputs and outputs to have an efficient classical description, i.e. be representable in $\mathcal{O}(\text{poly}(n))$ bits. When we require *quantum inputs* such as a state $|\psi\rangle$ or an observable O , we assume that they are given in terms of an efficient procedure that prepares $|\psi\rangle$ or samples O . All problems are formulated such that output is an estimate of a defined quantity, accurate to precision $\epsilon > 0$ with high probability. These problems are summarized in table 1.1.

A generic Hamiltonian cannot be described efficiently, as the definition of a Hermitian operator on an n -qubit Hilbert space \mathcal{H} requires 2^{2n-1} complex parameters. Natural Hamiltonians aren't however fully generic, as they inherit locality from the underlying physics. In local systems, the Hamiltonian can be described as a sum of local interaction terms h_j , each of which only acts non-trivially on a small Hilbert subspace \mathcal{H}_j and as the

Problem	Input	Target	Output
Dynamics	$ \psi\rangle, O, t$	$U(t)$	$\tilde{O} : \tilde{O} - \langle \psi(t) O \psi(t) \rangle < \epsilon$
Gibbs	O, T	$\rho_T \propto e^{-\frac{H}{T}}$	$\tilde{O} : \tilde{O} - \text{Tr}\{O\rho_T\} < \epsilon$
Eigenstate	O, j	$ E_j\rangle$	$\tilde{O} : \tilde{O} - \langle E_j O E_j \rangle < \epsilon$
Expectation	$ \psi\rangle$	Expval $\langle H \rangle$	$\tilde{E} : \tilde{E} - \langle \psi H \psi \rangle < \epsilon$
Eigenvalue sampling	$ \psi\rangle$	Projective meas. of H	$\tilde{E} : \tilde{E} - E_j < \epsilon,$ with prob. $p_j = \langle E_j \psi \rangle ^2$
Eigenvalue estimation	j	E_j	$\tilde{E} : \tilde{E} - E_j < \epsilon$

Table 1.1: Summary of some primitive problems in quantum simulation. The input of an efficient representation of H and of the required accuracy ϵ is implicit for every problem.

1 on the complementary subspace $\mathcal{H} \setminus \mathcal{H}_j$ ¹:

$$H = \sum_j^{J \sim \text{poly}(n)} h_j; \quad h_j := [h_j]_{\mathcal{H}_j} \otimes \mathbb{1}_{\mathcal{H} \setminus \mathcal{H}_j}. \quad (1.3)$$

This description is efficient as long as the number of terms in the sum scales polynomially $J \sim \mathcal{O}(\text{poly}(n))$ and the dimension of each local Hilbert space is constant $\dim[\mathcal{H}_j] \sim \mathcal{O}(1)$.

The two most common models of locality are illustrated in Fig. 1.1. A *geometrically local* system is composed of subsystem arranged in space, for example in a lattice. Each element can only couple to its geometric neighbors. As the number of neighbors of each subsystem is bounded by a constant, the number of couplings we need to describe the Hamiltonian can only scale linearly in the size of the system. Sometimes it is convenient to think of subsystems as particles, which can be geometrically delocalized and can interact with all other particles; in these cases, interactions typically involve no more than k -particles at a time. This defines a k -local model, the number of possible interactions scales as n^k (thus still polynomial). Note that geometric locality is more stringent than k -locality. In both

¹This description is only rigorous for local commuting systems (qudits). The action of a local fermionic Hamiltonian on the Hilbert space is more complex, as fermionic operators acting on different locations need anticommute. However, the number of terms that can appear in a fermionic local Hamiltonian remains the same.




Figure 1.1: Visualization of the Hamiltonian Eq. (1.3) in a 2-dimensional lattice model (geometric locality) and a two-body interacting particle model (k -locality).

cases, the Hilbert space on which each interaction acts \mathcal{H}_j is composed of up to a constant number k of subsystems, thus each h_j can be defined with $\mathcal{O}(2^k)$ parameters (constant in n). Other equivalently efficient descriptions of sparse Hamiltonians exist; we focus on local models as they mostly cover the natural systems of interest.

1.3.1 Hamiltonian dynamics

Reproducing the effect of time evolution of a state is the essence of simulation, in the stricter sense of the term. We refer to this problem to as *Hamiltonian simulation* or *dynamical simulation*. The inputs of the problem, along with H and ϵ , are the initial state $|\psi(0)\rangle := |\psi\rangle$, the observable O , and the evolution time t . The goal is producing an estimate \tilde{O} of the expectation value of O on the state $|\psi(t)\rangle$ solution of the time-dependent Schrödinger equation (1.2), up to the required accuracy $|\tilde{O} - \langle \psi(t) | O | \psi(t) \rangle| < \epsilon$.

Solving Hamiltonian simulation boils down to synthesizing the propagator $U(t)$, the unitary operator that sends $|\psi\rangle$ to $|\psi(t)\rangle$. More generally, the propagator solves the operator-Schrödinger equation, $-i\partial_t U(t) = H(t)U(t)$. A quantum algorithm for Hamiltonian simulation needs to define a procedure to construct an approximation of $U(t)$, starting from the description of the Hamiltonian. In the case of a time-independent Hamiltonian the target operator is $U(t) = e^{-iHt}$, which is not a sparse operator (unlike its generator H).

An example use case of Hamiltonian dynamics is studying the evolution of a known equilibrium state after an external field quenches (i.e. suddenly

changes) the system's Hamiltonian. The synthesis of the evolution operator $U(t)$ is also an key subroutine in many other simulation algorithms, such as simulated thermalization (a version of which is studied in Chapter 2), eigenvalue sampling (Section 1.4.5), quantum-assisted Hamiltonian learning and many others. Furthermore, any efficient quantum computation can be encoded in a Hamiltonian simulation problem [2, 20] – in complexity theory terms, Hamiltonian simulation is BQP-complete.

The first proposal for a procedure to approximately synthesize $U(t)$ on a universal quantum computer was put forward by Lloyd in 1996 [21], and it is based on Suzuki-Trotter product formulas [22, 23]. Through this procedure Lloyd showed that a universal quantum computer can solve the Hamiltonian dynamics problem for any k -local Hamiltonian efficiently. The last years saw many proposals of improved algorithms for Hamiltonian simulation, most of which lead to a significant reduction of the asymptotic costs. Finally, analog quantum simulators can implement natively the dynamics generated by some restricted class of Hamiltonian. We expand on these in Section 1.4.2.

1.3.2 State preparation

Thermal state preparation

A physical system weakly interacting with an environment tends to thermalize, that is to reach an equilibrium state ensemble with a fixed temperature. Analogously to the classical case, a quantum thermal ensemble distribution mirrors the Gibbs measure, with the probability of a state decaying exponentially with the state's energy divided by the temperature. As the energy is described by the Hamiltonian H , the density matrix for the *Gibbs state* at temperature T is defined by

$$\rho_T := \frac{1}{Z(T)} e^{-\frac{H}{T}}, \quad (1.4)$$

where $Z(T) := \text{Tr}[e^{-H/T}]$ is the partition function which normalizes the density matrix. As most systems in Nature are found in a thermal equilibrium state, studying the properties of thermal states has intrinsic interest. Characterizing thermal state properties as a function of temperature and other external parameters can lead to the identification of transitions between phases of matter, whose study is of special importance to material science.

The inputs to the problem of thermal state preparation, along with H and ϵ , are the observable O and the temperature T . The goal is producing

an estimate \tilde{O} of the the expectation value of O on the Gibbs state, up to the required accuracy accuracy $|\tilde{O} - \text{Tr}[\rho_T O]| < \epsilon$. To solve this problem, one can attempt to prepare an approximation of the Gibbs ensemble ρ_T on the quantum device.

Thermal state preparation algorithms include methods that simulate thermalization, by coupling the model system with a model bath. The algorithm we introduce in Chapter 2 fits in this class, although we specifically explore its application to ground state preparation. Other algorithms for thermal state simulation include a quantum version of metropolis sampling [24], variational approaches [25, 26], and filtering-based techniques [27].

Eigenstate preparation

The eigenstates of the Hamiltonian represent the equilibrium pure states of an isolated system, left invariant by time evolution. Each eigenstate $|E_j\rangle$ has a well-defined energy E_j , matching an eigenvalue of the Hamiltonian and satisfying the time-independent Schrödinger equation

$$H |E_j\rangle = E_j |E_j\rangle. \quad (1.5)$$

We assume an ordering of the eigenvalues $E_j < E_{j-1}$ without loss of generality.

The lowest-energy eigenstate denoted $|E_0\rangle$ or the *ground state*, holds particular significance, as it represent the zero-temperature limit of the thermal state Eq. (1.4). When the lowest excitation energy, also known as the "ground-state gap" of a system ($\Delta_{\text{GS}} = E_1 - E_0$), significantly exceeds the temperature of the environment T , thermal states closely resembling the ground state naturally emerge. For instance, this is the case for electrons within a molecule or semiconductor, where typical excitation energies are on the order of 1eV (equivalent to 10^4K , two orders of magnitude higher than room temperature). Simulating the electronic ground state in these systems allows for the prediction of molecular and material properties in their natural state. Moreover, low-lying excited eigenstates are relevant to the study of spectroscopic properties. Further insights into the quantum simulation of molecular ground states are provided in Section 1.5.1.

The inputs to the problem of eigenstate preparation, along with H and ϵ , are the observable O and the eigenvalue index j . The goal is to produce an estimate \tilde{O} of the the expectation value of O on the j -th eigenstate, up to the required accuracy accuracy $|\tilde{O} - \langle E_j | O | E_j \rangle| < \epsilon$. Solving this problem requires the quantum simulation to aim at preparing an approximation of the eigenstate $|E_j\rangle$.

Many algorithms for eigenstate preparation have been proposed, each employing very diverse techniques and heuristics, due to the intrinsic difficulty of the problem. The variational quantum eigensolver (Sec. 1.4.4) and adiabatic state preparation (Sec. 1.4.3) are important heuristic techniques for eigenstate preparation, discussed in more detail in the next section. Projective variants of quantum phase estimation (Sec. 1.4.5) can also be used for heralded preparation of eigenstates. Other ground state preparation techniques include among others simulated thermalization (a variant of which is studied in Chapter 2), and various flavors of quantum imaginary time evolution [28, 29].

1.3.3 Energy measurements

The Hamiltonian H serves as an observable representing the system's energy. In the preceding sections, we assumed the presence of an observable O as input, along with an efficient method for its sampling. This section addresses the challenge of implementing procedures to measure the Hamiltonian observable, starting from its efficient classical description. There are two distinct tasks, both referred to as energy measurement: *expectation value estimation* and *eigenvalue sampling*.

Expectation value estimation

Expectation value estimation involves an input state $|\psi\rangle$, along with the description of H and the target precision ϵ . The objective is to provide an estimate \tilde{E} of the expected energy of $|\psi\rangle$ to the required precision $|\tilde{E} - \langle\psi|H|\psi\rangle| < \epsilon$. Expectation value estimation is a key subroutine in the variational quantum eigensolver (discussed in Section 1.4.4). Additionally, the same technique can be applied to estimate the expectation value of observables other than the system Hamiltonian.

The most basic algorithm to estimate $\langle\psi|H|\psi\rangle$ for H defined in Eq. (1.3) relies on the linearity of the expectation value:

$$\langle\psi|H|\psi\rangle = \sum_j^J \langle\psi|h_j|\psi\rangle. \quad (1.6)$$

Given that h_j acts non-trivially on a subsystem of constant size, we can implement a quantum operation to sample it in constant time. In Chapter 4 we perform a thorough study of expectation value estimation under a restricted model of quantum computation, with applications in error-mitigated NISQ computing and metrology.

Eigenvalue sampling

The task of implementing projective measurement of H (up to target accuracy ϵ) on an input state $|\psi\rangle$ is more complex. The aim of eigenvalue sampling algorithm is to produce an estimate \tilde{E} of an eigenvalue E_j , sampled from the probability distribution described by the Born rule $p_j = |\langle E_j|\psi\rangle|^2$, up to the required accuracy $|\tilde{E} - E_j| < \epsilon$. Note that this task entails two objectives: sampling from the probability distribution p_j and estimating the eigenvalue E_j , which is not known in advance based on the description (1.3) of the Hamiltonian. The entanglement-based quantum phase estimation algorithm, utilizing Hamiltonian dynamics simulation as a subroutine, can be used to implement this measurement efficiently (see Section 1.4.5).

1.3.4 The ground state energy problem

The eigenvalue estimation problem revolves around the estimation of a specific eigenvalue E_j of H , with j being the sole input to the problem in addition to H and ϵ . Frequently, this problem specifically addresses the ground state, where it is known as the *ground state energy problem* or sometimes simply the *ground state problem*. Solving the ground state energy problem involves a combination of ground state preparation and energy measurement algorithms. A good approximation of the ground state energy can be achieved either by measuring $\langle H \rangle$ on a well-approximated $|E_0\rangle$, or by repeatedly sampling eigenvalues of H from an input state that demonstrates a substantial fidelity $|\langle E_0|\psi\rangle|^2$ with the ground state.

The ground state energy itself plays a pivotal role in the simulation of chemistry and material science, as exemplified in Section 1.5.1. Furthermore, any quantum computation that can be efficiently verified can be mapped into a ground state problem using Kitaev's history-state Hamiltonian construction [13, 30]; in complexity theory terms, the ground state problem is QMA-complete. The history state construction has been iteratively improved upon in literature, allowing for the encoding of any QMA problem into Hamiltonians that describe systems as straightforward as one-dimensional chains of nearest-neighbor-interacting subsystems, each with Hilbert subspace of constant size $N = 12$ [31]. Given that the quantum phase estimation algorithm implements eigenvalue sampling efficiently, the task of preparing a state $|\psi\rangle$ with a good fidelity to the ground state [i.e. $|\langle \psi|E_0\rangle|^2 \gtrsim 1/\text{poly}(n)$] must also be QMA-hard. Nevertheless, the statement on the complexity of ground state preparation is only true for worst-case. History state Hamiltonians have very specific forms, signifi-

cantly different from most natural system's Hamiltonians; for the latter, heuristic ground state preparation are more likely to succeed.

1.4 Algorithms for quantum simulation

In this section, we introduce some of several prominent algorithms designed to address the quantum simulation problems presented in the previous section. Since a primary objective of quantum algorithm research is to identify problems and methodologies that enable quantum devices to exhibit an advantage over classical numerical approaches, we will start by providing a concise conceptual introduction to prominent classical techniques used to solve quantum simulation problems. Subsequently, we will provide an overview of three fundamental quantum algorithms (or, more precisely, classes thereof), each of particular significance to the context of this thesis.

1.4.1 Classical algorithms

All the problems discussed in the preceding section can be framed in terms of classical input-output relationships, making them amenable to both classical and quantum algorithms. However, all the problems we described implicitly involve the storage and manipulation of states of the simulated quantum system. While quantum algorithms inherently handle this task, classical simulation algorithms must work with highly compressed state representations to maintain memory and time efficiency. Any compression of the quantum state introduces approximations, and tends to be accurate and efficient only in specific cases. Here, we highlight some successful classical algorithms used to address the ground state problem, emphasizing their primary assumptions.

Mean-field methods represent the state as a tensor product of linearly-many local subsystem states. By neglecting entanglement between different parts of the system, these methods offer an efficient yet rough representation suitable for approximating ground states in weakly-interacting systems. Even in such cases, mean-field methods often require to be aided by perturbation theory to achieve quantitative results. An example of a popular mean-field method is the **Hartree-Fock self-consistent field method**, frequently employed for Fermionic particle systems like molecular and material models. It is often used as a starting point for other methods, including ground state quantum simulations (see Section 1.5.1 for an example).

1

Similarly, **tensor-network methods** address ground state problems for local Hamiltonians by constraining the level of entanglement between local subsystems. This is accomplished by compressing the 2^n amplitudes of the state vector into a set of small tensors, along with a set of rules for contracting their indices structured to replicate the locality of the system. Tensor network techniques are particularly well-suited to treating geometrically local systems (see Fig. 1.1) when entanglement between distant parts of the system is small. One-dimensional tensor networks, also known as matrix product states, excel in describing ground states of gapped one-dimensional systems. Extending tensor network methods to higher-dimensional systems is more challenging due to the computational complexity of arbitrary index contractions. However, heuristic tensor network methods continue to improve, providing some of the most accurate results for correlated condensed matter systems [32] and chemistry applications [33, 34].

In the realm of Fermionic models of natural systems, **density functional theory (DFT)** offers a set of highly efficient algorithms for studying ground states. DFT achieves compression by storing only the local density of particles, a classical variable whose information content scales linearly with the system's size. The Hohenberg-Kohn theorems establish that such a description suffices to fully characterize a ground state for a two-body Fermionic Hamiltonian [35, 36]. Nevertheless, the computational complexity of exactly retrieving the energy from this compressed representation remains exponential. To simplify this computation, various mathematically and/or empirically motivated approximations have been developed, resulting in a diverse range of *approximate functionals*. Thanks to the compact nature of the compressed representation and the relatively straightforward computation of these approximate functionals, DFT proves to be exceptionally computationally efficient. Consequently, it stands out as the most practical method for handling exceedingly large systems. The accuracy of DFT outcomes is limited to systems where particle correlations remain relatively weak, rendering it particularly suitable for addressing many practical problems in fields such as organic chemistry.

Quantum Monte Carlo (QMC) encompasses a broad array of classical techniques that apply stochastic, or Monte Carlo, methods to quantum simulation problems. QMC methods typically circumvent the need to store an accurate representation of the ground state by recasting the ground state problem as a high-dimensional integral involving simple wave functions and operators. The complexity of solving this integral is managed through an approximate solution based on stochastic integration, facilitated by cleverly designed importance sampling. A poor choice in the

definition of the integral or of the importance sampling step can lead to a divergence in the variance of the result, thus requiring exponential sampling time to achieve the target accuracy. While this divergence often occurs in Fermionic systems due to the so-called *sign problem* [37, 38], techniques based on heuristic constraints enable QMC to obtain competitive results for complex problems in condensed matter physics and chemistry.

1.4.2 Hamiltonian simulation algorithms

The first approach to efficiently synthesize an evolution operator $U(t)$ on a quantum device, addressing the Hamiltonian simulation problem (see Section 1.3.1), was proposed by Lloyd in 1996 [21]. This technique, often referred to as Trotterization, adapts expansion methods for exponentials of operators developed by Trotter [22] and Suzuki [23]. The central concept revolved around recognizing that, for small time increments δt , the linearized evolution operator generated by the Hamiltonian $H(t)$ (decomposed as in Eq. (1.3)) can be approximated as

$$U(t, \delta t) = e^{-iH(t)\delta t} = \prod_{j \in [J]} e^{-ih_j(t)\delta t} + \mathcal{O}(\delta t^2). \quad (1.7)$$

The second order error in δt is due to the non-commuting terms $[h_j, h_{j'}] \neq 0$. Since each $e^{-ih_j\delta t}$ only acts non-trivially on a constant-size subsystem \mathcal{H}_j , it can typically be implemented in constant or linear time on a quantum device. To achieve the desired $U(t) = \prod_{\tau=1}^{t/\delta t} U(\tau\delta t, \delta t)$, the short-time evolution operator is synthesized with the desired accuracy and applied $\frac{t}{\delta t}$ times.

The approximation error in Trotterization can be controlled in various ways. This includes reducing δt , modifying the approximation (1.7) by increasing the order of the remainder, or randomizing the sequence of the local operators $e^{ih_j\delta t}$ using a technique known as quantum stochastic drift protocol (qDRIFT) [39]. While analytical bounds for the Trotter approximation error are less favorable than for more advanced methods, they can sometimes be tightened by leveraging the Hamiltonian's locality through Lieb-Robinson bounds [40–42]. Extensions of these methods that aim to counteract the approximation using locality provide highly efficient techniques for simulating geometrically local Hamiltonians.

Trotter-like methods represent the simplest algorithms for Hamiltonian simulation, and to date they are the only ones experimentally tested on quantum devices. In some of these experiments, the approximation error is embraced and included in the models as a periodic drive potential,

resulting in so-called *Floquet models*.

Over the past decade, various techniques have been developed that take a fundamentally different approach than Trotter expansion to synthesize $U(t)$. These methods are typically based on the concept of block encodings of a sub-unitary operator H/λ (with $\lambda \geq \|H\|$), constructed by expressing H as a linear combination of unitary operators (LCU) [43–45]. The evolution $U(t)$ is then synthesized by transforming the block-encoded operator through techniques derived from quantum walk theory [43] and nuclear magnetic resonance control. These techniques are known as **qubitization** [46], **quantum signal processing** [47], and the quantum singular value transform [48].

All the techniques mentioned above are designed for universal quantum computers and can be applied to simulate any physical model given a fault-tolerant quantum device of sufficiently large scale. In contrast, *analog quantum simulators* natively implement the evolution according to a restricted class of Hamiltonians, defined by the underlying physical implementation and the experimental control on the device. Due to these limitations, each quantum simulator can only address a subset of simulation problems, with its capabilities heavily dependent on its specific implementation.

1.4.3 Adiabatic state preparation

The adiabatic algorithm [49] is an early technique for eigenstate preparation (see Section 1.3.2), grounded in the adiabatic theorem introduced by Born and Fock in 1928 [50]. According to this theorem, a system evolving under a time-dependent Hamiltonian $H(t)$ will remain in its instantaneous eigenstate $|E_j(t)\rangle$ if the Hamiltonian changes gradually enough (i.e., *adiabatically*) and there is a persistent energy gap between $E_j(t)$ and its neighboring eigenvalues $E_{j+1}(t), E_{j-1}(t)$. For simplicity, hereon we focus on ground states, considering only the ground state gap $\Delta = E_1 - E_0$.

The adiabatic theorem can be harnessed by constructing a time-dependent Hamiltonian

$$H(t) = \left(1 - \frac{t}{t_{\max}}\right) H_{\text{init}} + \frac{t}{t_{\max}} H_{\text{target}}, \quad (1.8)$$

Here, H_{target} represents the Hamiltonian whose ground state $|E_0^{\text{target}}\rangle$ we aim to prepare, while H_{init} is a Hamiltonian whose ground state $|E_0^{\text{init}}\rangle$ is known and easy to prepare. After initializing the quantum device in the state $|E_0(0)\rangle = |E_0^{\text{init}}\rangle$, we can simulate the evolution dictated by




Figure 1.2: Example of nonadiabatic transitions. The plot shows the low-energy spectrum of an adiabatic algorithm Hamiltonian $H(t)$ [see Eq. (1.8)], for the same Ising chain model used in the simulations of Chapter 5. We plot the fidelity $|\langle E_j(t)|\psi(t)\rangle|$ of the adiabatic state $|\psi\rangle(t)$ with the instantaneous eigenstates $|E_j(t)\rangle$ (both in color and line size) throughout the progress of the quasi-adiabatic evolution. Note how around the gap closing at $\frac{t}{t_{\max}} \approx 0.5$ the fidelity with the instantaneous ground state is reduced in favor of the first excited state.

$H(t)$, ensuring T_{\max} is adequately large to ensure the Hamiltonian changes adiabatically. If the system remains in the instantaneous eigenstate $|E_0(t)\rangle$, we will have successfully prepared the target ground state at time $t = t_{\max}$.

This form of the adiabatic theorem holds exactly in the infinite t_{\max} limit, which is clearly unpractical. Finite rates of change in $H(t)$ result in *nonadiabatic transitions*, representing a transfer of some amplitude between instantaneous eigenstates, as depicted in Figure 1.2. The rate of transition depends on the rate of change of the Hamiltonian $\|\frac{d}{dt}H(t)\|$ (faster changes lead to larger transitions) and on the ground state gap $\Delta(t) = E_1(t) - E_0(t)$. The nonadiabatic transitions are stronger when the gap $\Delta(t)$ is small. While it is possible to enhance adiabatic algorithm outcomes by employing a schedule that adjusts the rate of change in $H(t)$ based on the gap, determining the gap a priori is challenging.

Rigorous bounds on the final state's approximation quality, dependent on the schedule and gap, have been established in a series of adiabatic theorems [51–53]. A popular version of the adiabatic theorem states that a time proportional to $[\min_t \Delta(t)]^{-2}$ is required to prepare the target ground state with fidelity $1 - \epsilon$.

The adiabatic algorithm is particularly suitable for analog quantum

simulators, which can natively implement dynamics generated by $H(t)$. On a universal quantum computer, implementing the adiabatic algorithm would entail simulating a time-dependent Hamiltonian, utilizing one of the algorithms discussed in the previous section and thus incurring additional computational overhead.

1.4.4 The variational quantum eigensolver

The variational quantum eigensolver (VQE) is a ground state preparation algorithm (see Section 1.3.2) tailored to the capabilities of NISQ devices, originally introduced in 2014 [19].

The VQE, illustrated in Fig. 1.3, is based on the interplay between a quantum device (or *quantum processing unit*, QPU) and a classical computer. The quantum device is used to run a subroutine that involves two main steps: 1. Preparation a parametrized *ansatz state* $|\psi(\boldsymbol{\theta})\rangle$; and 2. Estimation of the Hamiltonian expectation value $E(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle$ (as in Sec. 1.3.3). The ansatz state is generated through a set of parametrized operations on the quantum device, forming what is known as a *parametrized quantum circuit* (PQC). This subroutine is used by a classical algorithms, running an outer optimization loop with the objective of minimizing $E(\boldsymbol{\theta})$. Notably, since no state possesses lower energy than the ground state $|E_0\rangle$, the result of the VQE is *variationally bounded*, meaning $\min_{\boldsymbol{\theta}} E(\boldsymbol{\theta}) \geq E_0$.

An essential factor for the VQE's effectiveness is the design of the PQC preparing the ansatz state $|\psi(\boldsymbol{\theta})\rangle$. The set of ansatz states cannot cover the entire Hilbert space, since a classical variable $\boldsymbol{\theta}$ containing the same amount of information as the amplitudes characterizing the quantum state (which grows exponentially with the system size) would defeat the purpose of employing a quantum device. Instead, the choice of the PQC is guided by various heuristics. These can include extensions of perturbation theory (e.g., unitary coupled cluster [19]), constructions inspired by the problem's Hamiltonian (e.g., variational Hamiltonian ansatz [54], quantum alternating operator ansatz [55, 56]), or its symmetries (e.g., quantum-number preserving fabrics [57]). Adaptive ansätze have also been proposed, where the PQC is generated dynamically at runtime [58]. Given the importance of keeping the PQC compact, ansätze optimized for a specific hardware architecture are common in current proof-of-concept experiments [59].

While VQE provides an invaluable testing platform for quantum algorithm in the NISQ era, its potential to achieve practical quantum advantage is a topic of debate. Even with strong heuristics in the ansatz construction, low-dimensional ansatz manifolds may not accurately represent complex

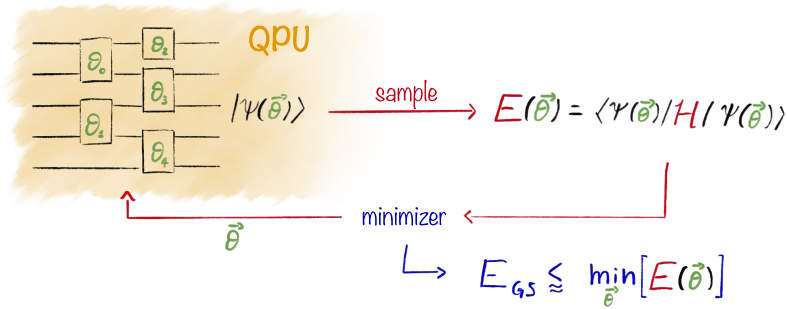


Figure 1.3: Block scheme of the variational quantum eigensolver. The subroutine run on the quantum device (labeled QPU) prepares an ansatz state $|\psi(\boldsymbol{\theta})\rangle$ through a quantum circuit whose gates are parameterized by the elements θ_j of $\boldsymbol{\theta}$. The expected energy $E(\boldsymbol{\theta})$ of $|\psi(\boldsymbol{\theta})\rangle$ is measured by sampling on many repetitions of the circuit, and minimized by a classical outer loop.

correlated systems; certifying the quality of the result is also complicated. Minimizing the cost function $E(\boldsymbol{\theta})$ is difficult due to the complex, multimodal nature of the optimization landscape, occasional vanishing gradients, and added challenges from hardware and sampling noise. Even when VQE has the potential to prepare the ground state of a system of interest, obtaining a high-accuracy estimate of the final energy estimate necessitates a substantial number of shots. This cost is further increased by the necessary application of error mitigation techniques.

1.4.5 Quantum phase estimation algorithms

The problem of quantum phase estimation is closely related to that of eigenvalue sampling described in Section 1.3.3. The objective of phase estimation is to sample eigenvalues φ_j of a unitary operator U with probabilities proportional to the fidelity $p_j = |\langle \varphi_j | \psi \rangle|^2$ with an input state $|\psi\rangle$. When considering $U = e^{-iHt}$, the sampled unit eigenvalues φ_j can be mapped to corresponding eigenvalues of the generator H as $E_j = \arg(\varphi_j)/t$, provided that t is chosen appropriately to ensure the argument function is single-valued (typically $t < \pi \|H\|$).

To gain insight into the relationship between the dynamics generated by H and the measurement of the observable H , we can draw parallels to an

early model of measurements introduced by Von Neumann [60]². In this model, a *pointer* representing the measurement apparatus is coupled to the system, with the aim to map the quantity we want to measure onto the pointer's state. This is represented in the left panel of Fig. 1.4. The pointer is described by a single continuous variable q , with generalized eigenstates $|x\rangle_{\text{P}}$, initially in the state $|0\rangle_{\text{P}}$. The displacement operator $D(\alpha) = e^{-i\alpha p}$ generated by the conjugate momentum p can be used to encode a number α on the pointer's state $D(\alpha)|0\rangle_{\text{P}} = |\alpha\rangle_{\text{P}}$. Correspondingly, the coupling the pointer and the system through the generator $p \otimes H$ yields

$$e^{-ip \otimes H t} |0\rangle_{\text{P}} \otimes |\psi\rangle = \sum_j |E_j\rangle_{\text{P}} \otimes |E_j\rangle \langle E_j | \psi \rangle \quad (1.9)$$

Tracing out the original system leaves the pointer in the ensemble state $p_j |E_j\rangle \langle E_j|_{\text{P}}$, yielding the desired result of eigenvalue sampling.

The classic quantum phase estimation algorithm (QPEA), introduced by Kitaev in 1995 [61], can be described as a Von Neumann measurement model with a digital pointer. The variable representing the pointer's position q is discretized and encoded in a binary register of m qubits, initialized to the all-zero state $|0\rangle_{P_{m-1}} \dots |0\rangle_{P_0}$. A quantum Fourier transform (QFT) [62] is applied to the pointer register, transforming its basis to represent a discrete counterpart of the conjugate momentum p . Subsequently, each qubit P_k in the pointer register is employed to control 2^k iterations of the operator $U = e^{-iH t}$, specifically using the operation $|1\rangle \langle P_k | U^{2^k} + |0\rangle \langle P_k | \mathbb{1}$. This effectively implements the discrete analog of the evolution $e^{-ip \otimes H t}$. Finally, an inverse QFT returns the pointer register to representing the pointer's position. The resulting state encodes an approximation of E_j , with precision depending on the size of the pointer register m and with probability proportional to the input state fidelity. A block diagram of the circuit for QPEA is shown in the top-right panel of Figure 1.4.

An extremely reduced version of the QPEA circuit can be created by modeling the pointer with a single qubit. The QFT on a single qubit coincides with a Hadamard gate, and the output of each run of the circuit is binary (0 or 1), hence the common name for this circuit is the **Hadamard Test (HT)**. The HT circuit is represented in the bottom-right panel of Figure 1.4. The Hadamard test is explored in depth as a model of a generalized measurement with binary output in Chapter 4. Although a single run of the HT is not particularly informative, the expectation value of

²The connection between Von Neumann measurement model and quantum phase estimation introduced in these pages was inspired by a presentation given by Seth Lloyd during his Lorentz professorship in Leiden, in 2019.

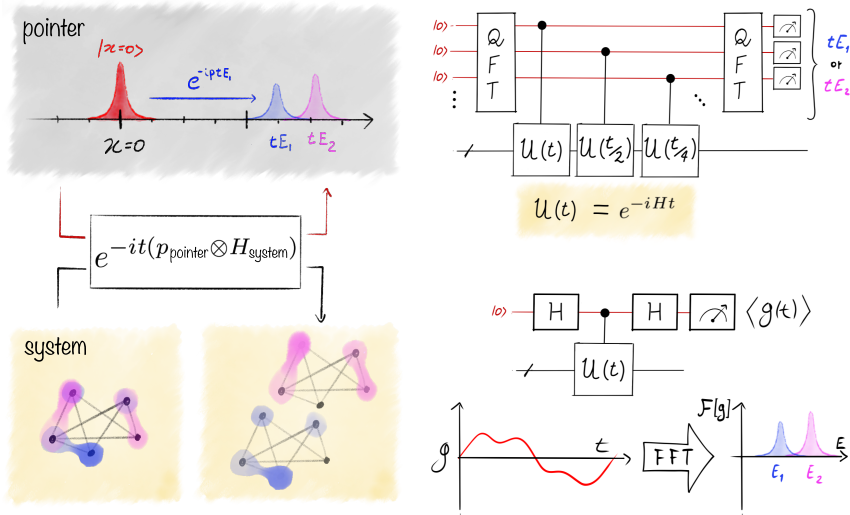


Figure 1.4: (Left) Von Neumann measurement of the system energy H_{system} using a continuous-variable pointer. (Top right) Entanglement-based quantum phase estimation algorithm, employing the quantum Fourier transform (QFT). The binary-encoded result is read out on the pointer qubits, marked in red. (c) Single-control quantum phase estimation. The Hadamard test circuit it used to measure a signal function $g(t)$. Processing the signal with a fast Fourier transform (FFT) or a similar method yields a spectral density function.

the results (interpreted as -1 and $+1$ instead of 0 and 1) matches $\langle \psi | U | \psi \rangle$, with U being the controlled operator. Measuring this for the evolution operator $U(t) = e^{-iHt}$ at different times t allows us to reconstruct a *signal function* $g(t) := \langle \psi | e^{-iHt} | \psi \rangle$. The frequencies of the Fourier components of $g(t)$, which can be extracted using various signal processing techniques, approximate the eigenvalues E_j , and their amplitudes approximate p_j . While this result is fundamentally different from eigenvalue sampling, both methods can be applied to the eigenvalue estimation problem described in Section 1.3.4. The algorithms based on this idea are commonly referred to as **single-control QPEAs**, and due to the shorter and simpler circuits they required they are more suited to the NISQ regime. In Chapter 3, we introduce an error mitigation technique specially tailored to single-control QPEAs.

1.5 Molecular simulation on a quantum computer

A quantum description of molecules – objects emerging from the interaction between atomic nuclei and electrons – stands as one of the early successes of quantum mechanics [63]. The theoretical study of molecules is one of the fields that most benefited from the development of classical computational simulation. Computational chemistry, a mature discipline, employs a diverse range of methods with applications spanning from fundamental research to industrial process development. Established classical techniques such as density functional theory, mean field and perturbation theory methods (described in Section 1.4.1) allow for efficient simulation of many large and complex molecules, particularly in organic chemistry. This empowers chemists to make predictions regarding molecular properties, geometries, reaction pathways and rates. Importantly, this bypasses the costly and sometimes unfeasible task of physically isolating or synthesizing the target molecule in the laboratory. Predictions from computational simulations can guide experiments, and provide valuable insight for chemical engineers, biochemists and other scientists in the development of industrial processes or the advancement of our understanding of the natural world.

Nonetheless, there are questions within the realm of quantum chemistry that pose challenges for classical techniques. These typically pertain the study of the electronic structure of *strongly correlated* molecules. In such systems, strong interactions disrupt perturbative assumptions, and electrons become entangled in complex ways that cannot be adequately described by the approximate models commonly used in computational simulations (like those described in Section 1.4.1). Quantum simulation holds the promise of providing a novel and distinct tool to tackle the unresolved questions in correlated quantum chemistry. The ability to precisely represent strongly-correlated quantum states renders quantum devices particularly attractive in this context.

While quantum hardware has not yet reached a point where it can support practical computations, ongoing research in quantum simulation is bringing the prospect of useful quantum advantage closer. This comprehensive research endeavor encompasses many stages: the identification of open problems that pose challenges to classical methods, the development of techniques to simplify and adapt these problems to quantum methods, the improvement of quantum algorithms tailored to address these specific challenges, and the establishment of benchmarks to test the algorithms and certify the results. The domain of computational quantum chemistry

offers a multitude of open questions with evident implications for both research and commercial applications.

1.5.1 The pipeline for electronic structure on a quantum computer

In this section, we introduce a key problem in quantum chemistry: the identification of the electronic structure ground state. Various quantum simulation methods aiming to address this have been proposed and explored at length. We show that this problem forms the basis for computing chemical reaction rates, probing molecular properties, and conducting molecular dynamics simulation. Following this, we outline the widely employed second-quantization approach, which simplifies the molecular model and facilitates its mapping onto a quantum device. The resulting quantum simulation problem can be addressed using some of the quantum algorithms introduced in Section 1.4.

The electronic structure ground state problem

A molecule is a system composed of electrons (here labeled as e) and atomic nuclei (labeled as N). The Hamiltonian of such a system can be decomposed in kinetic and electrostatic interaction terms³,

$$H = \underbrace{T_N + T_e}_{\text{Kinetic energy}} + \underbrace{U_{NN} + U_{ee} + U_{eN}}_{\text{Coulomb electrostatic energy}}. \quad (1.10)$$

The dynamics of these two groups of particles can be treated separately, following the approximation introduced by Born and Oppenheimer in 1927 [63], which is founded in large mass difference between nuclei and electrons. Initially, the heavier and slower nuclei are treated as immobile charges fixed at coordinates \mathbf{R} . Their dynamics can be reintroduced later, either classically (treating \mathbf{R} as a classical variable) or semiclassically (quantizing \mathbf{R} but neglecting its correlations with the electronic state). An *electronic structure Hamiltonian* can then be expressed as a function of \mathbf{R}

$$H_e(\mathbf{R}) = U_{NN}(\mathbf{R}) + U_{eN}(\mathbf{R}) + T_e + U_{ee}. \quad (1.11)$$

³This molecular structure Hamiltonian neglects relativistic effects (including spin-orbit) and nuclear spins, which is often a reasonable approximation. In areas of theoretical chemistry where these effects are significant, the treatment can be adjusted accordingly.

1

The first term $U_{\text{NN}}(\mathbf{R})$ is a nuclear repulsion constant which does not influence electronic dynamics, like the nuclear kinetic term T_{N} which we removed by fixing the nuclear positions. The term $U_{\text{eN}}(\mathbf{R})$ now represents the potential experienced by electrons due to nuclear charges. The remaining operators T_{e} and U_{ee} , are independent of \mathbf{R} and represent the kinetic energy of electrons and the electrostatic repulsion between electron pairs, respectively.

Solving the ground state problem for the electronic structure Hamiltonian Eq. (1.11) yields the energy $E_0(\mathbf{R})$ as a function of nuclear coordinates. The energy required to excite electrons to higher eigenstates is typically on the order of an electronvolt, which is two orders of magnitude larger than both the ambient temperature and the typical nuclear kinetic energy. This difference justifies the approximation of considering only electronic ground states, as well as the Born-Oppenheimer approximation. The local minima of $E_0(\mathbf{R})$ define the equilibrium geometries for the molecule. The shape of the potential around the equilibrium points defines the vibrational and rotational spectra of the molecule. More generally, nuclear motion can be reintroduced as governed by the potential $E_0(\mathbf{R})$, assuming that electrons follow their instantaneous ground state adiabatically. The \mathbf{R} -derivatives of $E_0(\mathbf{R})$ represent the forces acting on nuclei, which can be used in *molecular dynamics* simulations. Additionally, other molecular properties, such as polarizability, can be predicted based on the representation of the electronic structure ground state.

One relevant target of electronic structure calculations is the prediction of reaction rates. A chemical reaction can be viewed as a trajectory in nuclear coordinate space, starting from a stable state \mathbf{R}_{R} , with the nuclei arranged as reactant sub-molecules, and ending at another stable state \mathbf{R}_{P} representing the products. Along this reaction path connecting $\mathbf{R}_{\text{R}} \rightarrow \mathbf{R}_{\text{P}}$, a saddle point \mathbf{R}_{T} signifies a transition state: the highest energy point along the path. The activation energy of the reaction $\Delta E_{\text{A}} = E_0(\mathbf{R}_{\text{T}}) - E_0(\mathbf{R}_{\text{R}})$ is a key factor in determining the reaction rate constant k^4 . Both the (empirical) Arrhenius equation and the (first-principle) Eyring equation, reveal that k depends on the activation energy exponentially, $k \propto e^{-E_{\text{a}}}$. Therefore, resolving the activation energy to high precision is essential.

⁴Complex reactions may occur along multiple pathways, each consisting of several steps between metastable states. The overall reaction rate can be computed by applying the concepts described here to each individual step of the reaction and combining the results.

Discretization

Addressing the challenge just described amounts to solving the ground state problem, described in Section 1.3.4, for the electronic structure Hamiltonian Eq. (1.11). Tackling this with quantum simulation first requires to translate this Hamiltonian into the form of an operator acting on the state space of a quantum device. A common approach, assuming a second-quantized treatment of the problem, encompasses three steps: 1. discretization of the single-particle state space, 2. mean-field calculation, leading to the selection of single-particle orbitals relevant to correlations, and 3. mapping fermionic operators to qubit operators.

The first step, widely used in computational chemistry, involves discretizing space using a set of $2M_{\text{ao}}$ single-particle wavefunctions called *atomic orbitals*⁵ $\chi_\mu(\mathbf{x}, \sigma) \in \mathbb{C}$, with $\mathbf{x} \in \mathbb{R}^3$ representing the electron's position and $\sigma \in \mathbb{Z}_2$ its spin. Atomic orbitals are chosen to approximate the low-energy mean-field eigenstates for individual atoms, and to be numerically integrable; they need not be orthogonal to each other. Various methods exist for selecting atomic orbitals, and the resulting collections of orbitals are referred to as a *basis set*. By combining a large number of atomic orbitals from different atoms, more complex single-particle wavefunctions known as *molecular orbitals* ϕ_p can be constructed:

$$\phi_p(\mathbf{x}, \sigma) = \sum_{\mu \in [2M_{\text{ao}}]} C_{\mu p} \chi_\mu(\mathbf{x}, \sigma). \quad (1.12)$$

The combination of atomic orbitals to form molecular orbitals is illustrated in Fig. 1.5 (left and center). Molecular orbitals are typically chosen to be orthonormal.

An anti-symmetrized product of orthogonal molecular orbitals defines a many-electron state known as *Slater determinant*

$$\Phi(\mathbf{x}_1, \sigma_1; \dots; \mathbf{x}_{N_e}, \sigma_{N_e}) = \sum_{\pi \in \mathcal{P}_{N_e}} \text{sgn}(\pi) \bigotimes_{p=0}^{N_e} \phi_p(\mathbf{x}_{\pi(p)}, \sigma_{\pi(p)}), \quad (1.13)$$

where N_e is the number of electrons in the state and \mathcal{P}_m is the set of permutations of size m . The state depends on the molecular orbital coeffi-

⁵The orbitals we define here are often called *spin-orbitals*, as their definition includes the spinor component of the wavefunction. It is common to further factorize *spatial orbitals* and spinors. We avoid this for the sake of synthesis, but we add a factor 2 in the number of spin-orbitals to maintain the compatibility with the common definition of M as the number of spatial orbitals.




Figure 1.5: **Left:** artistic representation of the first few atomic orbitals and their energy levels. **Center:** atomic orbitals from two atoms combine to construct molecular orbitals. For the sake of clarity, we represent only a few of the first molecular orbitals generated by combining atomic orbitals from two atoms. **Right:** mean-field energy of each molecular orbital. One possible complete active space subdivision of orbitals is shown.

icients $C_{p\mu}$ through ϕ_p . A Slater determinant represents an uncorrelated many-body state, where each electron is not entangled to the others. Its energy can be computed efficiently in terms of integrals involving the atomic orbitals, the Coulomb operator $\frac{1}{x_1-x_2}$, the kinetic operator ∇_x^2 and the nuclear positions and charges.

Reduction and second quantization

The Hartree-Fock self-consistent-field method (HF-SCF) efficiently optimizes the coefficients $C_{\mu p}$ to construct the minimal-energy Slater determinant Φ^{HF} and a set of orthogonal molecular orbitals ϕ_p^{HF} , ordered by their mean field energy $\epsilon_p^{\text{HF}} \leq \epsilon_{p-1}^{\text{HF}}$. To introduce correlations on top of the mean field state Φ^{HF} , one can consider superpositions of Slater determinants generated by transferring some of the electrons from *occupied* molecular orbitals (indexed by $p < N_e$) to *unoccupied* molecular orbitals ($q \geq N_e$). Transferring an electron between orbitals $p \rightarrow q$ always results in a positive change in mean-field energy $\Delta_{p \rightarrow q}^{\text{HF}} = \epsilon_q^{\text{HF}} - \epsilon_p^{\text{HF}}$, but considering superpositions between such states can yield a lower, beyond-mean-field energy. However, when $\Delta_{p \rightarrow q}^{\text{HF}}$ is large, the large increase in mean-field energy contrasts the gain from adding the excited Slater determinant to the

superposition. Therefore, when constructing the superposition of Slater determinants, we can limit to considering orbitals with energies within a selected window around Fermi energy $\epsilon_F := \epsilon_{N_e}^{\text{HF}}$. The molecular orbitals are thus divided in three sets: a *core* set of fully-occupied low-energy orbitals for $p < 2M_C$, an set of *active* orbitals $2M_C \leq p < 2M_C + 2M_A$ used to describe correlations among $N_A = N_e - 2M_C$ electrons, and the remaining space of high-energy *virtual* orbitals which are discarded in the calculation of correlations. This subdivision is illustrated in Figure 1.5. Calculations that prescribe this subdivision are called *complete active space* methods, and denoted by the abbreviation $\text{CAS}(N_A, M_A)$.

The electronic structure Hamiltonian Eq. (1.11) can be constrained to the active space, accounting for the effect of frozen electrons in mean-field terms. The resulting Hamiltonian can be divided into terms that act on zero, one single, or pairs of electrons in the active space:

$$H_A = \overbrace{U_{\text{NN}} + U_{\text{NC}} + U_{\text{CC}} + T_C}^{\text{constant}} + \underbrace{T_A + U_{\text{CA}} + U_{\text{NA}}}_{\text{one active electron}} + \overbrace{U_{\text{AA}}}^{\text{two active electrons}}. \quad (1.14)$$

In turn, this can be written in terms of second-quantized fermionic operators acting on a single electron ($c_p^\dagger c_q$) and pairs of electrons ($c_p^\dagger c_q^\dagger c_r c_s$)

$$H_A = \text{const.} + \sum_{p,q \in [2M_A]} g_{pq} c_p^\dagger c_q c_r^\dagger c_s + \sum_{p,q,r,s \in [2M_A]} h_{pqrs} c_p^\dagger c_q^\dagger c_r c_s. \quad (1.15)$$

The coefficients g_{pq} and h_{pqrs} are obtained through integrals of the atomic orbitals with the kinetic and Coulomb operators and are known as *one-electron* and *two-electron integrals*, respectively. The number of Slater determinants that can be obtained by distributing N_A electrons over $2M_A$ orbitals is the binomial coefficient $\frac{(2M_A)!}{N_e!(2M_A - N_e)!}$, which grows exponentially in the size of the active space. The description of the Hamiltonian, instead, only requires the $\mathcal{O}(M_A^4)$ integrals.

Describing electronic structure accurately is most challenging in systems that necessitate large active spaces. The appropriate active space size is determined by chemical intuition, often guided by approximate classical simulations. The size of the molecule is not a reliable predictor of the required active space size; many large organic molecules are not particularly challenging because the electrons predominantly occupy bonding orbitals, well described by mean-field theory. There, transferring electrons to empty orbitals has a large energy cost and thus correlations remain small. Methods like density functional theory predict the ground state energy of such

systems efficiently and accurately. An example of a challenging problem is given by molecules containing transition metals; these provide many atomic *open-shell* orbitals, which lead to a number of complex molecular orbitals with similar energies which need to be considered in the active space construction. One such case is nitrogenase, an enzyme crucial for converting atmospheric nitrogen into ammonia, containing iron and molybdenum atoms embedded in an organic matrix (FeMo cofactor). A deeper understanding of this enzyme could potentially lead to improvements in industrial process of ammonia synthesis (critical, among others, for fertilizer production), currently achieved through the costly Haber-Bosch process.

Mapping to a quantum computer

The Slater determinants defined over active molecular orbitals constitute a set of orthogonal states. These can be mapped onto the states of a quantum device, serving as a basis to represent more general correlated states. Similarly, the Fermionic creation and annihilation operators c_p^\dagger, c_p that construct the Hamiltonian Eq. (1.15) need to be mapped onto operators on the device's Hilbert space. The Jordan-Wigner transformation, a common mapping between fermionic systems and qubit system, prescribes assigning one qubit to each molecular spin-orbital. An occupied spin-orbital is represented by the state $|1\rangle$ of the corresponding qubit, while an empty orbital corresponds to $|0\rangle$. The mapping of fermionic operators is defined as

$$c_p^\dagger = \bigotimes_{j=0}^{p-1} \sigma_j^z \otimes \sigma_j^+; \quad c_p = \bigotimes_{j=0}^{p-1} \sigma_j^z \otimes \sigma_j^-, \quad (1.16)$$

where σ_j^z is the Pauli Z operator on the j -th qubit and $\sigma_j^\pm = \frac{1}{2}(\sigma_j^x \pm i\sigma_j^y)$ are the raising/lowering operators. This mapping of fermionic operators ensures the correct anti-commutation relations, but sacrifices locality.

Solution by quantum simulation

With a model of the electronic structure problem represented in terms of states and operators on a quantum device, we can now explore the use of quantum simulation algorithms to solve it. If the goal is to solve the ground state problem on a NISQ device, one viable approach is to employ the Variational Quantum Eigensolver (VQE), as described in Section 1.4.4. Several VQE ansätze tailored to chemistry problems have been developed [64, 65], with the original proposal of the VQE being motivated by implementing a unitary version of the coupled-cluster ansatz

[19]. Various optimizations have been proposed for estimating expectation values (see Section 1.3.3) of a Hamiltonian of the form Eq. (1.15) [66–68]. An effective and practical method consists in measuring the *one-* and *two-electron reduced density matrices* of the input state $|\psi\rangle$, respectively

$$D_p^q = \langle \psi | c_p^\dagger c_q | \psi \rangle \quad \text{and} \quad D_{pq}^{rs} = \langle \psi | c_p^\dagger c_q^\dagger c_r c_s | \psi \rangle. \quad (1.17)$$

In practice, current quantum devices can only handle very small toy models of molecules using VQE. The largest VQE simulation for electronic structure with the final expectation value estimation step performed on a quantum computer utilized an active space of only 10 orbitals (with a further approximation simplifying the spin degree of freedom), and it relied heavily on error mitigation techniques [69], including the one introduced in Chapter 3.

Proposals for fault-tolerant quantum simulation of chemistry typically focus on the application of the quantum phase estimation algorithm (see Section 1.4.5) [70]. Using QPEA to solve the ground state problem requires a Hamiltonian simulation subroutine and a technique for preparing a state with a significant overlap with the ground state. Recent years have seen the proposal of numerous optimized Hamiltonian simulation algorithms for electronic structure, many of which employ qubitization and LCU techniques (see Section 1.4.2). These methods primarily concentrate on constructing efficient compressions of the highly structured information contained in the one- and two-electron integrals, making them more suitable for uploading onto quantum computers [71–74]. Preparing a state with a good overlap with the ground state remains a challenging problem. It has been suggested that this step could potentially hinder the achievement of exponential quantum advantages in simulating electronic structure ground states [75]. Since most molecules naturally approach their electronic ground state through processes like thermalization and adiabatic evolution along a reaction coordinate, we expect heuristic methods can be used to efficiently emulate these phenomena. Dissipative algorithms, such as the one proposed in Chapter 2, or adiabatic state preparation, could be optimized to approximate molecular ground states effectively.

1.5.2 The journey towards useful quantum advantage in chemistry

The example presented in the previous section illustrates the multitude of decisions involved in defining a pipeline for studying a physical phenomenon using quantum simulation. These decisions impact the level of

1

approximation, precision of results, and the practicality of the method. Research in simulation can address any step of the pipeline, from the very choice of the problem to study to the final algorithmic result.

Extensive research has been conducted on the problem of electronic structure ground states in quantum simulation, with efforts focused on optimizing molecule modeling, mapping, and various quantum algorithms – in the previous section, we only presented the most common approach. An important body of research today includes the estimation of the costs of running quantum simulation algorithms on prospective quantum computer, especially for fault-tolerant approaches. For instance, in [74], it was estimated that a fault-tolerant quantum computer could execute the Quantum Phase Estimation Algorithm (QPEA) for a model of the nitrogenase FeMo cofactor, with a precision sufficient for estimating reaction rates, in four days using approximately 4 million physical qubits (assuming reasonable but as yet unachieved physical error rates).

To attain useful quantum advantage, it is essential to identify problems that are easy to solve through quantum simulation but challenging for classical methods. The ground state problem for molecules containing transition metals is hard for classical methods, due to the large size of the active space in which correlations need to be accounted for non-perturbatively. However, classical approximate methods for electronic structure ground state continue evolving, with some of them being able to approximate more and more complex correlations [75], making the remaining challenging cases even more demanding for quantum solutions.

Another set of problems challenging for classical simulation is provided by photochemistry, where low-lying excited electronic eigenstates play a significant role in reactions initiated by light rather than thermal energy. Describing excited states on a quantum computer is not significantly more complex than describing ground states, while classical simulation algorithms often exploit ground-state-specific properties. Photochemical reactions also involve the conversion of light energy into phonons, which can lead to the breakdown of the Born-Oppenheimer approximation, especially near electronic structure spectrum degeneracy points known as *conical intersections* (detailed in Chapter 6). The numerical study of chemistry beyond the Born-Oppenheimer approximation is very limited due to the complex correlations between electrons and nuclei, and the difficulty of discretizing space in the absence of clear atomic orbitals. A concurrent simulation of electrons and nuclei could be in principle achieved through first-quantized quantum simulation methods [76]. The size of the quantum computers required to implement these methods is beyond the foreseeable future, but in the long term such quantum computers could be achieved.

Another example of shift in the problem focus includes the recent proposal of using Hamiltonian learning to reconstruct NMR spectra [77, 78]. This problem also pertains the field of computational quantum chemistry, but it completely bypasses the electronic structure Hamiltonian, focusing on the interactions between nuclear spins instead. Finally, quantum algorithms could be used to assist known classical methods in innovative ways. Finally, quantum algorithms can complement existing classical methods in innovative ways, as seen in the *quantum computing quantum Monte Carlo* (QC-QMC) proposal. There, sampling on a variational quantum state provides information about correlations in the electronic structure wavefunction, while high-precision energy calculations are performed classically.

1.6 Outline of this thesis

This thesis introduces a number of quantum algorithms tailored to the simulation of physical systems. Chapter 2 introduces a category of algorithms aiming to prepare ground states of natural systems through simulated cooling. Chapter 3 presents a novel error mitigation approach, *echo verification*, and investigates its application in single-control phase estimation algorithms. The two following chapters explore applications of echo verification in different contexts: Chapter 4 develops techniques for estimating expectation values within the measurement model induced by echo verification; Chapter 5 applies echo verification to mitigate the effect of non-adiabatic transitions in the adiabatic algorithm. Finally, Chapter 5 explores the challenge of detecting conical intersections in molecular models – an under-explored problem in quantum chemistry, well-suited to quantum simulation – and proposes a resilient quantum algorithm to solve it.

Chapter 2: Quantum Digital Cooling

In chapter 2, we explore the idea of simulating cooling by coupling the system to a single-qubit “fridge”. This auxiliary qubit is reset periodically to its low-energy state, allowing to extract energy and entropy from the system. The use of a single-qubit fridge differs from the natural thermalization process, where systems cool by releasing energy into extensive, cold, and ergodic baths. While simulating such baths is theoretically possible, it comes with a significant computational burden. Our investigation of single-qubit fridges leads us to introduce a category of algorithms designed for preparing the ground states of simulated Hamiltonians, which we name

quantum digital cooling. We establish two approaches to quantum digital cooling: a efficient strong-coupling approach which coarsely approximates the ground state, and a more expensive approach that can achieve arbitrary accurate approximations. We then assess the performance of both methods through numerical benchmarks.

Chapter 3: Error mitigation via verified phase estimation

Chapter 3 introduces a novel error mitigation technique that leverages the quantum information remaining in the device's state following a Hadamard-test-based measurement. This method involves checking if, after the measurement, the system state can be projected back onto the input state. This *verification* step has a high probability of failure in the presence of errors, enabling the rejection of erroneous results, and it does not introduce any bias to the measurement outcome when errors are absent. Our technique finds successful application within the framework of the single-control quantum phase estimation algorithm (introduced in Section 1.4.5), leading to the development of the *verified phase estimation (VPE)* algorithm. We demonstrate the application of VPE to both phase estimation and expectation value estimation problems for various models, showing improvements of several orders of magnitude over unmitigated estimation at near-term error rates.

Additionally, this chapter introduces a variant of VPE, known as *control-free VPE*, which eliminates the need for a control qubit. This simplifies the control circuitry significantly, bringing near-term implementations of our technique within reach.

In further literature, this error mitigation technique gained recognition under the name of *echo verification* [69, 79], a terms that draws parallel with the Loschmidt echo [80]. Control-free VPE has also been tested experimentally in a superconducting quantum processor, realizing simulations of systems of up to 10 qubits, and thus leading to the implementation of the largest variational algorithm for a correlated chemistry system to date [69].

Chapter 4: Optimizing the information extracted by a single-qubit measurement

In chapter 4, we study expectation value estimation (see Section 1.3.3) in a restricted model of quantum computation, where we are only allowed to extract a single bit of information per each n -qubit quantum state preparation. This restriction is motivated by echo verification, where all

but one qubit are used in the verification step to detect errors. We indeed show that, in echo-verification-like schemes, extracting more than one bit of information is counterproductive towards estimating expectation values.

Within this restricted model, we optimize expectation value estimation by decomposing the target observable into a sum of bitwise-measurable terms. We construct optimal decompositions analytically, and we propose a set of rules to improve on a given decomposition which can sometimes be applied even in presence of experimental constraints. We find the optimal decomposition of a fast-forwardable operator, and show a numerical improvement over a simple Pauli decomposition by a factor $n^{0.7}$.

Chapter 5: Virtual mitigation of coherent non-adiabatic transitions by echo verification

In chapter 5, we develop an extension of echo verification tailored to applications to the adiabatic state preparation algorithm (see Section 1.4.3). This technique, which we call *adiabatic echo verification*, mitigates both coherent and incoherent errors arising, respectively, from non-adiabatic transitions and hardware noise. This is an unconventional application of error mitigation, which is typically applied to hardware errors only. Even in the absence of hardware noise, the estimator bias of the observable is reduced when compared to standard adiabatic preparation, achieving up to a quadratic improvement.

Our method is based on two quasi-adiabatic evolutions with mirrored schedules implementing the echo. These are interleaved by a dephasing step by random-time evolution, and by the Hadamard-test-based measurement. The dephasing promotes coherent errors from non-adiabatic transitions into incoherent errors, making them amenable to verification. Our method requires positive-time dynamics only, making it more suitable to application in analog quantum simulators.

Chapter 6: A hybrid quantum algorithm to detect conical intersections

In chapter 6, we tackle a problem of interest in photochemistry: the detection of conical intersections in molecular models. The chapter introduces a quantum simulation algorithm tailored to the NISQ era to solve this problem.

Conical intersections are significant points in the geometry of a molecule, where the electronic potential energy surfaces cross in a topologically protected manner. Close to a conical intersection the Born-Oppenheimer

1

approximation breaks down, with crucial implications for chemical processes like photoisomerization and non-radiative relaxation. One prominent example is their role in the vision process, where a conical intersection facilitates the isomerization of retinal after absorbing a photon, which in turn triggers a cascade of chemical signal that result in the perception of light.

Conical intersections are characterized by a non-zero Berry phase, a topological invariant defined on a closed path in atomic coordinate space. The berry phase assumes the value of π when the path encircles the intersection manifold, and 0 otherwise. The algorithm we propose tracks the approximate ground state along the chosen path, using a parametrized quantum circuit ansatz updated by a fixed number of Newton-Raphson steps. At the end of the algorithm, a Hadamard test is used to measure a single bit of information, which determines whether the Berry phase is π or 0. Since the final result is discrete, our procedure succeeds even for a cumulative error bounded by a constant; this allows us to bound (analytically) the total sampling cost and to readily verify the success of the procedure. The application of our algorithm is demonstrated numerically on a small toy model of the formalimine molecule ($\text{H}_2\text{C}=\text{NH}$).