


UNIVERSITÀ DEGLI STUDI DI PAVIA
DOTTORATO DI RICERCA IN FISICA – XXXVIII CICLO

Efficient Quantum Simulations on Digital Quantum Computers

Davide Cugini



This work is licensed under a [Creative Commons “Attribution 4.0 International”](#) license.





Università degli Studi di Pavia
Dipartimento di Fisica “Alessandro Volta”

DOTTORATO DI RICERCA IN FISICA – XXXVIII CICLO

Efficient Quantum Simulations on Digital Quantum Computers

Davide Cugini

Submitted to the Graduate School of Physics of the University of
Pavia in partial fulfilment of the requirements for the degree of

DOTTORE DI RICERCA IN FISICA

DOCTOR OF PHILOSOPHY IN PHYSICS

Supervisors

Prof. Dario Gerace
Prof. Giacomo Guarnieri
Prof. Francesco Tacchino

*Alla mia famiglia e ai miei cari, con
amore e gratitudine.*

ACKNOWLEDGEMENTS

First and foremost, I want to thank my supervisors, Prof. Dario Gerace, Prof. Giacomo Guarnieri, and Prof. Francesco Tacchino, for their support and for always pushing me forward throughout this journey. My time at Los Alamos National Laboratory was one of the highlights of my PhD, and I owe that to Dr. Touheed Atif and Dr. Yigit Subasi who contributed to making that experience both extremely formative and genuinely enjoyable. I'm also grateful to the friends and colleagues who shared the day-to-day of research life at the University of Pavia with me: Dr. Davide Nigro, Dr. Francesco Scala, Francesco Ghisoni, Dr. Claudia Schrama, and Dr. Angela Rosy Morgillo. The conversations, collaborations, and everything in between made this a much richer experience. A special mention goes to Luca Spagnoli, the person whose friendship and scientific exchange I've found most enriching, in ways that go well beyond what ends up written in a thesis. Finally, thanks to Dr. Mario Motta, Dr. Torin Stetina, Prof. Nathan Wiebe, and Dr. Ignacio Loaiza for the ongoing collaborations and the many stimulating exchanges. I look forward to what comes next.

LIST OF PUBLICATIONS

The following manuscripts are part of the thesis:

1. **Davide Cugini**, Davide Nigro, Mattia Bruno, Dario Gerace, *Exponential optimization of adiabatic quantum-state preparation*, [Physical Review Research 7.1 \(2025\)](#) [1].
2. **Davide Cugini**, Giacomo Guarnieri, Mario Motta, Dario Gerace, *Symmetry-guided quantum state preparation: Branched-Subspaces Adiabatic Preparation (B-SAP)*, [Quantum Science and Technology 11.2 \(2026\)](#) [2].
3. **Davide Cugini**, Touheed Anwar Atif, Yiğit Subaşı, *Resource-Optimal Importance Sampling for Randomized Quantum Algorithms*, [arXiv preprint: 2603.13495 \(2026\)](#) [3].
4. **Davide Cugini**, Francesco Ghisoni, Angela Rosy Morgillo, Francesco Scala, *Spectral Gap Estimation via Adiabatic Preparation*, [Physica Scripta \(2026\)](#) [4].

The following manuscripts are not part of the thesis:

5. **Davide Cugini**, Dario Gerace, Pietro Govoni, Aurora Perego, Davide Valsecchi, *Comparing quantum and classical machine learning for Vector Boson Scattering background reduction at the Large Hadron Collider*, [Quantum Machine Intelligence 5.2 \(2023\)](#) [5].
6. **Davide Cugini**, André M. Timpanaro, Giacomo Livan, Giacomo Guarnieri, *Universal emergence of local Zipf–Mandelbrot law*, [Journal of Statistical Mechanics: Theory and Experiment 10 \(2025\)](#) [6].
7. **Davide Cugini**, Giacomo Guarnieri, *State convertibility and fluctuation theorems from a dynamical reference: majorization meets martingales*, [arXiv preprint: 2608.19391 \(2026\)](#) [7].
8. **Davide Cugini**, *Pathwise Random Hamiltonian Simulation*, [arXiv preprint: 2608.29756 \(2026\)](#) [8].

CONTENTS

Summary	7
Acknowledgements	7
List of Publications	8
Introduction	12
I Fundamentals of Digital Quantum Computing	16
1 Postulates of Quantum Mechanics	18
1.1 Pure-state description	18
1.2 Density-operator description	20
2 Basics of Quantum Computing	23
2.1 Qubits and the computational basis	23
2.2 Quantum registers	25
2.3 Elementary quantum gates	25
2.4 Universal quantum computation	27
2.5 Quantum circuits	28
3 Quantum Fourier Transform	29
3.1 The Discrete Fourier Transform	29
3.2 The Quantum Fourier Transform (QFT)	30
3.3 Computational complexity	31
3.4 Periodic states and the interpretation of the QFT	32
II Quantum-State Preparation	33
4 Overview of Quantum-State Preparation Methods	36
4.1 Variational Quantum Algorithms	37
4.2 Adiabatic Quantum-State Preparation	38

5	Optimization of Adiabatic Preparation	42
5.1	Explicit control of the adiabatic error bound	42
5.2	Initial Hamiltonian optimization	46
5.2.1	Example: XZ Heisenberg model	47
5.2.2	Example: Ising model	49
5.3	Periodic boundary conditions and site parity	51
6	Branched-Subspaces Adiabatic Preparation	54
6.1	Application: the XYZ Heisenberg chain	59
6.1.1	Example: target state for $n = 0$	62
6.1.2	Example: target state for $n = 1$	63
6.1.3	Benchmark with Conventional Adiabatic Preparation	64
6.2	Criteria for guaranteeing the satisfiability of B-SAP	65
III	Hamiltonian Evolution	70
7	Resource-Optimal Importance Sampling for Randomized Quantum Algorithms	75
7.1	Optimal Importance Sampling	76
7.1.1	Composite Quantum Channels	78
7.2	Importance Sampling for Imperfect Quantum Channels	79
7.3	Importance Sampling with Error Detection	81
7.3.1	Case Study: IID Parameter Model with Poissonian Error Statistics	84
7.4	Discussion	85
IV	Observable Measurement	88
8	Quantum Phase Estimation	92
9	Spectral Gap Estimation via Adiabatic State Preparation	94
9.1	Ising Model	96
9.2	H_2 and He_2 Molecules	99
9.3	Remarks and Outlook	99
	Conclusion	101
	Appendices	104
A	Exponential behavior of the Adiabatic error	105
B	Proof of Theorem 4	116
C	Proof of Theorem 5	119
D	Superposition states for Ising Hamiltonians	124

E	Observable choice	126
F	Quantum measurement error	128
G	Imperfect superposition state preparation	129
H	Details on IonQ Aria	132
	H.1 IonQ Native Gates	132
	H.2 IonQ Aria noise model	133
I	Additional Ising experiments	136
	Bibliography	137

INTRODUCTION

Contrary to common perception, the influence of quantum mechanics on computation predates the advent of quantum computers. Modern classical computers themselves rely on devices whose operation is fundamentally governed by quantum physics. In particular, semiconductor transistors, the building blocks of classical hardware, exploit quantum-mechanical effects to control charge transport at microscopic scales. Historically, the primary motivation for incorporating quantum physics into electronic devices was not to harness quantum information as such, but rather to enable the continued miniaturization of computational components. As transistors were made progressively smaller, their behavior could no longer be accurately described within a purely classical framework, necessitating a quantum-mechanical description of the underlying physical processes. For several decades, this strategy of scaling down device dimensions has been the dominant driver of technological progress in classical computing. Empirically, it has been captured by Moore's law, which observes an approximately exponential growth in the number of transistors that can be integrated on a single chip over time. While Moore's law is not a fundamental physical principle and has shown signs of slowing in recent years, continued miniaturization remains a central paradigm in semiconductor technology. However, this approach faces intrinsic limitations. As device dimensions approach the atomic scale, further shrinking of the p–n junctions on which conventional transistors are based becomes increasingly challenging due to quantum tunneling, fabrication constraints, power dissipation, and variability at the level of individual atoms. These considerations point to a fundamental limit of the traditional scaling paradigm.

This naturally raises the question of how further advances in computation may be achieved beyond simple miniaturization. Quantum computing represents a qualitatively different approach to exploiting quantum mechanics. Rather than merely using quantum theory to accurately model the behavior of increasingly small devices, quantum computers employ quantum mechanics as a computational resource. In this paradigm, information is encoded in quantum states of qubits, and computation proceeds through the coherent manipulation of superposition and entanglement. By operating directly within the exponentially large Hilbert space of a quantum system, quantum computers can manipulate information in ways that are intractable for classical devices. While not all of this information can be directly accessed due to the constraints of quantum measurement, it can be exploited algorithmically to perform computations that would be impossible or inefficient on classical machines.

In particular, quantum simulations lie at the heart of quantum computing, dat-

ing back to Feynman’s seminal proposal in 1982. The fundamental motivation stems from the exponential growth of the Hilbert space describing a quantum system as its size increases, which renders exact numerical simulation intractable on classical computers. In quantum chemistry, for instance, the dimension of the Hilbert space grows exponentially with the number of particles participating in a chemical reaction. In condensed matter physics, it scales with the number of spins, while in lattice formulations of quantum field theory it grows with the number of sites used to discretize spacetime. As a consequence, classical computational resources are rapidly overwhelmed, making exact simulation feasible only for small systems or those with particularly favorable structure. By directly manipulating qubits, quantum computers operate within the same Hilbert space as the system being simulated. This allows them, in principle, to process exponentially large state spaces using resources that scale only polynomially with system size, thereby circumventing the scaling limitations faced by classical devices. Consequently, quantum computers provide a promising platform for simulating complex quantum systems, including large molecules in quantum chemistry, strongly correlated materials in condensed matter physics, and discretized fields in lattice quantum field theory.

A quantum algorithm for simulating a physical system can be described as a sequence of three fundamental steps: state preparation, Hamiltonian evolution, and measurement. Let the quantum register be initialized in a density matrix ρ , representing the initial configuration of the physical system under study. This state evolves under a Hamiltonian H encoding the system dynamics and is ultimately probed through the measurement of an observable O .

The first step, state preparation, consists in initializing the quantum register in the state ρ . In practice, this state is typically prepared from a simple reference state, such as the all-zero computational basis state, by applying a suitable sequence of elementary quantum operations called quantum gates. The second step, Hamiltonian evolution, implements the time evolution dictated by the system Hamiltonian. Throughout this work, we adopt the common convention of setting the reduced Planck constant to unity, $\hbar = 1$. Under this convention, the Hamiltonian H generates the unitary time-evolution operator

$$U(t) = e^{-iHt}, \quad (0.1)$$

which evolves the initial state ρ over a time interval t , thereby encoding the system’s dynamics into the quantum algorithm. The third step, measurement, extracts physical information through the observable O . According to the postulates of quantum mechanics, an observable admits a spectral decomposition

$$O = \sum_i o_i \Pi_i, \quad (0.2)$$

where Π_i denotes the projector onto the eigenspace associated with the real eigenvalue o_i . A projective measurement of O performed on the evolved state $U(t)\rho U^\dagger(t)$ yields one of the eigenvalues o_i with probability

$$p_i = \text{Tr}[\Pi_i U(t)\rho U^\dagger(t)]. \quad (0.3)$$

By repeating measurements on identically prepared copies of the evolved state, one can estimate expectation values and higher-order statistical moments, e.g., the mean

$$\mathbb{E}[O] = \text{Tr}[O U(t) \rho U^\dagger(t)] \quad (0.4)$$

or, more generally, the k -th moment

$$\mathbb{E}[O^k] = \text{Tr}[O^k U(t) \rho U^\dagger(t)]. \quad (0.5)$$

The discussion above describes the standard model of quantum simulation at the algorithmic level, where state preparation, unitary evolution, and measurement are treated as ideal operations. This abstraction is useful for formulating quantum algorithms and analyzing their computational properties. In practice, however, the performance of a quantum simulation is ultimately determined by the capabilities of the underlying quantum hardware. Current devices operate in the Noisy Intermediate-Scale Quantum (NISQ) regime, characterized by tens to a few hundreds of physical qubits without full quantum error correction. Decoherence, control imperfections, and finite gate fidelities limit the depth of executable quantum circuits, thereby affecting the accuracy of state preparation, Hamiltonian evolution, and measurement outcomes. Consequently, the implementation of quantum simulation algorithms on near-term hardware requires a careful balance between algorithmic complexity and hardware constraints [9]. For this reason, many NISQ algorithms are designed around shallow circuit architectures and are often complemented by error-mitigation techniques [10], which suppress the effect of noise in measurement outcomes without providing full fault tolerance.

Achieving quantitatively accurate simulations of large, strongly correlated quantum systems is generally expected to require fault-tolerant quantum computation. Fault tolerance is based on the threshold theorem [11], which states that if the error rate of elementary quantum operations lies below a certain critical threshold, then quantum error correction schemes can be implemented to suppress logical errors arbitrarily well, effectively enabling exponential suppression of errors with computational time. Typically, fault tolerance is achieved by encoding logical qubits into multiple physical qubits using quantum error-correcting codes, enabling arbitrarily long computations with controlled error rates at the cost of substantial resource overheads. Within this framework, a critical distinction emerges between Clifford gates, which can be implemented fault-tolerantly with relatively low overhead, and non-Clifford gates, which introduce severe resource bottlenecks but are strictly necessary for universal quantum computation. Universal quantum computation, defined as the ability to approximate any unitary operation on a quantum system to arbitrary accuracy, requires the inclusion of at least one non-Clifford gate, most commonly the T -gate. Fault-tolerant implementation of the T -gate is significantly more resource-intensive than that of Clifford gates and typically relies on magic-state distillation [12]. As a result, the overall cost of fault-tolerant quantum algorithms is often dominated by the number of required T -gates. Consequently, the achievable rate of high-fidelity T -gate execution constitutes a central performance metric for large-scale quantum computers. Resource-estimation studies based on leading quantum error-correcting codes, most prominently the surface code [13, 14], indicate that the

overall execution rate of fault-tolerant quantum algorithms is primarily determined by the logical clock speed and magic state production rate rather than the underlying hardware architecture. In particular, analyses of surface-code implementations and magic state factories suggest sustained T -gate throughput in the range of 10^4 – 10^6 operations per second under realistic physical error rates and code distances [15, 16]. These estimates imply that algorithms requiring between 10^7 and 10^9 T -gates correspond to wall-clock runtimes ranging from minutes to hours on fully fault-tolerant architectures [14, 16].

The primary objective of this work is to advance the development and optimization of quantum algorithms for quantum simulation, addressing efficiency from both asymptotic algorithmic and hardware-resource perspectives. From the algorithmic viewpoint, the focus is on minimizing the fundamental query or step complexity required to achieve a target objective. From the hardware-oriented perspective, the analysis incorporates explicit constraints such as circuit depth and gate counts, critical for NISQ devices, as well as the logical T -gate cost, which dominates the resource profile of fault-tolerant architectures.

The thesis is organized as follows. Part I provides an overview of the main notions in quantum mechanics and quantum computing needed in the rest of this thesis. The remainder of the manuscript is divided into three parts that mirror the structure of a typical quantum simulation. Part II presents a comprehensive overview of quantum state-preparation techniques and introduces both optimized variants of existing algorithms and a novel method proposed by the author. Part III presents a new theoretical result originally developed to optimize Hamiltonian evolution, which is subsequently shown to apply to the broader class of all randomized quantum algorithms. Finally, Part IV focuses on observable measurement, reviewing standard techniques and describing the experimental implementation of a recent method on real digital quantum hardware.




Figure 1: Structural overview of the thesis organization across quantum simulation steps.

PART I

FUNDAMENTALS OF DIGITAL QUANTUM COMPUTING

The development of quantum computing relies on the combination of the principles of quantum mechanics with the concepts of information processing and computation. This part introduces the fundamental theoretical framework underlying digital quantum computation, starting from the mathematical formulation of quantum theory and progressively moving towards the circuit model of computation. In the first chapter, we review the postulates of quantum mechanics and the mathematical structures required to describe quantum systems. We begin with the pure-state formulation, where physical states are represented as rays in Hilbert spaces, and discuss the principles governing composition, evolution, observables, and measurements. We then extend this framework to the more general density-operator formalism, which provides the appropriate description of mixed states, statistical uncertainty, and quantum subsystems. Building upon this foundation, the subsequent chapter specializes these concepts to the framework of digital quantum computing. We introduce the basic elements of the gate-based model, including qubits, quantum registers, elementary quantum gates, quantum circuits, and the notion of universal quantum computation. These concepts establish the operational language used to describe quantum algorithms and quantum information processing. Finally, we examine the Quantum Fourier Transform (QFT), one of the central primitives of quantum computation. Beyond its role as a quantum analogue of the classical Fourier transform, the QFT provides a fundamental mechanism for extracting phase information and constitutes a key component of several important quantum algorithms, including quantum phase estimation and Shor's factoring algorithm. The material presented in this part establishes the conceptual and mathematical basis for the study of quantum algorithms and quantum simulation methods developed in the following sections of this thesis.

CHAPTER 1

POSTULATES OF QUANTUM MECHANICS

Quantum computing is formulated within the mathematical framework of quantum mechanics. In its most elementary description, the state of an isolated quantum system is represented by a ray in a complex Hilbert space. This pure-state formalism, however, is not sufficient to describe situations involving statistical uncertainty, open-system dynamics, or incomplete knowledge about state preparation. In such contexts, the state description must be generalized to a density operator. In this chapter, we review the standard postulates of quantum mechanics, starting from the pure-state formalism and subsequently extending the discussion to mixed states and the density-operator framework.

1.1 Pure-state description

We start by presenting the postulates of quantum mechanics in the pure-state formulation.

Postulate I To every quantum system there corresponds a Hilbert space $\mathcal{H}$. A physical state is represented by a ray in $\mathcal{H}$, defined as

$$|\psi\rangle := \{\phi \in \mathcal{H} \mid \exists \alpha \in \mathbb{C} \text{ such that } \phi = \alpha\psi\}, \quad (1.1)$$

i.e., the equivalence class of all vectors proportional to a given nonzero vector $\psi \in \mathcal{H}$. A representative of the ray is typically chosen to be normalized, such that

$$\| |\psi\rangle \| := \sqrt{\langle \psi | \psi \rangle} = 1. \quad (1.2)$$

Postulate II Given two quantum systems associated with Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, the composite system is described by the tensor product space $\mathcal{H}_A \otimes \mathcal{H}_B$. Pure states

of the composite system are rays in $\mathcal{H}_A \otimes \mathcal{H}_B$. In particular, product states are of the form

$$|\psi\rangle = |\psi\rangle_A \otimes |\phi\rangle_B, \quad (1.3)$$

with $|\psi\rangle_A \in \mathcal{H}_A$ and $|\phi\rangle_B \in \mathcal{H}_B$. More generally, there exist states in $\mathcal{H}_A \otimes \mathcal{H}_B$ that cannot be written in this form; such states are called “entangled”.

Postulate III The time evolution of a closed quantum system from time t_0 to time t is described by a unitary operator $U(t, t_0)$ such that

$$|\psi(t)\rangle = U(t, t_0) |\psi(t_0)\rangle. \quad (1.4)$$

Unitarity ensures preservation of the inner product and therefore of normalization. When the dynamics is generated by a (possibly time-dependent) Hamiltonian $H(t)$, the state satisfies the Schrödinger equation

$$i \frac{d}{dt} |\psi(t)\rangle = H(t) |\psi(t)\rangle, \quad (1.5)$$

and the corresponding evolution operator is given by

$$U(t, t_0) = \text{T exp} \left[-i \int_{t_0}^t dt' H(t') \right], \quad (1.6)$$

where T denotes the time-ordering operator.

Postulate IV Every physical observable is represented by a Hermitian operator $O : \mathcal{H} \rightarrow \mathcal{H}$. Consequently, O admits a spectral decomposition of the form

$$O = \sum_j o_j \Pi_j, \quad o_j \in \mathbb{R}, \quad (1.7)$$

where $\{\Pi_j\}$ is a set of orthogonal projectors satisfying

$$\Pi_j \geq 0 \quad \forall j, \quad \Pi_j \Pi_n = \delta_{jn} \Pi_j \quad \forall j, n, \quad \sum_j \Pi_j = I. \quad (1.8)$$

Postulate V The only possible outcomes of a measurement of an observable $O = \sum_j o_j \Pi_j$ on a state $|\psi\rangle$ are the eigenvalues $\{o_j\}$. The probability of obtaining outcome o_j is given by

$$p_j = \text{Tr} [|\psi\rangle\langle\psi| \Pi_j]. \quad (1.9)$$

Postulate VI Upon obtaining outcome o_j in a projective measurement of $O = \sum_j o_j \Pi_j$, the state updates according to

$$|\psi\rangle \longrightarrow \frac{\Pi_j |\psi\rangle}{\sqrt{\text{Tr} [|\psi\rangle\langle\psi| \Pi_j]}}. \quad (1.10)$$

1.2 Density-operator description

The pure-state formalism is insufficient to describe situations in which the preparation of a quantum system is not perfectly known or when the system under study is part of a larger composite system. In such cases, the state must be described in a more general way that accounts for classical uncertainty and correlations with external degrees of freedom. This leads to the density-operator formalism, which extends the notion of a quantum state while remaining consistent with the operational predictions of quantum mechanics.

Postulate I To every quantum system there corresponds a Hilbert space $\mathcal{H}$. A physical state is represented by a density operator ρ , i.e., a linear operator $\rho : \mathcal{H} \rightarrow \mathcal{H}$ such that

$$\rho = \rho^\dagger, \quad \rho \geq 0, \quad \text{Tr}(\rho) = 1. \quad (1.11)$$

Pure states are recovered as the special case $\rho = |\psi\rangle\langle\psi|$, with $\langle\psi|\psi\rangle = 1$.

A simple example of a situation where the state is not pure is given by an ensemble of preparations in which, at each run, the system is prepared in one of the pure states $|\psi_j\rangle$ with classical probabilities p_j , where $p_j \geq 0$ and $\sum_j p_j = 1$. The corresponding density operator is then defined by

$$\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|, \quad (1.12)$$

which describes statistical uncertainty over the preparation procedure rather than quantum superposition, and is called a mixed state. A particularly important case in thermodynamics is the Gibbs state, in which a system with Hamiltonian H at inverse temperature β is described by $\rho_\beta = e^{-\beta H}/Z$, with $Z = \text{Tr}[e^{-\beta H}]$, representing the thermal equilibrium mixture over energy eigenstates.

Remarkably, in the density-operator representation, the physical state is invariant under phase transformations of the underlying vectors, since

$$|\psi_j\rangle \longrightarrow e^{i\theta_j} |\psi_j\rangle \quad (1.13)$$

leaves the associated density operator unchanged,

$$\rho \longrightarrow \rho. \quad (1.14)$$

Thus, the density-operator formalism incorporates from the outset the equivalence of vectors differing by a global phase.

Postulate II Given two systems associated with $\mathcal{H}_A$ and $\mathcal{H}_B$, the composite system is described by $\mathcal{H}_A \otimes \mathcal{H}_B$. States of the composite system are density operators ρ acting on $\mathcal{H}_A \otimes \mathcal{H}_B$. Product states take the form

$$\rho = \rho_A \otimes \rho_B, \quad (1.15)$$

while more general states cannot be written in this form.

Postulate III The time evolution of a closed quantum system is described by a unitary transformation. If the state at time t_0 is $\rho(t_0)$, then at time t it is given by

$$\rho(t) = U(t, t_0) \rho(t_0) U^\dagger(t, t_0), \quad (1.16)$$

where $U(t, t_0)$ is unitary. If the evolution is generated by a Hamiltonian $H(t)$, the density operator satisfies

$$i \frac{d\rho}{dt} = [H(t), \rho(t)]. \quad (1.17)$$

Postulate IV Every physical observable is represented by a Hermitian operator $O : \mathcal{H} \rightarrow \mathcal{H}$. Consequently, O admits a spectral decomposition of the form

$$O = \sum_j o_j \Pi_j, \quad o_j \in \mathbb{R}, \quad (1.18)$$

where $\{\Pi_j\}$ is a set of orthogonal projectors satisfying

$$\Pi_j \geq 0 \quad \forall j, \quad \Pi_j \Pi_n = \delta_{jn} \Pi_j \quad \forall j, n, \quad \sum_j \Pi_j = I. \quad (1.19)$$

Postulate V A projective measurement associated with an observable $O = \sum_j o_j \Pi_j$ yields outcome o_j with probability

$$p_j = \text{Tr}[\rho \Pi_j]. \quad (1.20)$$

Postulate VI Conditioned on obtaining outcome o_j , the post-measurement state is given by

$$\rho \longrightarrow \frac{\Pi_j \rho \Pi_j}{\text{Tr}[\rho \Pi_j]}. \quad (1.21)$$

The postulates presented above describe closed, ideal quantum systems whose dynamics are strictly unitary. In practice, however, physical quantum computers are only approximately closed systems. Unavoidable interactions with the surrounding environment, control imperfections, and unwanted cross-talk introduce noise, decoherence, and dissipation. Consequently, describing real-world quantum hardware requires extending the concept of unitary time evolution to open quantum systems.

The most general physical transformation that a quantum state can undergo in an open system is modeled by a *quantum channel*. Formally, a quantum channel is a linear map $\mathcal{E} : \mathcal{L}(\mathcal{H}_A) \rightarrow \mathcal{L}(\mathcal{H}_B)$ that maps linear operators on an input Hilbert space $\mathcal{H}_A$ to linear operators on an output Hilbert space $\mathcal{H}_B$. To represent a physically valid state transformation, $\mathcal{E}$ must satisfy two fundamental properties:

1. **Trace preservation:** $\text{Tr}[\mathcal{E}(\rho)] = \text{Tr}[\rho] = 1$ for all $\rho \in \mathcal{L}(\mathcal{H}_A)$, ensuring that total probability is strictly conserved.

2. **Complete positivity:** $(\mathcal{E} \otimes I_R)(\rho_{AR}) \geq 0$ for any valid density operator ρ_{AR} defined on a composite space $\mathcal{H}_A \otimes \mathcal{H}_R$, where I_R is the identity map on an arbitrary auxiliary system R .

Physically, complete positivity ensures that $\mathcal{E}$ transforms valid states into valid, positive-semidefinite density operators even when system A is entangled with an untouched reference system R . Quantum channels provide the overarching mathematical framework for open-system dynamics, encompassing noisy gate operations, decoherence processes, and state-preparation errors resulting from tracing out an unobserved environment.

In this chapter, we have established the foundational mathematical framework governing quantum systems. Beginning with isolated systems in the pure-state formulation, we reviewed the postulates handling state representation, unitary dynamics, observables, and projective measurements. To account for incomplete knowledge, subsystems, and environmental interactions, we then extended this description to density operators and quantum channels. This unified formalism treats pure states, statistical mixtures, open-system noise, and non-classical correlations such as entanglement on an equal mathematical footing. Together, these principles establish the theoretical basis for quantum information processing, where physical states serve as information carriers, transformations represent computational operations, and measurements extract classical outcomes. In the next chapter, we specialize this abstract framework to finite-dimensional Hilbert spaces, introducing the operational building blocks of digital quantum computing: qubits, quantum gates, and circuit architectures.

CHAPTER 2

BASICS OF QUANTUM COMPUTING

The quantum computing paradigm most commonly adopted in practice is the digital, or gate-based, model of computation. In this framework, information is encoded in discrete quantum systems, typically qubits, and processed by applying sequences of elementary unitary operations. This approach is the direct quantum analogue of classical digital computation, in which bits are manipulated by logical gates arranged into circuits. The central difference is that, in the quantum case, the admissible transformations are constrained by unitarity and can exploit superposition, interference, and entanglement. This chapter introduces the basic building blocks of digital quantum computing and the notion of universal quantum computation.

2.1 Qubits and the computational basis

The fundamental unit of digital quantum information is the qubit. A qubit is a two-level quantum system whose state belongs to a two-dimensional Hilbert space $\mathcal{H}_2$. A convenient orthonormal basis of this space is the computational basis, denoted by

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (1.22)$$

Any pure qubit state can be written as a linear combination of these basis vectors:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \quad (1.23)$$

The coefficients α and β are complex amplitudes, and the normalization condition ensures that measurement outcomes are properly normalized.

Upon measuring the qubit in the computational basis, the outcome 0 is obtained with probability $|\alpha|^2$ and the outcome 1 with probability $|\beta|^2$. The post-measurement state collapses to the corresponding basis vector. This simple example already displays a central feature of quantum computation: the output of a quantum algorithm is not read out by directly inspecting the state, but by converting amplitudes into classical probabilities through measurement.




Figure 2.1: Bloch-sphere representation of a single-qubit pure state $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$. The angle θ denotes the polar angle measured with respect to the z axis, while φ denotes the azimuthal angle in the xy plane measured from the x axis. The north and south poles correspond to the computational basis states $|0\rangle$ and $|1\rangle$, which are obtained for $\theta = 0$ and $\theta = \pi$, respectively.

A useful geometric representation of single-qubit pure states is provided by the Bloch sphere. Any normalized pure state of a single qubit can be written, up to an irrelevant global phase, as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\varphi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \quad (1.24)$$

where $\theta \in [0, \pi]$ and $\varphi \in [0, 2\pi)$. The parameters (θ, φ) define a point on the unit sphere, known as the Bloch sphere, according to the correspondence

$$(x, y, z) = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta). \quad (1.25)$$

In this representation, the computational basis states $|0\rangle$ and $|1\rangle$ correspond to the north and south poles of the sphere, respectively, while superposition states occupy points on its surface. Single-qubit unitary gates generated by the Pauli operators admit a direct geometric interpretation as rotations of the Bloch vector around the corresponding axes. Figure 2.1 illustrates the Bloch-sphere representation of a single-qubit state and the angular parameters θ and φ .

2.2 Quantum registers

A quantum register of n qubits is described by the tensor product space

$$\mathcal{H}_2^{\otimes n}. \quad (1.26)$$

Its computational basis is formed by the 2^n vectors

$$|x\rangle = |x_1 x_2 \cdots x_n\rangle, \quad x_j \in \{0, 1\}, \quad (1.27)$$

where x is the binary string associated with the integer $0 \leq x < 2^n$. A general pure state of an n -qubit register is therefore

$$|\psi\rangle = \sum_{x=0}^{2^n-1} \alpha_x |x\rangle, \quad \sum_{x=0}^{2^n-1} |\alpha_x|^2 = 1. \quad (1.28)$$

The state space grows exponentially with the number of qubits, which is one of the reasons quantum systems can exhibit computational behavior that is difficult to simulate classically. The basis vectors $|x\rangle$ are often referred to as classical configurations of the register. However, unlike classical bits, a quantum register can exist in coherent superpositions of many configurations simultaneously, and the relative phases between amplitudes are physically meaningful. The ability to coherently manipulate relative phases is precisely the resource exploited by quantum algorithms that exhibit provable advantages over their classical counterparts, such as Grover's search algorithm, amplitude amplification techniques, and the Quantum Fourier Transform, which will be discussed in detail in the following chapter.

2.3 Elementary quantum gates

The processing of quantum information is implemented by quantum gates, namely unitary operators acting on one or more qubits. Since unitary operators preserve inner products, they preserve normalization and ensure reversible time evolution at the level of the computational model. Single-qubit gates can be represented by 2×2 unitary matrices. Common examples include the Pauli operators, whose action on the computational basis reads:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (1.29)$$

as well as the Hadamard gate,

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \quad (1.30)$$

The Pauli- X gate acts as a NOT gate on the computational basis, exchanging $|0\rangle$ and $|1\rangle$. The Hadamard gate creates and removes equal superpositions, since

$$H|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \quad H|1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}. \quad (1.31)$$

This operation is fundamental in quantum algorithms because it converts localized basis states into coherent superpositions.

A particularly important class of single-qubit operations consists of continuous rotations generated by the Pauli operators. Rotations around the Cartesian axes of the Bloch sphere are defined as

$$R_X(\theta) = e^{-i\theta X/2} = \begin{bmatrix} \cos(\theta/2) & -i \sin(\theta/2) \\ -i \sin(\theta/2) & \cos(\theta/2) \end{bmatrix}, \quad (1.32)$$

$$R_Y(\theta) = e^{-i\theta Y/2} = \begin{bmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{bmatrix}, \quad (1.33)$$

$$R_Z(\theta) = e^{-i\theta Z/2} = \begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}. \quad (1.34)$$

These gates admit a natural geometric interpretation as rotations of the qubit state on the Bloch sphere around the x , y , and z axes, respectively [17]. In particular, rotations around the z axis modify the relative phase between computational basis amplitudes. For this reason they are also called phase gates and they play a central role in interference-based quantum algorithms, such as those formalized within the architecture of quantum signal processing [18]. Analogous to classical 3D rotations, any arbitrary single-qubit unitary transformation can be decomposed into a sequence of three elementary rotations around orthogonal axes, typically referred to as the Euler decomposition of the $SU(2)$ group [17].

On the other hand, multi-qubit gates are those acting on two or more qubits and are essential for generating entanglement. The paradigmatic example is the controlled-NOT gate, or CNOT, acting on two qubits and defined by its action on the computational basis

$$|a, b\rangle \mapsto |a, a \oplus b\rangle, \quad a, b \in \{0, 1\}, \quad (1.35)$$

where $\oplus$ denotes addition modulo 2. The first qubit is the control and the second is the target. The CNOT gate flips the target if and only if the control qubit is in state $|1\rangle$, and its standard circuit representation is




Figure 2.2: Circuit diagram of the two-qubit controlled-NOT (CNOT) gate.

Controlled gates are indispensable because purely local operations on separate qubits cannot create entanglement. For instance, applying a Hadamard gate to the first qubit of the two-qubit state $|00\rangle$ and then a CNOT yields the Bell state

$$|\Phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}, \quad (1.36)$$

which cannot be factorized into a tensor product of single-qubit states. Entanglement is one of the central resources that distinguish quantum information processing from its classical counterpart.

2.4 Universal quantum computation

The notion of universality is central to quantum computing. A set of gates is called universal if any unitary transformation on an arbitrary number of qubits can be approximated to arbitrary precision by a finite sequence of gates from that set. Universality is therefore the quantum analogue of computational completeness in classical logic. A standard universal gate set is given by the collection $\{H, T, \text{CNOT}\}$, where T is the $\pi/8$ phase gate,

$$T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \quad (1.37)$$

Together with the CNOT gate, the gates H and T generate a dense subset of the unitary group on multiple qubits. Consequently, any quantum algorithm can be approximated with arbitrary accuracy by circuits built from these elementary operations.

More generally, any gate set that can approximate an arbitrary single-qubit unitary and contains at least one entangling two-qubit gate is universal. This statement reflects a fundamental structural result: single-qubit gates provide arbitrary local rotations on the Bloch sphere, while entangling gates supply the non-classical correlations required for genuinely quantum computation [17, 19, 20].

This structural universality can be understood through a two-tier conceptual decomposition. First, any arbitrary single-qubit unitary can be expressed, up to a global phase, as a sequence of rotations around two nonparallel axes. Second, any multi-qubit unitary can be mapped exactly to a circuit consisting entirely of single-qubit gates and CNOT gates. While exact decompositions may demand an impractical number of gates, fault-tolerant quantum computing requires only that we can construct approximations to within an arbitrarily small error ϵ . In this constructive setting, the landmark Solovay–Kitaev theorem demonstrated that any single-qubit unitary can be ϵ -approximated using a sequence of $\mathcal{O}(\ln^c(1/\epsilon))$ gates drawn from a finite, fault-tolerant set. While the original proof yielded an exponent of $c \approx 3.97$, subsequent synthesis algorithms optimized this scaling to the theoretical limit of $c = 1$ for prominent universal sets like Clifford+ T [21, 22].

Universality implies that quantum computers are programmable in a gate-by-gate fashion, independently of the particular algorithm to be executed. A physical device need only implement a small set of elementary operations with sufficient fidelity; more complicated transformations can then be synthesized from them. This modularity is the operational foundation of digital quantum computing. It is also the reason why many different physical platforms can realize the same computational model. Superconducting circuits, trapped ions, neutral atoms, photonic systems, and spin qubits may differ substantially at the hardware level, but they can all serve as universal quantum processors provided they support a suitable universal gate set with controllable measurements [17].




Figure 2.3: Schematic representation of a quantum circuit. Starting from an initial 4-qubit state, typically $|0\rangle^{\otimes 4}$, a sequence of unitary gates $U_1, U_2, \dots, U_m$ is applied, implementing the overall transformation $U = U_m \cdots U_2 U_1$. The computation ends with projective measurements in the computational basis, yielding the classical output of the algorithm.

2.5 Quantum circuits

A quantum algorithm is usually represented as a quantum circuit. In this representation, qubits are drawn as horizontal wires, and gates are applied sequentially along the wires. A circuit specifies a unitary transformation that acts on an initial state, typically chosen to be a computational basis state such as $|0\rangle^{\otimes n}$. If a circuit applies the gates $U_1, U_2, \dots, U_m$ in sequence, as illustrated in Figure 2.3, then the overall action is

$$U = U_m \cdots U_2 U_1. \quad (1.38)$$

Because matrix multiplication is noncommutative, the order of the gates matters. In this sense, a quantum circuit is a structured prescription for constructing a desired global unitary from simple local operations. After the circuit has been applied, one measures one or more qubits, typically in the computational basis. The resulting classical bit string is interpreted as the output of the algorithm. This separation between coherent unitary processing and final projective readout is a defining feature of the circuit model.

The next chapter is dedicated to one of the most fundamental quantum algorithms, the Quantum Fourier Transform (QFT). The QFT constitutes a central building block of quantum computation and appears as a core subroutine in several quantum algorithms exhibiting provable quantum advantage, including Shor's factoring algorithm, the amplitude-estimation routine, and the phase-estimation protocol, which we discuss in Chapter 8. Beyond its algorithmic relevance, the QFT also provides a paradigmatic example of how quantum interference and coherent phase manipulation can be exploited within the circuit model to process information in ways that are inaccessible to classical computation.

CHAPTER 3

QUANTUM FOURIER TRANSFORM

The quantum Fourier transform (QFT) is one of the most important and widely used subroutines in quantum computation. It is the quantum analogue of the Discrete Fourier Transform (DFT) and plays a central role in several landmark algorithms, including phase estimation, order finding, and Shor's factoring algorithm. Its importance stems from the fact that it converts phase information encoded in the amplitudes of quantum states into a form that can be accessed through measurement with high efficiency.

This chapter introduces the QFT in a systematic way. We begin with its mathematical definition, then discuss its action on basis states and on general superpositions. We next derive its circuit decomposition in terms of elementary gates, analyze its computational complexity, and finally explain its role in quantum algorithms.

3.1 The Discrete Fourier Transform

The DFT on $\mathbb{C}^N$ is the linear map that sends a vector of amplitudes $\{x_j\}_{j=0}^{N-1}$ to the vector $\{y_k\}_{k=0}^{N-1}$ defined by

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} e^{2\pi i j k / N} x_j. \quad (1.39)$$

In matrix form, the DFT is represented by the unitary matrix F_N with entries

$$[F_N]_{kj} = \frac{1}{\sqrt{N}} e^{2\pi i j k / N}. \quad (1.40)$$

The unitarity of F_N follows from the orthogonality of complex exponentials.

The quantum Fourier transform is precisely the action of this unitary map on the computational basis of an n -qubit register.

3.2 The Quantum Fourier Transform (QFT)

The QFT on $N = 2^n$ computational basis states is defined by

$$\mathcal{F}_N |x\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{2\pi i x k / N} |k\rangle, \quad x \in \{0, 1, \dots, N-1\}. \quad (1.41)$$

By linearity, its action on an arbitrary state

$$|\psi\rangle = \sum_{x=0}^{N-1} \alpha_x |x\rangle \quad (1.42)$$

is given by

$$\mathcal{F}_N |\psi\rangle = \sum_{k=0}^{N-1} \left(\frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} e^{2\pi i x k / N} \alpha_x \right) |k\rangle. \quad (1.43)$$

Thus the QFT maps the amplitude vector $\{\alpha_x\}$ to its Fourier-transformed version, up to normalization. To understand the structure of the QFT, it is convenient to write the integer x in binary form:

$$x = x_{n-1}2^{n-1} + x_{n-2}2^{n-2} + \dots + x_12 + x_0, \quad x_j \in \{0, 1\}. \quad (1.44)$$

Using $N = 2^n$, one can rewrite the QFT of a basis state $|x\rangle$ as

$$\mathcal{F}_N |x\rangle = \frac{1}{\sqrt{N}} \bigotimes_{j=0}^{n-1} (|0\rangle + e^{2\pi i 0.x_j x_{j-1} \dots x_0} |1\rangle), \quad (1.45)$$

where

$$0.x_j x_{j-1} \dots x_0 = \frac{x_j}{2} + \frac{x_{j-1}}{2^2} + \dots + \frac{x_0}{2^{j+1}} \quad (1.46)$$

is the binary fraction formed by the least significant bits of x . This factorized form is highly informative. It shows that the QFT maps a computational basis state to a tensor product of single-qubit superpositions with phases determined by the binary expansion of x . The appearance of these phases is the reason why the QFT is useful in problems where periodicity or hidden phase structure must be extracted. A subtle point is that the qubit order is reversed in the standard circuit decomposition. As a result, the output of the circuit implements the QFT up to a final bit-reversal permutation. Since the QFT is unitary, its inverse is simply its adjoint:

$$\mathcal{F}_N^{-1} = \mathcal{F}_N^\dagger. \quad (1.47)$$

Its action on basis states is

$$\mathcal{F}_N^{-1} |x\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{-2\pi i x k / N} |k\rangle. \quad (1.48)$$

The inverse QFT is frequently used in algorithms in which a quantum state acquires phase information during an evolution and that information must then be converted

into a measurable computational basis outcome. Quantum phase estimation, to which Chapter 8 is dedicated, is the canonical example.

The main practical value of the QFT is that it can be implemented efficiently. The standard decomposition uses Hadamard gates and controlled phase rotations.

Let us define the controlled phase gate

$$R_m = \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^m} \end{bmatrix}, \quad (1.49)$$

and the controlled version of R_m , which applies this phase to a target qubit when the control qubit is in state $|1\rangle$.

For a register of n qubits, the QFT circuit proceeds as follows. Starting from the most significant qubit, one applies a Hadamard gate and then a sequence of controlled phase rotations from the less significant qubits onto it. This pattern is repeated recursively on the remaining qubits. Finally, the order of the qubits is reversed.

For example, on three qubits, the circuit structure is

$$|x_2 x_1 x_0\rangle \mapsto (H \otimes I \otimes I) \cdot (\text{controlled phases}) \cdots \quad (1.50)$$

with the explicit controlled rotations arranged so that qubit 0 accumulates the phases contributed by qubits 1 and 2, qubit 1 accumulates the phase from qubit 2, and qubit 2 receives only a Hadamard transformation.

More explicitly, the QFT circuit on n qubits can be written schematically as

$$\mathcal{F}_{2^n} = \text{SWAP} \cdot \left(H_0 \prod_{j=1}^{n-1} \text{CR}_{j,0}(R_{j+1}) \right) \cdot \left(H_1 \prod_{j=2}^{n-1} \text{CR}_{j,1}(R_j) \right) \cdots H_{n-1}, \quad (1.51)$$

where $\text{CR}_{j,k}(R_m)$ denotes a controlled- R_m gate with control qubit j and target qubit k , and SWAP denotes the final reversal of qubit order. The exact indexing depends on convention, but the underlying structure is always the same.

3.3 Computational complexity

While a direct implementation of the DFT on a classical computer requires $\mathcal{O}(N^2) = \mathcal{O}(4^n)$ operations (and even the classical Fast Fourier Transform requires $\mathcal{O}(N \log N) = \mathcal{O}(n2^n)$ operations), the QFT can be implemented with a number of elementary gates that scales quadratically in n . Indeed, the standard QFT circuit uses one Hadamard gate per qubit and, for each qubit, a number of controlled phase gates proportional to the number of lower-order qubits. The total gate count is therefore $\mathcal{O}(n^2)$. This exponential speedup relative to classical Fourier transforms in terms of the register size is one of the reasons the QFT is so powerful. Moreover, in many practical applications, an approximate quantum Fourier transform suffices [23]. Because the controlled phase-shift angles decrease exponentially with the distance between qubits, the smallest controlled rotations become exponentially close to the identity operator. Consequently, one may truncate these small-angle gates below a given precision threshold with negligible impact on the overall fidelity of the final state [24]. For

a fixed target error, this truncation bounds the number of controlled rotations per qubit, effectively reducing the total circuit size to $\mathcal{O}(n \log n)$. This logarithmic gate overhead renders the approximate QFT highly advantageous in algorithmic settings where exact phase resolution is not strictly required, such as in certain instances of phase estimation or robust period-finding.

3.4 Periodic states and the interpretation of the QFT

The QFT is particularly effective when acting on states with periodic structure. Consider a state of the form

$$|\psi\rangle = \frac{1}{\sqrt{M}} \sum_{r=0}^{M-1} |x_0 + rT\rangle, \quad (1.52)$$

where the integer values $x_0 + rT$ form an arithmetic progression with period T . Applying the QFT to such a state produces a distribution sharply peaked at frequencies that are multiples of N/T , up to discretization effects.

This property is the core mechanism behind period-finding algorithms. The QFT transforms hidden periodicity in the computational basis into concentration of probability mass in the Fourier basis. Measurement of the transformed state then reveals information about the period. The most celebrated application of the QFT is Shor's factoring algorithm [25]. There, the problem of factoring an integer is reduced to finding the period of a modular exponentiation function. The QFT is used to extract this period efficiently from a superposition over function values. The algorithm proceeds by preparing a uniform superposition, evaluating the modular function in a second register, measuring the second register, and applying the inverse QFT to the first register. The resulting measurement statistics reveal the order of a number modulo N , from which the factors of N can be recovered with classical post-processing [17]. Without the QFT, this exponential quantum speedup would not be available. The QFT is therefore a central structural component of one of the most important quantum algorithms known.

PART II

QUANTUM-STATE PREPARATION

Quantum state preparation refers to the task of initializing a quantum register of L qubits in a target state of interest, starting from a simple and experimentally accessible reference state, typically the computational basis state $|0\rangle^{\otimes L}$. The target state may be, for instance, an eigenstate of the Hamiltonian of the system under study, a controlled superposition of several eigenstates, or even a mixed state such as a Gibbs state. A defining feature of quantum simulation is that the target state is not assumed to be known analytically. If an explicit classical description of the state were available, the motivation for performing a quantum simulation would largely vanish. Instead, the objective is to generate the desired state implicitly by exploiting the physical structure or the dynamics of the system under investigation.

Quantum state preparation is widely regarded as one of the principal bottlenecks in quantum simulation [17, 26]. For a broad class of applications, including ground-state energy estimation, equilibrium properties, and static correlation functions, the preparation of the target quantum state constitutes the dominant computational task, after which the desired observables can be obtained through direct measurements without further time evolution. Consequently, the overall efficiency of these quantum simulation algorithms is often determined primarily by the complexity of the state preparation procedure.

From the perspective of computational complexity theory, quantum state preparation is intractable in full generality. Closely related decision and construction problems, such as preparing the ground state of a local Hamiltonian or determining its properties, are known to be Quantum Merlin–Arthur (QMA)-complete. This was first established in the seminal work of Kitaev on the local Hamiltonian problem and later refined and extended by Kempe, Kitaev, and Regev, among others [27, 28]. These results imply that, under standard complexity-theoretic assumptions, no polynomial-time quantum algorithm is expected to efficiently prepare ground states or arbitrary target states for generic Hamiltonians. As a result, quantum simulations cannot be expected to be polynomially efficient for arbitrary instances.

However, the physically relevant problems encountered in practice are rarely generic. In areas such as quantum chemistry and condensed matter physics, Hamiltonians typically exhibit strong structure, including locality, symmetries, conserved quantities, or a restricted low-energy sector. In quantum field theories, symmetry principles play an even more fundamental role, as they constitute the starting point for the formulation of the theory itself. These structural features can significantly reduce the effective complexity of the state preparation problem and often enable efficient or approximate preparation strategies, despite negative worst-case complexity results.

Furthermore, the states of interest in quantum simulation are typically not arbitrary vectors in the exponentially large Hilbert space, but rather physically relevant states characterized by specific structural properties. In most applications, the goal is to prepare ground states or low-lying excited states of local Hamiltonians, which often exhibit strong constraints arising from locality, symmetries, conservation laws, and their low-energy nature. As an example, ground states of local gapped Hamiltonians are known to display restricted entanglement patterns, such as area-law scaling of the entanglement entropy in many relevant settings. This property reflects the fact

that physically meaningful states occupy only a small and highly structured subset of the full Hilbert space and has been exploited extensively in classical approaches, such as tensor network methods. Similarly, quantum algorithms for state preparation aim to identify and leverage these underlying structures to efficiently approximate relevant quantum states, overcoming the limitations implied by worst-case complexity results. Therefore, the central challenge in quantum simulation is not the preparation of arbitrary quantum states, but the efficient preparation of structured states arising from physically relevant Hamiltonians.

Having outlined the conceptual and complexity-theoretic landscape, we now discuss concrete quantum state preparation techniques, starting with methods for preparing specific eigenstates of the system Hamiltonian H . Current noisy intermediate-scale quantum (NISQ) devices motivated the development of strategies compatible with constrained quantum resources. Variational Quantum Algorithms (VQAs) provide one such approach, combining parametrized quantum circuits with classical optimization [29, 30]. While they can approximate complex target states with reduced circuit depth, VQAs face scalability limitations due to ansatz design requirements, high-dimensional optimization, and phenomena such as barren plateaus, which hinder convergence in large systems [31]. Adiabatic Preparation (AP) offers a conceptually different route [32], relying on slowly evolving a simple initial state under a time-dependent Hamiltonian toward the desired eigenstate of the target Hamiltonian. Despite its conceptual simplicity, AP is highly sensitive to the presence of small spectral gaps or level crossings, which are common in many-body systems. These crossings limit the ability of conventional adiabatic protocols to selectively prepare specific eigenstates, particularly excited states, thereby constraining the range of states that can be reliably accessed [33].

This part is structured into three chapters. Chapter 4 provides a detailed overview of the techniques mentioned above, with particular emphasis on their respective strengths and limitations. Chapter 5 presents an original improvement of the Adiabatic Preparation technique, developed within the framework of optimal control. Finally, Chapter 6 is devoted to the Branched-Subspaces Adiabatic Preparation (B-SAP) algorithm, a novel quantum-state preparation technique that combines adiabatic preparation and variational methods through group-theoretical considerations.

CHAPTER 4

OVERVIEW OF QUANTUM-STATE PREPARATION METHODS

A large fraction of quantum state preparation problems of practical and theoretical interest can be formulated as the preparation of eigenstates of a Hamiltonian H . This includes, in particular, ground states and low-lying excited states, which play a central role in quantum chemistry, condensed matter physics, and lattice field theories. The emphasis on eigenstates is not incidental: in many settings, physically relevant information is encoded directly in the spectral structure of H , rather than in arbitrary superpositions of its eigenbasis. More specifically, the ground state of H is associated with the minimum energy configuration and therefore encodes the equilibrium properties of interacting quantum systems at zero temperature. Excited eigenstates, on the other hand, govern spectral response, transport phenomena, and dynamical excitations. In many cases, physically meaningful regimes correspond not to arbitrary states in Hilbert space, but to structured subsets of the eigenbasis selected by symmetry, conservation laws, or energy constraints.

For these reasons, the eigenstate preparation problem serves as a canonical benchmark for quantum state preparation techniques. It captures both the algorithmic difficulty of accessing physically relevant quantum states and the structural features that can be exploited to reduce this difficulty in specific problem classes. In the following sections, we review the main approaches to this problem. We begin with Variational Quantum Algorithms (VQAs), which formulate eigenstate preparation as a parametrized optimization problem over quantum circuits. We then discuss Adiabatic Preparation (AP), which instead exploits the dynamical properties of slowly driven quantum systems to project onto eigenstates of interest.

4.1 Variational Quantum Algorithms

Variational quantum algorithms have emerged as a leading approach for approximating eigenstates of quantum many-body systems by combining parametrized quantum circuits with classical optimization [29]. A VQA consists of a parametrized quantum circuit, typically defined by a chosen ansatz structure, which prepares a trial state $|\psi(\theta)\rangle$ starting from a reference state, usually $|0\rangle^{\otimes L}$, as illustrated in Figure 4.1. These parameters are iteratively optimized to minimize a cost function, typically defined in terms of the expectation value of the target Hamiltonian or a related observable. The optimization is performed using classical algorithms, forming a hybrid quantum-classical loop.

For example, when targeting the ground state of a Hamiltonian H , the cost function is typically chosen as the energy expectation value $\langle\psi(\theta)|H|\psi(\theta)\rangle$, whose minimization variationally targets the lowest eigenstate. In contrast, when aiming to prepare a generic eigenstate, one may instead minimize the energy variance $\langle\psi(\theta)|(H - \langle H \rangle)^2|\psi(\theta)\rangle$, which vanishes if and only if the trial state is an exact




Figure 4.1: Sketch of the structure of a hybrid Variational Quantum Algorithm. The quantum register is initialized in the reference state $|0\rangle^{\otimes L}$, upon which ℓ layers of parametrized quantum operations are applied. At the end of the circuit, measurement outcomes are collected and processed by a classical optimizer. The optimizer updates the circuit parameters in order to minimize a given cost function, and the updated parameters are fed back into the quantum circuit. This hybrid quantum-classical optimization loop is repeated until convergence to a minimum of the cost function is achieved.

eigenstate of H . The primary advantage of VQAs is their ability to efficiently explore a subset of the Hilbert space, enabling approximation of complex quantum states with reduced circuit depth compared to exact methods. The flexibility of the circuit ansatz allows tailoring to the structure of the target state, improving performance for specific problems [34, 35].

However, VQAs exhibit several limitations. For instance, although the choice of the loss function is natural in eigenstate preparation tasks, identifying an appropriate objective for other settings can be non-trivial. Consider, for example, the case in which the target state is not a single eigenstate but a specific superposition of Hamiltonian eigenstates. In such situations, it is generally unclear how to define a loss function that both singles out the desired coherent superposition and remains experimentally accessible. Moreover, the choice of ansatz is critical: poorly designed circuits may fail to capture essential correlations of the target state, while circuits implementing fully generic unitary operations would require a number of parameters which is exponentially large in the system size, making optimization infeasible. Furthermore, variational quantum algorithms (VQAs) can suffer from barren plateaus, i.e., regions of the parameter space where gradients vanish exponentially with system size [31]. This can make classical optimization an exponentially hard task, thereby diminishing the advantage of quantum computing over classical computation. Finally, classical optimization challenges, including local minima and vanishing gradients, further hinder scalability and practical implementation on large systems [36].

4.2 Adiabatic Quantum-State Preparation

An alternative strategy for the preparation of eigenstates of a target Hamiltonian H is represented by adiabatic preparation, which relies on the adiabatic theorem [32, 37]. The protocol requires the choice of an auxiliary Hamiltonian H_0 whose eigenstates are easy to prepare. The quantum register, initialized in the default state $|0\rangle^{\otimes L}$, is first prepared in a selected eigenstate of H_0 and subsequently evolved under the parametrized Hamiltonian

$$H^{\text{AP}}(s) = H_0 + f(s)[H - H_0], \quad (2.1)$$

where $s \in [0, 1]$ is a dimensionless parameter and $f(s)$ is a smooth interpolation function satisfying $f(0) = 0$ and $f(1) = 1$ [33]. The evolution is governed by the operator

$$U_\tau(s) = \mathcal{T} \exp \left[-i\tau \int_0^s H^{\text{AP}}(s') ds' \right], \quad (2.2)$$

where τ denotes the total evolution time and $\mathcal{T}$ is the time-ordering operator. The effectiveness of AP depends critically on the spectral properties of $H^{\text{AP}}(s)$. Denoting the instantaneous spectrum by $\sigma(s)$, one may partition it into subsets $\{\sigma_n(s)\}_n$ and define the corresponding projectors $\{P_n(s)\}_n$ onto the corresponding subspaces.




Figure 4.2: Subfigure (a) illustrates the decomposition of the adiabatic spectrum $\sigma(s)$ into distinct subsets $\sigma_n(s)$, separated by finite spectral gaps. Subfigure (b) shows the poles of the resolvent $[H^{\text{AP}}(s) - z]^{-1}$ as a function of z , for a fixed value of s . Empty dots denote the eigenvalues belonging to $\sigma_n(s)$, while filled dots correspond to the remaining eigenvalues in $\sigma(s) \setminus \sigma_n(s)$. Throughout this discussion, we restrict attention to systems with discrete spectra.

Formally, such projectors can be defined as follows:

$$P_n(s) := \frac{1}{2\pi i} \oint_{\Gamma_n(s)} \frac{dz}{H^{\text{AP}}(s) - z}, \quad (2.3)$$

where $\Gamma_n(s)$ is a closed anti-clockwise curve surrounding only the eigenvalues in $\sigma_n(s)$ while keeping a distance of at least $\frac{\Delta_n(s)}{2}$ from the rest of the spectrum, as in Figure 4.2.

The central idea of AP is that a pure quantum state initialized within $\sigma_n(0)$ is approximately evolved by $U_\tau(s)$ in Eq. (2.2) into a state belonging to $\sigma_n(s)$. The error associated with this approximation can be quantified as

$$\mathcal{E}_n^{\text{AP}}(s, \tau) := \|(\mathbb{1} - P_n(s)) U_\tau(s) P_n(0)\|. \quad (2.4)$$

If the minimal spectral gap

$$\Delta_n(s) = \min_{\substack{x \in \sigma_n(s) \\ y \in \sigma(s) \setminus \sigma_n(s)}} |x - y| \quad (2.5)$$

between the subset $\sigma_n(s)$ and the remainder of the spectrum remains strictly positive for all s , the error decays at least as $1/\tau$ [33].

As a particular case, if the energy eigenvalues of $H^{\text{AP}}(s)$ are non-degenerate for all s , each subset $\sigma_n(s)$ corresponds to a single energy level. In this situation, the error defined in Eq. (2.4) coincides with the infidelity, commonly used in quantum computing to quantify the distance of the prepared state from the target state [17]. The bound above then guarantees that the eigenstate associated with the n -th energy level of H_0 is evolved by $U_\tau(s)$ into the corresponding eigenstate of $H^{\text{AP}}(s)$, up to an error that vanishes at least as $1/\tau$. In particular, for $s = 1$, the adiabatic evolution

yields direct access to the eigenstates of the target Hamiltonian, H , up to an error of the same order.

The scaling of the minimal spectral gap $\Delta_n(s)$ with system size L directly dictates the computational complexity of adiabatic preparation, revealing a fundamental distinction between problem classes. For regular, uniform models (such as translationally invariant spin chains), the interpolation path typically encounters continuous (second-order) quantum phase transitions where the gap closes polynomially with system size, $\Delta \propto L^{-z}$ [38, 39]. In these settings, adiabatic state preparation can generally be achieved in polynomial time $\tau \sim \text{poly}(L)$. In stark contrast, disordered or highly frustrated systems, such as quantum spin glasses, exhibit extremely rugged energy landscapes characterized by a multiplicity of local minima [40]. Along the adiabatic trajectory, such models frequently undergo first-order quantum phase transitions or suffer from non-geometric quantum localization, causing the spectral gap to close exponentially with system size, $\Delta \sim e^{-\alpha L}$ [41, 42]. Consequently, state preparation for disordered systems generically demands an exponential evolution time $\tau \sim e^{\alpha L}$, illustrating how physical disorder transforms a computationally tractable adiabatic path into an exponentially hard task. More broadly, both exact level crossings at finite system size and spectral gaps that vanish in the large- L limit represent regimes where individual energy levels can no longer be tracked in isolation. In either case, the affected energy levels must be grouped into a joint subspace to maintain a well-defined gap, resulting in an adiabatic evolution that maps an initial eigenstate to an unknown superposition within that manifold.

From the considerations above, it is clear that maintaining finite spectral gaps around the target energy level is essential for the effectiveness of adiabatic state preparation. In this context, symmetry considerations play a central role, as the presence of symmetries directly constrains the structure of the spectrum. Specifically, if the Hamiltonian $H^{\text{AP}}(s)$ is invariant under a nontrivial symmetry group at a given value of the evolution parameter $s \in [0, 1]$, its eigenspaces decompose into irreducible representations (irreps) of that group [43]. The dimension of each irrep then fixes the multiplicity of the associated energy level, implying that symmetries generically enforce spectral degeneracies. As a consequence, when symmetries are sufficiently large, multiple energy levels may become degenerate or cross as s is varied, leading to an enhanced degeneracy structure in the instantaneous spectrum of $H^{\text{AP}}(s)$.

For this reason, a natural strategy to suppress level crossings during the adiabatic evolution is to engineer the auxiliary Hamiltonian H_0 so as to minimize the symmetry group of $H^{\text{AP}}(s)$. Reducing the symmetry group generally lowers the dimensionality of the irreps, thereby limiting the multiplicity of degenerate eigenspaces that can give rise to crossings along the adiabatic path. In practice, this is often accomplished by leveraging the arbitrariness of the auxiliary Hamiltonian, specifically by adding symmetry-breaking terms to H_0 that explicitly remove the evident symmetries of $H^{\text{AP}}(s)$ [44–46].

Although the latter strategy is widely used in the literature, the procedure is highly delicate. The introduction of additional terms in H_0 can inadvertently generate accidental symmetries along the adiabatic path, i.e., symmetries that are not manifest

in the geometric or physical structure of the Hamiltonian but emerge from specific features of the model or of the added perturbations. Such accidental symmetries can induce unexpected degeneracies or level crossings in the instantaneous spectrum of $H^{\text{AP}}(s)$.

A paradigmatic classical example of emerging accidental symmetries is provided by the Laplace–Runge–Lenz vector in the Kepler problem: despite the apparent $\text{SO}(3)$ rotational symmetry, the $1/r$ potential endows the system with a hidden $\text{SO}(4)$ symmetry, yielding additional conserved quantities and ensuring that all bound orbits are closed [47]. The quantum analogue arises in the hydrogen atom, where the same $\text{SO}(4)$ symmetry leads to degeneracies between states with different angular momentum but identical principal quantum number [48]. Analogously, in the adiabatic setting considered here, accidental symmetries may give rise to unintended degeneracies that complicate the adiabatic evolution. Consequently, while explicit symmetry breaking in H_0 can remove trivial degeneracies, hidden symmetries may persist and reintroduce spectral degeneracies, thereby posing a nontrivial obstacle to the construction of a smooth and well-gapped adiabatic path.

These considerations highlight the need for a more refined strategy to enhance the robustness of adiabatic state preparation. The Branched-Subspace Adiabatic Algorithm that will be introduced in Chapter 6 is specifically designed to mitigate the difficulties arising from both explicit and accidental symmetries, thereby improving control over degeneracies encountered along the adiabatic evolution.

As a final remark, AP exhibits an often-overlooked advantage compared to the VQA approach. Specifically, if AP independently evolves a set of eigenstates of H_0 into the corresponding eigenstates of H , then by linearity of $U_\tau(1)$ any linear combination of eigenstates of H_0 is mapped to the same linear combination of eigenstates of H . In other words, solving the state preparation problem for each eigenstate automatically provides a solution for any pure state.

CHAPTER 5

OPTIMIZATION OF ADIABATIC PREPARATION

The discussion of the previous chapter highlighted both the conceptual simplicity and the practical limitations of adiabatic quantum-state preparation. While the adiabatic theorem guarantees that the subspace associated with $\sigma_n(0)$ can be evolved into the one associated with $\sigma_n(s)$ with an arbitrarily small error by increasing the evolution time τ , the actual performance of the protocol depends sensitively on the spectral properties of the interpolating Hamiltonian $H(s)$ and the specific choice of the auxiliary Hamiltonian H_0 . In particular, small spectral gaps or unfavorable spectral structures can significantly increase the characteristic timescale required to reach a desired fidelity.

These observations naturally raise the question of whether the adiabatic protocol itself can be systematically optimized. More precisely, one may ask to what extent the choice of the auxiliary Hamiltonian H_0 and of the interpolation path between H_0 and the target Hamiltonian H_1 can be engineered so as to accelerate the convergence of the adiabatic evolution. In this chapter, we report the first original contribution of this thesis; namely, we address this question by deriving an explicit analytical expression for the characteristic timescale governing the exponential suppression of the adiabatic error, and by exploiting this result to construct an optimized adiabatic preparation protocol.

5.1 Explicit control of the adiabatic error bound

Consider the interpolating Hamiltonian

$$H^{\text{AP}}(s) = H_0 + f(s)(H - H_0), \quad (2.6)$$

where $s \in [0, 1]$, with $f(0) = 0$ and $f(1) = 1$. Let $\sigma(s)$ denote the spectrum of $H^{\text{AP}}(s)$, and assume that it can be partitioned into spectral subsets $\sigma_n(s)$ separated by finite gaps. In the following, we focus on the complete adiabatic evolution, corresponding to $s = 1$, and omit the explicit dependence on s whenever no ambiguity arises. The

error associated with the adiabatic evolution of the n -th eigenspace is quantified by $\mathcal{E}_n^{\text{AP}}(\tau)$, defined in Eq. (2.4).

In [49], G. Nenciu proved that there exist positive functions $g_n(s)$ and positive constants C_n such that

$$\mathcal{E}_n^{\text{AP}}(\tau) \leq C_n \int_0^1 ds' \exp(-\lfloor \tau / g_n(s') \rfloor) + A(\tau), \quad (2.7)$$

where $A(\tau) = \mathcal{O}(1/\tau)$ and vanishes identically whenever $H^{\text{AP}}(s)$ admits a holomorphic extension to a strip around the real axis. This is precisely the regime targeted by the recently proposed adiabatic quantum phase estimation protocol of Ref. [50]. Despite establishing exponential suppression of the adiabatic error with the evolution time τ , Nenciu's theorem does not provide an explicit analytical expression for the characteristic timescale $g_n(s)$. Consequently, while the theorem guarantees favorable asymptotic scaling, it does not reveal how the interpolation path should be designed to minimize the adiabatic runtime.

In Ref. [1], we addressed this limitation by deriving an explicit analytical expression for $g_n(s)$ in terms of spectral properties of the interpolating Hamiltonian $H^{\text{AP}}(s)$. In particular, $g_n(s)$ is expressed in terms of four quantities:

i) the operator norm

$$\|V\| := \|H - H_0\|, \quad (2.8)$$

ii) the instantaneous spectral gap $\Delta_n(s)$ defined in Eq. (2.5),

iii) the bandwidth

$$D_n(s) := \max_{x, y \in \sigma_n(s)} |x - y| \quad (2.9)$$

of the target spectral sector, and

iv) the interpolation function $f(s)$.

It is important to note that these quantities can typically be estimated for systems of moderate size using analytical or classical numerical methods, e.g., through iterative eigensolvers such as the Lanczos algorithm [51]. The availability of an explicit expression for $g_n(s)$ provides a concrete way to control the characteristic timescale of the adiabatic process and therefore enables the construction of a purely classical optimization strategy for the choice of H_0 and $f(s)$. Crucially, this result does not relax any of the conventional conditions required for adiabatic preparation to hold—such as the presence of a non-zero spectral gap—but rather provides a quantitative route to optimize the protocol's runtime whenever these conditions are met. The expression for $g_n(s)$ obtained in [1] reads

$$g_n(s) = 4 \frac{W(\Delta_n(s), D_n(s))}{\varrho\left(s, \frac{\Delta_n(s)}{\|V\|}, f(s)\right)}, \quad (2.10)$$

where

$$W(x, y) = \left[2^8 \left(1 + \frac{2y}{\pi x} \right)^2 \left(\frac{\pi^2}{3} + \frac{4}{x} \right) \right]^{7/3}, \quad (2.11)$$

and

$$\varrho(x, y, f) = \sup_{\rho \in \mathbb{C}} \left\{ |\rho| \text{ such that } |f(x + \rho) - f(x)| < \frac{y}{4} \right\}. \quad (2.12)$$

The calculations leading to the previous result are reported in App. A. A direct inspection of Eq. (2.10) shows that $g(s)$ decreases when the spectral gap $\Delta_n(s)$ increases and when the bandwidth $D_n(s)$ decreases. Consequently, defining

$$\Delta_n^{\min} = \inf_{s \in [0,1]} \Delta_n(s), \quad D_n^{\max} = \sup_{s \in [0,1]} D_n(s), \quad (2.13)$$

one obtains the bound

$$g_n(s) \leq 4 \frac{W(\Delta_n^{\min}, D_n^{\max})}{\inf_{s \in [0,1]} \varrho\left(s, \frac{\Delta_n^{\min}}{\|V\|}, f(s)\right)} =: \tilde{g}_n \quad (2.14)$$

such that

$$\mathcal{E}_n^{\text{AP}}(\tau) \leq C_n \exp(-\lfloor \tau / \tilde{g}_n \rfloor). \quad (2.15)$$

One may then use the expression of $\tilde{g}_n$ in Eq. (2.14) as a loss function to be minimized by searching for an optimal initial Hamiltonian H_0 and/or an optimal interpolation function $f(s)$ within a given class of admissible choices. For instance, if the interpolation is fixed to the commonly employed linear schedule $f(s) = s$, the expression reduces to

$$\tilde{g}_n = 16 \frac{\|V\|}{\Delta_n^{\min}} W(\Delta_n^{\min}, D_n^{\max}), \quad (2.16)$$

which can then be minimized through an appropriate choice of H_0 . Previous bounds in the literature already indicated a dependence on the operator distance $\|V\|$ [52], although in those results the adiabatic error exhibits only a power-law dependence on the inverse evolution time. Equations (2.15) and (2.16) instead show that the error bound of the adiabatic protocol is exponentially small in $\|V\|^{-1}$.

Equation (2.16) allows one to approach the problem of reducing $\tilde{g}_n$ through the design of the auxiliary Hamiltonian H_0 in two different ways. A first, coarse strategy consists of minimizing the operator distance $\|V\| = \|H - H_0\|$. Although this choice is not necessarily optimal, it has the advantage of being transversal to all the spectral sectors labeled by n and therefore improves the adiabatic timescale simultaneously for all eigenspaces. Moreover, the quantity $\|V\|$ is typically the easiest to estimate or control in practice, making this approach less demanding in terms of classical computational resources. A more refined strategy instead aims at minimizing the full expression appearing in Eq. (2.16). In this case, the design of H_0 must take

into account not only the operator distance $\|V\|$ but also the spectral quantities $\Delta_n^{\min}$ and $D_n^{\max}$, which characterize respectively the minimum gap separating the target eigenspace from the rest of the spectrum and the bandwidth of the sector itself along the interpolation path. This second approach can potentially lead to a significantly smaller value of $\tilde{g}_n$, at the price of requiring a more detailed spectral analysis of the interpolating Hamiltonian.

5.2 Initial Hamiltonian optimization

The results presented in the previous section naturally suggest a strategy for optimizing adiabatic state preparation. In particular, we henceforth consider the setting in which the interpolation function $f(s)$ is fixed, while the initial Hamiltonian H_0 is optimized over a prescribed class of experimentally accessible Hamiltonians. The objective is to identify the choice of H_0 that minimizes the characteristic timescale governing the adiabatic evolution, thereby reducing the total runtime required to achieve a target accuracy. More generally, however, the analytical framework developed in Ref. [1] is not restricted to the optimization of the initial Hamiltonian. The same formalism can be employed to optimize the interpolation path itself, by designing the schedule $f(s)$ or, more broadly, the Hamiltonian trajectory so as to minimize the relevant characteristic timescale. This provides a principled and quantitative alternative to the heuristic interpolation-design strategies commonly adopted in the literature, such as in Ref. [53].

Given a target Hamiltonian H , a common choice for the auxiliary Hamiltonian H_0 consists of the diagonal part of H in the computational basis, whose eigenstates are trivially prepared on a quantum register [54, 55]. Here we extend this choice to a more general construction aimed at reducing $\tilde{g}_n$ while preserving the simplicity of the initial state preparation. In particular, we write the auxiliary Hamiltonian as

$$\langle i | H_0 | j \rangle = \delta_{ij} \langle i | H | j \rangle + \langle i | M | j \rangle, \quad (2.17)$$

where M is a Hermitian operator acting as a preconditioning term.

In this explanatory example, we require that M satisfy several practical constraints. First, H_0 should remain diagonal in the computational basis so that its eigenstates coincide with computational basis states. Second, the implementation of the corresponding unitary operator

$$\mathcal{U}_M(t) = \exp(-itM) \quad (2.18)$$

should not introduce a circuit depth that offsets the advantage obtained from reducing $\tilde{g}$. Finally, the number of parameters required to specify M should scale at most polynomially with the number of qubits. A particularly simple choice satisfying these requirements is

$$\langle i | M | i \rangle = \sum_{j=0}^{L-1} \alpha_j \text{bin}_j(i), \quad (2.19)$$

where $|i\rangle$ denotes the i -th computational basis state and $\text{bin}_j(i)$ is the j -th digit of i in binary representation. With this construction the corresponding unitary operator factorizes as

$$\mathcal{U}_M(t) = \bigotimes_{j=0}^{L-1} P_j(\alpha_j t), \quad (2.20)$$

where P_j denotes a single-qubit phase gate acting on the j -th qubit. The resulting circuit depth is therefore independent of the system size. The parameters $\{\alpha_j\}$ can

then be optimized classically so as to minimize either $\|V\|$ or the full characteristic timescale expression $\tilde{g}_n$. The resulting auxiliary Hamiltonian defines a preconditioned adiabatic path that significantly accelerates the convergence of the adiabatic evolution, leading to a faster preparation of the target quantum state.

5.2.1 Example: XZ Heisenberg model

Here we apply the AP protocol to the preparation of the ground state of the Heisenberg model on one- and two-dimensional (1D and 2D) square lattices with periodic boundary conditions. The purpose of this analysis is to numerically validate the theoretical predictions derived in the previous sections on models of direct relevance for quantum simulation. The corresponding quantum circuits are implemented and executed on the Qiskit *qasm* simulator [56]. To isolate the contribution of the adiabatic error $\mathcal{E}_n^{\text{AP}}$, we employ sufficiently small Trotter time steps, so that the systematic error introduced by the Trotter–Suzuki approximation becomes negligible. In the following, the state preparation is performed using three different strategies, all adopting the same linear interpolation function $f(s) = s$.

The first corresponds to the standard AP protocol, where M is set to zero in Eq. (2.17). We then consider a second strategy in which the parameters α_j are chosen so as to minimize $\|V\|$, which requires only a single application of the Lanczos algorithm (or an equivalent iterative eigensolver) and does not depend on the target n -th subspace. Although this choice does not guarantee that $\tilde{g}_n$ is minimized, it already incorporates part of the information entering the bound. Finally, we analyze a third strategy in which the quantity $\tilde{g}_n$ in Eq. (2.16) is evaluated explicitly by estimating both $\Delta_n^{\min}$ and $D_n^{\max}$. While this approach leads to the best scaling in the examples




Figure 5.1: Ground-state preparation for the 9-site one-dimensional (a) and two-dimensional (b) Heisenberg models with periodic boundary conditions, whose lattice geometries are shown schematically in the corresponding panels. The adiabatic error is plotted as a function of the preparation timescale τ for $J^x = 5J^z$ for three different protocols: the unoptimized one ($M = 0$), a partially optimized one obtained by minimizing $\|V\|$, and the fully optimized one obtained by minimizing the complete expression for $\tilde{g}_n$.

considered here, its application to more complex models may be less practical, since it requires the estimation of the spectrum at several intermediate values of s , each involving a separate application of the Lanczos algorithm.

The XZ Heisenberg model is a prototypical many-body system describing interacting spin-1/2 particles on a lattice, defined by the Hamiltonian

$$H = -\frac{1}{2} \sum_{\langle i,j \rangle} (J^z Z_i Z_j + J^x X_i X_j), \quad (2.21)$$

where the sum runs over nearest-neighbor sites, and X_i and Z_i denote Pauli operators acting on the i -th qubit. In this work we consider real and positive coupling constants J^z and J^x . The adiabatic interpolation is then defined by

$$H^{\text{AP}}(s) = -\frac{1}{2} \sum_{\langle i,j \rangle} (J^z Z_i Z_j + s J^x X_i X_j) + (1-s)M. \quad (2.22)$$

At $s = 0$ the Hamiltonian is diagonal and, in the unpreconditioned case ($M = 0$), possesses a degenerate ground space. The corresponding eigenstates $|0\rangle^{\otimes L}$ and $|1\rangle^{\otimes L}$ are orthonormal and invariant under lattice translations. To preserve translational symmetry along the adiabatic path, we impose the constraint $\alpha_i = \alpha$ ($\forall i$) in Eq. (2.20). This choice reduces the number of free parameters to a single one while maintaining the symmetry of the model. Physically, the preconditioner M then corresponds to the introduction of a uniform magnetic field along the z direction with strength α .

The results for the ground-state preparation are reported in Fig. 5.1 for both the 1D and 2D lattices. Independently of the dimensionality and of the parameter regime considered, the AP protocol yields an error that decreases approximately




Figure 5.2: Characteristic times g_{fit} obtained by fitting the numerical curves $\epsilon_{\text{AP}}(\tau)$ (as in Fig. 5.1) with a single exponential for different values of the ratio J^x/J^z of the 1D Heisenberg model. Standard deviations from the fit are reported as error bars (often smaller than the markers).

exponentially with the preparation time τ . Moreover, comparing the three strategies described above reveals a clear reduction of the required timescale to reach a fixed target accuracy, with the third (fully optimized) approach consistently providing the shortest preparation time. The differences between the approaches become more pronounced as the ratio J^x/J^z increases, as illustrated in Fig. 5.2. In this plot the characteristic time g_{fit} is extracted by fitting $\ln\{\epsilon_{\text{AP}}(\tau)\}$ with the linear function $C_{\text{fit}} - \tau/g_{\text{fit}}$. For $J^x \ll J^z$ the values obtained with the three methods are nearly indistinguishable, whereas a clear advantage of the preconditioning strategy emerges when $J^x \gg J^z$, i.e., when the target Hamiltonian differs significantly from the initial one. The peak observed around $J^x/J^z \simeq 1$ corresponds to the region where, in the thermodynamic limit, the system undergoes a quantum phase transition. In this regime the preparation of the ground state becomes intrinsically more demanding and the effectiveness of the optimization is correspondingly reduced. Finally, we remark that, for the specific spin models considered here, the improvement introduced by the preconditioning term occurs only for an odd number of sites with periodic boundary conditions. A detailed explanation of this feature is provided in Section 5.3.

5.2.2 Example: Ising model

In this section we report the numerical results obtained for the Ising model by performing the same numerical tests presented for the Heisenberg model. The Ising model in a uniform external magnetic field is described by the Hamiltonian

$$H = -\frac{h}{2} \sum_i Z_i - \frac{J^x}{2} \sum_{\langle i,j \rangle} X_i X_j. \quad (2.23)$$




Figure 5.3: Lowest energy levels of the 9-site Ising chain in a transverse field plotted as functions of the dimensionless interpolation parameter s (in units of h/J^x).




Figure 5.4: Ground-state preparation via the AP protocol for the 9-site (a) 1D and (b) 2D Ising models with periodic boundary conditions and a transverse magnetic field along z . The adiabatic error is plotted as a function of the preparation time τ (in units of the transverse-field energy h) for $J^x = 5h$.

To apply the AP protocol we define the interpolating Hamiltonian

$$H^{\text{AP}}(s) = -\frac{h}{2} \sum_i Z_i - \frac{sJ^x}{2} \sum_{\langle i,j \rangle} X_i X_j + (1-s)M. \quad (2.24)$$

In this model, the ground state of H_0 is non-degenerate, whereas the ground-state level of H becomes doubly degenerate in the limit $J^x/h \rightarrow \infty$. Consequently, the spectral gap between the ground state and the first excited state closes during the adiabatic evolution. Nevertheless, the AP protocol can still be shown to remain valid by decomposing the evolution into two separate stages. The first stage starts at $s = 0$ and stops at an intermediate value s^* . For concreteness we choose $s^* = 0.5J^x/h$, although the precise value is not critical for the argument. In this first part of the evolution the set $\sigma_0(s)$ contains only the ground-state energy, and the corresponding projector is denoted by $P_0(s)$. In the second stage of the evolution, i.e., for $s > s^*$, both the ground and first excited energy levels are included in $\sigma_0(s)$, as schematically illustrated in Fig. 5.3. The associated projector onto this two-dimensional subspace is denoted by $P'_0(s)$. Since the spectral gap separating $\sigma_0(s)$ from the rest of the spectrum remains strictly positive in each of the two intervals, both

$$\epsilon_0^{\text{AP}}(\tau) = \|(\mathbb{1} - P_0(s^*))U_\tau(s^*, 0)P_0(0)\| \quad (2.25)$$

and

$$\epsilon_0^{\text{AP}'}(\tau) = \|(\mathbb{1} - P'_0(1))U_\tau(1, s^*)P'_0(s^*)\| \quad (2.26)$$

are exponentially suppressed in the large- τ limit. Moreover, since the image of $P_0(s^*)$ is contained in the image of $P'_0(s^*)$, the total adiabatic error satisfies

$$\begin{aligned} \mathcal{E}_0^{\text{AP}}(\tau) &= \|(\mathbb{1} - P'_0(1))U_\tau(1, s^*)U_\tau(s^*, 0)P_0(0)\| \\ &\leq \|(\mathbb{1} - P'_0(1))U_\tau(1, s^*)P_0(s^*)\| + \|(\mathbb{1} - P_0(s^*))U_\tau(s^*, 0)P_0(0)\| \\ &\leq \epsilon_0^{\text{AP}'}(\tau) + \epsilon_0^{\text{AP}}(\tau), \end{aligned} \quad (2.27)$$




Figure 5.5: Characteristic times g_{fit} obtained by fitting the numerical curves $\epsilon_{\text{AP}}(\tau)$ with a single exponential for different values of the ratio J^x/h in the 1D Ising model. Standard deviations of the fit are shown as error bars (often smaller than the markers).

which implies that the overall preparation error also exhibits exponential suppression. The main numerical results are reported in Fig. 5.4 and Fig. 5.5.

5.3 Periodic boundary conditions and site parity

From numerical simulations, we noticed that the optimal choice for the preconditioner is strictly dependent on the parity in the number of sites. In fact, we found that in the case of an even number of qubits and periodic boundary conditions, the optimal preconditioner that minimizes $\tilde{g}$ is actually $M = 0$, both for the Ising and the Heisenberg models. In other words, the standard AP process, i.e., letting just the original Hamiltonian evolve in time to find its ground state, is already the optimal one for an even number of sites. Here we report an interpretation of this peculiar phenomenon. Consider the Hamiltonian term

$$V = -\frac{J^x}{2} \sum_{i=0}^{L-1} X_i X_{i+1}, \quad (2.28)$$

which is the same for both the Ising and Heisenberg models before the preconditioner term M is introduced. In the presence of periodic boundary conditions, if the number of qubits, L , is even, there is a clear distinction between even and odd sites in the lattice, as schematically illustrated in Fig. 5.6.

In fact, in this case it is possible to rewrite

$$V = -\frac{J^x}{2} \sum_{i \text{ even}}^{L-1} (X_i X_{i+1} + X_{i-1} X_i). \quad (2.29)$$




Figure 5.6: Schematic illustration of the different parity in the number of sites that can occur in the presence of periodic boundary conditions: (a) for even L , the odd (dark green) and even (light green) sites present a clear difference, since the last site (even) is mapped to the first one (odd); (b) in the case of odd L , the last and first site of the chain share the same parity, and no parity issue is present.

Moreover, it is possible to introduce the well-defined operator

$$S \equiv \bigotimes_{i \text{ even}} Z_i = S^{-1}, \quad (2.30)$$

from which we notice that

$$SVS = -V. \quad (2.31)$$

As a consequence, for each eigenvector $|\lambda\rangle$ of V such that

$$V|\lambda\rangle = \lambda|\lambda\rangle, \quad (2.32)$$

there exists a corresponding eigenvector $|\lambda\rangle \equiv S|\lambda\rangle$, with opposite eigenvalue

$$V|\lambda\rangle = -\lambda|\lambda\rangle. \quad (2.33)$$

This implies that the spectrum of V is completely symmetric in the case of even L . In particular, let $\lambda_{\max}$ be the maximum eigenvalue; then the minimum one is necessarily $-\lambda_{\max}$, and therefore

$$\|V\| = \lambda_{\max}. \quad (2.34)$$

The additional preconditioning term M is then introduced with the intention to reduce the previous norm. However, in the even L case, this attempt turns out to be ineffective with our proposal of M , and the reason is that

$$SMS = M, \quad (2.35)$$

i.e., M commutes with S . To show why, we now consider $-M$ to be a perturbation of V , and we apply perturbation theory to find the maximum and minimum eigenvalues of $V' := V - M$ as

$$\begin{cases} \Lambda_{\max} = \lambda_{\max} + \lambda_1 + \lambda_2 + \lambda_3 + \mathcal{O}(\|M\|^3) \\ \Lambda_{\min} = -\lambda_{\max} + \lambda_1 - \lambda_2 + \lambda_3 + \mathcal{O}(\|M\|^3), \end{cases} \quad (2.36)$$




Figure 5.7: Error dependence on the adiabatic process timescale for the ground state preparation of the periodic XZ Heisenberg chain, assuming $L = 10$, $J^x = 5J^z$, and an external magnetic field on the x axis of magnitude $h^x = 0.4J^z$.

where $\lambda_i \sim \|M\|^i$. The odd order corrections, such as λ_1 and λ_3 , shift the unperturbed maximum and minimum eigenvalues by the same contribution, thus leading to an overall increase of the operator norm of V' . On the other hand, the second-order term gives opposite contributions to the energy shifting, but it is always positive for the higher energy level. Hence, it also leads to an increasing norm of V . The final consequence is that, for this specific problem, any perturbative preconditioner $M \neq 0$ would increase the operator norm. Notice that we never used our specific choice expressed by Eq. (2.19), but only the fact that M is diagonal, and it hence commutes with S .

On the other hand, a slightly more complicated model is sufficient to break the condition expressed in Eq. (2.31). As a paradigmatic example, we have tested our method in the case of even-site chains for the XZ Heisenberg model in an external x-oriented magnetic field of magnitude h^x . In this case, we find a clear improvement in performance provided by our protocol with respect to the standard adiabatic preparation, as can be appreciated in Fig. 5.7 for the ground-state preparation of a 10-site chain.

CHAPTER 6

BRANCHED-SUBSPACES ADIABATIC PREPARATION

This chapter is dedicated to the Branched-Subspaces Adiabatic Preparation (B-SAP) method [2], devised to address the limitations of conventional adiabatic preparation (AP) and variational quantum algorithms (VQAs) discussed previously. The approach combines structural features of both paradigms in order to compensate for their respective shortcomings. The construction begins with a reformulation of the adiabatic state preparation framework. In standard AP, the initial Hamiltonian is typically chosen to have a symmetry group that is as small as possible, thereby minimizing spectral degeneracies. In contrast, in B-SAP the initial Hamiltonian H_0 is selected such that its symmetry group contains the symmetry group of the target Hamiltonian H as a subgroup. Under this choice, symmetries may be broken along the adiabatic path but are never created. Consequently, symmetry does not induce mergers of distinct energy levels, although it may cause level splitting during the evolution (see, e.g., Fig. 6.1). At first glance, this strategy may appear counterintuitive, since it leads to equal or higher degeneracy at the beginning of the adiabatic process compared to its final stage. In other words, the degeneracies intrinsic to the problem are not eliminated, but effectively shifted to the initial stage of the process. This does not constitute a practical obstruction, however, because the spectrum of H_0 is known a priori and its eigenstates are assumed to be efficiently preparable. The ambiguity associated with the high degeneracy of H_0 can therefore be resolved at negligible cost. In practice, one can initialize well-characterized basis states using relatively shallow quantum circuits before performing the adiabatic sweep, instead of attempting to construct complicated target eigenstates directly.

More precisely, let $\mathcal{B}$ be an orthonormal eigenbasis of H_0 , and let $\mathcal{B}_n \subset \mathcal{B}$ denote a subset of d_n eigenstates spanning the degenerate subspace $\mathcal{S}(\mathcal{B}_n)$ associated with the spectral cluster $\sigma_n(0)$. In what follows, we assume H_0 is chosen such that $d_n \in \text{poly}(L)$. By construction, the spectral cluster $\sigma_n(s)$ remains isolated from the rest of the spectrum along the adiabatic process and the adiabatic evolution operator $U_\tau(s)$ unitarily maps the basis $\mathcal{B}_n$ into $U_\tau(s)\mathcal{B}_n$ and therefore $\mathcal{S}(\mathcal{B}_n)$ into the instantaneous spectral subspace $\mathcal{S}(U_\tau(s)\mathcal{B}_n)$ associated with $\sigma_n(s)$. If $\sigma_n(s)$ consists of a




Figure 6.1: Illustrations of adiabatic energy spectra for two processes sharing the same final Hamiltonian but differing in their initial ones. The left case 6.1a, related to conventional AP, starts with a non-degenerate spectrum, but it features several level crossings. In contrast, the right panel illustrates the evolution performed through B-SAP 6.1b, which displays a forked structure with degeneracies that never increase throughout the evolution. Notice that the time evolution is plotted from 0 to 1 in 6.1a, and mirrored (i.e., from 1 to 0) in 6.1b, highlighting the differences in the final spectra.

single eigenvalue for all s , i.e., no level splitting occurs, the transported subspace remains irreducible and no additional structure needs to be resolved. In general, however, the degeneracy may be partially or completely lifted during the evolution. The problem then becomes to select a basis of $S(\mathcal{B}_n)$ whose vectors are transported by $U_\tau(s)$ into instantaneous eigenstates of $H^{\text{AP}}(s)$ within the evolved subspace. Operationally, one seeks a set of linearly independent eigenstates in $\mathcal{B}_n$ such that each state is adiabatically connected to a single eigenbranch of $H^{\text{AP}}(s)$.

To implement this selection procedure, we employ the Multistate-Contracted Variational Quantum Eigensolver (MC-VQE) method [57]. Let $\{|\Psi_\mu\rangle\}_{\mu=1}^{d_n}$ denote the elements of the arbitrarily chosen basis $\mathcal{B}_n$. For all μ and $\nu \neq \mu$, and for $p \in \{0, 1, 2, 3\}$, we prepare the overlap states

$$|\Psi_{\mu\nu}^p\rangle = \frac{1}{\sqrt{2}} (|\Psi_\mu\rangle + i^p |\Psi_\nu\rangle), \quad (2.37)$$

initialize them on the quantum register, evolve them under $U_\tau(s)$, and estimate the corresponding energy expectation values

$$\langle \Psi_{\mu\nu}^p(s) | H^{\text{AP}}(s) | \Psi_{\mu\nu}^p(s) \rangle := \langle \Psi_{\mu\nu}^p | U_\tau^\dagger(s) H^{\text{AP}}(s) U_\tau(s) | \Psi_{\mu\nu}^p \rangle, \quad (2.38)$$

where, for any state $|\psi\rangle$, we use the shorthand $|\psi(s)\rangle := U_\tau(s)|\psi\rangle$. From this set of measured quantities, the matrix elements of $H^{\text{AP}}(s)$ in the evolved basis $U_\tau(s)\mathcal{B}_n$ can

be reconstructed classically in post-processing as

$$[H^{\text{AP}}(s)]_{\mu\nu}^{(n)} := \langle \Psi_\mu(s) | H^{\text{AP}}(s) | \Psi_\nu(s) \rangle \quad (2.39)$$

$$= \frac{1}{2} \left[\langle \Psi_{\mu\nu}^0(s) | H^{\text{AP}}(s) | \Psi_{\mu\nu}^0(s) \rangle - \langle \Psi_{\mu\nu}^2(s) | H^{\text{AP}}(s) | \Psi_{\mu\nu}^2(s) \rangle \right. \\ \left. - i \langle \Psi_{\mu\nu}^1(s) | H^{\text{AP}}(s) | \Psi_{\mu\nu}^1(s) \rangle + i \langle \Psi_{\mu\nu}^3(s) | H^{\text{AP}}(s) | \Psi_{\mu\nu}^3(s) \rangle \right]. \quad (2.40)$$

Since $d_n \in \text{poly}(L)$, the resulting d_n -dimensional matrix can be diagonalized efficiently on a classical computer [51]. This yields approximate eigenvalues $E_\mu^{(n)}(s)$ and corresponding eigenvectors expressed in the basis $U_\tau(s)\mathcal{B}_n$ as

$$|E_\mu^{(n)}(s)\rangle = \sum_\nu c_{\mu\nu}(s) |\Psi_\nu(s)\rangle. \quad (2.41)$$

If the sole objective of the computation is the determination of the energy eigenvalues, the procedure terminates at this stage, since the classical diagonalization already provides the desired spectrum. On the other hand, one may wish to prepare one of the eigenstates obtained from the classical diagonalization, or an arbitrary superposition thereof, on the quantum register in order to estimate additional observables. Suppose, for instance, that the objective is to initialize

$$|\Psi_{\text{target}}\rangle = \sum_\nu \tilde{c}_\nu |\Psi_\nu(s)\rangle, \quad \sum_\nu |\tilde{c}_\nu|^2 = 1. \quad (2.42)$$

By linearity of $U_\tau(s)$, one has

$$|\Psi_{\text{target}}\rangle = \sum_\nu \tilde{c}_\nu |\Psi_\nu(s)\rangle = U_\tau(s) \sum_\nu \tilde{c}_\nu |\Psi_\nu\rangle. \quad (2.43)$$

Therefore, it suffices to prepare the desired linear combination of the initial basis states $\mathcal{B}_n$ and subsequently apply the adiabatic evolution to obtain the target state. To achieve this, one could employ the linear combination of unitaries technique within a fault-tolerant quantum computing framework [58]. In what follows, we focus on the procedure introduced in Ref. [2], which is particularly suitable for NISQ and early fault-tolerant architectures, as it does not require ancillary qubits or controlled applications of oracle operations. Namely, we first select a fixed reference state in $\mathcal{B}_n$, e.g., $|\Psi_1\rangle$, and then construct a parametrized unitary gate $\mathcal{G}(\underline{\alpha})$ acting as an endomorphism on $\mathcal{S}(\mathcal{B}_n)$. Its action on $|\Psi_1\rangle$ should allow one to explore the entire subspace $\mathcal{S}(\mathcal{B}_n)$ by varying the N_α real parameters $\underline{\alpha}$. Equivalently, the gate should define a mapping

$$\mathcal{G} : \mathbb{R}^{N_\alpha} \longrightarrow \mathcal{S}(\mathcal{B}_n), \quad \underline{\alpha} \mapsto \mathcal{G}(\underline{\alpha}) |\Psi_1\rangle. \quad (2.44)$$

The unitary endomorphisms of $\mathcal{S}(\mathcal{B}_n)$ form the group $U(d_n)$, whose Lie algebra $\mathfrak{u}(d_n)$ has dimension d_n^2 . Each generator corresponds to an infinitesimal rotation or phase transformation within the degenerate manifold. The following result holds [59]:

Theorem 1. Let $\{t_j\}_j$ be a set of generators of $\mathfrak{u}(d_n)$. Then, for every $V \in U(d_n)$, there exists a set of real parameters $\{\alpha_j\}_j$ such that

$$V = \prod_{j=1}^{d_n^2} \exp[i\alpha_j t_j]. \quad (2.45)$$

Each generator appears only once in the product, so $\mathcal{G}$ can be implemented as a sequence of parametrized elementary gates

$$G_j(\alpha_j) = \exp[i\alpha_j t_j], \quad j \in \{1, \dots, d_n^2\}. \quad (2.46)$$

However, when starting from a fixed reference state $|\Psi_1\rangle$, not all elements of $U(d_n)$ are required. The subgroup $U(d_n - 1) \subset U(d_n)$ leaves $|\Psi_1\rangle$ invariant,

$$V |\Psi_1\rangle = |\Psi_1\rangle, \quad \forall V \in U(d_n - 1), \quad (2.47)$$

so only the coset space $U(d_n)/U(d_n - 1)$ acts nontrivially on $|\Psi_1\rangle$. Its dimension is

$$d_n^2 - (d_n - 1)^2 = 2d_n - 1, \quad (2.48)$$

which is the number of real parameters needed to specify a normalized state in $\mathcal{S}(\mathcal{B}_n)$ starting from $|\Psi_1\rangle$. Accordingly, the mapping $\mathcal{G}$ can be made bijective by restricting to $2d_n - 1$ generators,

$$\mathcal{G}(\underline{\alpha}) |\Psi_1\rangle = \prod_{j=1}^{2d_n-1} G_j(\alpha_j) |\Psi_1\rangle. \quad (2.49)$$

If the target state is defined only up to a global phase, one additional generator can be removed, reducing the count to $2(d_n - 1)$. When d_n scales polynomially with the system size L , the number of trainable parameters also scales polynomially. Notably, the parametrized circuit constructed in this manner is not a heuristic ansatz, but a systematic realization of all physically relevant transformations within $\mathcal{S}(\mathcal{B}_n)$. While the choice of the generator basis $\{t_j\}$ of $\mathfrak{u}(d_n)$, and therefore of the corresponding set of gates $\{G_j\}$ in Eq. (2.49), is not unique, a convenient construction can be given explicitly. For each $\mu \in \{2, \dots, d_n\}$, define two generators of the algebra acting nontrivially only on the two-dimensional subspace $\text{span}\{|\Psi_1\rangle, |\Psi_\mu\rangle\}$:

$$t_\mu^{(X)} = |\Psi_1\rangle\langle\Psi_\mu| + |\Psi_\mu\rangle\langle\Psi_1|, \quad (2.50)$$

$$t_\mu^{(Y)} = -i|\Psi_1\rangle\langle\Psi_\mu| + i|\Psi_\mu\rangle\langle\Psi_1|. \quad (2.51)$$

Restricted to the ordered basis $\{|\Psi_1\rangle, |\Psi_\mu\rangle\}$, these operators act respectively as Pauli X and Pauli Y matrices, while acting trivially on the orthogonal complement. Consequently, for each $\mu \neq 1$, the associated parametrized gates

$$G_\mu^{(X)}(\alpha_\mu^{(X)}) = \exp[i\alpha_\mu^{(X)} t_\mu^{(X)}], \quad (2.52)$$

$$G_\mu^{(Y)}(\alpha_\mu^{(Y)}) = \exp[i\alpha_\mu^{(Y)} t_\mu^{(Y)}], \quad (2.53)$$

implement controlled rotations within the two-dimensional subspace spanned by $|\Psi_1\rangle$ and $|\Psi_\mu\rangle$, leaving all other basis states invariant. This construction provides $2(d_n - 1)$ generators, which are sufficient to explore all physically distinct normalized states in $\mathcal{S}(\mathcal{B}_n)$ up to a global phase. If global phases are to be retained explicitly, an additional diagonal generator,

$$t^{(Z)} = |\Psi_1\rangle\langle\Psi_1|, \quad (2.54)$$

may be included to control the relative phase between $|\Psi_1\rangle$ and the remaining components. Operationally, the parametrized circuit thus consists of sequential X - and Y -type rotations coupling the reference state $|\Psi_1\rangle$ to each $|\Psi_\mu\rangle$, generating arbitrary amplitude redistribution and phase structure within the degenerate subspace. Importantly, the structure of the operator $\mathcal{G}(\underline{\alpha})$ is completely independent of the target Hamiltonian H and depends only on the structure of the initial Hamiltonian H_0 . At this point we can write

$$\prod_j G_j(\alpha_j) |\Psi_1\rangle = \sum_\nu a_\nu(\underline{\alpha}) |\Psi_\nu\rangle, \quad (2.55)$$

where the coefficients $a_\nu(\underline{\alpha})$ are completely determined by the parameter vector $\underline{\alpha}$. As a final step, the parameters $\{\alpha_j\}$ must be chosen such that the prepared state coincides with $|\Psi_{\text{target}}\rangle$ up to an unobservable global phase. This condition is equivalent to enforcing unit magnitude overlap,

$$\left| \langle \Psi_{\text{target}} | \prod_j G_j(\alpha_j) |\Psi_1\rangle \right| = \left| \sum_\nu \tilde{c}_\nu^* a_\nu(\underline{\alpha}) \right| = 1, \quad (2.56)$$

or, equivalently, one may formulate the problem as the minimization of the real-valued classical loss function

$$\mathcal{L}(\underline{\alpha}) = 1 - \left| \sum_\nu \tilde{c}_\nu^*(s) a_\nu(\underline{\alpha}) \right|^2. \quad (2.57)$$

Importantly, the optimization can be entirely classical: no quantum measurements are required at this stage. This is because the loss function in Eq. (2.57) can be computed directly from the coefficients $a_\nu(\underline{\alpha})$, which in turn are obtained by evaluating the product of the matrices G_j . Although G_j act on the full Hilbert space, their non-trivial action is restricted to the degenerate subspace $\mathcal{S}(\mathcal{B}_n)$, so the effective size of the matrices is only $d_n \times d_n$. Since d_n scales polynomially with the system size, the computation of the coefficients $a_\nu(\underline{\alpha})$ is efficient. Moreover, the global minimum of the loss function is known a priori and equals 0, attained uniquely when

$$a_\nu(\underline{\alpha}) = \tilde{c}_\nu(s) \quad \forall \nu, \quad (2.58)$$

up to an irrelevant global phase. Consequently, no spurious local minima can be mistaken for the solution. Note that while the target physical state is uniquely defined, the optimal parameter vector $\underline{\alpha}^{\text{opt}}$ itself may not be unique due to parameterization redundancies; nevertheless, any parameter vector achieving the global minimum $\mathcal{L} = 0$ is equally valid, as it prepares the correct physical state up to a global




Figure 6.2: Schematic illustration of the quantum circuit implementing the B-SAP algorithm. The initial quantum state $|\Psi_1\rangle \in \mathcal{S}(\mathcal{B}_n)$ is first prepared by applying a low-depth circuit to the default $|0\rangle^{\otimes L}$ state. The parametrized unitary $G(\underline{\alpha}^{\text{opt}})$ then allows exploration of the full subspace $\mathcal{S}(\mathcal{B}_n)$, by varying the parameters $\underline{\alpha}^{\text{opt}}$. Finally, the state undergoes adiabatic evolution via $U_\tau(1)$.

phase. Once an optimal set of parameters $\underline{\alpha}^{\text{opt}}$ has been determined, one can initialize the target state $|\Psi_{\text{target}}\rangle$ on the quantum register as follows: i) initialize $|\Psi_1\rangle$, ii) act with the optimized quantum circuit $G(\underline{\alpha}^{\text{opt}})$, iii) adiabatically evolve with $U_\tau(s)$. The structure of the algorithm is represented in Fig. 6.2.

6.1 Application: the XYZ Heisenberg chain

As an explicit application of the B-SAP protocol, we consider the one-dimensional XYZ Heisenberg model on a spin chain with an even number of sites L and periodic boundary conditions. The target Hamiltonian is

$$H = - \sum_{i=0}^{L-1} [J^x X_i X_{i+1} + J^y Y_i Y_{i+1} + J^z Z_i Z_{i+1}], \quad (2.59)$$

where the index $i+1$ is understood modulo L . Without loss of generality, we assume

$$|J^z| \geq |J^x| \geq |J^y|, \quad (2.60)$$

since any different ordering can be restored by a global rotation of the spin axes. The Hamiltonian is invariant under lattice translations and reflections, generating the so-called dihedral group D_L . In addition, when J^x , J^y , and J^z are all distinct, the remaining spin symmetry consists of global π rotations around the coordinate axes. In particular, the operator

$$\mathcal{X} = \bigotimes_{i=0}^{L-1} X_i \quad (2.61)$$

commutes with H and defines a parity quantum number that will be conserved along the adiabatic evolution. Following the general construction introduced previously,

we choose the initial Hamiltonian

$$H_0 = -J^z \sum_{i=0}^{L-1} Z_i Z_{i+1}, \quad (2.62)$$

whose eigenbasis coincides with the computational basis

$$\mathcal{B} = \left\{ |\underline{b}\rangle = \bigotimes_{j=0}^{L-1} |b_j\rangle \mid b_j \in \{0, 1\} \right\}. \quad (2.63)$$

These states can be prepared trivially on a quantum register, and the symmetry group of H_0 contains that of H , ensuring that no additional symmetries are generated during the interpolation. For a computational basis state $|\underline{b}\rangle$ one finds

$$\langle \underline{b} | Z_i Z_{i+1} | \underline{b} \rangle = 1 - 2(b_i \oplus b_{i+1}), \quad (2.64)$$

where $\oplus$ denotes addition modulo two. It is therefore convenient to introduce the linear transformation

$$\Phi |\underline{b}\rangle = \bigotimes_{j=0}^{L-1} |b_j \oplus b_{j+1}\rangle, \quad (2.65)$$

which maps a bit string to the pattern of domain walls between neighboring spins. Notice that Φ is not bijective, as it stands. Indeed, for any $|\underline{b}\rangle$ it holds

$$\bigoplus_i \langle \underline{b} | \Phi^\dagger \frac{1 - Z_i}{2} \Phi | \underline{b} \rangle = 0, \quad (2.66)$$

meaning that the Hamming weight

$$h(\Phi |\underline{b}\rangle) = \sum_i \langle \underline{b} | \Phi^\dagger \frac{1 - Z_i}{2} \Phi | \underline{b} \rangle \quad (2.67)$$

is always even, and therefore the image of Φ is at most half of $\mathcal{S}(\mathcal{B})$. In order to make Φ invertible, we partition $\mathcal{B}$ into

$$\mathcal{B}^{(0)} = \{ |\underline{b}\rangle \mid b_0 = 0 \} \quad \mathcal{B}^{(1)} = \{ |\underline{b}\rangle \mid b_0 = 1 \},$$

splitting the bit strings according to their value on the first bit. The restriction of the domain of Φ to either $\mathcal{S}(\mathcal{B}^{(0)})$ or $\mathcal{S}(\mathcal{B}^{(1)})$ leads to a bijective map. For the case of $\Phi : \mathcal{S}(\mathcal{B}^{(0)}) \mapsto \mathcal{S}(\mathcal{B})$, the inverse operation is

$$\Phi_0^{-1} |\underline{b}\rangle = \bigotimes_{j=0}^{N-1} \left| \bigoplus_{k=j}^{L-1} b_k \right\rangle, \quad (2.68)$$

while for $\Phi : \mathcal{S}(\mathcal{B}^{(1)}) \mapsto \mathcal{S}(\mathcal{B})$ its inverse is

$$\Phi_1^{-1} |\underline{b}\rangle = \mathcal{X} \Phi_0^{-1} |\underline{b}\rangle. \quad (2.69)$$

Making use of these expressions, and recalling the Hamming weight definition in Eq. (2.67), the eigenvalue equation for H_0 can now be recast into the more informative form

$$H_0 |\Phi(\underline{b})\rangle = -J^z [L - 2h(\Phi(\underline{b}))] |\Phi(\underline{b})\rangle. \quad (2.70)$$

This equation has important consequences. First, we notice that Φ bijectively maps the sets $\mathcal{B}_n^{(0,1)} = \mathcal{B}_n \cap \mathcal{B}^{(0,1)}$ into the set $\mathcal{W}_{2n} \subset \mathcal{B}$ of bit strings whose Hamming weight is equal to $2n$. This in turn implies that there exist $L/2 + 1$ distinct energy levels, $\{E_n(0)\}$, all separated by the same gap $\Delta E = 4J^z$, with a degeneracy equal to twice the size of $\mathcal{W}_{2n}$, i.e.,

$$d_n = 2 \binom{L}{2n}. \quad (2.71)$$

Remarkably, Eq. (2.71) allows one to conclude that $d_n = \mathcal{O}(L^{2n})$, such that the number of trainable parameters in the final circuit will only grow polynomially with L . If one manages to explore $\mathcal{S}(\mathcal{W}_{2n})$, one can then use $\Phi_{0,1}^{-1}$ to map the states back into $\mathcal{S}(\mathcal{B}_n^{(0,1)})$. However, this only allows one to separately prepare the generic vector either in $\mathcal{S}(\mathcal{B}_n^{(0)})$ or in $\mathcal{S}(\mathcal{B}_n^{(1)})$, while the goal is to prepare any $|\psi_n\rangle \in \mathcal{S}(\mathcal{B}_n)$ that is expressed as

$$|\psi_n\rangle = \cos(\theta) |\psi_n^{(0)}\rangle + \sin(\theta) e^{i\varphi} |\psi_n^{(1)}\rangle, \quad (2.72)$$

for proper $|\psi_n^{(0,1)}\rangle \in \mathcal{S}(\mathcal{B}_n^{(0,1)})$, $\theta \in [0, \pi/2]$, and $\varphi \in [0, 2\pi)$. The pivotal observation to overcome this problem is that $\mathcal{X}$, in light of Eq. (2.69), is an operator that bijectively maps $\mathcal{B}_n^{(0)}$ into $\mathcal{B}_n^{(1)}$. Moreover, $\mathcal{X}$ and $H(s)$ commute for any s , hence it is always possible to simultaneously diagonalize them such that the expectation value of $\mathcal{X}$ remains constant throughout the adiabatic procedure. This is enough to gain insight into the form of Eq. (2.72). Let $|\psi_n^\pm\rangle$ be an eigenstate with energy E_n and eigenvalue ± 1 with respect to $\mathcal{X}$. Then

$$\langle \psi_n^\pm | \mathcal{X} | \psi_n^\pm \rangle = \sin(2\theta) \operatorname{Re} [e^{i\varphi} \langle \psi_n^{(0)} | \mathcal{X} | \psi_n^{(1)} \rangle] = \pm 1, \quad (2.73)$$

in which orthogonality between $\mathcal{S}(\mathcal{B}_n^{(0)})$ and $\mathcal{S}(\mathcal{B}_n^{(1)})$ has been used. It follows that

$$\begin{cases} \theta = \pi/4 \\ e^{i\varphi} |\psi_n^{(1)}\rangle = \pm \mathcal{X} |\psi_n^{(0)}\rangle. \end{cases} \quad (2.74)$$

Therefore

$$|\psi_n^\pm\rangle = \frac{1}{\sqrt{2}} [\mathbb{1} \pm \mathcal{X}] |\psi_n^{(0)}\rangle, \quad (2.75)$$

thus ensuring that, once we are able to explore $\mathcal{S}(\mathcal{B}_n^{(0)})$, exploration of $\mathcal{S}(\mathcal{B}_n)$ automatically follows. Finally, we can constructively determine the unitary gates $G_j(\alpha_j)$ of Eq. (2.56) to explore $\mathcal{S}(\mathcal{B}_n^{(0)})$, specifically focusing on the $n = 0$ and $n = 1$ cases as examples, up to global phases.

To illustrate the performance of the protocol, we performed numerical simulations using the Qiskit framework for system sizes up to $L = 10$ qubits. The quality of

the prepared state is quantified through

$$\mathcal{E} = \left\| \left(1 - P_{E_v}(1) \right) U_\tau(1) \prod_j G_j(\alpha_j) |\Psi\rangle \right\|, \quad (2.76)$$

where $P_{E_v}(1)$ denotes the projector onto the eigenspace associated with the target eigenvalue $E_v(1)$. For nondegenerate levels, $\mathcal{E}$ coincides with the standard infidelity. This quantity incorporates all sources of error, including finite-time adiabatic evolution and the Trotterization used in the digital implementation.

6.1.1 Example: target state for $n = 0$

In this case $d_0 = 1$; the only state in $\mathcal{W}_0$ is $|\Psi\rangle = |0\rangle^{\otimes L}$, i.e., the default state of a digital quantum computer. Hence no parameterization is required and the operators Φ^{-1} , $\frac{1}{\sqrt{2}}(\mathbb{1} \pm \mathcal{X})$, and $U_\tau(1)$ can be applied directly. The error in the preparation of the ground state of H , defined in Eq. (2.76), obtained using a simulator of a digital quantum computer for a Heisenberg chain of length $L = 10$, is reported in Fig. 6.3 in logarithmic scale for both $\mathcal{X} = 1$ and $\mathcal{X} = -1$. In both cases, the numerical results obtained using only $L/2 = 5$ Trotter steps, which is the minimal number required to enable entanglement between any pair of sites in the periodic chain, show that $\mathcal{E}$ is negligible compared to 2^{-L} , i.e., the average overlap between a random state in the Hilbert space and the target one. As expected, this is particularly evident for $|J^x/J^z|, |J^y/J^x| \ll 1$, i.e., in parameter regions where the target Hamiltonian H is not too different from H_0 . However, even for values close to critical regions such




Figure 6.3: Error between the target ground state and the actually prepared state in the two parity sectors $\mathcal{X} = \pm 1$. Such an error is quantified by $\mathcal{E}$ (see Eq. (2.76)) for all possible ratios among the coupling constants of the XYZ Heisenberg Hamiltonian H in Eq. (2.59). Both panels refer to the ferromagnetic model with $L = 10$. The adiabatic procedure has been performed for $L/2 = 5$ Trotter steps, equivalent to half of the chain length, each with $0.25 \cdot [J^z]^{-1}$ duration.




Figure 6.4: Error between the target ground state and the actually prepared state in the $\mathcal{X} = +1$ sector, calculated for different numbers of Trotter steps in the adiabatic protocol of the B-SAP method. The error, $\mathcal{E}$ (see Eq. (2.76)), is reported for possible ratios among the coupling constants of the ferromagnetic XYZ Heisenberg Hamiltonian H in Eq. (2.59), with $L = 10$. The adiabatic procedures have been performed with (a) 5, (b) 10, and (c) 20 Trotter steps of duration $0.25 \cdot [J^z]^{-1}$.

as $J^x = J^z$ or $J^y = J^x$, the error achieved in state preparation by using the B-SAP protocol never exceeds 0.3. Finally, it is important to stress that this error can be significantly reduced further, for all parameters, by simply increasing the number of Trotter steps employed in the digital implementation of the adiabatic procedure. This effect is explicitly illustrated in Fig. 6.4, where the number of Trotter steps is increased from $L/2$ to L and $2L$, respectively.

6.1.2 Example: target state for $n = 1$

In this case, the dimension of the relevant subspace becomes

$$d_1 = \frac{1}{2}L(L-1), \quad (2.77)$$

corresponding to configurations with Hamming weight equal to two. Since each state in $\mathcal{W}_2$ is uniquely determined by the positions of its two nonzero entries in the bit string, it is convenient to denote the basis states as $|a, b\rangle$, where a and b label the positions of such entries along the chain and $0 \leq a < b \leq L-1$. The set of operators $\{G_\mu^{(X)}\}$ and $\{G_\mu^{(Y)}\}$ of Eq. (2.52) coincide with the well-known X and Y Givens rotations. Fixing the reference state in Eq. (2.44) to be $|\Psi\rangle_1 = |0, 1\rangle$, one restricts to the subset of $2(d_n - 1)$ Givens rotations acting nontrivially on the initial state.

Further reduction in the number of trainable parameters can be obtained by exploiting the fact that the Hamiltonian (2.59) is purely real. Consequently, there exists an orthogonal change of basis that diagonalizes it, and it is therefore sufficient to reproduce the action of the orthogonal group $O(d_1)$ rather than the full unitary group $U(d_1)$. This is achieved by restricting to the subset of gates $\{G_\mu^{(X)}\}$. Accordingly, the total number of trainable parameters is reduced to $d_1 = L(L-1)/2$.




Figure 6.5: Error between the lowest-energy target excited states and the actually prepared states in both parity sectors, $\mathcal{X} = \pm 1$. The error, $\mathcal{E}$ (see Eq. (2.76)), is reported for all possible ratios among the coupling constants of the ferromagnetic XYZ Heisenberg Hamiltonian H in Eq. (2.59). Calculations reported in this Figure have been performed for $L = 10$. The adiabatic procedure has been realized with a number of Trotter steps equivalent to $L/2 = 5$, each with $0.25 \cdot [J^z]^{-1}$ duration.

We apply this procedure to the preparation of the first four excited eigenstates of the Heisenberg Hamiltonian for both parity sectors of the operator $\mathcal{X}$ using quantum registers of up to $L = 10$ qubits and $L/2 = 5$ Trotter steps in the adiabatic protocol. The resulting preparation error $\mathcal{E}$ defined in Eq. (2.76) is shown in Fig. 6.5 as a logarithmic color map over the parameter space of the Hamiltonian. The results indicate that the error remains either very small, typically below the percent level, or limited to values below 10% over most of the parameter region. As the excitation level increases, the regions of larger error become more extended at fixed Trotter step size, reflecting the fact that nearby excited energy levels are closer in energy and therefore more difficult to distinguish through classical diagonalization of the matrix $[H^{\text{AP}}(1)]_{\mu\nu}^{(1)}$. Nevertheless, the results confirm that the B-SAP protocol remains effective also for the preparation of excited eigenstates.

6.1.3 Benchmark with Conventional Adiabatic Preparation

We here compare the B-SAP protocol with conventional adiabatic preparation (AP), which serves as a benchmark. Figure 6.6 reports the errors obtained in the preparation of the ground and first excited states of the same Hamiltonian using standard AP under identical simulation conditions, namely $L = 10$ qubits and the same number of Trotter steps. The initial Hamiltonian is chosen as

$$H_0 = -J^z \sum_i \frac{Z_i}{2^i}, \quad (2.78)$$

which is diagonal in the computational basis and has no degeneracies. The ground-state errors remain moderately small over a broad region of the parameter space and




Figure 6.6: Error in preparing the ground (a) and first excited (b) states of the ferromagnetic XYZ Heisenberg Hamiltonian H of Eq. (2.59) using standard adiabatic preparation. Both panels show results for $L = 10$, with $L/2 = 5$ Trotter steps of duration $0.25 \cdot [J^z]^{-1}$ each.

are comparable to those obtained with the B-SAP protocol for the same number of qubits and Trotter steps. In contrast, the preparation of the first excited state exhibits a clear qualitative difference. Except in the trivial limit $J^x/J^z = 0$, standard AP fails due to unavoidable level crossings during the adiabatic evolution, as illustrated by the spectrum shown in Fig. 6.1a. This comparison highlights the key advantage of the B-SAP strategy, which enables reliable preparation not only of ground states but also of excited eigenstates of the target Hamiltonian.

6.2 Criteria for guaranteeing the satisfiability of B-SAP

In this section, we analyze different scenarios that may arise when attempting to guarantee that the spectral-gap conditions required by B-SAP are satisfied. As a representative testbed, we consider the transverse-field Ising model on a closed chain of L sites. The target Hamiltonian is

$$H_T = -J^x \sum_{i=1}^L X_i X_{i+1} + B^z \sum_{i=1}^L Z_i, \quad (2.79)$$

with periodic boundary conditions $X_{L+1} \equiv X_1$. The Hamiltonian H_T is invariant under translations of the chain, reflections, and global π -rotations of all spins about the z axis. Two natural choices for the initial Hamiltonian are

$$H_0 = -J^x \sum_{i=1}^L X_i X_{i+1}, \quad \text{if } |J^x| \geq |B^z|, \quad (2.80)$$




Figure 6.7: Spectrum of the transverse-field Ising chain with $L = 4$ sites as a function of the interpolation parameter s . The thick lines show the exact eigenvalues of $H(s)$. The shaded areas around the initial Hamming-weight energies $\sigma_n(0)$ represent the estimated maximal shifts $\pm s \|\Pi_n V \Pi_n\|$ for each Hamming-weight sector n . Panel (a) corresponds to a small ratio $J^x/B^z = 1/L$, where the gaps between sectors are large and the bound provides a clear separation. Panels (b) and (c) correspond to larger ratios J^x/B^z , in which case one should also evolve the closest Hamming-weight eigenspaces that may cross or become nearly degenerate, with an additional polynomial overhead, to ensure the correct ordering of eigenstates.

and

$$H_0 = B^z \sum_{i=1}^L Z_i, \quad \text{if } |B^z| \geq |J^x|. \quad (2.81)$$

In the following, we focus on the latter choice and use $|B^z|$ as the energy unit. By a suitable rotation of the reference frame, we can always take $B^z \geq 0$, which we assume throughout. We adopt the standard convention

$$Z|0\rangle = +|0\rangle, \quad Z|1\rangle = -|1\rangle. \quad (2.82)$$

With this convention, the eigenstates of

$$H_0 = B^z \sum_{i=1}^L Z_i \quad (2.83)$$

are the computational basis states $|\underline{b}\rangle$, whose energies depend only on their Hamming weight

$$h(|\underline{b}\rangle) = n, \quad (2.84)$$

namely

$$H_0 |\underline{b}\rangle = B^z (L - 2n) |\underline{b}\rangle. \quad (2.85)$$

Therefore, the eigenspace $\mathcal{S}(\mathcal{B}_n)$ associated with Hamming weight n has degeneracy

$$d_n = \binom{L}{n}, \quad (2.86)$$

and adjacent Hamming-weight sectors are separated at $s = 0$ by an energy gap $2B^z$. We consider the linear interpolation

$$H^{\text{AP}}(s) = H_0 + s(H_T - H_0), \quad s \in [0, 1], \quad (2.87)$$

where the perturbation driving the evolution away from the initial spectrum is

$$V := H_T - H_0 = -J^x \sum_{i=1}^L X_i X_{i+1}. \quad (2.88)$$

B-SAP evolves the initial eigenspace $\mathcal{S}(\mathcal{B}_n)$ with exponentially suppressed mixing provided that, throughout the interpolation, the spectral branch $\sigma_n(s)$ originating from this eigenspace remains separated from all other spectral branches $\sigma_m(s)$ by a finite gap. To derive sufficient conditions ensuring this separation, we bound the maximal displacement of the spectrum associated with a fixed Hamming-weight sector under the perturbation sV . The projection onto individual Hamming-weight sectors provides a sector-resolved estimate of the spectral motion. A fully rigorous alternative would be to bound the motion of the complete spectrum using Weyl's theorem; however, such an approach generally produces substantially looser bounds. We define the projected Hamiltonian within the Hamming-weight- n sector as

$$H_{\text{eff}}^{(n)} := \Pi_n H_T \Pi_n \quad (2.89)$$

$$= \Pi_n (H_0 + V) \Pi_n \quad (2.90)$$

$$= H_0|_{\mathcal{S}(\mathcal{B}_n)} + \Pi_n V \Pi_n, \quad (2.91)$$

where Π_n is the projector onto $\mathcal{S}(\mathcal{B}_n)$. The operator $X_i X_{i+1}$ flips the pair of bits (b_i, b_{i+1}) . Such a flip preserves the Hamming weight only when acting on a boundary of type 01 or 10; conversely, flips of 00 or 11 change the Hamming weight and are removed by the projection. For a bit string $\underline{b}$ of weight n , let $q(\underline{b})$ denote the number of boundaries of type 01 or 10. Under periodic boundary conditions, the maximum possible number of such boundaries is obtained by arranging the bits as alternately as possible, yielding

$$q_{\text{max}}(n) = 2 \min(n, L - n). \quad (2.92)$$

Within $\mathcal{S}(\mathcal{B}_n)$, each allowed boundary flip gives rise to an off-diagonal matrix element of magnitude $|J^x|$. Hence, each row of $\Pi_n V \Pi_n$ contains at most $q_{\text{max}}(n)$ nonzero elements of magnitude $|J^x|$. By the Gershgorin circle theorem, or equivalently by a row-sum bound, the operator norm satisfies

$$\|\Pi_n V \Pi_n\| \leq |J^x| q_{\text{max}}(n) = 2|J^x| \min(n, L - n). \quad (2.93)$$

Therefore, the perturbation $s\Pi_n V \Pi_n$ can shift the eigenvalues originating from the n -th Hamming-weight sector by at most

$$\Delta\sigma_n(s) := \max_{x \in \sigma_n(s), y \in \sigma_n(0)} |x - y| \quad (2.94)$$

$$\leq 2s|J^x| \min(n, L - n). \quad (2.95)$$

A sufficient condition preventing overlap between two spectral branches $\sigma_n(s)$ and $\sigma_m(s)$ with $n \neq m$ is that the sum of their maximal possible displacements remains smaller than their initial separation:

$$\Delta\sigma_n(s) + \Delta\sigma_m(s) \leq |\sigma_n(0) - \sigma_m(0)| = 2B^z |n - m|. \quad (2.96)$$

Using Eq. (2.94) and considering the worst-case value $s = 1$, we obtain the sufficient condition

$$2|J^x| [\min(n, L - n) + \min(m, L - m)] \leq 2B^z |n - m|. \quad (2.97)$$

Restricting to sectors with $n, m \leq L/2$ (the complementary cases follow from particle-hole symmetry), Eq. (2.97) becomes

$$|J^x|(n + m) \leq B^z |n - m|. \quad (2.98)$$

This represents a conservative sufficient condition: whenever it is satisfied, the spectral branches originating from the two Hamming-weight sectors cannot overlap during the interpolation. The most stringent requirement generally corresponds to neighboring sectors, $m = n \pm 1$. A simple sufficient condition guaranteeing Eq. (2.98) for every Hamming-weight sector is therefore

$$\frac{|J^x|}{B^z} \leq \frac{1}{L}. \quad (2.99)$$

In this regime, the B-SAP gap conditions are guaranteed to hold a priori for all sectors. Outside this strict regime, however, efficient state preparation remains possible. Indeed, for generic $|J^x| \leq B^z$ and fixed n , only a finite number of Hamming-weight sectors can potentially overlap with $\sigma_n(s)$. These sectors satisfy

$$m \in [n - \ell_{\max}, n + k_{\max}], \quad (2.100)$$

where

$$\ell_{\max} = \left\lfloor \frac{2n|J^x|}{B^z + |J^x|} \right\rfloor, \quad (2.101)$$

$$k_{\max} = \left\lfloor \frac{2n|J^x|}{B^z - |J^x|} \right\rfloor. \quad (2.102)$$

Importantly, the number of such sectors is independent of the system size L . Each potentially relevant sector $\mathcal{S}(\mathcal{B}_m)$ has dimension

$$d_m = \binom{L}{m}. \quad (2.103)$$

For fixed m , this scales polynomially with L ,

$$d_m \sim \frac{L^m}{m!}. \quad (2.104)$$

Consequently, the total number of states involved in certifying possible crossings is bounded by

$$\mathcal{O}(L^{n+k_{\max}}) \quad (2.105)$$

up to prefactors depending only on n and $k_{\max}$. Hence, even when the simple sufficient condition ensuring global gap preservation is not satisfied, the additional overhead required to identify and resolve possible crossings remains polynomial in the system size. By adiabatically evolving all potentially relevant sectors, one can verify whether crossings occur and, if necessary, reconstruct the correct ordering of the evolved eigenstates a posteriori.

PART III

HAMILTONIAN EVOLUTION

Hamiltonian simulation is one of the central primitives of quantum computing. In its simplest form, Hamiltonian simulation is the task of implementing an approximation to the time-evolution operator $U(t) = e^{-iHt}$ generated by a Hamiltonian H . Besides its obvious relevance for simulating quantum systems, Hamiltonian simulation constitutes the algorithmic backbone of numerous quantum algorithms, including phase estimation, quantum chemistry, optimization, linear-system solvers, and matrix-function evaluation. Over the last three decades, two complementary paradigms have emerged for addressing this problem. The first is based on *Quantum Singular Value Transformation* (QSVT), which has become the method of choice in the fault-tolerant setting owing to its nearly optimal asymptotic complexity. The second consists of *operator-splitting* techniques, in which the evolution generated by a complicated Hamiltonian is approximated through compositions of evolutions generated by simpler Hamiltonian terms. This family includes the well-known Trotter–Suzuki formulas and more recent randomized approaches such as qDRIFT. Although operator-splitting methods generally exhibit less favorable asymptotic scaling with the target precision, they are considerably simpler to implement and remain the dominant approach for near-term quantum devices and many practical fault-tolerant applications.

The QSVT framework relies on the concept of a *block encoding*. Rather than representing an operator directly, one embeds it as a sub-block of a larger unitary acting on an extended Hilbert space.

Definition 1 (Block encoding). Let A be an operator acting on n qubits. A unitary U_A acting on $n + a$ qubits is called an (α, a, ϵ) block encoding of A if

$$|A - \alpha(\langle 0|^{\otimes a} \otimes I)U_A(|0\rangle^{\otimes a} \otimes I)| \leq \epsilon, \quad (3.1)$$

where $\alpha > 0$ is a normalization factor and a denotes the number of ancillary qubits. We denote the set of such block encodings as $\text{BE}_{\alpha,a}(A, \epsilon)$. When $\epsilon = 0$, the block encoding is said to be exact and one simply writes $A \in \text{BE}_{\alpha,a}(A)$.

Intuitively, the upper-left block of U_A reproduces the target operator A , up to normalization and approximation error. This construction allows non-unitary operators to be manipulated coherently using unitary quantum circuits and forms the fundamental building block underlying modern quantum linear algebra. Starting from a block encoding of the Hamiltonian, $U_H \in \text{BE}_{\alpha,a}(H)$, QSVT provides a systematic procedure for implementing polynomial transformations of the singular values of H by alternating applications of U_H with carefully chosen single-qubit rotations [60]. In the context of Hamiltonian simulation, one approximates the functions $\cos(Ht)$ and $\sin(Ht)$ through polynomial expansions of appropriate parity, which are synthesized independently using QSVT. The resulting block encodings are subsequently combined through the Linear Combination of Unitaries (LCU) framework [61] to obtain a block encoding of the full propagator

$$U(t) = e^{-iHt}. \quad (3.2)$$

The principal advantage of QSVT is its asymptotic efficiency. Hamiltonian simulation can be achieved with query complexity

$$\mathcal{O}\left(\alpha t + \frac{\ln(1/\epsilon)}{\ln(e + \ln(1/\epsilon)/(\alpha t))}\right), \quad (3.3)$$

where α denotes the normalization factor of the block encoding, t the simulation time, and ϵ the desired precision. This nearly optimal scaling explains why QSVT-based algorithms are generally preferred in the fault-tolerant regime, particularly for long simulation times and high-accuracy applications [62]. These favorable asymptotics come at the price of a relatively sophisticated implementation. Besides requiring a block encoding of the Hamiltonian, QSVT introduces ancillary qubits and relies on the synthesis of long sequences of precisely optimized single-qubit phases, whose computation requires dedicated classical algorithms [63]. Consequently, despite its asymptotic optimality, QSVT often carries substantial constant overheads and is comparatively demanding to realize in practice.

Throughout the remainder of this thesis, we instead focus on the second paradigm, namely operator splitting. Operator-splitting methods assume that the Hamiltonian admits a decomposition

$$H = \sum_{\gamma=1}^{\Gamma} H_{\gamma}, \quad (3.4)$$

where each elementary propagator

$$U_{\gamma}(t) = e^{-iH_{\gamma}t} \quad (3.5)$$

can be efficiently implemented. The central idea is to approximate the evolution generated by H through suitable compositions of these elementary evolutions. The simplest example is the first-order Trotter formula [64],

$$S_1(t) = \left(\prod_{\gamma=1}^{\Gamma} e^{-iH_{\gamma}t/s} \right)^s, \quad (3.6)$$

where s denotes the number of Trotter steps. Higher-order Suzuki formulas systematically improve the approximation by recursively constructing symmetric product formulas, while randomized variants such as qDRIFT [65] replace deterministic sequences by randomly sampled Hamiltonian terms, as discussed in the following subsection. Compared with QSVT, operator-splitting methods exhibit a less favorable dependence on the target precision. For any fixed Suzuki order, the query complexity scales polynomially with $1/\epsilon$, whereas QSVT achieves an essentially polylogarithmic dependence. Nevertheless, operator splitting offers several practical advantages. It requires only the ability to implement the individual propagators $e^{-iH_{\gamma}t}$, introduces no algorithmic ancilla overhead, and naturally exploits locality, sparsity, and commutation relations among the Hamiltonian terms. As a result, it remains the most widely adopted framework for Hamiltonian simulation on near-term quantum devices and continues to play an important role in fault-tolerant quantum computing.

As a final remark, it is relevant to highlight that Hamiltonian simulation naturally extends to time-dependent Hamiltonians $H(t)$, where the objective is to approximate the corresponding time-ordered exponential. In this setting, operator-splitting methods generalize straightforwardly by discretizing time into sufficiently short intervals and applying product formulas to the instantaneous Hamiltonians. By contrast, extending QSVT to time-dependent dynamics is considerably more involved, since the standard construction assumes access to a single fixed block-encoded operator. Accommodating a continuously varying Hamiltonian generally requires additional techniques, such as piecewise-constant approximations or linear combinations of distinct block encodings, while preserving the correct time ordering of the evolution.

The QDRIFT protocol Here we introduce the QDRIFT protocol, proposed by Campbell in [65], which provides a stochastic approach to approximate the time evolution generated by a Hamiltonian H . Given an initial quantum state ρ , the goal is to simulate its evolution under H for a time t . The QDRIFT protocol achieves this by constructing a quantum channel $\mathcal{C}^\ell$, defined as the composition of ℓ identical channels $\mathcal{C}$, each approximating a short-time evolution step

$$\mathcal{U}[\rho] = e^{-iHt/\ell} \rho e^{iHt/\ell} = e^{t\mathcal{L}/\ell}[\rho], \quad (3.7)$$

where the Liouvillian superoperator is given by

$$\mathcal{L}[\rho] = -i[H, \rho]. \quad (3.8)$$

Let the Hamiltonian admit a decomposition

$$H = \sum_{\gamma=1}^{\Gamma} \lambda_{\gamma} H_{\gamma}, \quad \lambda_{\gamma} \in \mathbb{R}, \quad \|H_{\gamma}\| = 1, \quad (3.9)$$

where each term H_{γ} generates a unitary evolution that can be efficiently implemented on quantum hardware. The single-step QDRIFT channel is defined as

$$\mathcal{C}[\rho] = \sum_{\gamma=1}^{\Gamma} p(\gamma) e^{-i\lambda_{\gamma} H_{\gamma} t/\ell p(\gamma)} \rho e^{iH_{\gamma} t/\ell p(\gamma)} \quad (3.10)$$

$$= \sum_{\gamma=1}^{\Gamma} p(\gamma) e^{t\mathcal{L}_{\gamma}/\ell p(\gamma)}[\rho], \quad (3.11)$$

where

$$\mathcal{L}_{\gamma}[\rho] = -i\lambda_{\gamma}[H_{\gamma}, \rho], \quad (3.12)$$

and

$$p(\gamma) := \frac{|\lambda_{\gamma}|}{\Lambda}, \quad \Lambda := \sum_{\gamma} |\lambda_{\gamma}| \quad (3.13)$$

is a sampling distribution. Such a channel can be implemented on quantum hardware by first classically sampling an index γ according to the probability distribution $p(\gamma)$, and then applying the unitary $e^{-i\lambda_\gamma H_\gamma t / (\ell p(\gamma))}$ to the quantum state. Repeating this procedure independently for each of the ℓ steps realizes the channel $\mathcal{C}^\ell$. When using the QDRIFT approximation to estimate the expectation value of an observable O , the resulting bias satisfies the bound (see [65])

$$|\text{Tr}[O \mathcal{C}^\ell[\rho]] - \text{Tr}[O \mathcal{U}^\ell[\rho]]| \leq \|O\| \frac{(2t\Lambda)^2}{\ell} e^{\frac{2t\Lambda}{\ell}}. \quad (3.14)$$

Therefore, in the limit of $\ell \rightarrow \infty$, the estimation bias goes to zero as $\mathcal{O}(\ell^{-1})$, matching the asymptotic scaling of the first-order Trotter formula. Unlike deterministic product formulas, however, QDRIFT does not depend on any specific ordering of the Hamiltonian terms, and its error bound is independent of commutation properties of the decomposition. This yields a more uniform and, in many cases, tighter control of the approximation error [65]. Moreover, at each time step the protocol requires the implementation of only a single Hamiltonian term H_γ , in contrast to Trotterized schemes, which apply every term in the Hamiltonian decomposition at each step. Consequently, QDRIFT is particularly well suited to Hamiltonians containing a large number of terms Γ .

Despite these advantages, standard QDRIFT is intrinsically limited to first-order accuracy in time, inheriting an inverse-linear query complexity with respect to the target precision ϵ . Recently, this bottleneck was resolved by Pathwise Random Hamiltonian Simulation (PRHS) [8], which generalizes randomized product formulas to arbitrary order. By dividing each time step into correlated slices governed by a uniquely constructed quasi-probability distribution, PRHS achieves a factorial bias decay in the expansion order. This enables query costs scaling sub-polynomially with $1/\epsilon$ in high-precision regimes, bridging the gap between low-overhead near-term randomized methods and fault-tolerant accuracy targets without requiring ancillary qubits.

CHAPTER 7

RESOURCE-OPTIMAL IMPORTANCE SAMPLING FOR RANDOMIZED QUANTUM ALGORITHMS

The QDRIFT protocol provides a paradigmatic example of a randomized quantum algorithm, where the stochastic selection of unitary evolutions enables an efficient approximation of Hamiltonian dynamics. More generally, a wide class of quantum algorithms relies on probabilistic choices over a set of quantum operations, giving rise to estimators whose performance is inherently governed by both the underlying sampling distribution and the statistical fluctuations induced by measurement. This observation naturally raises the question of whether the classical sampling distribution defining a randomized quantum protocol can be optimized in order to reduce the overall computational cost. In particular, while standard constructions such as QDRIFT fix the sampling distribution according to structural properties of the problem (see e.g. Eq. (3.13)), this choice is not necessarily optimal when accounting for implementation-dependent costs associated with different circuit realizations.

In this chapter, we present the general and problem-agnostic framework of [3] for applying classical importance sampling to randomized quantum protocols. Given an original sampling distribution p and a cost function c associated with the implementation of the corresponding quantum operations, we consider alternative sampling distributions q together with appropriate reweighting factors, ensuring that the resulting estimator remains unbiased. The central objective is to determine the optimal distribution q^* that minimizes a net-cost figure of merit defined as the product between the average cost per circuit execution and the variance of the estimator. Within this framework, importance sampling enables a systematic tradeoff between execution cost and statistical fluctuations: sampling expensive operations less frequently reduces the cost per run, at the expense of increasing the worse-case variance, and vice versa. We show how to optimally balance these competing contributions without requiring prior knowledge of the target quantity. We further extend the analysis to realistic scenarios in which the quantum computation is affected by algorithmic approximations or physical noise, proving that importance sampling preserves the es-

estimator bias despite modifying the sampling distribution. Additionally, we consider settings where error-detection schemes are employed, and characterize how these mechanisms modify both the optimal sampling strategy and the achievable reduction in computational resources.

Remarkably, while the approach was initially developed in the context of the QDRIFT algorithm, its generality enables direct application to all randomized quantum protocols, providing a principled method to reduce implementation costs while preserving the same asymptotic accuracy guarantees.

7.1 Optimal Importance Sampling

We consider the problem of estimating the expectation value $\text{Tr}[O \rho_{\text{target}}]$ of an observable O on a target quantum state ρ_{target} . We focus on scenarios in which ρ_{target} admits a representation as an ensemble average over a family of more easily implementable quantum states, obtained by applying quantum channels $\mathcal{E}_{\underline{\theta}}$ to a fixed reference state ρ , with sampling probabilities given by a distribution $p(\underline{\theta})$. Namely,

$$\rho_{\text{target}} = \int d\underline{\theta} p(\underline{\theta}) \mathcal{E}_{\underline{\theta}}[\rho] = \mathbb{E}_p[\mathcal{E}_{\underline{\theta}}[\rho]]. \quad (3.15)$$

In the case of the QDRIFT protocol, the parameter $\underline{\theta}$ reduces to the set of parameters $\{\gamma_j\}_{j=1}^s$ of the different steps, while the Liouvillian evolutions

$$\prod_{j=1}^s e^{\frac{t}{sp(\gamma_j)} \mathcal{L}_{\gamma_j}} \quad (3.16)$$

play the role of the channels $\mathcal{E}_{\underline{\theta}}$. Let

$$O = \sum_x x \Pi_x \quad (3.17)$$

be the spectral decomposition of the observable, where Π_x is the projector onto the eigenspace associated with eigenvalue x . An unbiased estimator of $\text{Tr}[O \rho_{\text{target}}]$ can be obtained by sampling $\underline{\theta} \sim p$, preparing $\mathcal{E}_{\underline{\theta}}[\rho]$, and measuring in the eigenbasis of O . Indeed,

$$\begin{aligned} \text{Tr}[O \rho_{\text{target}}] &= \int d\underline{\theta} p(\underline{\theta}) \sum_x \text{Tr}[\Pi_x \mathcal{E}_{\underline{\theta}}[\rho]] x \\ &= \sum_x \int d\underline{\theta} p(\underline{\theta}) r(x|\underline{\theta}) x \\ &= \mathbb{E}_{pr}[x], \end{aligned} \quad (3.18)$$

where $r(x|\underline{\theta}) := \text{Tr}[\Pi_x \mathcal{E}_{\underline{\theta}}[\rho]]$ denotes the conditional measurement distribution, and $\mathbb{E}_{pr}$ is the expectation over the joint distribution $p(\underline{\theta})r(x|\underline{\theta})$. We assume $\|O\| \leq 1$, which can always be enforced by rescaling. The total computational cost is determined by both the number of samples required to reach precision δ and the cost of

generating each sample. We denote by $c(\underline{\theta})$ the cost associated with implementing $\mathcal{E}_{\underline{\theta}}[\rho]$ and performing the measurement. The expected cost per sample depends on the sampling distribution. To optimize the trade-off between cost and statistical fluctuations, we introduce importance sampling (IS). Specifically, we sample $\underline{\theta} \sim q(\underline{\theta})$ and reweight outcomes using

$$w(\underline{\theta}) := \frac{p(\underline{\theta})}{q(\underline{\theta})}, \quad (3.19)$$

so that the resulting estimator remains unbiased

$$\mathbb{E}_{qr}[w(\underline{\theta})x] = \mathbb{E}_{pr}[x]. \quad (3.20)$$

Under the modified sampling, the expected cost per run is

$$\mathbb{E}_{qr}[c(\underline{\theta})] = \mathbb{E}_q[c(\underline{\theta})]. \quad (3.21)$$

To characterize the statistical overhead, we consider the worst-case variance over all observables with $\|O\| \leq 1$:

$$\begin{aligned} \text{Var}_{qr}^{\max} &= \max_{|x| \leq 1} \{ \mathbb{E}_{qr}[(wx)^2] - \mathbb{E}_{qr}[wx]^2 \} \\ &\leq \mathbb{E}_{qr}[w^2 x^2] \\ &= \mathbb{E}_q[w^2(\underline{\theta})]. \end{aligned} \quad (3.22)$$

By the Central Limit Theorem, the number of samples required to achieve precision δ scales as $\mathbb{E}_q[w^2] \delta^{-2}$. This motivates the definition of the net cost

$$NC_q := \mathbb{E}_q[c(\underline{\theta})] \mathbb{E}_q[w^2(\underline{\theta})]. \quad (3.23)$$

With this definition, the bound to the total computational cost scales as $\delta^{-2} NC_q$, independently of the observable. In the absence of IS ($q = p$), one has

$$NC_p = \mathbb{E}_p[c(\underline{\theta})]. \quad (3.24)$$

The optimal choice of q is characterized in the following Theorem.

Theorem 2 (Optimal sampling distribution). Let $p(\underline{\theta})$ be a probability distribution and $c(\underline{\theta}) > 0$ a cost function. Then, for any sampling distribution q ,

$$NC_q \geq (\mathbb{E}_p[c^{1/2}(\underline{\theta})])^2, \quad (3.25)$$

with equality if and only if

$$q^*(\underline{\theta}) = \frac{p(\underline{\theta}) c^{-1/2}(\underline{\theta})}{\mathbb{E}_p[c^{-1/2}(\underline{\theta})]}. \quad (3.26)$$

Proof. Using $w = p/q$, we rewrite

$$\begin{aligned} NC_q &= \mathbb{E}_q[c] \mathbb{E}_q[w^2] \\ &= \mathbb{E}_p[c w^{-1}] \mathbb{E}_p[w]. \end{aligned} \quad (3.27)$$

Applying the Cauchy–Schwarz inequality,

$$\mathbb{E}_p[c^{1/2}] = \mathbb{E}_p[c^{1/2} w^{-1/2} \cdot w^{1/2}] \leq (\mathbb{E}_p[c w^{-1}])^{1/2} (\mathbb{E}_p[w])^{1/2}. \quad (3.28)$$

Squaring both sides yields the claim. Equality holds iff $w \propto c^{1/2}$. $\square$

The relative advantage of optimal IS is quantified by

$$\frac{NC_{q^*}}{NC_p} = \frac{(\mathbb{E}_p[c^{1/2}])^2}{\mathbb{E}_p[c]} = \left(1 + \frac{\text{Var}_p[c^{1/2}]}{\mathbb{E}_p[c^{1/2}]^2} \right)^{-1}. \quad (3.29)$$

Thus, $NC_{q^*} \leq NC_p$, with larger gains when $c^{1/2}$ exhibits large relative fluctuations under p . Importantly, optimal IS does not modify the underlying quantum operations: the same channels $\mathcal{E}_{\underline{\theta}}$ are implemented, and only the classical sampling distribution is altered.

7.1.1 Composite Quantum Channels

In the following, we focus on settings characterized by the following assumptions:

1. The parameter vector $\underline{\theta} = (\theta_1, \dots, \theta_s)$ has dimension s , and its components are drawn independently according to a product distribution.
2. The cost function is additive:

$$c(\underline{\theta}) = \sum_{i=1}^s c_i(\theta_i), \quad (3.30)$$

where each c_i depends only on θ_i .

Each variable θ_i may represent either a single scalar parameter or a collection of parameters. This structure naturally arises when the channel $\mathcal{E}_{\underline{\theta}}$ is given by a sequential composition

$$\mathcal{E}_{\underline{\theta}} = \mathcal{E}_{\theta_s} \circ \mathcal{E}_{\theta_{s-1}} \circ \dots \circ \mathcal{E}_{\theta_1}, \quad (3.31)$$

where each channel $\mathcal{E}_{\theta_i}$ depends on an independent parameter θ_i . This includes QDRIFT-type constructions [65], as well as Eigenpath Traversal protocols [66], in ideal settings. Under the independence assumption, the random variables $c_i := c_i(\theta_i)$ are independent, with mean μ_i and variance σ_i^2 . The total cost $c(\underline{\theta})$ is therefore a sum

of independent random variables. For large s , its distribution approaches a Gaussian by the Central Limit Theorem, with

$$\mu = \sum_{i=1}^s \mu_i := s \bar{\mu}, \quad (3.32)$$

$$\sigma^2 = \sum_{i=1}^s \sigma_i^2 := s \bar{\sigma}^2, \quad (3.33)$$

where $\bar{\mu}$ and $\bar{\sigma}^2$ denote the average mean and variance. Assuming that $\bar{\mu}$ and $\bar{\sigma}^2$ remain finite as $s \rightarrow \infty$, one finds (see Appendix C) that

$$\mathbb{E}_p[c(\underline{\theta})] = s \bar{\mu}, \quad (3.34)$$

$$(\mathbb{E}_p[c^{1/2}(\underline{\theta})])^2 = s \bar{\mu} - \frac{\bar{\sigma}^2}{4\bar{\mu}} + \mathcal{O}\left(\frac{1}{s}\right). \quad (3.35)$$

Substituting into Eq. (3.29), we obtain

$$\frac{NC_{q^*}}{NC_p} = 1 - \frac{1}{s} \left(\frac{\bar{\sigma}}{2\bar{\mu}} \right)^2 + \mathcal{O}\left(\frac{1}{s^2}\right). \quad (3.36)$$

This expression shows that the advantage of optimal importance sampling vanishes as $s \rightarrow \infty$. In other words, in multilayer settings with independent components and additive costs, the relative reduction in net cost becomes asymptotically negligible. This behavior admits a simple interpretation: as the number of layers increases, the total cost concentrates around its mean, and the relative fluctuations of $c^{1/2}(\underline{\theta})$ decrease, thereby reducing the potential gain from importance sampling. However, relaxing either independence or additivity can lead to qualitatively different regimes. In particular, correlations between layers or non-additive cost structures may sustain a non-vanishing advantage. Furthermore, distinct behaviors arise in the presence of error-detection schemes, as will be discussed in Section 7.3.

7.2 Importance Sampling for Imperfect Quantum Channels

In the previous section we showed how the net computational cost can be minimized by IS. Since the cost function can be chosen to quantify circuit depth or CNOT gate count, this might suggest that the overall noise-induced error could also change, and perhaps even be reduced. In this section, we show that this intuition is misleading and, in fact, the bias of the final estimator is unaffected by IS.

In practical scenarios the quantum channel $\mathcal{E}_{\underline{\theta}}$ introduced in Eq. (3.15) may only be approximated by a noisy channel $\mathcal{N}_{\underline{\theta}}$, both because of algorithmic approximation errors and hardware imperfections, including gate noise and decoherence. Let

$$\varepsilon_{\underline{\theta}}(\rho) := \mathcal{N}_{\underline{\theta}}(\rho) - \mathcal{E}_{\underline{\theta}}(\rho) \quad (3.37)$$

denote the error in the implementation of the state $\mathcal{E}_{\underline{\theta}}(\rho)$. Although IS changes the sampling distribution from $p(\underline{\theta})$ to $q(\underline{\theta})$, and thus changes the expected error

in the prepared state from $\mathbb{E}_p[\varepsilon_{\underline{\theta}}(\rho)]$ to $\mathbb{E}_q[\varepsilon_{\underline{\theta}}(\rho)]$, this does not affect the bias of the final observable estimator. Indeed, the reweighting factor $w(\underline{\theta}) = p(\underline{\theta})/q(\underline{\theta})$ exactly compensates for the change in sampling distribution, so that the bias is fully determined by the original distribution $p(\underline{\theta})$ and the implementation error $\varepsilon_{\underline{\theta}}(\rho)$. This is stated more formally in the following theorem.

Theorem 3. Let

$$\mathcal{N}_{\underline{\theta}}(\rho) = \mathcal{E}_{\underline{\theta}}(\rho) + \varepsilon_{\underline{\theta}}(\rho) \quad (3.38)$$

be an approximate implementation of the channel $\mathcal{E}_{\underline{\theta}}(\rho)$, sampled according to $p(\underline{\theta})$. Let $q(\underline{\theta})$ be an arbitrary importance-sampling distribution, with weight $w(\underline{\theta}) = p(\underline{\theta})/q(\underline{\theta})$. Then, for any observable $O = \sum_x x \Pi_x$, the bias induced by the channel approximation in the importance-sampled estimator is independent of q and is given by

$$\mathbb{E}_{q_{r'}}[w(\underline{\theta})x] - \mathbb{E}_{p_r}[x] = \mathbb{E}_p[\text{Tr}(O \varepsilon_{\underline{\theta}}(\rho))]. \quad (3.39)$$

Proof. Define the ideal and noisy conditional outcome distributions for the POVM $\{\Pi_x\}_x$ as

$$r(x|\underline{\theta}) := \text{Tr}[\Pi_x \mathcal{E}_{\underline{\theta}}(\rho)], \quad (3.40)$$

$$r'(x|\underline{\theta}) := \text{Tr}[\Pi_x \mathcal{N}_{\underline{\theta}}(\rho)]. \quad (3.41)$$

The expectation value of the importance-sampled estimator built from the noisy channel is

$$\begin{aligned} \mathbb{E}_{q_{r'}}[w(\underline{\theta})x] &= \int d\underline{\theta} q(\underline{\theta}) w(\underline{\theta}) \sum_x r'(x|\underline{\theta}) x \\ &= \int d\underline{\theta} p(\underline{\theta}) \sum_x r'(x|\underline{\theta}) x \\ &= \mathbb{E}_{p_{r'}}[x]. \end{aligned} \quad (3.42)$$

Therefore,

$$\begin{aligned} \mathbb{E}_{q_{r'}}[w(\underline{\theta})x] - \mathbb{E}_{p_r}[x] &= \mathbb{E}_{p_{r'}}[x] - \mathbb{E}_{p_r}[x] \\ &= \int d\underline{\theta} p(\underline{\theta}) \sum_x (r'(x|\underline{\theta}) - r(x|\underline{\theta})) x \\ &= \int d\underline{\theta} p(\underline{\theta}) \text{Tr}[O (\mathcal{N}_{\underline{\theta}}(\rho) - \mathcal{E}_{\underline{\theta}}(\rho))] \\ &= \mathbb{E}_p[\text{Tr}(O \varepsilon_{\underline{\theta}}(\rho))]. \end{aligned} \quad (3.43)$$

This quantity depends on p and on the channel approximation error, but not on the IS distribution q . $\square$

In summary, IS can reduce the computational resources required to reach a target precision, but it neither alleviates nor exacerbates the bias induced by imperfect state preparation or hardware noise. The expectation value of any observable remains governed by the original sampling distribution $p(\underline{\theta})$ and by the implemented noisy channel $\mathcal{N}_{\underline{\theta}}$.

7.3 Importance Sampling with Error Detection

In this section, we consider a scenario in which a device, after completing a circuit execution, outputs a classical *flag* $F \in \{\text{ok}, \text{err}\}$ indicating whether a detectable fault occurred. We refer to runs with $F = \text{ok}$ as successful and runs with $F = \text{err}$ as failed. We assume that the flag is perfectly reliable, with no false positives or negatives. If an error occurs, two situations can arise. In some cases, the error can be corrected, allowing the final outcome to be treated as if no error occurred. In these cases, the additional resources required for the error-correction procedure should be accounted for in the overall cost function. In other cases, errors can be detected but not corrected. This situation is common in current quantum devices, for example, in trapped-ion systems, where certain faults can be reliably flagged but not actively corrected [67–69]. Alternatively, errors may be detected through violations of conservation laws or symmetries that are preserved by the ideal, error-free evolution [70–73]. In the case of errors that can be detected but not corrected, various strategies can be employed. First we describe how these situations can be handled without IS. Consider a randomized algorithm for estimating $\text{Tr}[O \rho_{\text{target}}]$, where ρ_{target} can be decomposed as in Eq. (3.15) using a probability distribution $p(\underline{\theta})$. The simplest procedure is to sample a value of $p(\underline{\theta})$, apply the corresponding parametrized quantum channel $\mathcal{E}_{\underline{\theta}}$ on the initial state ρ , and measure O . If the flag is ok, the outcome is recorded; otherwise, the whole process is repeated with the same $\underline{\theta}$ until success. This is necessary in order for the samples $\underline{\theta}$ to come from the target distribution $p(\underline{\theta})$ and for the final state to be ρ_{target} . This in turn implies that the expected cost to record one successful outcome depends on both the cost associated with $\mathcal{E}_{\underline{\theta}}$ and the success probability $f(\underline{\theta}) > 0$ of the flag indicating ok. Since we are not using IS at this stage, the expected cost coincides with the net cost and reads

$$NC_p = \mathbb{E}_p \left[\frac{c(\underline{\theta})}{f(\underline{\theta})} \right], \quad (3.44)$$

as we will prove later in this section. When $f(\underline{\theta})$ is small for some $\underline{\theta}$, the corresponding circuit must be run many times on average before it does not flag an error. This increases the average cost of the algorithm. One may therefore impose a limit L on the number of repetitions, after which a new $\underline{\theta}$ is sampled from $p(\underline{\theta})$. However, this effectively alters the sampling distribution $p(\underline{\theta})$, and IS is required to maintain an unbiased estimator. In what follows we present two approaches that allow recording independent outcomes whose expectation values exactly correspond to $\text{Tr}[O \rho_{\text{target}}]$, differing in how they handle unsuccessful runs. We then analyze these methods and derive the optimal IS distributions that minimize the net cost for each protocol for a given target estimation error.

ZeroFill(L). In the ZeroFill(L) protocol, one first samples a parameter configuration $\underline{\theta}$ according to the distribution $q(\underline{\theta})$. The corresponding quantum circuit is then executed, incurring a cost $c(\underline{\theta})$ and subject to an error-detection mechanism. If no error is detected, which occurs with probability $f(\underline{\theta})$, the observable O is measured and the outcome is multiplied by a weight factor before being recorded. If instead

an error is detected, the execution is repeated using the same value of $\underline{\theta}$. This retry procedure is performed for at most L attempts. If all L executions fail, the protocol records the value 0 (Zero-Fill) and proceeds to a new parameter sample. The goal is to estimate $\mathbb{E}_p[O \mathcal{E}_{\underline{\theta}}(\rho)]$ from repeated samples. Unbiasedness is ensured by choosing the weight

$$w_z(\underline{\theta}; L) = \frac{p(\underline{\theta})}{q(\underline{\theta})k(\underline{\theta}; L)}, \quad (3.45)$$

where

$$k(\underline{\theta}; L) = 1 - (1 - f(\underline{\theta}))^L. \quad (3.46)$$

Discard(L). The Discard(L) protocol follows a similar structure, allowing up to L repeated executions for a fixed parameter configuration $\underline{\theta}$. Whenever an execution succeeds, the corresponding measurement outcome is recorded after reweighting. However, if all L attempts fail, no outcome is stored and the parameter configuration is discarded entirely before a new sample is drawn. In this case, unbiasedness requires the weight

$$w_D(\underline{\theta}; L) = \mathbb{E}_q[k(\underline{\theta}; L)] \frac{p(\underline{\theta})}{q(\underline{\theta})k(\underline{\theta}; L)}. \quad (3.47)$$

Equivalently,

$$w_D(\underline{\theta}; L) = \mathbb{E}_q[k(\underline{\theta}; L)] w_z(\underline{\theta}; L). \quad (3.48)$$

The performance of the two preceding algorithms can be characterized in terms of the expected cost per recorded outcome and the corresponding variance. As shown in the following theorem, both quantities admit the same formal expressions for the ZeroFill(L) and Discard(L) protocols when written in terms of their respective weight functions $w_z(\underline{\theta}; L)$ and $w_D(\underline{\theta}; L)$.

Theorem 4. Consider a target distribution $p(\underline{\theta})$, a cost function $c(\underline{\theta})$, and one of the error-handling schemes introduced above. Let $f(\underline{\theta})$ denote the probability that the circuit associated with $\underline{\theta}$ executes without detectable error, and let $q(\underline{\theta})$ be the sampling distribution.

Then, the expected cost required to obtain a single recorded outcome is

$$\mathbb{E}_p \left[\frac{c(\underline{\theta})}{w_{z,D}(\underline{\theta}; L) f(\underline{\theta})} \right], \quad (3.49)$$

while the corresponding worst-case variance, maximized over all observables satisfying $\|O\| \leq 1$, is

$$\mathbb{E}_p[w_{z,D}(\underline{\theta}; L)]. \quad (3.50)$$

Proof. A proof is provided in Appendix B. □

Although Eqs. (3.49) and (3.50) hold for both $\text{ZeroFill}(L)$ and $\text{Discard}(L)$, the two protocols generally exhibit different performance because their associated weight functions are not identical. In particular, the expected cost of $\text{ZeroFill}(L)$ equals that of $\text{Discard}(L)$ multiplied by the factor $\mathbb{E}_q[k(\underline{\theta}; L)]$, whereas the variance bound of $\text{ZeroFill}(L)$ is reduced by the inverse of the same factor. We also emphasize that Eq. (3.49) quantifies the expected cost required to obtain a successfully recorded outcome. If instead one is interested only in the cost of executing a single run of the algorithm, irrespective of whether the outcome is ultimately retained, the relevant quantity is simply

$$\mathbb{E}_q[c(\underline{\theta})]. \quad (3.51)$$

Corollary 1. For both the $\text{ZeroFill}(L)$ and $\text{Discard}(L)$ protocols, characterized by success probability $f(\underline{\theta})$ and sampling distribution $q(\underline{\theta})$, the net cost satisfies

$$NC_q \geq \left(\mathbb{E}_p \left[\sqrt{\frac{c(\underline{\theta})}{f(\underline{\theta})}} \right] \right)^2. \quad (3.52)$$

Equality holds if and only if the sampling distribution is given by

$$q_L^*(\underline{\theta}) \propto \frac{p(\underline{\theta})}{k(\underline{\theta}; L)} \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}}. \quad (3.53)$$

Proof. Consider the net cost

$$NC_q \geq \mathbb{E}_p \left[\frac{c(\underline{\theta})}{w_{Z,D}(\underline{\theta}; L) f(\underline{\theta})} \right] \mathbb{E}_p [w_{Z,D}(\underline{\theta}; L)],$$

obtained by multiplying the expected cost in Eq. (3.49) by the variance bound in Eq. (3.50). Applying the Cauchy–Schwarz inequality yields Eq. (3.52), with equality if and only if

$$w_{Z,D}(\underline{\theta}; L) \propto \sqrt{\frac{c(\underline{\theta})}{f(\underline{\theta})}}. \quad (3.54)$$

Finally, since both $\text{ZeroFill}(L)$ and $\text{Discard}(L)$ satisfy

$$q(\underline{\theta}) \propto \frac{p(\underline{\theta})}{k(\underline{\theta}; L) w_{Z,D}(\underline{\theta}; L)}, \quad (3.55)$$

the optimal sampling distribution in Eq. (3.53) follows immediately. $\square$

Remarkably, both the optimal net cost and the associated optimal sampling distribution q_L^* coincide for the $\text{ZeroFill}(L)$ and $\text{Discard}(L)$ protocols. Moreover, although the optimal sampling distribution depends explicitly on L , the resulting optimal net cost does not. In the limit $L \rightarrow \infty$, one has

$$k(\underline{\theta}; L) \rightarrow 1, \quad \mathbb{E}_q[k(\underline{\theta}; L)] \rightarrow 1, \quad (3.56)$$

and the weight functions reduce to the standard importance-sampling form

$$w(\underline{\theta}) = \frac{p(\underline{\theta})}{q(\underline{\theta})}. \quad (3.57)$$

In this regime, Theorem 4 implies that sampling directly from $p(\underline{\theta})$ yields a worst-case variance bound equal to 1, together with the expected cost given in Eq. (3.44), which is strictly larger than the optimal net cost achievable through importance sampling for any positive cost function. Finally, taking the limit $L \rightarrow \infty$ in Corollary 1 yields the optimal sampling distribution

$$q_{\infty}^*(\underline{\theta}) \propto p(\underline{\theta}) \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}}. \quad (3.58)$$

7.3.1 Case Study: IID Parameter Model with Poissonian Error Statistics

In this section, we present an analytically tractable instance of the general framework introduced in Section 7.1.1. We consider a collection of parameters $\{\theta_1, \dots, \theta_s\}$ modeled as independent and identically distributed random variables, and assume that the total cost function $c(\underline{\theta})$ is additive, with identical single-parameter contributions associated with each θ_i :

$$c(\underline{\theta}) = \sum_{i=1}^s c_1(\theta_i). \quad (3.59)$$

We further assume that the probability of observing no detectable errors follows a Poissonian ansatz depending on the total cost,

$$f(\underline{\theta}) = e^{-\lambda c(\underline{\theta})}, \quad \lambda > 0. \quad (3.60)$$

This model naturally describes situations in which the cost function faithfully captures the total circuit execution time and errors originate from stationary noise processes continuously acting on the quantum hardware. Under these assumptions, the probability of remaining error-free decays exponentially with the accumulated cost.

Theorem 5. Given a target distribution $p(\underline{\theta})$, let $q(\underline{\theta})$ be an arbitrary sampling distribution. Consider i.i.d. random variables $\underline{\theta} = (\theta_1, \dots, \theta_s)$ with additive cost

$$c(\underline{\theta}) = \sum_{i=1}^s c_1(\theta_i). \quad (3.61)$$

Assume that the detectable-error success probability is given by

$$f(\underline{\theta}) = e^{-\lambda c(\underline{\theta})}, \quad (3.62)$$

for some $\lambda \geq 0$. Define

$$\mu_1 = \mathbb{E}_p[c_1(\theta)], \quad \sigma_1^2 = \text{Var}_p[c_1(\theta)], \quad (3.63)$$

as the mean and variance of the single-variable cost function $c_1(\theta)$.

Using one of the protocols described in Section 7.3 with $L = \infty$, the ratio between the net cost obtained under the optimal sampling distribution $q^*(\theta)$ and that obtained under the original distribution $p(\theta)$ admits the asymptotic expansion

$$\begin{aligned} \frac{NC_{q^*}}{NC_p}(s) &= \left(\frac{\mu_1 + \frac{\lambda\sigma_1^2}{2}}{\mu_1 + \lambda\sigma_1^2} \right) e^{-s\left(\frac{\lambda\sigma_1}{2}\right)^2} \\ &\times \left[1 - \frac{1}{4s} \left(\frac{\sigma_1}{\mu_1 + \frac{\lambda\sigma_1^2}{2}} \right)^2 + \mathcal{O}(s^{-3/2}) \right]. \end{aligned} \quad (3.64)$$

Proof. A proof is provided in Appendix C. $\square$

In the absence of detectable errors, corresponding to $\lambda = 0$, one finds

$$\frac{NC_{q^*}}{NC_p}(s) \sim 1 - \frac{1}{s}, \quad s \rightarrow \infty, \quad (3.65)$$

so that the ratio converges to unity asymptotically. In this regime, importance sampling does not provide an asymptotic advantage. By contrast, for any $\lambda > 0$, the ratio decays exponentially with the system size:

$$\frac{NC_{q^*}}{NC_p}(s) \sim \exp\left[-s\left(\frac{\lambda\sigma_1}{2}\right)^2\right], \quad s \rightarrow \infty. \quad (3.66)$$

Therefore, whenever $\sigma_1 \neq 0$, the optimal importance-sampling strategy yields an exponentially vanishing net-cost ratio. Moreover, the quantity $NC_{q^*}/NC_p(s)$ admits a unique global maximum at a finite value of s , approximately given by

$$s^* \approx \frac{1}{8} \left(\frac{\sigma_1^2}{\mu_1 + \frac{\lambda\sigma_1^2}{2}} \right)^2 \left[1 + \sqrt{1 + \frac{64}{\lambda^2\sigma_1^4} \left(\mu_1 + \frac{\lambda\sigma_1^2}{2} \right)^2} \right]. \quad (3.67)$$

This value identifies the characteristic size of θ at which the relative net-cost reduction induced by importance sampling is weakest before the asymptotic exponential suppression becomes dominant.

7.4 Discussion

In this chapter, we developed a general framework that uses classical importance sampling as an outer optimization layer for randomized quantum algorithms. The main result is an analytical characterization of the optimal sampling distribution that

minimizes the net cost associated with estimating a target quantity, thereby promoting importance sampling from a heuristic variance-reduction technique to a principled resource-allocation tool for quantum computation. Within this formulation, the optimization acts solely on the classical sampling measure, while the underlying quantum channels and measurement procedures remain unchanged.

A central outcome of the analysis is that importance sampling preserves the bias induced by imperfect quantum implementations. In particular, when the ideal channels are replaced by noisy or approximate ones, the reweighting procedure compensates only for the change in the sampling distribution, not for the physical or algorithmic errors themselves. As a consequence, the framework can substantially reduce the expected resource consumption required to reach a target precision, but it cannot be used to mitigate coherent approximation errors or hardware noise. This distinction is essential for interpreting the practical scope of the method.

The framework is deliberately flexible with respect to the cost function. Depending on the platform and the computational regime, the relevant resource may be circuit depth, two-qubit gate count, non-Clifford cost, runtime, or another implementation-dependent metric. This makes the method adaptable to both near-term and fault-tolerant settings. Although the formulation is observable-independent at the level of the optimization problem, it naturally suggests observable-aware extensions in which prior information about the target quantity is used to refine the sampling strategy and improve the combined cost–variance trade-off.

The magnitude of the attainable improvement depends on the statistical structure of the randomized protocol. When the cost is additive across many independent layers, the total resource consumption concentrates around its mean and the relative gain from importance sampling vanishes asymptotically, as quantified by the large- s analysis in Section 7.1.1. By contrast, in the presence of error-detection schemes, the induced success probabilities can produce a much broader effective cost distribution, and the advantage of optimal sampling may become exponentially large in the system size. This contrast highlights a key limitation of the basic framework and, at the same time, points to a route for overcoming it through structured randomness, correlated sampling across layers, or other designs that avoid concentration effects.

The applications discussed in [3] illustrate the breadth of the framework, including randomized product formulas, mixed-state simulation, classical shadow estimation, and probabilistic error cancellation. Across these examples, the same principle applies: whenever a randomized quantum protocol admits a tunable classical sampling law and a meaningful cost model, importance sampling provides a systematic way to reduce the expected computational burden without altering the algorithmic target. For this reason, performance comparisons between deterministic and randomized formulations should always take the optimized randomized version into account, rather than the unoptimized one.

Overall, the results of this chapter establish importance sampling as a general resource-optimization mechanism for randomized quantum computation. The framework clarifies when cost reductions are possible, how the optimal distribution should be chosen, and which limitations are intrinsic to the method. These conclusions provide a foundation for further developments in resource-aware quantum algorithm

design, especially in settings where randomness is already part of the computational structure.

PART IV

OBSERVABLE MEASUREMENT

In the previous sections, we described how to initialize a quantum state ρ on a quantum register and unitarily evolve it to the target state ρ_{target} according to a prescribed Hamiltonian. The final step consists of performing measurements to extract physically relevant information, such as the expectation values of observables of interest. More precisely, let O be an observable with spectral decomposition

$$O = \sum_x x \Pi_x, \quad (4.1)$$

where Π_x denotes the projector onto the eigenspace associated with eigenvalue x . Our goal is to estimate its expectation value

$$\langle O \rangle := \text{Tr}[\rho_{\text{target}} O]. \quad (4.2)$$

If direct measurements of O are available, an estimator for $\langle O \rangle$ can be obtained by repeatedly executing the quantum circuit and averaging the measurement outcomes. However, in most practical settings, quantum devices are restricted to performing projective measurements in a limited set of bases, typically corresponding to local operators such as single-qubit Pauli Z measurements in the computational basis. Consequently, estimating the expectation value of a generic observable requires reducing the measurement of O to operations compatible with the available quantum resources. Several strategies have been developed to estimate expectation values on a quantum computer, and the most appropriate one depends on the available oracle access, the structure of the observable O , and the dominant resource constraint. In what follows, we assume $\delta > 0$ is the target accuracy to be achieved in the estimation of $\langle O \rangle$. As a first step, we survey the main approaches to estimating $\langle O \rangle$ in the various oracle settings encountered in practice. For each setting, we highlight the algorithms that are most naturally suited to the available access model.

i) A first approach is based on basis rotation. If an efficiently implementable unitary U is known such that $UOU^\dagger = D$ is diagonal in the computational basis, one can prepare ρ_{target} on a quantum register, apply U , and perform a projective measurement in the computational basis. The final estimator is obtained as a weighted average of the eigenvalues of D , where the weights are given by the empirical frequencies of the measurement outcomes. This strategy requires no ancilla qubits and introduces only the additional circuit depth associated with the implementation of U . Assuming bounded eigenvalue samples, the number of repetitions required to reach an accuracy δ scales as $\mathcal{O}(\|O\|^2/\delta^2)$ according to Hoeffding's inequality [74]. Its main limitation is that finding an efficient basis transformation is generally intractable for a generic observable, and even when such a transformation is known, its implementation may require a substantial circuit depth.

ii) A second class of approaches relies on the availability of a block encoding of the observable. Suppose that an efficient block encoding $U_O \in \text{BE}_{\alpha,a}(O)$ can be synthesized. In this case, U_O can be used as a primitive in a larger quantum circuit whose measurement statistics are related to the desired expectation value. The simplest example of this strategy is the Hadamard test [75–77], which estimates the relevant overlap, in the case of unitary O , through the success probability of a block encoding

of

$$\frac{1}{2}(1 + U_O) \quad (4.3)$$

acting on the state $\rho_{\text{target}} \otimes |0^a\rangle \langle 0^a|$, that is

$$p_{\text{success}} = \frac{1}{4} \text{Tr}[(1 + O)\rho_{\text{target}}(1 + O)^\dagger] = \frac{1}{2}(1 + \langle O \rangle). \quad (4.4)$$

The method requires one additional ancilla qubit, besides the internal ancillas involved in the block encoding, together with controlled access to U_O . Its sampling complexity scales as $\mathcal{O}(\alpha^2/\delta^2)$, where α is the block-encoding factor. The main advantage of this approach is its conceptual simplicity and its applicability whenever a block encoding can be constructed. However, controlled operations increase the circuit depth, the normalization factor α directly affects the sampling cost, and the construction of U_O may itself introduce significant ancilla and synthesis overhead. A more advanced technique is amplitude estimation [78], which improves the dependence on the target precision from $\mathcal{O}(1/\delta^2)$ to $\mathcal{O}(1/\delta)$ by embedding Hadamard-test-type circuits into a quantum phase estimation procedure. In this setting, one typically requires $\mathcal{O}(\log(\alpha/\delta))$ ancilla qubits for phase readout and $\mathcal{O}(\alpha/\delta)$ queries to the block encoding. The total circuit depth scales as

$$\mathcal{O}(\alpha/\delta) \times \text{depth}(U_O). \quad (4.5)$$

The main advantage of amplitude estimation is the quadratic improvement in query complexity with respect to sampling-based approaches. However, this improvement comes at the cost of substantially deeper circuits and the requirement of coherent access to many sequential applications of U_O .

iii) Finally, if an efficient block encoding of the time-evolution operator e^{-iOt} is available, quantum phase estimation (QPE) [79] provides a natural alternative. QPE samples from the spectral distribution of O induced by the state $|\psi\rangle$: the algorithm returns an eigenvalue λ_k with probability p_k , and the expectation value is recovered as the empirical mean of the sampled eigenvalues. This approach requires $m = \mathcal{O}(\log(1/\delta))$ ancilla qubits and coherent controlled simulations of e^{-iOt} up to evolution times $t = \mathcal{O}(1/\delta)$, leading to an overall query complexity of order $\mathcal{O}(1/\delta)$. Its main advantage is that it provides access to the full spectral distribution of O with respect to $|\psi\rangle$, allowing the estimation of higher moments and other spectral quantities without additional asymptotic cost. Its main limitation is the substantial circuit depth required, as it relies on long coherent time evolutions. A detailed explanation of the QPE routine can be found in the following chapter.

So far we have been discussing a general framework for selecting an estimation strategy according to the available oracle access. In practical applications, however, additional structure of the observable or the quantum state can be exploited. Suppose, for instance, that O admits a decomposition

$$O = \sum_i c_i O_i, \quad c_i \in \mathbb{R}. \quad (4.6)$$

By linearity, the expectation value can be obtained by separately estimating the expectation values of the individual operators O_i using any of the previously described methods, and subsequently combining the results according to

$$\langle O \rangle = \sum_i c_i \langle O_i \rangle. \quad (4.7)$$

An important example of this approach is the Pauli decomposition, in which the operators O_i correspond to Pauli strings. Such decompositions are widely used in quantum simulation and variational quantum algorithms because Pauli operators can often be measured efficiently. However, the decomposition is not always readily available, and constructing it may be challenging or impractical for certain observables. Furthermore, if the number of terms grows exponentially with the system size, estimating each contribution separately becomes computationally prohibitive. The decomposition in Eq. (4.6) also provides a natural framework for optimizing the allocation of measurement resources. In particular, the available measurement budget can be distributed among the different terms according to importance-sampling schemes proportional to $|c_i|$, or through more refined strategies that account for the variance of each contribution and the specific computational constraints of the problem [3, 80]. Moreover, when several expectation values $\langle O_i \rangle$ are required simultaneously, classical shadow techniques [81] can be employed to reconstruct multiple observables from the same measurement data, especially when the operators share locality properties or admit efficient simultaneous reconstruction. A particularly simple case arises when the observable is proportional to a projector, namely $O = c\Pi$ with $c \in \mathbb{R}$. In this situation, the expectation value can be rewritten as

$$\langle \psi | O | \psi \rangle = c \langle \psi | \Pi | \psi \rangle \quad (4.8)$$

$$= c^{-1} \langle \psi | O^\dagger O | \psi \rangle \quad (4.9)$$

$$= c^{-1} \|O|\psi\rangle\|^2. \quad (4.10)$$

This expression shows that the expectation value is directly related to the success probability associated with a block encoding of O . Consequently, when such a block encoding is available, the estimation problem can be addressed even more directly than through a Hadamard-test construction.

In the next two chapters, we focus on estimation problems in which the quantity of interest is either the expectation value of the system Hamiltonian or the energy difference between two eigenstates. Chapter 8 is devoted to the standard Quantum Phase Estimation (QPE) algorithm and its application to energy estimation. Chapter 9 instead introduces a NISQ-friendly protocol for estimating energy differences directly, without requiring the individual energy eigenvalues to be computed.

CHAPTER 8

QUANTUM PHASE ESTIMATION

The Quantum Phase Estimation (QPE) algorithm, originally introduced by Kitaev [79] and subsequently formalized in the quantum circuit model by Cleve *et al.*, is a fundamental primitive of quantum computation. The task of the protocol is to estimate the phase ϕ associated with an eigenvalue of a unitary operator U , defined through the eigenvalue equation

$$U |\psi\rangle = e^{2\pi i \phi} |\psi\rangle, \quad (4.11)$$

where $|\psi\rangle$ is an eigenstate of U and $\phi \in [0, 1)$. The objective is to extract an n -bit approximation of ϕ with high probability using a quantum circuit whose complexity scales polynomially in n . The QPE algorithm exploits quantum parallelism and interference: an ancillary register of n qubits is initialized in a uniform superposition via Hadamard gates, while a second register contains the eigenstate $|\psi\rangle$, assumed to be available as an input resource. Controlled applications of successive powers of the unitary operator U^{2^k} are then performed, as illustrated in Fig. 8.1, thereby encoding the binary expansion of ϕ into the relative phases of the computational basis states of the ancillary register. This step emphasizes a crucial aspect of the protocol, namely that the implementation of controlled-unitary operations, often referred to as oracle access to U , constitutes a fundamental resource alongside the preparation of the eigenstate. Following this phase kickback procedure, one is left with the quantum state

$$\frac{1}{2^{n/2}} \sum_{k=0}^{2^n-1} |k\rangle \otimes U^k |\psi\rangle = \frac{1}{2^{n/2}} \sum_{k=0}^{2^n-1} e^{2\pi i k \phi} |k\rangle \otimes |\psi\rangle, \quad (4.12)$$

where the second equality follows directly from the eigenvalue equation. At this stage, it is sufficient to observe that the structure of the ancillary register in Eq. (4.12) matches that of Eq. (1.41). One can therefore apply the inverse Quantum Fourier Transform (QFT⁻¹) to the ancillary register, which decodes the accumulated phase information into a binary representation accessible via measurement. The measurement outcome yields an estimate $\tilde{\phi}$ such that $|\tilde{\phi} - \phi| \leq 2^{-n}$ with probability bounded




Figure 8.1: Quantum circuit for the Quantum Phase Estimation algorithm. The top n qubits constitute the ancillary register. After preparation in a uniform superposition, they control successive powers of the unitary operator U acting on the eigenstate register. Finally, the inverse Quantum Fourier Transform is applied to the ancillary register before measurement.

away from zero, which can be further amplified through standard repetition techniques. From the perspective of computational complexity, QPE provides an exponential advantage over classical phase estimation methods in regimes where $|\psi\rangle$ can be efficiently prepared and U can be efficiently simulated on a quantum device, since classical approaches typically require resources that scale exponentially with the desired precision. This exponential improvement underlies a broad class of quantum algorithms, most notably Shor's factoring algorithm, as well as quantum algorithms for eigenvalue problems in chemistry and physics. In the latter applications, the unitary oracle U is commonly realized as the time-evolution operator

$$U := e^{-iH\delta t}, \quad (4.13)$$

generated by a target Hamiltonian H over a fixed time interval δt such that

$$\|H\|\delta t < 2\pi. \quad (4.14)$$

The eigenstate $|\psi\rangle$ is then identified with an eigenvector $|E\rangle$ of H , associated with an energy eigenvalue E to be estimated. As emphasized above, efficient preparation of $|E\rangle$ is required to estimate E via QPE, together with an efficient implementation of the unitary evolution U , highlighting the relevance of the topics discussed in Part II and Part III. In other words, the performance of QPE critically depends on both the efficient realization of controlled powers of the unitary operator and the availability of its eigenstate, conditions that are generally nontrivial and often constitute the primary bottleneck in practical implementations. Nevertheless, within these constraints, Quantum Phase Estimation remains one of the central algorithmic tools demonstrating the computational power of quantum information processing.

CHAPTER 9

SPECTRAL GAP ESTIMATION VIA ADIABATIC STATE PREPARATION

The spectral gap, defined as the energy difference between two eigenstates (typically the ground state and the first excited state) of a quantum system, is a central quantity in both condensed matter physics and quantum chemistry. In condensed matter systems, the spectral gap governs low-temperature properties, stability, and phase transitions [38, 82]. In particular, it plays a key role in determining transport properties: systems with a large gap typically behave as insulators, whereas metallic systems are characterized by a vanishing gap at the Fermi level [83]. Near critical points, small or closing gaps are associated with drastic changes in physical behavior and the emergence of critical phenomena [84, 85]. In quantum chemistry, energy gaps between molecular orbitals, such as the highest occupied molecular orbital and the lowest unoccupied molecular orbital, determine molecular stability, reactivity, and optical properties [86, 87]. These quantities directly influence chemical processes including catalysis, bonding, and photochemical transitions [88]. Despite their importance, computing spectral gaps for large or strongly correlated systems remains a challenging task on classical computers. The exponential growth of the Hilbert space with system size makes exact diagonalization infeasible, while additional complications arise in fermionic systems due to the sign problem [89].

Quantum simulation offers a potential route to overcome these limitations. In a fault-tolerant setting, a natural strategy for estimating the spectral gap consists in independently applying QPE to the relevant eigenstates and subsequently taking the difference between the measured eigenvalues. Owing to the use of the QFT, this approach achieves an exponential improvement in precision scaling compared to classical diagonalization methods. Nevertheless, it still presents several conceptual and practical limitations. From a conceptual perspective, if the quantity of interest is solely the spectral gap, estimating the two individual eigenvalues may constitute an unnecessary use of computational resources. This suggests the possibility of alternative strategies capable of accessing the gap more directly, without reconstructing the corresponding energies separately. Moreover, from a practical perspective, QPE requires repeated controlled applications of the unitary evolution operator. Although

the query complexity scales only logarithmically with the target precision, the overall implementation cost can still become prohibitively large, particularly in many-body systems where each realization of the time evolution is itself computationally demanding. Consequently, despite its asymptotic efficiency, the direct application of QPE may remain impractical in physically relevant regimes. In particular, during the NISQ era [90], available quantum devices are constrained by limited qubit counts and coherence times, typically rendering QPE infeasible. To address these limitations, several methods for estimating spectral gaps on NISQ devices have been proposed, including imaginary-time evolution [91], robust phase estimation [92], and Bayesian phase-difference estimation [93]. In Ref. [4], we develop an alternative approach that avoids the explicit estimation of energy eigenvalues. The method is particularly well suited to NISQ devices, as it replaces deep quantum circuits with a larger number of shallow circuit executions, effectively trading circuit depth for measurement repetitions. The method is based on the temporal fluctuations of expectation values of observables and consists of two steps.

i) First, the coherent superposition

$$|\Psi(i, j)\rangle = \frac{1}{\sqrt{2}}(|E_i\rangle + |E_j\rangle), \quad (4.15)$$

where $|E_i\rangle$ and $|E_j\rangle$ are eigenstates of the system Hamiltonian H with eigenvalues E_i and E_j , is prepared on quantum hardware, for instance using B-SAP [2]. The prepared state is then evolved under H for a time t , and the expectation value of an arbitrary observable O is measured. By repeating this procedure for different evolution times, one obtains an estimate of

$$\langle O(t) \rangle = \frac{1}{2} (\langle E_i | O | E_i \rangle + \langle E_j | O | E_j \rangle) + \mathcal{A} \cos(\Delta E_{ji} t + \theta), \quad (4.16)$$

where $\Delta E_{ji} := E_i - E_j$, while $\mathcal{A} \geq 0$ and $\theta \in [0, 2\pi)$ correspond to the polar representation of the complex transition matrix element $\langle E_j | O | E_i \rangle$. Operationally, this first quantum step consists of preparing the superposition state, evolving it under H for a set of times $\{t_r\}$, and measuring the corresponding expectation values $\langle O(t_r) \rangle$, thereby generating a dataset.

ii) The energy gap is subsequently extracted from this dataset, either through Fourier analysis or, more efficiently in the NISQ regime, by fitting the measured signal to a sinusoidal model. The choice of observable is not restrictive, as any operator satisfying $\mathcal{A} \neq 0$ allows the energy gap to be accessed. However, selecting observables that maximize the oscillation amplitude improves the signal-to-noise ratio and enhances the robustness of the estimation.

This approach enables, in principle, the estimation of energy differences between arbitrary eigenstates without requiring full spectral reconstruction. Moreover, it is particularly advantageous for systems where direct energy measurements are costly due to the complexity of the Hamiltonian, e.g., its number of terms in a decomposition of the form of Eq. (4.6). In what follows we report numerical simulations and experiments on quantum hardware which demonstrate the feasibility of the methods within current technological constraints.

9.1 Ising Model

We consider the Ising model on a square lattice in d dimensions with periodic boundary conditions (PBC). The system is described by the Hamiltonian

$$H = -\frac{J_1}{2} \sum_{\langle ij \rangle} \sigma_i^1 \sigma_j^1 - \frac{h_3}{2} \sum_i \sigma_i^3, \quad (4.17)$$

where σ_i denotes the Pauli operators acting on site i , and $\langle ij \rangle$ indicates a sum over nearest neighbors. We focus on the ferromagnetic regime $J_1 > 0$ and aim to estimate the spectral gap between the ground state and the first excited state.

As a first step, we need to construct the superposition state $|\Psi(0, 1)\rangle$. To this end, we introduce the auxiliary Hamiltonian

$$H_0 = -\frac{J_1}{2} \sum_{\langle ij \rangle} \sigma_i^1 \sigma_j^1, \quad (4.18)$$

which corresponds to Eq. (4.17) with $h_3 = 0$. The ground-state subspace of H_0 is twofold degenerate and is spanned by $\{|+\rangle^{\otimes L}, |-\rangle^{\otimes L}\}$, where

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle). \quad (4.19)$$

We can employ B-SAP to determine coefficients $c_{i,+}$ and $c_{i,-}$ such that the states

$$|E_i(0)\rangle := c_{i,+} |+\rangle^{\otimes L} + c_{i,-} |-\rangle^{\otimes L} \quad (4.20)$$

are mapped, under adiabatic evolution, to the eigenstates

$$|E_i\rangle = \lim_{\tau \rightarrow \infty} U_\tau(1) |E_i(0)\rangle. \quad (4.21)$$

By linearity, this implies

$$|\Psi(0, 1)\rangle = \lim_{\tau \rightarrow \infty} U_\tau(1) \frac{1}{\sqrt{2}} [(c_{0,+} + c_{1,+}) |+\rangle^{\otimes L} + (c_{0,-} + c_{1,-}) |-\rangle^{\otimes L}], \quad (4.22)$$

so that $|\Psi(0, 1)\rangle$ can be prepared by initializing the appropriate superposition in the degenerate subspace and subsequently performing adiabatic evolution. For the present problem, symmetry considerations allow one to determine the coefficients $\{c_{i,\pm}\}$ without resorting to MC-VQE in the B-SAP procedure. Specifically, consider the operator

$$\mathcal{R} = \bigotimes_{j=0}^{L-1} R_z^j(\pi), \quad (4.23)$$

where $R_z^j(\theta)$ denotes a rotation about the z -axis by an angle θ acting on qubit j . This operator commutes with the Ising Hamiltonian for all values of the coupling parameters, and its expectation value is therefore conserved during the adiabatic evolution. The states $|+\rangle^{\otimes L}$ and $|-\rangle^{\otimes L}$ do not diagonalize $\mathcal{R}$, since

$$\mathcal{R} |\pm\rangle^{\otimes L} = (-i)^L |\mp\rangle^{\otimes L}. \quad (4.24)$$

Introducing instead

$$|\Phi^\pm\rangle = \frac{1}{\sqrt{2}}(|+\rangle^{\otimes L} \pm |-\rangle^{\otimes L}), \quad (4.25)$$

one finds

$$\mathcal{R}|\Phi^\pm\rangle = \pm(-i)^L |\Phi^\pm\rangle. \quad (4.26)$$

Thus, the ground-state subspace of H_0 is spanned by $|\Phi^+\rangle$ and $|\Phi^-\rangle$, which are eigenstates of $\mathcal{R}$ with eigenvalues $(-i)^L$ and $-(-i)^L$, respectively. Let's now consider the ground and first excited states of H , arising from the lifting of the degeneracy in the ground-state subspace of H_0 . In the limit $h_3/J_1 \rightarrow \infty$, the ground state of H is given by

$$|\Omega_0\rangle \equiv |0\rangle^{\otimes L}, \quad (4.27)$$

which is an eigenstate of $\mathcal{R}$ with eigenvalue $(-i)^L$. Since $\mathcal{R}$ is conserved under the adiabatic evolution, it follows that $|\Omega_0\rangle$ must originate from $|\Phi^+\rangle$, i.e.,

$$\lim_{\tau \rightarrow \infty} U_\tau |\Phi^+\rangle = |\Omega_0\rangle. \quad (4.28)$$

By orthogonality and conservation of $\mathcal{R}$, the complementary state satisfies

$$\lim_{\tau \rightarrow \infty} U_\tau |\Phi^-\rangle = |\Omega_1\rangle, \quad (4.29)$$

where $|\Omega_1\rangle$ denotes the first excited state of H within the same symmetry sector. This identification extends to any finite value of h_3/J_1 . To see this, note that the adiabatic evolution targeting a Hamiltonian with finite h_3/J_1 can be regarded as an initial segment of a longer adiabatic path that asymptotically reaches the limit $h_3/J_1 \rightarrow \infty$. Since $\mathcal{R}$ commutes with the Hamiltonian at all times along this path, its eigenvalues label invariant symmetry sectors that cannot mix during the evolution. In the asymptotic limit $h_3/J_1 \rightarrow \infty$, the ground state is uniquely identified as $|\Omega_0\rangle = |0\rangle^{\otimes L}$ and carries eigenvalue $(-i)^L$ under $\mathcal{R}$. Therefore, the adiabatic continuation of the initial state $|\Phi^+\rangle$ must remain within the same symmetry sector and evolve into the instantaneous ground state throughout the entire path. Since no level crossings are expected between different symmetry sectors, this correspondence holds already at any finite value of h_3/J_1 . In other words, the state originating from $|\Phi^+\rangle$ continuously tracks the ground state of H for all $h_3/J_1 > 0$, while $|\Phi^-\rangle$ analogously evolves into the first excited state within the complementary symmetry sector. Finally, using linearity and the relation $|+\rangle^{\otimes L} = \frac{1}{\sqrt{2}}(|\Phi^+\rangle + |\Phi^-\rangle)$, we obtain

$$|\Psi(0, 1)\rangle = \lim_{\tau \rightarrow \infty} U_\tau |+\rangle^{\otimes L}. \quad (4.30)$$

Now that we know how to prepare the quantum state on the register, we are left with the choice of the observable. Based on numerical tests we employ

$$O = \sigma^1 \otimes (\sigma^0)^{\otimes(L-1)}, \quad (4.31)$$

which appears to maximize $\mathcal{A}$ in Eq. (4.16) (Appendix E). The experiments involve 4-site and 20-site Ising chains with PBC. The circuit depth per Trotter step is independent of L , reaching a maximum of 80 two-qubit gates when decomposed into IonQ




Figure 9.1: Spectral gap as a function of h_3/J_1 for (a) $L = 4$ and (b) $L = 20$ Ising chains with PBC. Simulations include classical noiseless results and experiments on IonQ Aria 1. Red-shaded regions indicate adiabatic and infrared limits where circuit depth is insufficient. Error bars represent a single standard deviation from the fit.

native gates (Appendix H). We use 40 Trotter steps, divided into 15 adiabatic and 25 time-evolution steps. To evaluate $O(t)$, the 25 times t are chosen as Chebyshev nodes [94, 95] to accurately capture oscillations. Each circuit is simulated classically with 8192 shots, incorporating intrinsic measurement errors (Appendix F), and compared against exact diagonalization of H .

Figure 9.1 shows strong agreement between noiseless simulations and numerical benchmarks, with relative errors of order 10^{-2} . The infrared and adiabatic limits correspond to regimes where our approach becomes unstable due to Trotter approximation constraints or long adiabatic evolution times. In the infrared limit, the $O(t)$ frequency approaches zero, requiring many time steps. In the adiabatic limit, fidelity decreases with increasing transverse field, sometimes exhibiting unexplained oscillations (Appendix G). The method demonstrates robustness against IonQ-inspired noise, with thermal relaxation and readout contributing minimally, and two-qubit gate fidelities being the main source of error (Appendix H). The average relative error under this noise model is of order 10^{-1} , and performance remains stable with reduced shot counts. Analogous results for 4-site 2D Ising lattices are presented in Appendix I. Experiments on the IonQ Aria 1 device for 4- and 20-site chains were performed under hardware constraints, limiting measurements to three h_3/J_1 values for 4 qubits and two for 20 qubits, avoiding infrared and adiabatic limits. Only 10 of the 25 originally chosen times t were used for fitting $O(t)$, and 2500 shots (with error mitigation) replaced the 8192-shot classical simulations. Results (Figure 9.1) agree well with noiseless and noisy simulations, demonstrating the method’s resilience. Integrated error mitigation on the device further improved accuracy [96].




Figure 9.2: Spectral gap for (a) H_2 and (b) He_2 as a function of bond length. Noiseless simulations match benchmark values closely. Error bars represent a single standard deviation from the fit.

9.2 H_2 and He_2 Molecules

The spectral gap estimation was applied to the diatomic molecules H_2 and He_2 . The electronic Hamiltonian in second quantization is

$$\mathcal{H} = \sum_{pq} h_{pq} c_p^\dagger c_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} c_p^\dagger c_q^\dagger c_r c_s, \quad (4.32)$$

with coefficients h_{pq} and h_{pqrs} obtained via Hartree-Fock [97, 98]. Mapping to qubits via Jordan-Wigner gives

$$H = \sum_{i_1, \dots, i_L} a_{i_1, \dots, i_L} \bigotimes_{j=1}^L \sigma^{i_j}, \quad i_j \in \{0, 1, 2, 3\}. \quad (4.33)$$

The auxiliary Hamiltonian H_0 is taken as the diagonal of H , making eigenstates computational basis states. The superposition $|\Psi(0, 1)\rangle$ is constructed from the ground and first-excited states of H_0 . Observables for these systems are tensor products of identities and σ^1 matrices, satisfying $|\langle \omega_1 | O | \omega_0 \rangle| = 1$ (Appendix E). Numerical simulations used 40 Trotter steps (5 adiabatic, 35 evolution). H_2 requires 4 qubits; He_2 requires 8. Energy gaps as functions of bond length are shown in Figure 9.2. Noiseless simulations closely match benchmarks, with deviations at large bond lengths for H_2 due to imperfect state preparation. He_2 results remain robust across the tested range. On the other hand, circuit depth considerations make noisy simulations on current hardware impractical: 2-qubit gate counts reach 180 for H_2 and 6640 for He_2 , yielding low fidelity.

9.3 Remarks and Outlook

The protocol presented in this chapter was primarily conceived for the NISQ and early fault-tolerant regimes, where coherent circuit depth constitutes the principal computational bottleneck. Its main advantage lies in replacing long coherent quantum

evolutions with a larger number of independent shallow circuit executions, making it considerably more resilient to decoherence and gate imperfections than standard implementations of quantum phase estimation [99, 100].

This practical advantage, however, comes at the price of a less favorable asymptotic scaling with the target precision. To resolve the spectral gap $\Delta = E_1 - E_0$ with accuracy ϵ , the Fourier uncertainty principle dictates a maximum coherent evolution time per circuit execution of

$$T_{\max} = \mathcal{O}(\epsilon^{-1}). \quad (4.34)$$

However, because the spectral gap is extracted via classical signal processing of noisy expectation values $\langle O(t) \rangle$ [101, 102], statistical sampling noise (shot noise) propagates into the frequency fitting via the Cramér-Rao bound. Consequently, the total accumulated runtime across all circuit executions scales according to $\mathcal{O}(\epsilon^{-2})$. In contrast, Quantum Phase Estimation (QPE) [99] processes phase information coherently using an ancilla register of $m = \mathcal{O}(\log \epsilon^{-1})$ qubits. By avoiding the shot-noise penalty associated with repeated expectation-value measurements, QPE achieves the optimal Heisenberg-limited total runtime scaling of $\mathcal{O}(\epsilon^{-1})$, making it asymptotically superior in the fully fault-tolerant regime.

These considerations suggest that the present approach and quantum phase estimation should not be regarded as competing algorithms, but rather as complementary techniques tailored to different hardware regimes. While the former is particularly attractive for near-term and early fault-tolerant quantum devices, the latter ultimately becomes preferable once sufficiently large, fully fault-tolerant quantum computers become available.

An important direction of my current research is the development of a fully fault-tolerant protocol for spectral gap estimation that retains the Heisenberg-limited asymptotic scaling of quantum phase estimation while avoiding the unnecessary estimation of individual eigenvalues. The central observation is that the spectral gap is itself the quantity of interest, whereas conventional QPE reconstructs the two eigenenergies independently before taking their difference. By exploiting this additional structure, the goal is to directly estimate the energy difference, thereby combining the optimal precision scaling with a more resource-efficient use of quantum memory and gate depth. Preliminary investigations indicate that this approach may substantially reduce the overhead associated with conventional phase-estimation-based gap evaluation.

CONCLUSIONS

This thesis has addressed the problem of designing and optimizing quantum algorithms for the simulation of many-body quantum systems, with a particular focus on the efficient use of computational resources across the full quantum simulation pipeline, from state preparation to Hamiltonian evolution and measurement. The starting point of this work is the observation that quantum computation does not merely extend classical computational paradigms, but reorganizes the way in which information is represented and processed. While classical hardware relies on increasingly sophisticated but fundamentally classical control of physical systems, quantum computation exploits coherent quantum dynamics directly. In this setting, quantum simulation emerges as one of the most natural and impactful applications, since the exponential complexity of quantum many-body systems makes their classical description intractable beyond modest sizes. Within this framework, a generic quantum simulation protocol can be decomposed into three fundamental stages: state preparation, Hamiltonian evolution, and measurement. Each of these stages introduces distinct algorithmic and physical constraints, and each has been the subject of dedicated developments in this thesis.

Part II of the thesis focused on quantum state preparation. After establishing its central role in quantum simulation and its intrinsic computational hardness in general settings, we analyzed existing approaches based on Variational Quantum Algorithms and Adiabatic Preparation. While VQAs provide flexible low-depth ansätze at the cost of difficult optimization landscapes, adiabatic methods offer a conceptually simple dynamical route to eigenstate preparation but are severely limited by small spectral gaps and long evolution times. Building on these observations, two main contributions were developed. First, an optimal-control-inspired modification of adiabatic state preparation was introduced, in which a preconditioning term is designed to reshape the adiabatic path and reduce the characteristic preparation time. This approach provides a systematic route to accelerating adiabatic protocols and demonstrates exponential improvements in representative many-body models. Second, a hybrid strategy combining variational and adiabatic ideas was developed. By incorporating group-theoretic structure into the variational ansatz and constraining the evolution within symmetry-resolved subspaces, the resulting Branched-Subspaces Adiabatic Preparation (B-SAP) algorithm achieves efficient preparation of both ground and excited states with improved scalability compared to standard methods.

Part III addressed Hamiltonian evolution, focusing on operator splitting tech-

niques as a central paradigm for quantum simulation. Product formulas, including Trotter and higher-order Trotter–Suzuki decompositions, provide a systematic way of approximating complex unitary dynamics by sequences of simpler exponentials. Beyond deterministic constructions, this part introduced a broader perspective based on randomized protocols and classical optimization of quantum resources. In particular, a general framework for importance sampling in randomized quantum algorithms was developed. Within this framework, the cost of a quantum protocol is quantified through a net-cost functional combining circuit execution cost and statistical variance. It was shown that optimal sampling strategies reducing the computational overhead while preserving the estimator bias can be derived analytically. This approach was applied to a variety of settings, including Hamiltonian simulation protocols such as QDrift, as well as more general tasks involving noisy channels and measurement-based estimators. The resulting formalism provides a unified perspective on how classical sampling theory can be used to systematically optimize quantum algorithm performance.

Part **IV** focused on observable measurement, with particular emphasis on the estimation of spectral properties of quantum systems. While standard approaches to spectral gap estimation rely on independent computation of low-lying energy eigenvalues via methods such as quantum phase estimation, these techniques can be resource-intensive and conceptually indirect when only energy differences are of interest. To address this limitation, a direct gap estimation method based on adiabatic state preparation was introduced. The key idea is to construct a specific superposition state whose time-dependent measurement statistics encode the energy gap through characteristic oscillatory behavior. By analyzing these temporal fluctuations, the spectral gap can be extracted via classical post-processing and fitting procedures. The method, which relies on the accurate initialization of special quantum states on the quantum register, was validated numerically on prototypical spin systems and small molecular Hamiltonians and experimentally on real quantum hardware for systems of moderate size. These results indicate that meaningful spectral information can be extracted using shallow circuits, suggesting a viable path toward gap estimation in noisy intermediate-scale quantum devices and, in the long term, more complex applications in fault-tolerant architectures.

Taken together, the results of this thesis highlight a unifying theme: the efficiency of quantum simulation is not determined solely by the availability of quantum hardware, but crucially by the structure of the algorithms used to exploit it. Across state preparation, evolution, and measurement, the central challenge is the same: reducing the effective resource cost while preserving the physical information encoded in quantum states. From this perspective, several directions emerge naturally for future research. On the state preparation side, further exploration of symmetry-aware and control-theoretic methods may lead to more systematic constructions of low-depth circuits for physically relevant classes of Hamiltonians. For Hamiltonian evolution, extending importance sampling techniques to more sophisticated fully fault-tolerant resource models, where T -gate counts dominate complexity, appears particularly promising. Finally, in the context of measurement, a natural direction for future work is the development of more general frameworks for extracting spectral and

dynamical information directly from time-dependent observables. In particular, extending these approaches to fault-tolerant architectures is an active line of research that I am currently pursuing, with the goal of making such methods compatible with large-scale quantum computation.

Overall, this work contributes to the broader effort of transforming quantum simulation from a primarily conceptual advantage into a practically efficient computational tool, by systematically addressing algorithmic bottlenecks across all stages of the computation.

APPENDICES

APPENDIX A

EXPONENTIAL BEHAVIOR OF THE ADIABATIC ERROR

In this section we review the proof of the exponential bound provided in the main text. We begin by extending the domain of $H(s)$ to $s < 0$ and $s > 1$

$$H(s) = \begin{cases} H_0 & \text{for } s < 0 \\ [1 - f(s)]H_0 + f(s)H_1 & \text{for } 0 \leq s \leq 1 \\ H_1 & \text{for } s > 1 \end{cases} \quad (5.1)$$

and by defining the (adiabatic) time evolution operator $U_\tau(s_1, s_0)$ from the differential equation

$$i \frac{dU_\tau}{ds}(s_1, s_0) = \tau H(s_1) U_\tau(s_1, s_0). \quad (5.2)$$

Integrating both sides and using recursively Eq. (5.2) one obtains

$$U_\tau(s_1, s_0) = T \exp \left[-i\tau \int_{s_0}^{s_1} ds H(s) \right], \quad (5.3)$$

where T is the time ordering operator. In what follows we assume that $\forall s \in \mathbb{R}$, the subset $\sigma_0(s)$ of the full spectrum $\sigma(s) = \sigma_0(s) \cup \sigma_1(s)$ of $H(s)$ is isolated and bounded, i.e.

$$\begin{cases} d(s) = \inf_{x \in \sigma_0(s), y \in \sigma_1(s)} |x - y| \geq d_{\min} > 0 \\ D(s) = \sup_{x, y \in \sigma_0(s)} |x - y| \leq D_{\max} < \infty. \end{cases} \quad (5.4)$$

Let $P(s)$ be the projector on the eigenspace associated with $\sigma_0(s)$. The goal of this section is to establish a bound for the adiabatic error

$$\epsilon_{\text{AP}}(\tau) \equiv \|(\mathbb{1} - P(s_1)) U_\tau(s_1, s_0) P(s_0)\|, \quad (5.5)$$

with an exponential falloff in the large τ regime. If we sum and subtract a generic projector $G(s; \tau)$ and we repeatedly use the triangular inequality together with the

unitarity of U_τ , we obtain

$$\begin{aligned} \epsilon_{\text{AP}}(\tau) &\leq \|P(s_1) - G(s_1; \tau)\| + \|P(s_0) - G(s_0; \tau)\| + \|U_\tau^\dagger(s_1, s_0)(1 - G(s_1; \tau))U_\tau(s_1, s_0)\| \\ &\leq \|P(s_1) - G(s_1; \tau)\| + \|P(s_0) - G(s_0; \tau)\| + \int_{s_0}^{s_1} ds \delta(s; \tau), \end{aligned} \quad (5.6)$$

with the auxiliary function

$$\delta(s; \tau) \equiv \left\| \frac{dG}{ds}(s; \tau) + i\tau[H(s), G(s; \tau)] \right\|. \quad (5.7)$$

It is evident that if one can prove the existence of an operator $G(s; \tau)$ satisfying

$$G(s; \tau) = P(s) \quad \forall \tau, \quad \forall s \notin [0, 1], \quad (5.8)$$

the first two terms in Eq. (5.6) cancel out from the interval $[0, 1]$. Additionally, if $\delta(s; \tau)$ falls off exponentially in the limit $\tau \rightarrow \infty$ it follows that, for any $s_0 < 0$ and $s_1 > 1$, also $\epsilon_{\text{AP}}(\tau)$ decreases at least exponentially in τ .

We now turn to the explicit construction of $G(s; \tau)$ subject to condition (5.8) and to the requirement that $\delta(s; \tau)$ has an exponential falloff. At this scope we introduce the index of a generic operator A

$$\text{Ind}_\gamma(A) \equiv \frac{1}{2\pi i} \oint_\gamma \frac{dz}{A - z}, \quad (5.9)$$

where γ is a closed curve in the complex plane. Notice that this definition is also valid for a complex number. Using the definition above and starting from a generic operator $T(s; \tau)$, we may define its projector $G(s; \tau) \equiv \text{Ind}_\gamma(T(s; \tau))$ from an anti-clockwise curve γ such that $\text{Ind}_\gamma(1) = 1$ and $\text{Ind}_\gamma(0) = 0$, as represented in Fig. A.1. Clearly the previous expression may be ill-defined when $T(s; \tau)$ has eigenvalues laying on the curve γ and this eventuality is discussed and excluded below. Notice also that the relation between T and G is not bijective, in fact there exist infinitely many operators T giving the same projector G . In what follows we rewrite conditions (5.8) and (5.7) in terms of T . From now on we denote the k -th derivative of a generic operator A with respect to the variable s with

$$A^{(k)}(s) \equiv \frac{d^k A}{ds^k}(s), \quad (5.10)$$

with the convention that $A^{(0)}(s) = A(s)$.

I Eq. (5.8) implies that

$$G(s; \tau) - P(s) = \frac{1}{2\pi i} \oint_\gamma dz \left(\frac{1}{T(s; \tau) - z} - \frac{1}{P(s) - z} \right) \quad (5.11)$$




Figure A.1: Gray dots in the figure represent the eigenvalues of T . The thick circumference is a possible choice for the curve γ .

should vanish for $s \notin [0, 1]$. In other words

$$\boxed{\text{Ind}_\gamma(T(s; \tau)) = \text{Ind}_\gamma(P(s)) \quad \forall s \notin [0, 1].} \quad (5.12)$$

II We hereby tackle eq. (5.7) by expressing $\delta(s; \tau)$ in terms T

$$\begin{aligned} \delta(s; \tau) &= \frac{1}{2\pi} \left\| \oint_\gamma dz \frac{1}{T(s; \tau) - z} (T^{(1)}(s; \tau) + i\tau [H(s), T(s; \tau)]) \frac{1}{T(s; \tau) - z} \right\| \\ &\leq \frac{L_\gamma}{2\pi} \left[\sup_{z \in \gamma} \left\| \frac{1}{T(s; \tau) - z} \right\| \right]^2 \|T^{(1)}(s; \tau) + i\tau [H(s), T(s; \tau)]\| \end{aligned} \quad (5.13)$$

where L_γ is the length of the curve γ . $\delta(s; \tau)$ is guaranteed to decrease exponentially in τ if there exists an s -independent constant K such that, for τ large enough,

$$\boxed{\sup_{z \in \gamma} \left\| \frac{1}{T(s; \tau) - z} \right\| \leq K < \infty} \quad (5.14)$$

and

$$\boxed{\lim_{\tau \rightarrow \infty} \tau^{-1} \ln \|T^{(1)}(s; \tau) + i\tau [H(s), T(s; \tau)]\| < 0.} \quad (5.15)$$

Notice that condition (5.14) also implies that the definition $G(s; \tau)$ is well-defined.

So far we have obtained three conditions for $T(s; \tau)$. In particular, Eq. (5.14) is equivalent to claiming the existence of $0 \leq K < \infty$ s.t. $K^{-1} \leq \inf_{z \in \gamma} \|T(s; \tau) - z\|$ for $\tau \gg 0$. To simplify the derivation, we choose γ to be a circumference centred in $z = 1$ with radius $1/2$, as represented in Fig. A.2; it follows that

$$\inf_{z \in \gamma} \|T(s; \tau) - z\| \geq \inf_{z \in \gamma} \| \|P(s) - z\| - \|T(s; \tau) - P(s)\| \| = \left| \frac{1}{2} - \|T(s; \tau) - P(s)\| \right|, \quad (5.16)$$




Figure A.2: Gray dots in the figure represent the eigenvalues of T . The thick circumference, is a possible choice for the curve γ . The erased region is made up from those points whose distance from γ is lower than K^{-1} . Condition in eq. (5.14) requires, for τ large enough, that T has no eigenvalues in this region.

and hence condition (5.14) is guaranteed if

$$\sup_s \|T(s; \tau) - P(s)\| < \frac{1}{2} \quad \forall \tau \quad \cup \quad \inf_s \|T(s; \tau) - P(s)\| > \frac{1}{2} \quad \forall \tau. \quad (5.17)$$

We now focus on the condition in Eq. (5.15) and look for an operator $T(s; \tau)$ whose dependence on s and τ can be factorized as

$$T(s; \tau) \equiv \sum_{j=0}^{N(s; \tau)} \frac{O_j(s)}{\tau^j}. \quad (5.18)$$

One could just set N to infinity, but this choice would require to deal with the norm of an infinite series in eq. (5.17). We postpone the choice of $N(s; \tau)$, instead, and treat it as a generic step function of s parametrized by τ . Outside the support of $N^{(1)}(s; \tau)$, i.e. almost everywhere, by means of the triangular inequality one obtains

$$\begin{aligned} \|T^{(1)}(s; \tau) + i\tau [H(s), T(s; \tau)]\| &\leq \tau \| [H(s), O_0(s)] \| + \sum_{j=0}^{N(s; \tau)} \frac{\|O_j^{(1)}(s) + i[H(s), O_{j+1}]\|}{\tau^j} \\ &\quad + \frac{\|O_{N(s; \tau)}^{(1)}\|}{\tau^{N(s; \tau)}}. \end{aligned} \quad (5.19)$$

In order to have an exponential convergence to zero, we set any term on the r.h.s but the last one to zero, and exploit the arbitrariness on $N(s; \tau)$ to exponentially suppress it. At this point we impose the system of equations

$$\begin{cases} [H(s), O_0(s)] = 0 \\ O_j^{(1)}(s) + i[H(s), O_{j+1}] = 0 \quad \forall j \in \mathbb{N}, \end{cases}$$




Figure A.3: The figure represents the poles of $R(s, z)$, as a function of z , for a fixed value of s . In particular, gray dots correspond to the eigenvalues in $\sigma_0(H(s))$, while black dots correspond to those in $\sigma_1(H(s))$. Here we consider systems with a discrete spectrum, e.g. field theories defined in a finite spatial box.

admitting the following solutions

$$\begin{cases} O_0(s) = P(s) \\ O_j = Q(s)S_j(s)Q(s) - P(s)S_j(s)P(s) + \frac{1}{2\pi} \oint_{\Gamma(s)} \left(R(s, z) [O_{j-1}^{(1)}(s), P(s)] R(s, z) \right) \end{cases} \quad (5.20)$$

with

$$Q(s) = \mathbb{1} - P(s), \quad S_j(s) \equiv \sum_{m=1}^{j-1} O_m(s) O_{j-m}(s), \quad R(s, z) \equiv [H(s) - z]^{-1}, \quad (5.21)$$

and $\Gamma(s)$ a closed anti-clockwise curve surrounding only the eigenvalues in $\sigma_0(s)$ while keeping a distance of at least $\frac{d(s)}{2}$ from the spectrum (see Fig. A.3), formally

$$\begin{cases} \text{Ind}_{\Gamma(s)}(E) = 1 & \text{if } E \in \sigma_0(s), \\ \text{Ind}_{\Gamma(s)}(E) = 0 & \text{if } E \in \sigma_1(s), \\ |z - E| \geq \frac{d(s)}{2} & \forall z \in \Gamma(s), \forall E \in \sigma(s). \end{cases} \quad (5.22)$$

Using these results Eq. (5.15) takes the form

$$\lim_{\tau \rightarrow \infty} \tau^{-1} \ln \left\| \frac{O_{N(s;\tau)}^{(1)}(s)}{\tau^{N(s;\tau)}} \right\| < 0 \quad (5.23)$$

subject to the constraints

$$\sup_s \left\| \sum_{j=1}^{N(s;\tau)} \frac{O_j(s)}{\tau^j} \right\| < \frac{1}{2} \quad \forall \tau \quad \cup \quad \inf_s \left\| \sum_{j=1}^{N(s;\tau)} \frac{O_j(s)}{\tau^j} \right\| > \frac{1}{2} \quad \forall \tau \quad (5.24)$$

valid for large τ and derived from Eq. (5.19). Our goal is to find bounds for the operators $O_j^{(k)}$ for Hamiltonians $H(s)$ of the form in Eq. (5.1), that we will use to demonstrate that both Eq. (5.23) and Eq. (5.24) hold, with a proper choice of $N(s; \tau)$. Since $P(s)$ can be expressed in terms of $R(s, z)$ as

$$P(s) = \frac{1}{2\pi i} \oint_{\Gamma(s)} dz R^{-1}(s, z) \quad (5.25)$$

and the O_j operators can be calculated from Eq. (5.20), we start from evaluating

$$r_k(s) \equiv \sup_{z \in \Gamma(s)} \|R^{(k)}(s, z)\| \leq \frac{k!}{2\pi} \oint_{C(s)} \sup_{z \in \Gamma(s)} \left\| \frac{R(\xi, z) d\xi}{(\xi - s)^{k+1}} \right\|, \quad (5.26)$$

where we have used the Cauchy theorem for the k -th derivative. In particular, $C(s)$ can be chosen as the circumference in the complex plane centred in s with radius $\rho > 0$. Using the fact the $f(s)$ is holomorphic and switching to polar coordinates one obtains

$$\begin{aligned} r_k(s) &\leq \frac{k!}{2\pi} \left(\frac{1}{\rho} \right)^k \int_0^{2\pi} \sup_{z \in \Gamma(s)} \|R(s + \rho e^{i\theta}, z)\| d\theta = \\ &= \frac{k!}{2\pi} \left(\frac{1}{\rho} \right)^k \int_0^{2\pi} \sup_{z \in \Gamma(s)} \left\| [H(s) - z + (H(s + \rho e^{i\theta}) - H(s))]^{-1} \right\| d\theta = \\ &= \frac{k!}{2\pi} \left(\frac{1}{\rho} \right)^k \int_0^{2\pi} \sup_{z \in \Gamma(s)} \left\| [H(s) - z + \Delta(f(s + \rho e^{i\theta}) - f(s))]^{-1} \right\| d\theta. \end{aligned} \quad (5.27)$$

Since $f(s)$ is continuous we can take ρ small enough so that

$$\|H(s) - z\| \geq |f(s + \rho e^{i\theta}) - f(s)| \|\Delta\| \quad \forall \theta \in [0, 2\pi), \forall z \in \Gamma(s). \quad (5.28)$$

Hence, denoting with $|\Psi\rangle$ the normalized vectors of the Hilbert space, we obtain

$$\begin{aligned} &\left\| [H(s) - z + \Delta(f(s + \rho e^{i\theta}) - f(s))]^{-1} \right\| = \\ &= \left(\inf_{|\Psi\rangle} |\langle \Psi | H(s) - z + \Delta(f(s + \rho e^{i\theta}) - f(s)) | \Psi \rangle| \right)^{-1} \leq \\ &\leq \left(\inf_{|\Psi\rangle} |\langle \Psi | H(s) - z | \Psi \rangle| - \sup_{|\Psi\rangle} |\langle \Psi | \Delta(f(s + \rho e^{i\theta}) - f(s)) | \Psi \rangle| \right)^{-1} = \\ &= \left(\inf_{|\Psi\rangle} |\langle \Psi | H(s) - z | \Psi \rangle| - |f(s + \rho e^{i\theta}) - f(s)| \|\Delta\| \right)^{-1}, \end{aligned} \quad (5.29)$$

that, together with the hypothesis that the spectrum of $H(s)$ has a distance of at least $\frac{d(s)}{2}$ from any z along the curve $\Gamma(s)$, gives

$$r_k(s) \leq \frac{k!}{2\pi} \left(\frac{1}{\rho} \right)^k \int_0^{2\pi} \left(\frac{d(s)}{2} - |f(s + \rho e^{i\theta}) - f(s)| \|\Delta\| \right)^{-1} d\theta. \quad (5.30)$$

We use again the continuity of $f(s)$ to impose a tighter condition on ρ , i.e.

$$|f(s + \rho e^{i\theta}) - f(s)| \|\Delta\| < \frac{d(s)}{4} \quad \forall \theta, \quad (5.31)$$

that is compatible with the previous one in Eq. (5.28). Finally, we define

$$\varrho(s) \equiv \sup \left\{ \rho > 0 \text{ s.t. } |f(s + \rho e^{i\theta}) - f(s)| \|\Delta\| < \frac{d(s)}{4} \quad \forall \theta \right\}. \quad (5.32)$$

and use it as the radius for $C(s)$, obtaining

$$\boxed{r_k(s) \leq k! \left(\frac{1}{\varrho(s)} \right)^k \frac{4}{d(s)}}. \quad (5.33)$$

Knowing the bounds for r_k , we proceed to the study of $\|O_j^{(k)}\|$, for which we introduce three important Lemmas. $\forall k \in \mathbb{N} \quad \exists F_k \leq \frac{4\pi^2}{3}$ such that

$$\sum_{l=0}^k \frac{1}{(l+1)^2(k-l+1)^2} \leq \frac{F_k}{(k+1)^2}. \quad (5.34)$$

Proof. The proof is obtained by the straightforward calculation

$$\begin{aligned} \sum_{l=0}^k \frac{1}{(l+1)^2} \frac{1}{(k-l+1)^2} &= \left(\sum_{l=0}^{\lfloor k/2 \rfloor} + \sum_{l=\lfloor k/2 \rfloor+1}^k \right) \frac{1}{(l+1)^2} \frac{1}{(k-l+1)^2} \\ &\leq \frac{2}{(\lfloor k/2 \rfloor + 1)^2} \sum_{l=0}^{\infty} \frac{1}{(l+1)^2} \leq \frac{8}{(k+1)^2} \frac{\pi^2}{6}. \end{aligned} \quad (5.35)$$

In particular, one can find explicitly that $F_0 = 1$ and $F_1 = 2$. $\square$

Let $M(s), N(s) \in \mathcal{C}^\infty$ be operators, $p, q \in \mathbb{N}$, and the functions $m(s), n(s), a(s) \in \mathbb{R}$ such that

$$\|M^{(k)}(s)\| \leq m(s) a^{p+k}(s) \frac{(p+k)!}{(k+1)^2}, \quad (5.36)$$

$$\|N^{(k)}(s)\| \leq n(s) a^{q+k}(s) \frac{(q+k)!}{(k+1)^2}. \quad (5.37)$$

for any integer k . Using the constants F_k introduced above, it follows that

$$\|(MN)^{(k)}(s)\| \leq F_k m(s) n(s) a^{p+q+k}(s) \frac{(p+q+k)!}{(k+1)^2}. \quad (5.38)$$

Proof. Using the standard notation for the binomial coefficients

$$\binom{k}{l} \equiv \frac{k!}{(k-l)!l!}, \quad (5.39)$$

and the hypothesis of this Lemma, from the Leibnitz formula we find

$$\begin{aligned} \|(MN)^{(k)}(s)\| &\leq \sum_{l=0}^k \binom{k}{l} \|M^{(l)}(s) N^{(k-l)}(s)\| \leq \\ &\leq \sum_{l=0}^k \binom{k}{l} m(s) n(s) a^{p+q+k} \frac{(p+l)!(q+k-l)!}{(l+1)^2(k-l+1)^2} = \\ &= m(s) n(s) a^{p+q+k} (p+q+k)! \sum_{l=0}^k \binom{k}{l} \frac{(p+l)!(q+k-l)!}{(p+q+k)!} \frac{1}{(l+1)^2(k-l+1)^2}. \end{aligned} \quad (5.40)$$

Considering that

$$\binom{k}{l} \frac{(p+l)!(q+k-l)!}{(p+q+k)!} = \frac{\binom{k}{l}}{\binom{p+q+k}{p+l}} \leq 1, \quad (5.41)$$

and using the previous Lemma we conclude the proof. $\square$

The derivation described above, despite some formal differences, is substantially following the original proof provided by Nenciu in [49]. The novelty of our result, stated in the lemma and in the theorem below, stands in establishing the connection between the bound for $\|O_j^{(k)}\|$ with the function $g(s)$ defined in Eq. (10) of the main text and therefore with the explicit properties of the Hamiltonian $H(s)$.

Consider $\varrho(s)$ defined by Eq. (5.32), the energy gap $d(s)$ defined in Eq. (5.4) and let L_Γ be the length of the curve Γ , used in the contour integral of Eq. (5.20). Then, $\forall j \in \mathbb{N} \setminus \{0\}$ and $k \in \{0, 1\}$, the bound

$$\|O_j^{(k)}(s)\| \leq X_j(s) \left(\frac{4}{\varrho(s)} \right)^{k+j} \frac{(k+j)!}{(k+1)^2} \quad (5.42)$$

holds for a set of coefficients $\{X_j(s)\}_{j \in \mathbb{N} \setminus \{0\}}$ satisfying the recursive condition

$$\begin{cases} X_1 \geq 2^2 \left(\frac{L_\Gamma}{2\pi} \right)^3 \left(\frac{4}{d} \right)^4 \\ X_j \geq 2^4 \left(\frac{L_\Gamma}{2\pi} \right)^2 \left(\frac{4}{d} \right)^2 \left[\left(\sum_{m=1}^{j-1} X_m X_{j-m} \right) + \left(\frac{4}{d} \right) X_{j-1} \right]. \end{cases} \quad (5.43)$$

Proof. We proceed by induction starting from $j = 1$, a simpler case since S_1 does not contribute. Using typical inequalities, namely

$$\|O_1^{(k)}(s)\| \leq \frac{1}{2\pi} \oint \left\| (R(s, z) [P^{(1)}, P] R(s, z))^{(k)} \right\| dz, \quad (5.44)$$

and exploiting Lemma A, from the following three considerations (for simplicity we omit the s -dependence)

- $\|R^{(k)}\| \leq r_k \leq \left(\frac{4}{d} \right) \left(\frac{1}{\varrho} \right)^k k! = \left(\frac{4}{d} \right) \left(\frac{4}{\varrho} \right)^k \frac{k!}{(k+1)^2}$
- $\|P^{(k)}\| \leq \left(\frac{L_\Gamma}{2\pi} \right) r_k \leq \left(\frac{L_\Gamma}{2\pi} \right) \left(\frac{4}{d} \right) \left(\frac{4}{\varrho} \right)^k \frac{k!}{(k+1)^2}$
- $\|P^{(k+1)}\| = \left(\frac{L_\Gamma}{2\pi} \right) r_{k+1} \leq \frac{1}{4} \left(\frac{L_\Gamma}{2\pi} \right) \left(\frac{4}{d} \right) \left(\frac{4}{\varrho} \right)^{k+1} \frac{(k+1)!}{(k+1)^2},$

one gets

$$\|O_1^{(k)}(s)\| \leq 2^2 \left(\frac{L_\Gamma}{2\pi} \right)^3 \left(\frac{4}{d} \right)^4 \left(\frac{4}{\varrho} \right)^{k+1} \frac{(k+1)!}{(k+1)^2}. \quad (5.45)$$

Then, substituting the first relation in Eq. (5.43) inside the above inequality, the bound in Eq. (5.42) is immediately satisfied, proving the Lemma for $j = 1$. Let us now suppose that the same property holds for any positive integer lower than j , then from Eq. (5.20)

$$\|O_j^{(k)}\| \leq \|[QS_j Q]^{(k)}\| + \|[PS_j P]^{(k)}\| + \frac{1}{2\pi} \oint_\Gamma \left\| (R(s, z) [O_{j-1}^{(1)}, P] R(s, z))^{(k)} \right\| dz. \quad (5.46)$$

Once again taking Lemma A together with the following three observations

- $\|R^{(k)}\| \leq \left(\frac{4}{d}\right) \left(\frac{4}{\varrho}\right)^k \frac{k!}{(k+1)^2}$
- $\|Q^{(k)}\| = \|P^{(k)}\| \leq \left(\frac{L_\Gamma}{2\pi}\right) r_k \leq \left(\frac{L_\Gamma}{2\pi}\right) \left(\frac{4}{d}\right) \left(\frac{4}{\varrho}\right)^k \frac{k!}{(k+1)^2}$
- $\|O_{j-1}^{(k+1)}\| \leq X_{j-1} \left(\frac{4}{\varrho}\right)^{k+j} \frac{(k+j)!}{(k+1)^2}$

leads to the system of inequalities

$$\begin{cases} \left\| [QS_jQ]^{(k)} \right\| \leq 2^3 \left(\frac{L_\Gamma}{2\pi}\right)^2 \left(\frac{4}{d}\right)^2 \left(\sum_{m=1}^{j-1} X_m X_{j-m}\right) \left(\frac{4}{\varrho}\right)^{k+j} \frac{(k+j)!}{(k+1)^2} \\ \left\| [PS_jP]^{(k)} \right\| \leq 2^3 \left(\frac{L_\Gamma}{2\pi}\right)^2 \left(\frac{4}{d}\right)^2 \left(\sum_{m=1}^{j-1} X_m X_{j-m}\right) \left(\frac{4}{\varrho}\right)^{k+j} \frac{(k+j)!}{(k+1)^2} \\ \frac{1}{2\pi} \oint_\Gamma \left\| \left(R(s, z) [O_{j-1}^{(1)}, P] R(s, z) \right)^{(k)} \right\| dz \leq 2^4 \left(\frac{4}{d}\right)^3 \left(\frac{L_\Gamma}{2\pi}\right)^2 X_{j-1} \left(\frac{4}{\varrho}\right)^{k+j} \frac{(k+j)!}{(k+1)^2} \end{cases} \quad (5.47)$$

Now Eq. (5.46) simplifies to

$$\|O_j^{(k)}\| \leq 2^4 \left(\frac{L_\Gamma}{2\pi}\right)^2 \left(\frac{4}{d}\right)^2 \left[\left(\sum_{m=1}^{j-1} X_m X_{j-m}\right) + \left(\frac{4}{d}\right) X_{j-1} \right] \left(\frac{4}{\varrho}\right)^{k+j} \frac{(k+j)!}{(k+1)^2}, \quad (5.48)$$

concluding the proof thanks to the constraints satisfied by hypothesis by the X_j . $\square$

Theorem 6. $\forall j \geq 1$ and $k \in \{0, 1\}$,

$$\|O_j^{(k)}(s)\| \leq (j+k)! g^{j+k}(s) \quad (5.49)$$

where

$$g(s) = 4 \frac{W(d(s), D(s))}{\varrho(s, d(s), \|\Delta\|)} \quad (5.50)$$

with $d(s)$ and $D(s)$ defined in eq. (5.4)

Proof. In order to prove the theorem we start from the result of Lemma A. We define

$$\frac{4}{\varrho(s)} \equiv \frac{g(s)}{W(s)} \quad (5.51)$$

with $W(s) \geq 1$ and large enough so that, $\forall i \geq 1$,

$$\frac{X_i(s)}{(W(s))^i} \leq 1. \quad (5.52)$$

We look for coefficients $\{X_i\}_{i \in \mathbb{N}}$, of the form $X_i(s) = W^{i-\alpha}(s) Y_i(s)$, with $\alpha \in (0, 1)$ and $0 \leq Y_i(s) \leq 1$, respecting the hypothesis of Lemma A and such that Eq. (5.52) is satisfied. We start from verifying if this choice is compatible with the second condition in (5.43), which becomes

$$W^{-\alpha} 2^4 \left(\frac{L_\Gamma}{2\pi}\right)^2 \left(\frac{4}{d}\right)^2 \left[\frac{\sum_{m=1}^{j-1} Y_m Y_{j-m}}{Y_j} + \left(\frac{4}{d}\right) \frac{Y_{j-1}}{Y_j} \right] \leq 1. \quad (5.53)$$

Clearly for fixed j it is always possible to choose W large enough so that the inequality holds. Therefore we choose

$$Y_i = \frac{1}{(i+1)^2}, \quad \text{satisfying } \frac{Y_{j-1}}{Y_j} \leq 4, \quad (5.54)$$

such that Lemma A guarantees that $\sum_{m=1}^{j-1} Y_m Y_{j-m} \leq \frac{8\pi^2}{6} Y_{j-1}$. With little algebra we bound the term in the square brackets of eq. (5.53) and solve it for W , arriving at the condition

$$W \geq \left[2^6 \left(\frac{L_\Gamma}{2\pi} \right)^2 \left(\frac{4}{d} \right)^2 \left(\frac{\pi^2}{3} + \frac{4}{d} \right) \right]^{1/\alpha}. \quad (5.55)$$

Similarly, our ansatz for the coefficients X_j must obey the condition for X_1 used in the proof of Lemma A, which translates into

$$W^{1-\alpha} \geq 2^4 \left(\frac{L_\Gamma}{2\pi} \right)^3 \left(\frac{4}{d} \right)^4. \quad (5.56)$$

Putting everything together, we find that all constraints are satisfied by

$$W = \max \left(\left[2^6 \left(\frac{L_\Gamma}{2\pi} \right)^2 \left(\frac{4}{d} \right)^2 \left(\frac{\pi^2}{3} + \frac{4}{d} \right) \right]^{1/\alpha}, \left[2^4 \left(\frac{L_\Gamma}{2\pi} \right)^3 \left(\frac{4}{d} \right)^4 \right]^{1/(1-\alpha)}, 1 \right). \quad (5.57)$$

In particular, if we fix $\alpha = \frac{3}{7}$, the first quantity is always greater than the second one. Moreover we can choose $L_\Gamma = \pi d + 2D$, and obtain one of our main results

$$W = \left[2^8 \left(1 + 2 \frac{D}{\pi d} \right)^2 \left(\frac{\pi^2}{3} + \frac{4}{d} \right) \right]^{7/3}. \quad (5.58)$$

which concludes the proof since

$$\|O_j^{(k)}\| \leq X_j(s) \left(\frac{4}{\varrho(s)} \right)^{k+j} \frac{(k+j)!}{(k+1)^2} \leq \frac{X_j}{W^j} g^{j+k} \frac{(j+k)!}{(k+1)^2} \leq g^{j+k} (j+k)!. \quad (5.59)$$

□

In what follows we use the result of Theorem 6 to prove that conditions eE. (5.24) and Eq. (5.23) hold in the interval $s \in [0, 1]$ for any holomorphic $f(s)$. We start from applying it to

$$\begin{aligned} \frac{\|O_{N(s;\tau)}^1\|}{\tau^{N(s;\tau)}} &\leq (N(s;\tau) + 1)! \frac{g^{N(s;\tau)+1}(s)}{\tau^{N(s;\tau)}} \\ &= (N(s;\tau) + 1)! \exp[-N(s;\tau) \ln \tau + (N(s;\tau) + 1) \ln g(s)] \\ &= \exp \left[N(s;\tau) \ln \frac{g(s)(N(s;\tau) + 1)}{\tau} + \ln g(s) + \ln(N(s;\tau) + 1) - (N(s;\tau) + 1) \right] \times \\ &\quad \times \sqrt{2\pi N(s;\tau)} (1 + O(N^{-1}(s;\tau))), \end{aligned} \quad (5.60)$$

where we have used Stirling's approximation. The exponential convergence is achieved in the $\tau \rightarrow \infty$ limit if

$$\ln \frac{g(s)(N(s; \tau) + 1)}{\tau} \leq 0, \quad \text{i.e.} \quad N(s; \tau) \leq \frac{\tau}{g(s)} - 1. \quad (5.61)$$

Neglecting higher order terms, under this condition the bound above simplifies to

$$\frac{\|O_{N(s; \tau)}^1\|}{\tau^{N(s; \tau)}} \leq \sqrt{2\pi N(s; \tau)} \exp[\ln g(s) + \ln(N(s; \tau) + 1) - (N(s; \tau) + 1)] \quad (5.62)$$

whose convergence may be accelerated by taking the greatest step function $N(s; \tau)$ subject to (5.61), i.e.

$$N(s; \tau) = \left\lfloor \frac{\tau}{g(s)} \right\rfloor - 1. \quad (5.63)$$

This choice results in a bound for Eq. (5.13) of the form

$$\delta(s, \tau) \leq \frac{K^2}{2} \exp\left[-\left\lfloor \frac{\tau}{g(s)} \right\rfloor\right] \tau \sqrt{\frac{2\pi\tau}{g(s)}} \left(1 + O\left(\frac{g(s)}{\tau}\right)\right), \quad (5.64)$$

that directly leads to

$$\epsilon_{\text{AP}}(\tau) \leq C \int_0^1 ds \exp(-\lfloor \tau/g(s) \rfloor), \quad C \in \mathbb{R}, \quad (5.65)$$

namely eq. (5) of the main text. Clearly $g(s)$, as defined in Theorem 6, plays the role of a characteristic time. The final check consists in verifying that our choice for $N(s; \tau)$ also satisfies Eq. (5.24) for large τ . From Theorem 6 we get

$$\left\| \sum_{j=1}^{N(s; \tau)} \frac{O_j(s)}{\tau^j} \right\| \leq \sum_{j=1}^{N(s; \tau)} \frac{\|O_j(s)\|}{\tau^j} \leq \frac{g}{\tau} \sum_{j=1}^{N(s; \tau)} j! \left(\frac{g(s)}{\tau}\right)^{j-1}. \quad (5.66)$$

Relying once more on Stirling's approximation, and recalling that $j \leq N(s; \tau) \leq \tau/g(s) - 1$, we conclude that the series above is convergent

$$\sum_{j=1}^{N(s; \tau)} j! \left(\frac{g(s)}{\tau}\right)^{j-1} \leq \sum_{j=1}^{N(s; \tau)} \exp\left[j(\ln j - 1) - (j-1) \ln \frac{\tau}{g(s)}\right] \leq \sum_{j=1}^{N(s; \tau)} j e^{-j} \leq \sum_{j=1}^{\infty} j e^{-j} < +\infty. \quad (5.67)$$

By demonstrating its convergence we automatically guarantee that for large τ , the l.h.s. of Eq. (5.66) can be made arbitrarily small and, in particular, smaller than $\frac{1}{2}$.

APPENDIX B

PROOF OF THEOREM 4

We begin with a proof for the ZeroFill(L) algorithm. For a fixed value of $\underline{\theta}$, each circuit attempt succeeds with probability $f(\underline{\theta})$. Since the circuit is repeated at most L times, the probability that all L runs fail is $(1 - f(\underline{\theta}))^L$. Therefore, the probability that $\underline{\theta}$ yields at least one successful measurement is

$$k(\underline{\theta}; L) := 1 - (1 - f(\underline{\theta}))^L.$$

Assume that $p(\underline{\theta})$ is absolutely continuous with respect to $q(\underline{\theta})k(\underline{\theta}; L)$, the estimator is unbiased for every observable O if

$$\mathbb{E}_{qr} [w_z(\underline{\theta})k(\underline{\theta}; L)x] = \mathbb{E}_{pr} [x].$$

This holds (a.e.) if and only if

$$w_z(\underline{\theta}) = \frac{p(\underline{\theta})}{q(\underline{\theta})k(\underline{\theta}; L)} = \frac{p(\underline{\theta})}{q(\underline{\theta})[1 - (1 - f(\underline{\theta}))^L]}. \quad (5.68)$$

To characterize the average cost per run, we note the following. Given a $\underline{\theta}$, the number of attempts required to record an outcome is $\min(J, L)$, where $J \sim \text{Geom}(f(\underline{\theta}))$ counts the number of trials until the first success, and L is the fixed positive integer limiting the maximum number of allowed trials for each $\underline{\theta}$. Using the tail-sum formula,

$$\mathbb{E}[\min(J, L)|\underline{\theta}] = \sum_{j=1}^{\infty} \Pr[\min(J, L) \geq j|\underline{\theta}]. \quad (5.69)$$

For $1 \leq j \leq L$, the event $\min(J, L) \geq j$ is equivalent to $J \geq j$. Hence

$$\Pr[\min(J, L) \geq j|\underline{\theta}] = \Pr[J \geq j|\underline{\theta}] = (1 - f(\underline{\theta}))^{j-1}. \quad (5.70)$$

For $j > L$, we cannot have $\min(J, L) \geq j$, so the probability is zero. Thus, the expected number of circuit runs given a $\underline{\theta}$ is

$$\mathbb{E}[\min(J, L)|\underline{\theta}] = \sum_{j=1}^L (1 - f(\underline{\theta}))^{j-1} = \frac{1 - (1 - f(\underline{\theta}))^L}{f(\underline{\theta})} = \frac{k(\underline{\theta}; L)}{f(\underline{\theta})}, \quad (5.71)$$

and the expected cost is $\frac{c(\underline{\theta})k(\underline{\theta};L)}{f(\underline{\theta})}$. Averaging over the sampling distribution $q(\underline{\theta})$ gives the expected cost per run as

$$\mathbb{E}_q \left[c(\underline{\theta}) \frac{k(\underline{\theta};L)}{f(\underline{\theta})} \right]. \quad (5.72)$$

For the Variance, we have

$$\text{Var}_{qr}^{\max} \leq \max_{|x| \leq 1} \mathbb{E}_{qr} [k(\underline{\theta};L)w_z^2(\underline{\theta})x^2] = \mathbb{E}_q [k(\underline{\theta};L)w_z^2(\underline{\theta})]. \quad (5.73)$$

Therefore,

$$NC_q = \mathbb{E}_q \left[c(\underline{\theta}) \frac{k(\underline{\theta};L)}{f(\underline{\theta})} \right] \mathbb{E}_q [k(\underline{\theta};L)w_z^2(\underline{\theta})]. \quad (5.74)$$

Using the condition (5.68), we perform the change of measure to obtain

$$NC_q = \mathbb{E}_p \left[\frac{c(\underline{\theta})}{w_z(\underline{\theta})f(\underline{\theta})} \right] \mathbb{E}_p [w_z(\underline{\theta})] \geq \left(\mathbb{E}_p \left[\sqrt{\frac{c(\underline{\theta})}{f(\underline{\theta})}} \right] \right)^2, \quad (5.75)$$

where the last inequality follows from the Cauchy-Schwarz inequality, with equality if and only if $w_z(\underline{\theta}) \propto \sqrt{c(\underline{\theta})/f(\underline{\theta})}$ which implies:

$$q(\underline{\theta}) \propto \frac{p(\underline{\theta})}{k(\underline{\theta};L)} \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}} = \frac{p(\underline{\theta})}{1 - (1 - f(\underline{\theta}))^L} \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}}. \quad (5.76)$$

We now provide a proof for the Discard(L) algorithm. In this algorithm $\underline{\theta}$ is sampled from $q(\underline{\theta})$ and if errors are detected in all L runs of the corresponding quantum circuit, a new parameter $\underline{\theta}$ is sampled from $q(\underline{\theta})$. This is in fact rejection sampling and results in $\underline{\theta}$ begin sampled from a new distribution given by:

$$h_L(\underline{\theta}) := \frac{q(\underline{\theta})k(\underline{\theta},L)}{\mathbb{E}_q[k(\underline{\theta};L)]}, \quad (5.77)$$

where $\mathbb{E}_q[k(\underline{\theta};L)] = \int q(\underline{\theta})k(\underline{\theta};L)d\underline{\theta}$ is the acceptance rate. Now we can think of the procedure as the standard IS algorithm with $q(\underline{\theta}) \rightarrow h_L(\underline{\theta})$ and for an unbiased estimator for any observable O , we choose

$$w_D(\underline{\theta}) = p(\underline{\theta})/h_L(\underline{\theta}). \quad (5.78)$$

Because of the rejection sampling step, we need to sample $\underline{\theta}$ from $q(\underline{\theta})$ on average $1/\mathbb{E}_q[k(\underline{\theta};L)]$ times in order to get a single sample from $h_L(\underline{\theta})$. From (5.71), for a given $\underline{\theta}$ the expected number of circuit runs is $k(\underline{\theta};L)/f(\underline{\theta})$. Hence the expected cost per sample is

$$\frac{1}{\mathbb{E}_q[k(\underline{\theta};L)]} \mathbb{E}_q \left[c(\underline{\theta}) \frac{k(\underline{\theta};L)}{f(\underline{\theta})} \right] = \mathbb{E}_{h_L} \left[\frac{c(\underline{\theta})}{f(\underline{\theta})} \right]. \quad (5.79)$$

As for the Variance, we have

$$\text{Var}_{qr}^{\max} \leq \mathbb{E}_{h_L} [w_D^2(\underline{\theta})]. \quad (5.80)$$

This gives the bound on the net cost as

$$NC_q = \mathbb{E}_{h_L} \left[\frac{c(\underline{\theta})}{f(\underline{\theta})} \right] \mathbb{E}_{h_L} [w_D^2(\underline{\theta})] = \mathbb{E}_p \left[\frac{c(\underline{\theta})}{w_D(\underline{\theta})f(\underline{\theta})} \right] \mathbb{E}_p [w_D(\underline{\theta})] \geq \left(\mathbb{E}_p \left[\sqrt{\frac{c(\underline{\theta})}{f(\underline{\theta})}} \right] \right)^2, \quad (5.81)$$

where the last inequality follows from the Cauchy-Schwarz inequality, with equality if and only if $w(\underline{\theta}) \propto \sqrt{c(\underline{\theta})/f(\underline{\theta})}$ which implies:

$$q(\underline{\theta}) \propto \frac{p(\underline{\theta})}{k(\underline{\theta}; L)} \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}} = \frac{p(\underline{\theta})}{1 - (1 - f(\underline{\theta}))^L} \sqrt{\frac{f(\underline{\theta})}{c(\underline{\theta})}}. \quad (5.82)$$

This completes the proof.

APPENDIX C

PROOF OF THEOREM 5

In the limit of large s , the Central Limit Theorem applies, implying that the distribution of $c(\underline{\theta})$ becomes close to a normal distribution $\mathcal{N}(c; s\mu_1, s\sigma_1^2)$, where μ_1 and σ_1^2 are the first two cumulants of each variable $c_1(\theta_i)$. By the Berry–Esseen theorem, if the third absolute moment of $c_1(\theta_i)$ is finite, then the cumulative distribution function of the normalized sum differs from the Gaussian CDF Φ by at most a constant times $1/\sqrt{s}$. Concretely, there exists a finite constant C (depending only on the third absolute moment) such that for all real y

$$\left| \mathbb{P}\left(\frac{c - s\mu_1}{\sqrt{s}\sigma_1} \leq y\right) - \Phi(y) \right| \leq \frac{C}{\sqrt{s}}.$$

To propagate this CDF bound to expectations of the (unbounded) test function $g(c) = c^\alpha e^{\lambda ac}$ we fix a cutoff $M > 0$ and write

$$\mathbb{E}[g(c)] = \mathbb{E}[g(c)\mathbf{1}_{c \leq M}] + \mathbb{E}[g(c)\mathbf{1}_{c > M}].$$

We apply Berry–Esseen to the bounded function $g\mathbf{1}$ to obtain an additive $O(s^{-1/2})$ error for the truncated expectation; then choose M large enough (using the assumed moment bound) so that the tail $\mathbb{E}[g(c)\mathbf{1}_{c > M}]$ is uniformly small. Letting $M \rightarrow \infty$ after the Berry–Esseen step implies that replacing the exact distribution of c by the Gaussian $\mathcal{N}(c; s\mu_1, s\sigma_1^2)$ incurs an additive error of order $O(s^{-1/2})$ in $\mathbb{E}[g(c)]$. Thus the Gaussian approximation step contributes an $O(s^{-1/2})$ additive error in the final asymptotic evaluation, i.e.

$$\mathbb{E}_p \left[\left(\frac{c(\underline{\theta})}{f(\underline{\theta})} \right)^\alpha \right] = \frac{\int_0^\infty dc \, c^\alpha e^{\lambda ac} \exp\left(-\frac{(c-s\mu_1)^2}{2s\sigma_1^2}\right)}{\int_0^\infty dc \, \exp\left(-\frac{(c-s\mu_1)^2}{2s\sigma_1^2}\right)} + \mathcal{O}\left(\frac{1}{\sqrt{s}}\right), \quad (5.83)$$

where the normal density is truncated at $c = 0$ and the denominator enforces the normalization on $[0, \infty)$. Here we introduce the parameter α : setting $\alpha = 0.5$ evaluates the net cost under the sampling distribution $q(\underline{\theta})$ (see (3.52)), whereas $\alpha = 1$ evaluates the net cost under the target distribution $p(\underline{\theta})$ (see (3.44)). Since we are

studying the asymptotic behavior for large s , we first control the denominator with a Mills-type expansion. Writing $c = s\mu_1 + \sqrt{s}\sigma_1 x$ gives

$$\begin{aligned} \int_0^\infty dc \exp\left(-\frac{(c-s\mu_1)^2}{2s\sigma_1^2}\right) &= \sqrt{s}\sigma_1 \int_{-\sqrt{s}\mu_1/\sigma_1}^\infty dx e^{-x^2/2} \\ &= \sqrt{2\pi s\sigma_1^2} \Phi\left(\frac{\sqrt{s}\mu_1}{\sigma_1}\right) \\ &= \sqrt{2\pi s\sigma_1^2} \left[1 + o\left(e^{-s\mu_1^2/(2\sigma_1^2)}\right)\right], \end{aligned}$$

Hence the denominator tends to $\sqrt{2\pi s\sigma_1^2}$ up to an exponentially small correction. First, we focus on the numerator and complete the square in the exponent:

$$\begin{aligned} &\int_0^\infty dc c^\alpha e^{\lambda\alpha c} e^{-(c-s\mu_1)^2/(2s\sigma_1^2)} \\ &= \int_0^\infty dc c^\alpha \exp\left(-\frac{1}{2s\sigma_1^2}\left(c - [s\mu_1 + s\lambda\alpha\sigma_1^2]\right)^2\right) \exp\left(\frac{1}{2s\sigma_1^2}\left([s\mu_1 + s\lambda\alpha\sigma_1^2]^2 - s^2\mu_1^2\right)\right) \\ &= \exp\left(s\alpha\left(\lambda\mu_1 + \frac{\alpha}{2}\lambda^2\sigma_1^2\right)\right) \int_0^\infty dc c^\alpha \exp\left(-\frac{1}{2s\sigma_1^2}\left(c - [s\mu_1 + s\lambda\alpha\sigma_1^2]\right)^2\right). \end{aligned}$$

We now put numerator and denominator together, obtaining

$$\mathbb{E}_p\left[\left(\frac{c(\theta)}{f(\theta)}\right)^\alpha\right] = \exp\left(s\alpha\left(\lambda\mu_1 + \frac{\alpha}{2}\lambda^2\sigma_1^2\right)\right) \left[1 + o\left(e^{-s\mu_1^2/(2\sigma_1^2)}\right)\right] \int_0^\infty dc c^\alpha \mathcal{N}(c; s\mu_1 + s\lambda\alpha\sigma_1^2, s\sigma_1^2). \quad (5.84)$$

The remaining term to be evaluated is the integral

$$\int_0^\infty dc c^\alpha \mathcal{N}(c; s\mu_1 + s\lambda\alpha\sigma_1^2, s\sigma_1^2) = \sigma_1^\alpha s^{\alpha/2} \int_{-\sqrt{s}[\mu_1 + \lambda\alpha\sigma_1^2]/\sigma_1}^\infty dx \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} + x\right)^\alpha \frac{e^{-x^2/2}}{\sqrt{2\pi}}. \quad (5.85)$$

The lower limit $-\sqrt{s}[\mu_1 + \lambda\alpha\sigma_1^2]/\sigma_1$ tends to $-\infty$ as $s \rightarrow \infty$ when $\mu_1 > 0$. Extending the integral down to $-\infty$ introduces an error equal to the discarded left tail

$$T(s) := \left| \int_{-\infty}^{-\sqrt{s}[\mu_1 + \lambda\alpha\sigma_1^2]/\sigma_1} dx \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} + x\right)^\alpha \frac{e^{-x^2/2}}{\sqrt{2\pi}} \right|. \quad (5.86)$$

We bound $T(s)$ by Cauchy–Schwarz as follows. We define

$$x_0 := \frac{\sqrt{s}[\mu_1 + \lambda\alpha\sigma_1^2]}{\sigma_1} > 0, \quad \phi(x) := \frac{e^{-x^2/2}}{\sqrt{2\pi}}, \quad h(x) := \left[\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} + x\right]^\alpha. \quad (5.87)$$

Then by Cauchy–Schwarz,

$$T(s) = \left| \int_{-\infty}^{-x_0} h(x) \phi(x) dx \right| \leq \left(\int_{-\infty}^{-x_0} h(x)^2 \phi(x) dx \right)^{1/2} \left(\int_{-\infty}^{-x_0} \phi(x) dx \right)^{1/2}. \quad (5.88)$$

We henceforth estimate each factor: (i) Tail probability factor (second factor). Using the standard bound for the Gaussian tail,

$$\int_{-\infty}^{-x_0} \phi(x) dx = \int_{x_0}^{\infty} \phi(x) dx \leq \frac{\phi(x_0)}{x_0} = \frac{1}{x_0 \sqrt{2\pi}} e^{-x_0^2/2}. \quad (5.89)$$

Hence its square root satisfies

$$\left(\int_{-\infty}^{-x_0} \phi(x) dx \right)^{1/2} \leq \frac{1}{(2\pi)^{1/4}} x_0^{-1/2} e^{-x_0^2/4}.$$

(ii) First factor. For $x \leq -x_0$ we may bound $h(x)^2$ by a polynomial times s^α . Indeed there exist constants (depending only on $\alpha, \mu_1, \lambda, \sigma_1$) such that

$$h(x)^2 \leq C s^\alpha (1 + x^{2\alpha}).$$

Hence

$$\int_{-\infty}^{-x_0} h(x)^2 \phi(x) dx \leq C s^\alpha \int_{x_0}^{\infty} (1 + x^{2\alpha}) \phi(x) dx \leq C s^\alpha \int_{-\infty}^{\infty} (1 + x^{2\alpha}) \phi(x) dx. \quad (5.90)$$

For α integer or half integer, the r.h.s. integral is a sum of moments from a normal distributions, hence it's finite. Combining (i) and (ii) in the Cauchy–Schwarz inequality yields

$$T(s) = O\left(s^{(2\alpha-1)/2} \exp\left(-\frac{s[\mu_1 + \lambda\alpha\sigma_1^2]^2}{2\sigma_1^2}\right)\right).$$

Thus the error from extending the lower limit in Eq. (5.85) to $-\infty$ is exponentially small in s and negligible compared to any algebraic powers of s . Hence, the final step is to tackle the integral

$$\int_{-\infty}^{\infty} dx \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} + x \right)^\alpha \frac{e^{-x^2/2}}{\sqrt{2\pi}}.$$

At this scope we apply Laplace's method by Taylor-expanding the $h(x)$ function around $x = 0$, that is around the maximum point of the Gaussian. Its Taylor expansion up to second order about $x = 0$ reads

$$h(x) = \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} \right)^\alpha \times \left[1 + \alpha \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} \right)^{-1} x + \frac{\alpha(\alpha-1)}{2} \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} \right)^{-2} x^2 \right] + R(x),$$

where $R(x)$ is the remainder. If we neglect the contribution of $R(x)$, the consequent error to the final result will be

$$\int_{-\infty}^{\infty} dx R(x) \frac{e^{-x^2/2}}{\sqrt{2\pi}}.$$

We henceforth bound the integral of the remainder against the Gaussian density by splitting the integral into the two regions

$$(A) \quad |x| \leq s^{1/4}, \quad (B) \quad |x| > s^{1/4}.$$

Region A: small x . For $|x| \leq s^{1/4}$ we have

$$\left| \frac{x}{\sqrt{s}(\mu_1 + \lambda\alpha\sigma_1^2)/\sigma_1} \right| = O(s^{-1/4}),$$

so for sufficiently large s the expansion parameter is uniformly small. By the Lagrange form of Taylor's remainder, $\exists \xi(x) \in [0, x] \subseteq [0, s^{1/4}]$

$$|R(x)| = \left| \frac{h^{(3)}(\xi(x))}{3!} x^3 \right| \leq C_\alpha \left| \frac{x}{\sqrt{s}(\mu_1 + \lambda\alpha\sigma_1^2)/\sigma_1} \right|^3,$$

with a constant C_α depending only on α and the fixed combination $\mu_1 + \lambda\alpha\sigma_1^2$. Therefore, integrating $|x|^3 e^{-x^2/2}$ for $|x| \leq s^{1/4}$ yields

$$\left| \int_{|x| \leq s^{1/4}} \frac{R(x)}{\sqrt{2\pi}} e^{-x^2/2} dx \right| = O(s^{-3/2}).$$

Region B: large x . On $|x| > s^{1/4}$ we bound the contribution by Cauchy–Schwarz:

$$\left| \int_{|x| > s^{1/4}} \frac{R(x)}{\sqrt{2\pi}} e^{-x^2/2} dx \right| \leq \left(\int_{|x| > s^{1/4}} \frac{R(x)^2}{\sqrt{2\pi}} e^{-x^2/2} dx \right)^{1/2} \left(\int_{|x| > s^{1/4}} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \right)^{1/2}. \quad (5.91)$$

The tail probability factor is $O(e^{-s^{1/2}/2})$. On the other hand, since $R(x)$ is a difference of polynomials, it is also a polynomial, and hence the first integral can be bounded with a finite value. Hence the overall Region B contribution is $O(e^{-s^{1/2}/2})$.

Total remainder. Combining Region A and Region B bounds we obtain the explicit remainder estimate for the Laplace step:

$$\int_{-\infty}^{\infty} \frac{R(x)}{\sqrt{2\pi}} e^{-x^2/2} dx = O(s^{-3/2}),$$

since the algebraic term dominates the exponentially small term for large s . Thus the integral over the whole real line satisfies the Laplace expansion

$$\int_{-\infty}^{\infty} dx \frac{h(x)}{\sqrt{2\pi}} e^{-x^2/2} = \left(\frac{\sqrt{s}\mu_1 + \sqrt{s}\lambda\alpha\sigma_1^2}{\sigma_1} \right)^\alpha \left[1 + \frac{\alpha(\alpha-1)}{2} \left(\frac{\mu_1 + \lambda\alpha\sigma_1^2}{\sigma_1} \right)^{-2} \frac{1}{s} + O(s^{-3/2}) \right]. \quad (5.92)$$

The final asymptotic expansion is

$$\begin{aligned} \mathbb{E}_p \left[\left(\frac{c(\underline{\theta})}{f(\underline{\theta})} \right)^\alpha \right] &= \exp \left(s \alpha \left(\lambda \mu_1 + \frac{\alpha}{2} \lambda^2 \sigma_1^2 \right) \right) s^\alpha \left(\mu_1 + \lambda \alpha \sigma_1^2 \right)^\alpha \times \\ &\times \left[1 + \frac{\alpha(\alpha-1)}{2} \left(\frac{\mu_1 + \lambda \alpha \sigma_1^2}{\sigma_1} \right)^{-2} \frac{1}{s} + O(s^{-3/2}) \right]. \end{aligned} \quad (5.93)$$

Finally, taking the $1/\alpha$ power and simplifying yields the expression

$$\begin{aligned} \left(\mathbb{E}_p \left[\left(\frac{c(\underline{\theta})}{f(\underline{\theta})} \right)^\alpha \right] \right)^{1/\alpha} &= e^{s \left(\lambda \mu_1 + \alpha \frac{\lambda^2 \sigma_1^2}{2} \right)} s \left(\mu_1 + \lambda \alpha \sigma_1^2 \right) \times \\ &\times \left[1 + \frac{\alpha(\alpha-1)}{2} \left(\frac{\mu_1 + \lambda \alpha \sigma_1^2}{\sigma_1} \right)^{-2} \frac{1}{s} + O(s^{-3/2}) \right]^{1/\alpha}. \end{aligned} \quad (5.94)$$

First, we notice that, as long as $\lambda \neq 0$ and either $\mu_1 \neq 0$ or $\sigma_1 \neq 0$, the expression showcases an exponential behavior as a function of s . Moreover, in the $L = \infty$ case, we can readily use it to compute the ratio

$$\frac{NC_{q^*}}{NC_p}(s) = e^{-s \left(\frac{\lambda \sigma_1}{2} \right)^2} \frac{\mu_1 + \frac{\lambda \sigma_1^2}{2}}{\mu_1 + \lambda \sigma_1^2} \left[1 - \frac{1}{4s} \left(\frac{\sigma_1}{\mu_1 + \frac{\lambda \sigma_1^2}{2}} \right)^2 + O(s^{-3/2}) \right]. \quad (5.96)$$

In particular, if $\lambda = 0$ then $NC_{q^*}/NC_p \rightarrow 1$, while for all $\lambda > 0$ one obtains $NC_{q^*}/NC_p \rightarrow 0$. Observe that NC_{q^*}/NC_p always has a single minimum, precisely at

$$s^*(L = \infty) = \frac{1}{8} \left(\frac{\sigma_1^2}{\mu_1 + \frac{\lambda \sigma_1^2}{2}} \right)^2 \left[1 + \sqrt{1 + \frac{64}{\lambda^2 \sigma_1^4} \left(\mu_1 + \frac{\lambda \sigma_1^2}{2} \right)^2} \right]. \quad (5.97)$$

APPENDIX D

SUPERPOSITION STATES FOR ISING HAMILTONIANS

We provide details on preparing the superposition state for the Ising model with periodic boundary conditions (PBC). Our goal is to identify an initial state for the adiabatic preparation (AP) that evolves into the desired superposition state $|\Psi(0, 1)\rangle$ defined in Equation (4.15) as the final outcome. For our choice of H_0 in Eq. (4.18) the ground state is degenerate, and therefore the adiabatic theorem does not hold. To circumvent this issue we employ the Branched-Subspace Adiabatic Preparation (B-SAP) algorithm for state preparation [2], which extends the applicability of the adiabatic theorem to scenarios in which the initial Hamiltonian H_0 possesses a degenerate spectrum. A direct implementation of B-SAP would require the application of a parametrized circuit prior to the evolution under U_τ . However, for the specific case of the Ising model, the training of these parameters can be avoided by exploiting symmetry considerations. To achieve this, we introduce the operator:

$$\mathcal{R} = \bigotimes_{j=0}^{L-1} R_z^j(\pi), \quad (5.98)$$

where $R_z^j(\theta)$ represents a rotation around the z -axis by angle θ on the j -th qubit. This operator commutes with the Hamiltonian for any value of J_1 , ensuring that its expectation value remains conserved throughout the AP. The ground states of the auxiliary Hamiltonian H_0 , $|+\rangle^{\otimes L}$ and $|-\rangle^{\otimes L}$, do not diagonalize $\mathcal{R}$, as:

$$\mathcal{R}|\pm\rangle^{\otimes L} = (-i)^L |\mp\rangle^{\otimes L}. \quad (5.99)$$

However, by defining the new basis:

$$|\Phi^\pm\rangle = \frac{1}{\sqrt{2}} (|+\rangle^{\otimes L} \pm |-\rangle^{\otimes L}), \quad (5.100)$$

we obtain:

$$\mathcal{R}|\Phi^\pm\rangle = \pm(-i)^L |\Phi^\pm\rangle. \quad (5.101)$$

Thus, the initial ground eigenspace is spanned by $|\Phi^+\rangle$ and $|\Phi^-\rangle$, which are eigenstates of $\mathcal{R}$ with eigenvalues $(-i)^L$ and $-(-i)^L$, respectively.

We now consider the state $|\Psi(0, 1)\rangle$ defined in Equation (4.22), which should be the equal superposition of the ground and first excited states of H , originating from the splitting of the initially degenerate ground eigenspace of H_0 . In the limit $h_3/J_1 \rightarrow \infty$, the final ground state is:

$$|\Omega_0\rangle \equiv |0\rangle^{\otimes L}, \quad (5.102)$$

with eigenvalue $(-i)^L$ with respect to $\mathcal{R}$. Since $\mathcal{R}$ is a conserved quantity, it follows that $|\Omega_0\rangle$ must evolve from $|\Phi^+\rangle$:

$$\lim_{\tau \rightarrow \infty} U_\tau |\Phi^+\rangle = |\Omega_0\rangle. \quad (5.103)$$

This also holds for any finite h_3/J_1 , where a shorter adiabatic process is sufficient. By complementarity, we also obtain:

$$\lim_{\tau \rightarrow \infty} U_\tau |\Phi^-\rangle = |\Omega_1\rangle. \quad (5.104)$$

Finally, by linearity, combining Eqs. (5.103) and (5.104), we arrive at:

$$\lim_{\tau \rightarrow \infty} U_\tau |+\rangle^{\otimes L} = \frac{1}{\sqrt{2}} \lim_{\tau \rightarrow \infty} U_\tau (|\Phi^+\rangle + |\Phi^-\rangle) = |\Psi(0, 1)\rangle. \quad (5.105)$$

APPENDIX E

OBSERVABLE CHOICE

For our procedure to work, the objective is to have an observable O that maximizes $\mathcal{A}$, as described in Eq. (4.16). This facilitates a more straightforward estimation of the period of the sinusoidal function, which in turn simplifies the fitting process for the energy gap.

Ising Model

For the Ising model our goal is to find an operator O for which $\mathcal{A}$ is substantially different from zero, independently from the chain/lattice size. At this scope we study the two regimes in which the eigenstates are analytically known. In particular, for $h_3/J_1 \rightarrow 0$, one has

$$\mathcal{A} = |\langle \Phi^+ | O | \Phi^- \rangle|. \quad (5.106)$$

Notably, any operator obtained with a tensor product of an odd number of σ^1 Pauli operations gives $\mathcal{A} = 1$. In the opposite regime, the one with $h_3/J_1 \rightarrow \infty$, one has

$$\mathcal{A} = |\langle \Omega_0 | O | \Omega_1 \rangle| = \frac{1}{\sqrt{L}} \left| \sum_j \langle 0 |^{\otimes L} O | 0 \rangle^j \otimes | 1 \rangle \otimes | 0 \rangle^{\otimes L-1-j} \right|. \quad (5.107)$$

We therefore choose $O = \sigma_j^1 \otimes \mathbb{1}^{\otimes L-1}$, where j can be any site because of translational invariance, so that $\mathcal{A} = 1/\sqrt{L} > 0$. These limiting behaviors are illustrated in Figure E.1, where we plot $\mathcal{A}$ as a function of h_3/J_1 for different chain lengths.

Moreover, we also test the values for $\mathcal{A}$ for all the possible choices of the observable O with a Pauli-string shape on Ising chains with PBC and a length up to 10 sites. Based on these considerations, we use the σ_j^1 operator to estimate the energy gap of all our experiments involving the Ising model.

Molecules

For small molecules, the choice of the observable O is based on the following considerations. Since both ground and excited state of the auxiliary Hamiltonian are in the




Figure E.1: Numerical estimations of the expectation value $\mathcal{A} = |\langle \Phi^+ | \sigma^1 \otimes (\mathbb{1}^{\otimes L-1}) | \Phi^- \rangle|$ as a function of the Ising's coupling constants ratio J_1/h_3 . For all chain lengths considered, $L = \{4, 6, 8, 10, 12\}$, $\mathcal{A}$ remains larger than $L^{-1/2}$ for any value of h_3/J_1 , ensuring that the amplitude is strictly positive across the entire range of coupling ratios.

computational basis, it is always possible to find a tensor product P of identities and σ^1 Pauli matrices such that

$$\langle \omega_1 | P | \omega_0 \rangle = \langle \omega_1 | \bigotimes_{j=1}^L \sigma_j^{i_j} | \omega_0 \rangle = 1, \quad i_j \in \{0, 1\}. \quad (5.108)$$

Assuming that the off-diagonal terms of H are small with respect to the diagonal terms, then choosing P as the operator for our procedure will satisfy $\mathcal{A} \neq 0$. In particular, for both H_2 and He_2 molecules, we choose the operator P such that:

$$\mathcal{A} = |\langle \Omega_1 | P | \Omega_0 \rangle| = |\langle \omega_1 | U_\tau^\dagger P U_\tau | \omega_0 \rangle|. \quad (5.109)$$

APPENDIX F

QUANTUM MEASUREMENT ERROR

Let

$$\langle O \rangle = \langle \psi | O | \psi \rangle \quad (5.110)$$

be the expectation value of an operator O on a quantum state $|\psi\rangle$, and let

$$\langle (O - \langle O \rangle)^2 \rangle \quad (5.111)$$

be its variance. Recalling that O should be unitary and Hermitian in order to be measured on a quantum computer

$$O^2 = O^\dagger O = \mathbb{1}, \quad (5.112)$$

so that its eigenvalues can only be $\{+1, -1\}$. In the non trivial case in which the eigenvalues are different, one can then decompose the state onto two eigenstates of O as

$$|\psi\rangle = \alpha|\psi_{+1}\rangle + \beta|\psi_{-1}\rangle, \quad (5.113)$$

with

$$\begin{cases} O|\psi_{\pm 1}\rangle = \pm|\psi_{\pm 1}\rangle \\ |\alpha|^2 + |\beta|^2 = 1. \end{cases} \quad (5.114)$$

Then

$$\begin{cases} \langle O \rangle = |\alpha|^2 - |\beta|^2 \\ \langle (O - \langle O \rangle)^2 \rangle = 4|\alpha|^2|\beta|^2. \end{cases} \quad (5.115)$$

APPENDIX G

IMPERFECT SUPERPOSITION STATE PREPARATION

In this Appendix, we provide insights into the robustness of the proposed method with respect to imperfect superposition state preparation. We henceforth only consider the effects of the finite time for both Adiabatic Preparation (AP) and of the Trotterization process, without taking into account those coming from quantum noise. In particular, we will now show analytically that if the error in the preparation is of order ϵ then the estimation of the expectation value will experience an error of the same order.

Let

$$|\nu_{ij}(\epsilon)\rangle = |\Psi_{ij}\rangle + \mathcal{O}(\epsilon) = \frac{1}{\sqrt{2}}(|\Omega_i\rangle + |\Omega_j\rangle) + \mathcal{O}(\epsilon) \quad (5.116)$$

be an imperfectly prepared superposition state, where ϵ denotes the error magnitude. This state can be expressed on the basis of the Hamiltonian eigenstates as

$$|\nu_{ij}(\epsilon)\rangle = \sum_n c_n(\epsilon) |\Omega_n\rangle. \quad (5.117)$$

We expand the coefficients in powers of ϵ :

$$c_n = \sum_{m=0}^{\infty} c_n^{(m)} \epsilon^m. \quad (5.118)$$

The coefficients c_n can be uniquely determined by imposing the same conditions typically used in perturbation theory (see, e.g., [43]). First, we require that for $\epsilon = 0$, the ideal state is recovered, i.e. $|\nu_{ij}(0)\rangle = |\Psi_{ij}\rangle$. This condition holds if and only if

$$c_i^{(0)} = \frac{1}{\sqrt{2}}, \quad c_j^{(0)} = \frac{1}{\sqrt{2}}, \quad c_{n \neq i,j}^{(0)} = 0, \quad (5.119)$$

or, more compactly,

$$c_n^{(0)} = \frac{1}{\sqrt{2}}(\delta_{ni} + \delta_{nj}). \quad (5.120)$$




Figure G.1: Fidelity of the approximate superposition state $|\Psi(0, 1)\rangle$, in the noiseless and noisy case ($|\nu\rangle$, ρ_ν), with the ground and first excited state ($|\psi_0\rangle$, $|\psi_1\rangle$) for an Ising chain of $L = 8$ qubits. The fidelity is found to be sub-optimal, decaying and oscillating for increasing h_3 . The discontinuities are given by the change in $\delta\tau$ (at $h_3 = \{3.6, 5\}$) imposed in the adiabatic process.

As a second condition, we require that the first-order truncated state be normalized up to second-order corrections, namely

$$1 - \langle \nu_{ij}(\epsilon) | \nu_{ij}(\epsilon) \rangle = \mathcal{O}(\epsilon^2). \quad (5.121)$$

Substituting Eq. (5.117) and Eq. (5.118) one obtains

$$1 - \left[\frac{1}{\sqrt{2}} (\langle \Omega_i | + \langle \Omega_j |) + \epsilon \sum_n c_n^{(1)*} \langle \Omega_n | + \mathcal{O}(\epsilon^2) \right] \times \quad (5.122)$$

$$\times \left[\frac{1}{\sqrt{2}} (|\Omega_i\rangle + |\Omega_j\rangle) + \epsilon \sum_n c_n^{(1)} |\Omega_n\rangle + \mathcal{O}(\epsilon^2) \right] = \mathcal{O}(\epsilon^2), \quad (5.123)$$

which yields $c_i^{(1)} = c_j^{(1)} = 0$. Therefore,

$$|\nu_{ij}(\epsilon)\rangle = \frac{1}{\sqrt{2}} (|\Omega_i\rangle + |\Omega_j\rangle) + \epsilon \sum_{n \neq i,j} c_n^{(1)} |\Omega_n\rangle + \mathcal{O}(\epsilon^2). \quad (5.124)$$

The expectation value of a time-dependent operator $O(t)$ on $|\nu_{ij}(\epsilon)\rangle$ is then

$$\begin{aligned} \langle \nu_{ij}(\epsilon) | O(t) | \nu_{ij}(\epsilon) \rangle &= \langle \Psi_{ij} | O(t) | \Psi_{ij} \rangle + \epsilon \sqrt{2} \sum_{n \neq i,j} \left[|c_n^{(1)} \langle \Omega_i | O | \Omega_n \rangle| \cos(\Delta_{i,n} t + \theta_{i,n}) \right. \\ &\quad \left. + |c_n^{(1)} \langle \Omega_j | O | \Omega_n \rangle| \cos(\Delta_{j,n} t + \theta_{j,n}) \right] + \mathcal{O}(\epsilon^2). \end{aligned} \quad (5.125)$$

Finally, we relate the error magnitude ϵ to the infidelity, defined as

$$\mathcal{I} = 1 - |\langle \nu_{ij}(\epsilon) | \Psi_{ij} \rangle|^2 = \mathcal{O}(\epsilon^2). \quad (5.126)$$

From this relation, we conclude that if the target state $|\Psi_{ij}\rangle$ is prepared with a given infidelity $\mathcal{I}$, the estimated expectation value of a generic time-dependent operator $O(t)$ is affected by an error proportional to $\sqrt{\mathcal{I}}$.

In particular, we numerically verify the fidelity of the prepared superposition state with both the ground and first excited states of the one-dimensional Ising chain as a function of the transverse field, as shown in Figure G.1. The theoretically expected optimal fidelities of the superposition state with the ground and first excited states are

$$F(|\Psi\rangle, |\psi_0\rangle) = F(|\Psi\rangle, |\psi_1\rangle) = 0.5. \quad (5.127)$$

As illustrated in Figure G.1, the fidelity does not reach the optimal theoretical value and decreases with increasing h_3 , both in the noiseless and in the noisy simulations. Consistently with the theoretical analysis presented in this Appendix, even though the prepared state $|\Psi\rangle$ is imperfect, the proposed method remains robust: it yields accurate estimates of expectation values with errors on the order of 10^{-2} and provides a reliable approximation of the spectral gap of the system. Finally, we note the presence of oscillations in the fidelity, which may originate from intrinsic features of the adiabatic preparation process, even though we were not able to identify them theoretically.

APPENDIX H

DETAILS ON IONQ ARIA

H.1 IonQ Native Gates

Since single qubit-gates are typically implemented on ion-trapped hardware with high average gate fidelity, potential issues in result's fidelity (when running a circuit on NISQ hardware) primarily stem from 2-qubit gates. For this reason, we simulate the Ising model with interactions along the x -axis and transverse field along the z -axis. This allows us to use only the native entangling gates available for IonQ Aria device, i.e. Mølmer-Sørensen gate (MS) [103], and to reduce the number of entangling gates only to the essential ones. More precisely, in a single Trotter step, one only needs two layers of entangling gates (see Figure H.1), implying 80 2-qubit gates depth when fixing the number of steps to 40.

The native gates are

$$\text{GPI2}(\phi) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -ie^{-i\phi} \\ -ie^{i\phi} & 1 \end{pmatrix}, \quad (5.128)$$

$$\text{VirtualZ}(\theta) = \begin{pmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{pmatrix}, \quad (5.129)$$




Figure H.1: Quantum circuit for a single Trotter step showing the evolution of the superposition state $|\Psi\rangle$. The circuit involves the application of $\text{GPI2}(\phi)$, $\text{RZ}(\theta)$, and $\text{IsingXX}(\theta)$ gates to prepare $|\Psi\rangle$ for spectral gap estimation.

and

$$MS(\phi_0, \phi_1, \theta) = \begin{pmatrix} \cos \frac{\theta}{2} & 0 & 0 & -ie^{-i(\phi_0+\phi_1)} \sin \frac{\theta}{2} \\ 0 & \cos \frac{\theta}{2} & -ie^{-i(\phi_0-\phi_1)} \sin \frac{\theta}{2} & 0 \\ 0 & -ie^{i(\phi_0-\phi_1)} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} & 0 \\ -ie^{i(\phi_0+\phi_1)} \sin \frac{\theta}{2} & 0 & 0 & \cos \frac{\theta}{2} \end{pmatrix}.$$

The Mølmer-Sørensen gate can be further simplified to obtain the Ising interaction $\sigma_i^1 \sigma_j^1$ by setting ϕ_0 and ϕ_1 equal to 0.

H.2 IonQ Aria noise model

To assess the resilience of our approach to noise, we incorporated a realistic noise model taking into account the technical features of IonQ's Aria device available online [104]. The noise sources considered include:

- *Thermal Relaxation*: This source of noise accounts for the effects of the thermal relaxation times T_1 and T_2 . Notably, IonQ Aria has $T_1 = 100s$ and $T_2 = 1s$, and while we took them into account in our simulations, their effects were found to be negligible and did not significantly impact the results.
- *Gate Fidelity*: We simulated errors arising from finite gate fidelity, modeled as depolarizing noise. The relationship between gate fidelity and the depolarizing parameter is derived from what reported in the Appendix of [105]. The two-qubit gate time is $600\mu s$, while single qubit gates are executed in $135\mu s$.
- *Readout Noise*: A readout noise of 0.39% was introduced to emulate imperfections in the measurement process.




Figure H.2: Observable oscillations in time for Ising 1D chain with $L = 4$ at $h_3/J_1 = 7.257$, measured (points) and fitted (dashed lines) for noiseless, noisy and real hardware (with error mitigation). For reference, it is also reported the expected oscillatory from diagonalization (dotted line). Error bars represent a single standard deviation computed from the fit.




Figure H.3: Robustness analysis as a function of the number of shots for the 1D Ising chain with $L = 4$. Scatter points show the measured spectral gap in units of J_1 , and error bars represent one standard deviation obtained from the fit.

Upon detailed examination, we have determined that the gate fidelity plays a significant role in the estimation of the spectral gap. This noise source has been implemented by considering a single qubit depolarizing channel, which acts on a density matrix ϱ

$$\mathcal{D}_p(\varrho) = (1-p)\varrho + p\frac{\mathbb{1}}{2}\varrho. \quad (5.130)$$

To take into account gate imperfections and finite qubit relaxation times we define the probability p as

$$p = 1 + 3\frac{2\epsilon - 1}{d}. \quad (5.131)$$

Here, $d = \exp\{-T_g/T_1\} + 2\exp\{-T_g/T_2\}$ and $\epsilon = 1 - F$, where F is the average gate fidelity. The term T_g represents the gate time, which is the duration required to perform a single-qubit or two-qubit gate operation. We define two depolarizing channels one per each type of gate (single/two-qubit gate). Our approach demonstrates robustness against gate fidelity errors and readout noise, suggesting its viability for simulation on real quantum hardware.

The results of the noise resilience analysis are presented in Figure 9.1. The robustness of our approach can be appreciated also in the time oscillations of the observable reported in Figure H.2. There, the damping of oscillations due to noise (manifested as a reduction in $\mathcal{A}$) is highlighted, compared to the expected behaviour $\mathcal{A}\cos(\Delta E + \theta)$. The amplitude $\mathcal{A}$ is taken from our preliminary study of the observable, and the phase θ is set as in the noiseless simulations. Interestingly the noise mitigation allows for having wider oscillations on real hardware compared to noisy simulations.

In particular, the IonQ Aria device reduces the impact of systematic errors by employing a specific error mitigation method called *debiasing* [96]. This technique maps a circuit into multiple variants, employing different qubit permutations or gate decompositions. The effectiveness of debiasing error mitigation is likely improved by our approach's robustness to a reduced number of shots. In fact, as reported in

Figure H.3, we assess the robustness of our approach over a reduced number of shots by reducing this value from 10^4 up to 100 shots for chains of 4 qubits. As witnessed by Figure H.3, the spectral gap estimate stays approximately constant over different number of shots.

APPENDIX I

ADDITIONAL ISING EXPERIMENTS

Here, we extend our study to two additional Ising configurations with periodic boundary conditions—a 10-spin chain and a 4-site lattice. Their spectral-gap dependence on h_3/J_1 is presented in Figure I.1. All previous conditions for the Ising chain experiment remain in this study. In particular, the 40 Trotter steps are divided in 15 adiabatic steps and 25 time evolution steps.

The numerical results for the 10-site Ising chain show excellent agreement between the benchmark values and the noiseless estimate of the spectral gap. We do not report experiments on real hardware because the limited availability of the IonQ Aria device prevented additional tests. Simulations with a realistic noise model were also not performed in this case, as we consider the proof-of-principle demonstrations with 4 and 20 qubits presented in the main text to be sufficient. In the 2D simulations, the number of entangling gates required per Trotter step is doubled with respect to the 1D model. Due to the prohibitive depth implied by the extra interactions, simulations with real hardware noise model lead to inconclusive results. It would be interesting to see if IonQ Aria noise mitigation is able to lead to measurable results.




Figure I.1: Spectral gap as a function of h_3/J_1 for a $L = 10$ Ising 1D (left) and a $L = 4$ Ising 2D lattice (right) with periodic boundary conditions (PBC). Our approach demonstrates optimal agreement with the benchmark values, in the noiseless case. Error bars represent a single standard deviation computed from the fit.

BIBLIOGRAPHY

- [1] D. Cugini, D. Nigro, M. Bruno, and D. Gerace, Exponential optimization of adiabatic quantum-state preparation, *Physical Review Research* **7**, L012074 (2025).
- [2] D. Cugini, G. Guarnieri, M. Motta, and D. Gerace, Symmetry-guided quantum state preparation: Branched-subspaces adiabatic preparation (b-sap), *Quantum Science and Technology* (2026).
- [3] D. Cugini, T. A. Atif, and Y. Subasi, Resource-optimal importance sampling for randomized quantum algorithms, *arXiv preprint arXiv:2603.13495* (2026).
- [4] D. Cugini, F. Ghisoni, A. R. Morgillo, and F. Scala, Spectral gap estimation via adiabatic preparation, *arXiv preprint arXiv:2512.19288* (2025).
- [5] D. Cugini, D. Gerace, P. Govoni, A. Perego, and D. Valsecchi, Comparing quantum and classical machine learning for vector boson scattering background reduction at the large hadron collider, *Quantum Machine Intelligence* **5**, 35 (2023).
- [6] D. Cugini, A. M. Timpanaro, G. Livan, and G. Guarnieri, Universal emergence of local zipf–mandelbrot law, *Journal of Statistical Mechanics: Theory and Experiment* **2025**, 103403 (2025).
- [7] D. Cugini and G. Guarnieri, State convertibility and fluctuation theorems from a dynamical reference: majorization meets martingales, *arXiv preprint arXiv:2608.19391* (2026).
- [8] D. Cugini, Pathwise random hamiltonian simulation, *arXiv preprint arXiv:2608.29756* (2026).
- [9] J. Preskill, Quantum computing in the nisq era and beyond, *Quantum* **2**, 79 (2018).
- [10] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Physical Review Letters* **119**, 180509 (2017).
- [11] D. Aharonov and M. Ben-Or, in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing* (1997), pp. 176–188.

- [12] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, *Physical Review A—Atomic, Molecular, and Optical Physics* **71**, 022316 (2005).
- [13] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Physical Review A* **86**, 032324 (2012).
- [14] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
- [15] D. Litinski, Magic state distillation: Not as costly as you think, *Quantum* **3**, 205 (2019).
- [16] C. Gidney and M. Ekerå, Factoring with 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [17] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, vol. 2 (Cambridge university press Cambridge, 2001).
- [18] G. H. Low and I. L. Chuang, Optimal hamiltonian simulation by quantum signal processing, *Physical review letters* **118**, 010501 (2017).
- [19] J.-L. Brylinski and R. Brylinski, in *Mathematics of quantum computation* (Chapman and Hall/CRC, 2002), pp. 117–134.
- [20] M. J. Bremner, C. M. Dawson, J. L. Dodd, A. Gilchrist, A. W. Harrow, D. Mortimer, M. A. Nielsen, and T. J. Osborne, Practical scheme for quantum computation with any two-qubit entangling gate, *Physical review letters* **89**, 247902 (2002).
- [21] V. Kliuchnikov, D. Maslov, and M. Mosca, Fast and efficient exact synthesis of single qubit unitaries generated by clifford and t gates, arXiv preprint arXiv:1206.5236 (2012).
- [22] N. J. Ross and P. Selinger, Optimal ancilla-free clifford+ t approximation of z-rotations., *Quantum Inf. Comput.* **16**, 901 (2016).
- [23] D. Coppersmith, Tech. Rep., IBM Research Report RC19642 (1994), <https://arxiv.org/abs/quant-ph/0201067>.
- [24] A. Barenco, A. Ekert, K.-A. Suominen, and P. Tóth, Approximate quantum fourier transforms and decoherence, *Physical Review A* **54**, 139 (1996).
- [25] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM review* **41**, 303 (1999).
- [26] P. Schleich, L. M. Calderón, C. Sun, M. Bagherimehrab, A. Aldossary, J. S. Kottmann, and A. Aspuru-Guzik, *Quantum Computing for Quantum Chemistry* (American Chemical Society, 2025).

- [27] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*, vol. 47 of *Graduate Studies in Mathematics* (American Mathematical Society, 2002).
- [28] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local hamiltonian problem, *SIAM Journal on Computing* **35**, 1070 (2006).
- [29] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, et al., Variational quantum algorithms, *Nature Reviews Physics* **3**, 625 (2021).
- [30] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [31] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, *Nature Reviews Physics* pp. 1–16 (2025).
- [32] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem, *Science* **292**, 472 (2001).
- [33] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Reviews of Modern Physics* **90**, 015002 (2018).
- [34] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028* (2014).
- [35] A. Peruzzo, J. McClean, P. Shadbolt, et al., A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
- [36] N. Moll, P. Barkoutsos, L. S. Bishop, et al., Quantum optimization using variational algorithms on near-term quantum devices, *Quantum Science and Technology* **3**, 030503 (2018).
- [37] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
- [38] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, Cambridge, 2011), ISBN 9780511973765, <https://doi.org/10.1017/CBO9780511973765>.
- [39] J. Dziarmaga, Dynamics of a quantum phase transition and relaxation to a steady state, *Advances in Physics* **59**, 1063 (2010).
- [40] K. Binder and A. P. Young, Spin glasses: Experimental facts, theoretical concepts, and open questions, *Reviews of Modern Physics* **58**, 801 (1986).
- [41] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, First-order quantum phase transitions in adiabatic quantum computation, *EPL (Europhysics Letters)* **89**, 20004 (2010).

- [42] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and L. Zdeborová, The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective, *Physics Reports* **523**, 127 (2013).
- [43] L. D. Landau and E. M. Lifshitz, *Quantum mechanics: non-relativistic theory*, vol. 3 (Elsevier, 2013).
- [44] B. Chakraborty, M. Honda, T. Izubuchi, Y. Kikuchi, and A. Tomiya, Classically emulated digital quantum simulation of the schwinger model with a topological term via adiabatic state preparation, *Physical Review D* **105**, 094503 (2022).
- [45] A. Hamma and D. A. Lidar, Adiabatic preparation of topological order, *Physical review letters* **100**, 030502 (2008).
- [46] A. N. Ciavarella, S. Caspar, M. Illa, and M. J. Savage, State preparation in the heisenberg model through adiabatic spiraling, *Quantum* **7**, 970 (2023).
- [47] H. Goldstein, C. P. Poole, J. Safko, et al., *Classical mechanics*, vol. 2 (Addison-wesley Reading, MA, 1950).
- [48] J. J. Sakurai and J. Napolitano, *Modern quantum mechanics* (Cambridge university press, 2020).
- [49] G. Nenciu, Linear adiabatic theory. exponential estimates, *Comm. Math. Phys.* **152**, 479 (1993).
- [50] A. Schmidhuber and S. Lloyd, Adiabatic quantum phase estimation, arXiv preprint arXiv:2605.22770 (2026).
- [51] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators (1950).
- [52] S. Jansen, M.-B. Ruskai, and R. Seiler, Bounds for the adiabatic approximation with applications to quantum computation, *Journal of Mathematical Physics* **48**, 102111 (2007).
- [53] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Physical Review A* **65**, 042308 (2002).
- [54] C. Mc Keever and M. Lubasch, Towards adiabatic quantum computing using compressed quantum circuits, *PRX Quantum* **5**, 020362 (2024).
- [55] E. Granet and H. Dreyer, Continuous hamiltonian dynamics on noisy digital quantum computers without trotter error, arXiv preprint arXiv:2308.03694 (2023).
- [56] Qiskit contributors, Qiskit: An open-source framework for quantum computing (2023).

- [57] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martínez, Quantum computation of electronic transitions using a variational quantum eigensolver, *Physical review letters* **122**, 230401 (2019).
- [58] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Information Computation* **13**, 901 (2013).
- [59] F. D. Murnaghan, *The unitary and rotation groups* (1962).
- [60] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing* (2019), pp. 193–204.
- [61] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *arXiv preprint arXiv:1202.5822* (2012).
- [62] M. L. Rhodes, M. Kreshchuk, and S. Pathak, Exponential improvements in the simulation of lattice gauge theories using near-optimal techniques, *PRX Quantum* **5**, 040347 (2024).
- [63] Y. Dong, L. Lin, H. Ni, and J. Wang, Infinite quantum signal processing, *Quantum* **8**, 1558 (2024).
- [64] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [65] E. Campbell, Random compiler for fast hamiltonian simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [66] S. Boixo, E. Knill, and R. D. Somma, Eigenpath traversal by phase randomization, *arXiv:0903.1652* (2009), <https://arxiv.org/abs/0903.1652>.
- [67] J. Hilder, D. Pijn, O. Onishchenko, A. Stahl, M. Orth, B. Lekitsch, A. Rodriguez-Blanco, M. Müller, F. Schmidt-Kaler, and U. Poschinger, Fault-tolerant parity readout on a shuttling-based trapped-ion quantum computer, *Physical Review X* **12**, 011032 (2022).
- [68] C. N. Self, M. Benedetti, and D. Amaro, Protecting expressive circuits with a quantum error detection code, *Nature Physics* **20**, 219 (2024).
- [69] A. Ransford, M. Allman, J. Arkinstall, J. Campora III, S. F. Cooper, R. D. Delaney, J. M. Dreiling, B. Estey, C. Figgatt, A. Hall, et al., Helios: A 98-qubit trapped-ion quantum computer, *arXiv preprint arXiv:2511.05465* (2025).
- [70] N. M. Linke, S. Johri, C. Figgatt, K. A. Landsman, A. Y. Matsuura, and C. Monroe, Measuring the rényi entropy of a two-site fermi-hubbard model on a trapped ion quantum computer, *Physical Review A* **98**, 052334 (2018).
- [71] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. O’Brien, Low-cost error mitigation by symmetry verification, *Physical Review A* **98**, 062339 (2018).
- [72] S. McArdle, X. Yuan, and S. Benjamin, Error-mitigated digital quantum simulation, *Physical review letters* **122**, 180501 (2019).

- [73] Z. Cai, Quantum error mitigation using symmetry expansion, *Quantum* **5**, 548 (2021).
- [74] W. Hoeffding, Probability inequalities for sums of bounded random variables, *Journal of the American statistical association* **58**, 13 (1963).
- [75] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
- [76] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, 339 (1998).
- [77] F. Tacchino, A. Chiesa, S. Carretta, and D. Gerace, Quantum computers as universal quantum simulators: State-of-the-art and perspectives, *Adv. Quantum Technol.* **3**, 1900052 (2020).
- [78] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *Contemporary Mathematics* **305**, 53 (2002).
- [79] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem, arXiv preprint quant-ph/9511026 (1995).
- [80] A. Arrasmith, L. Cincio, R. D. Somma, and P. J. Coles, Operator sampling for shot-frugal optimization in variational algorithms, arXiv preprint arXiv:2004.06252 (2020).
- [81] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nature Physics* **16**, 1050 (2020).
- [82] M. B. Hastings and T. Koma, Spectral gap and exponential decay of correlations, *Commun. Math. Phys.* **265**, 781 (2006).
- [83] W. Kohn, Theory of the insulating state, *Phys. Rev.* **133**, A171 (1964).
- [84] K. Szymański and K. Życzkowski, Universal witnesses of vanishing energy gap, *EPL* **136**, 30003 (2022).
- [85] J.-H. Chung and J. Shapiro, Topological classification of insulators: I. non-interacting spectrally-gapped one-dimensional systems, *Adv. Math.* **480**, 110486 (2025), ISSN 0001-8708.
- [86] J. Griffith and L. Orgel, Ligand-field theory, *Q. Rev., Chem. Soc.* **11**, 381 (1957).
- [87] M. Lupo Pasini, K. Mehta, P. Yoo, and S. Irle, Two excited-state datasets for quantum chemical uv-vis spectra of organic molecules, *Sci. Data* **10**, 546 (2023).
- [88] J.-L. Bredas, Mind the gap!, *Mater. Horiz.* **1**, 17 (2014).

- [89] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations, *Phys. Rev. Lett.* **94**, 170201 (2005).
- [90] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018), ISSN 2521-327X, <https://doi.org/10.22331/q-2018-08-06-79>.
- [91] J. M. Leamer, A. B. Magann, A. D. Baczewski, G. McCaul, and D. I. Bondar, Spectral gaps via imaginary time (2023), [2303.02124](https://arxiv.org/abs/2303.02124).
- [92] A. E. Russo, K. M. Rudinger, B. C. A. Morrison, and A. D. Baczewski, Evaluating energy differences on a quantum computer with robust phase estimation, *Phys. Rev. Lett.* **126**, 210501 (2021).
- [93] K. Gnatenko, H. Laba, and V. Tkachuk, Energy levels estimation on a quantum computer by evolution of a physical quantity, *Phys. Lett. A* **424**, 127843 (2022), ISSN 0375-9601.
- [94] P.L. Chebyshev, *Théorie des mécanismes connus sous le nom de parallélogrammes* (Imprimerie de l'Académie impériale des sciences, 1853).
- [95] G. Rendon, J. Watkins, and N. Wiebe, Improved Accuracy for Trotter Simulations Using Chebyshev Interpolation, *Quantum* **8**, 1266 (2024), ISSN 2521-327X, <https://doi.org/10.22331/q-2024-02-26-1266>.
- [96] A. Maksymov, J. Nguyen, Y. Nam, and I. Markov, Enhancing quantum computer performance via symmetrization (2023), [2301.07233](https://arxiv.org/abs/2301.07233).
- [97] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Simulation of electronic structure hamiltonians using quantum computers, *Mol. Phys.* **109**, 735 (2011).
- [98] N. Moll, A. Fuhrer, P. Staar, and I. Tavernelli, Optimizing qubit resources for quantum chemistry simulations in second quantization on a quantum computer, *J. Phys. A* **49**, 295301 (2016).
- [99] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem, arXiv preprint quant-ph/9511026 (1995).
- [100] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2010).
- [101] R. D. Somma, Quantum eigenvalue estimation via time series analysis, *New Journal of Physics* **21**, 123025 (2019).
- [102] A. Dutkiewicz, B. M. Terhal, and T. E. O'Brien, Heisenberg-limited quantum phase estimation of multiple eigenvalues with few control qubits, *Quantum* **6**, 830 (2022).

-
- [103] A. Sørensen and K. Mølmer, Quantum computation with ions in thermal motion, *Phys. Rev. Lett.* **82**, 1971 (1999).
 - [104] IonQ, Ionq aria: practical performance (2024), <https://ionq.com/resources/ionq-aria-practical-performance>.
 - [105] C. Blank, D. K. Park, J.-K. K. Rhee, and F. Petruccione, Quantum classifier with tailored quantum kernel, *npj Quantum Inf.* **6**, 41 (2020).