OUKA
The University of Osaka
Institutional Knowledge Archive

Title	Developing Efficient Quantum Algorithms for State Preparation and Post-Processing in Quantum Chemistry
Author(s)	Khinevich, Viktor
Citation	大阪大学, 2026, 博士論文
Version Type	VoR
URL	https://doi.org/10.18910/105546
rights	
Note	

The University of Osaka Institutional Knowledge Archive : OUKA

https://ir.library.osaka-u.ac.jp/

The University of Osaka

Developing Efficient Quantum Algorithms for State Preparation and Post-Processing in Quantum Chemistry

VIKTOR KHINEVICH

March 2026

Developing Efficient Quantum Algorithms for State Preparation and Post-Processing in Quantum Chemistry

A dissertation submitted to
THE GRADUATE SCHOOL OF ENGINEERING SCIENCE
The University of Osaka
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY IN SCIENCE

BY

VIKTOR KHINEVICH

March 2026

Abstract

Quantum computers are anticipated to provide computational advantages over classical devices for specific classes of problems. In quantum chemistry, such an advantage can be achieved on fault-tolerant quantum computers (FTQCs) using quantum phase estimation (QPE), which can, in principle, provide accurate electronic energies with polynomial scaling. However, the practical success of QPE and related algorithms critically depends on the availability of high-quality initial quantum states with sufficient overlap with the target eigenstates. Moreover, even FTQCs are expected to be limited in size and total computational resources for the foreseeable future, restricting quantum calculations to reduced subspaces and allowing only a partial treatment of electronic correlation. For chemically meaningful predictions, the remaining correlation effects must therefore be recovered through post-processing.

Motivated by these limitations, this dissertation develops two complementary strategies for the preparation of chemically relevant quantum states on FTQCs and a classical post-processing scheme. The first strategy introduces a unified framework for power-iteration-like state preparation methods based on generalized quantum signal processing (GQSP). This framework encompasses and generalizes several quantum algorithms, including quantum power iteration, quantum power Lanczos, quantum inverse iteration, and quantum folded spectrum method, enabling efficient energy filtering. The second strategy enforces fundamental physical symmetries of electronic Hamiltonians, such as particle number and spin. Symmetry projectors are constructed using linear combination of unitaries, GQSP and generalized quantum singular value transformation. They ensure symmetry-adapted state preparation and improve overlap with target eigenstates. A classical post-processing framework was also designed to recover the electronic correlation not captured by quantum calculations in restricted subspaces. This method, referred to as QC-CBT-AFQMC (quantum-classical computational basis tomography auxiliary-field quantum Monte Carlo), enables the reconstruction of quantum-prepared multiconfigurational wave functions and their subsequent refinement.

Together, the proposed quantum state preparation algorithms and the QC-CBT–AFQMC post-processing pipeline establish a concrete hybrid quantum–classical workflow. These approaches extend the reach of FTQC toward quantitatively accurate and chemically relevant simulations of strongly correlated molecular systems.

Abbreviations

Acronym	Description
AA	Amplitude Amplification
AFQMC	Auxiliary-Field Quantum Monte Carlo
BO	Born–Oppenheimer approximation
CAS	Complete Active Space
CASCI	Complete Active Space Configuration Interaction
CASPT2	Complete Active Space Second-Order Perturbation Theory
CASSCF	Complete Active Space Self-Consistent Field
CBS	Complete Basis Set
CBT	Computational Basis Tomography
CC	Coupled Cluster
CCSD	Coupled Cluster with Singles and Doubles
CCSD(T)	CCSD with Perturbative Triples
CCSDT	Coupled Cluster with Singles, Doubles, and Triples
CI	Configuration Interaction
CISD	Configuration Interaction with Singles and Doubles
CISDT	Configuration Interaction with Singles, Doubles, and Triples
COM	Center of Mass
DB-QITE	Double-Bracket Quantum Imaginary-Time Evolution
DMC	Diffusion Monte Carlo
FCI	Full Configuration Interaction
FTQC	Fault-Tolerant Quantum Computation
GQSP	Generalized Quantum Signal Processing
GQSVT	Generalized Quantum Singular Value Transformation
GTO	Gaussian-Type Orbitals
HEA	Hardware-Efficient Ansatz
HF	Hartree–Fock
ITE	Imaginary-Time Evolution
JW	Jordan–Wigner mapping
LCAO	Linear Combination of Atomic Orbitals
LCU	Linear Combination of Unitaries
MCSCF	Multiconfigurational Self-Consistent Field
MO	Molecular Orbitals
MPn	Møller–Plesset Perturbation Theory

Acronym	**Description**
MRCC	Multireference Coupled Cluster
MRCISD	Multireference Configuration Interaction with Singles and Doubles
MRMP2	Multireference Møller–Plesset Second-Order Perturbation Theory
MRPT	Multireference Perturbation Theory
NCOM	Nuclear Center of Mass
NEVPT2	n-Electron Valence State Second-Order Perturbation Theory
NP	Nondeterministic Polynomial Time
PEC	Potential Energy Curve
PES	Potential Energy Surface
ph-AFQMC	Phaseless AFQMC
PL	Power Lanczos Method
PT	Perturbation Theory
PVM	Projection-valued measure
QFD	Quantum Filter Diagonalization
QFSM	Quantum Folded Spectrum Method
QII	Quantum Inverse Iteration Method
QITE	Quantum Imaginary-Time Evolution
QMA	Quantum Merlin–Arthur
QMC	Quantum Monte Carlo
QPE	Quantum Phase Estimation
QPI	Quantum Power Iteration Method
QPL	Quantum Power Lanczos Method
QROM	Quantum Read-Only Memory
QSCI	Quantum Selected Configuration Interaction
QSE	Quantum Subspace Expansion
QSP	Quantum Signal Processing
QSVT	Quantum Singular Value Transformation
QC-AFQMC	Quantum-Classical Hybrid AFQMC
QC-CBT-TCC	Quantum-Classical Hybrid Tailored Coupled Cluster with CBT
RHF	Restricted Hartree–Fock
ROHF	Restricted Open-Shell Hartree–Fock
SA-CASSCF	State-Averaged CASSCF
SCF	Self-Consistent Field
SQD	Sample-Based Quantum Diagonalization
STO	Slater-Type Orbitals
THC	Tensor Hypercontraction
UCC	Unitary Coupled Cluster
UCCD	Unitary Coupled Cluster with Doubles
UCCSD	Unitary Coupled Cluster with Singles and Doubles
VITE	Variational Imaginary-Time Evolution
VMC	Variational Monte Carlo
VQE	Variational Quantum Eigensolver
(X)MCQDPT2	(Extended) Multiconfigurational Quasi-Degenerate Perturbation Theory of the Second Order

Contents

1 Introduction **1**

2 Preliminaries **7**

2.1 Fundamentals of Quantum Chemistry 8

2.1.1 Electronic Structure Problem 8

2.1.1.1 Quantum Mechanical Description of Molecules . 8

2.1.1.2 Separation of Center of Mass 9

2.1.1.3 Separation of Electronic and Nuclear Motion . . 10

2.1.2 Electronic Wave Function 12

2.1.3 Second Quantization 15

2.1.4 Mean-Field Foundations 18

2.1.5 Capturing Correlation 21

2.1.5.1 Perturbation Theory 21

2.1.5.2 Configuration Interaction 22

2.1.5.3 Coupled Cluster Theory 24

2.1.5.4 Multiconfigurational Self-Consistent Field 25

2.1.5.5 Multireference Perturbation, Configuration Interaction, and Coupled Cluster 27

2.1.6 Quantum Monte Carlo Approaches 29

2.1.6.1 Variational Monte Carlo 29

2.1.6.2 Diffusion Monte Carlo 30

2.1.6.3 Auxiliary-Field QMC 31

2.1.7 Computational and Conceptual Problems of Modern Quantum Chemistry 32

2.2 Fundamentals of Quantum Computing 34

2.2.1 Emergence of Quantum Computing 34

2.2.2 Deterministic Evolution 35

2.2.2.1 State of a Quantum Computer and Its Evolution 35

2.2.2.2 Quantum Gates 36

2.2.2.3 Universal Gate Set 38

2.2.2.4 Circuit Cost Models and Compilation 40

2.2.3 Measurements 40

2.2.3.1 Projective Measurements and the Born Rule . . 41

2.2.3.2 What Do We Measure in Quantum Computing? 41

2.2.3.3 Hadamard Test . . . 43
2.2.3.4 State Tomography . . . 44
2.2.3.5 Pauli Shadows . . . 45
2.2.3.6 Clifford Shadows . . . 45
2.2.3.7 Matchgate Shadows . . . 46
2.2.4 Quantum Algorithms . . . 46
2.2.4.1 Fermion to Qubit Mapping . . . 47
2.2.4.2 Linear Combination of Unitaries . . . 47
2.2.4.3 Qubitization . . . 49
2.2.4.4 Grover Search and Amplitude Amplification . . . 49
2.2.4.5 Quantum Signal Processing . . . 50
2.2.4.6 Generalized Quantum Signal Processing . . . 52
2.2.4.7 Quantum Phase Estimation . . . 52
2.2.4.8 Quantum Selected Configuration Interaction . . . 54
2.2.4.9 Quantum Krylov Subspace Methods . . . 54
2.2.5 Quantum State Preparation . . . 55
2.2.5.1 Variational Quantum Eigensolver . . . 55
2.2.5.2 Imaginary-Time Evolution . . . 56
2.2.5.3 State Filters . . . 56
2.2.6 Classical Post-Processing . . . 57
2.2.6.1 Quantum-Classical Hybrid NEVPT . . . 57
2.2.6.2 Quantum-Classical Hybrid AFQMC . . . 58
2.2.6.3 Quantum-Classical Hybrid Tailored Coupled Cluster . . . 59
2.2.7 Open Problems of Quantum Computing for Quantum Chemistry . . . 59

3 Quantum Power Iteration Unified Using Generalized Quantum Signal Processing 61
3.1 Introduction . . . 61
3.2 Theory . . . 63
3.2.1 Polynomial Transformation of a Hamiltonian . . . 63
3.2.2 Quantum Power Iteration Method . . . 65
3.2.3 Quantum Power Lanczos Method . . . 66
3.2.4 Quantum Inverse Iteration Method . . . 68
3.2.5 Quantum Folded Spectrum Method . . . 70
3.3 Computational Details . . . 71
3.4 Results and Discussion . . . 72
3.4.1 Theoretical Scaling . . . 72
3.4.2 Numerical Comparison of QPI and QPL . . . 79
3.4.3 Numerical Results for QII . . . 83
3.4.4 Excited States Calculations Using QFSM . . . 87
3.5 Conclusion . . . 89

4 Symmetry-Adapted State Preparation for Quantum Chemistry on Fault-Tolerant Quantum Computers 91
4.1 Introduction . . . 92
4.2 Theory . . . 93
4.2.1 Symmetry Operators . . . 93
4.2.2 Symmetry Projectors . . . 95
4.2.2.1 Integral Projectors . . . 95
4.2.2.2 Lagrange Interpolation Projectors . . . 96
4.2.3 Projectors Implementation . . . 96
4.2.3.1 Linear Combination of Unitaries . . . 96
4.2.3.2 Generalized Quantum Signal Processing . . . 99
4.2.3.3 Generalized Quantum Singular Value Transformation . . . 99
4.3 Computational Details . . . 100
4.4 Results and Discussion . . . 101
4.4.1 Theoretical Scaling . . . 101
4.4.2 Success Probability . . . 102
4.4.3 Numerical Accuracy of Projectors . . . 104
4.4.4 Gate Precision Sensitivity . . . 106
4.4.5 Resource Estimates . . . 109
4.4.6 Projectors as a QPE Filter . . . 112
4.4.7 Limitations and Outlook . . . 115
4.5 Conclusion . . . 119

5 Enhancing Quantum Computations with the Synergy of Auxiliary-Field Quantum Monte Carlo and Computational Basis Tomography 121
5.1 Introduction . . . 122
5.2 Methods . . . 124
5.2.1 Computational Basis Tomography . . . 124
5.2.2 Auxiliary-Field QMC . . . 127
5.2.3 QC-CBT-AFQMC . . . 129
5.3 Results and Discussion . . . 130
5.3.1 Potential Energy Curves . . . 131
5.3.2 Influence of the Trial Wave Function on AFQMC Convergence . . . 135
5.3.3 Effect of Measurement Quantity on CBT . . . 138
5.3.4 Analysis of Reaction Barriers . . . 138
5.4 Conclusions . . . 142

6 Conclusions 143

Appendix 149
Appendix A: Supporting Information for Chapter 3 . . . 149
A1 Angle Finding . . . 149

A2 Comparison of QII with Q-Inv and I-Iter for Different Truncation . . . 152
A3 Pseudocodes . . . 154
Appendix B: Supporting Information for Chapter 4 . . . 158
B1 Generalized Quantum Singular-Value Transformation . . . 158
Appendix C: Supporting Information for Chapter 5 . . . 159
C1 Dependency Analysis of QC-CBT-AFQMC on the Number of Measurements for Ethylene and Nitrogen . . . 159
C2 Complete Basis Set Extrapolation . . . 159

Bibliography **167**

Acknowledgments **183**

List of Publications **185**

Chapter 1

Introduction

Quantum computers are devices that exploit quantum mechanical phenomena, such as superposition, entanglement, and contextuality, to perform computations. While any physical realization of a classical computer ultimately obeys the laws of quantum mechanics, the microscopic degrees of freedom in a classical device interact uncontrollably with the environment, leading to rapid decoherence. As a result, classical computers can be accurately described in terms of classical bits. In contrast, a quantum computer is engineered to isolate and control its state such that quantum coherence can be preserved for a sufficiently long time to execute quantum algorithms. This opens up the possibility of solving certain computational problems more efficiently than is believed to be possible on classical hardware.

Although a rigorous and general separation between problems efficiently solvable by quantum and classical computers remains elusive [1, 2], several quantum algorithms exhibit asymptotic performance superior to all currently known classical counterparts. Notable examples include Shor's algorithm for integer factorization [3, 4] and Grover's algorithm for unstructured search [5]. More generally, computational primitives related to these algorithms, such as quantum phase estimation (QPE) [6, 7], amplitude amplification (AA) [8, 9, 10, 11, 12], and quantum amplitude estimation [8] form the foundation of many modern quantum algorithms.

One of the earliest and most natural applications envisioned for quantum computers is the simulation of quantum many-body systems [13, 14, 15]. Quantum chemistry represents a particularly important and challenging class of such problems. The central difficulty in *ab initio* electronic structure theory is often described as the curse of dimensionality: the dimension of the relevant Hilbert space grows combinatorially or exponentially with system size. For example, even a relatively small molecule such as propane (C_3H_8) yields a full configuration interaction (FCI) space of dimension on the order of 1.3×10^{12} in a minimal STO-3G basis [16], making exact classical treatment infeasible except on large-scale high-performance computing platforms.

Beyond this explosion in dimension, quantum chemistry also runs into formal

limits on what is efficiently computable. Approximating the ground state energy of a general local Hamiltonian is **QMA**-complete [17, 18], thus it requires exponential time even on a quantum computer in the worst case. Moreover, the prediction whether the given Hamiltonian represents insulator or conductor is undecidable in general [19]. Even Hartree–Fock (HF) is a **NP**-complete problem in the worst case [20, 18], meaning it cannot be solved in polynomial time on a classical computer (and presumably on a quantum as well). Trying to use randomness to circumvent the problems is also doomed, since quantum Monte Carlo (QMC) methods face the sign problem for fermions, which is **NP**-hard in general [21].

Despite all of these theoretical results, practical quantum chemistry is still possible because molecular Hamiltonians have structure and symmetries, and because heuristic chemical reasoning can be applied. Thus, quantum computing is not expected to eliminate worst-case hardness or undecidability; rather, it provides new heuristics and new computational primitives that may extend the repertoire of feasible approximations.

Quantum computers offer a different way to represent and manipulate many-body states. N qubits naturally store 2^N complex amplitudes, and quantum circuits can implement structured transformations of these amplitudes. This motivates a broad family of quantum algorithms for quantum chemistry, including methods for estimating ground and excited state energies [22, 7, 23, 24, 25], simulating dynamics [26, 27, 28], and sampling many-body quantum states [29, 30].

A central practical obstacle, common to essentially all quantum algorithms, is state preparation. The performance of many algorithms deteriorates rapidly if the initial state has a small overlap with the target eigenstate, leading to prohibitive circuit depth or measurement cost [7, 29, 30, 31, 32, 33, 23]. If an exact state could be prepared efficiently, QPE would, in principle, enable the extraction of ground and excited state energies with exponential speedup to any known classical method. Moreover, accurate state preparation would allow efficient sampling in quantum-selected configuration interaction (QSCI) [29, 30] and the construction of accurate effective Hamiltonians in Krylov subspace methods [34, 35, 36, 31]. Therefore, state preparation is a primary bottleneck for turning quantum algorithms for chemistry into useful computational tools.

At the same time, even fault-tolerant quantum computers (FTQCs) are expected to remain severely limited in size and computational resources [37, 38, 39, 40]. Quantum processors operate at clock speeds orders of magnitude slower than classical processors, implying that practical limitations arise not only from the number of available qubits but also from the total achievable computation time. These constraints restrict quantum simulations to reduced active spaces and allow only for a partial treatment of electronic correlation effects. However, for chemically accurate predictions, correlation outside the active space often plays a crucial role and must be accounted for by additional means. A natural approach is therefore to employ hybrid strategies in which quantum hardware is used to generate multiconfigurational states in restricted subspaces, while the remaining correlation is recovered using classical electronic structure methods.

State tomography methods are widely implemented to measure quantum states for use in classical post-processing schemes. However, state tomography methods are often prohibitively expensive, both in terms of measurement requirements and the classical computational cost [41, 42, 43, 44].

In this dissertation, we address these two closely connected challenges: the preparation of chemically relevant quantum states and their efficient post-processing using classical algorithms. On the quantum side, we develop and analyze a family of power-iteration-like methods adapted for quantum computing with generalized quantum signal processing (GQSP). The approach is dictated by the fact that the Hamiltonian is the only fundamental source of information about the quantum system. Hence, it drives the evolution toward desired energy eigenstates similarly to imaginary-time evolution (ITE), but more versatile.

In addition, molecular electronic states obey exact symmetries, including particle number and spin. Conservation of these symmetries is important for a good overlap of the initial state with the exact one, but some quantum state preparation schemes may violate these symmetries. Thus, we construct efficient symmetry projectors that enforce fundamental conserved quantities of molecular Hamiltonians, such as particle number and spin, using the linear combination of unitaries (LCU), GQSP, and generalized quantum singular value transformation (GQSVT). This ensures that the prepared states reside within physically meaningful subspaces.

Finally, since it is unlikely that quantum computers will directly treat the full molecular Hamiltonian of realistic systems in the near term, we develop a post-processing pipeline based on computational basis tomography (CBT) to reconstruct quantum-prepared many-body states. These reconstructed states are then employed as trial wavefunctions in auxiliary-field QMC (AFQMC), a state-of-the-art classical method for strongly correlated electronic systems. In this way, quantum hardware is used to generate high-quality multiconfigurational states in restricted active spaces, which can subsequently enhance the performance of advanced classical algorithms, thereby extending the practical applicability of quantum computing to chemically relevant problems.

The remainder of this dissertation is organized to mirror the computational workflow of quantum chemistry simulations using quantum computers. After establishing the necessary theoretical background in Chap. 2, we proceed starting from the quantum state preparation via power-iteration-based methods (Chap. 3) and symmetry projection (Chap. 4) — to the classical aspects, namely efficient readout and post-processing using QC-CBT-AFQMC (Chap. 5). This progression reflects the sequence of tasks required in practice: preparing a quantum state, enforcing physical constraints, and extracting chemically relevant information for subsequent classical treatment.

Dissertation Overview

Chapter 2: Preliminaries

We provide the foundational concepts required for the remainder of the dissertation. It begins with a compact formulation of the molecular electronic structure problem and a brief overview of the principal classical methods used to study low-lying electronic states of molecules (Sec. 2.1). The second half of the chapter provides the basic elements of quantum computing theory, followed by a discussion focused on quantum computing approaches relevant to quantum chemistry (Sec. 2.2).

Chapter 3: Quantum Power Iteration Unified Using Generalized Quantum Signal Processing

We introduce several power-method-based techniques for quantum state preparation that can be implemented within a single framework using GQSP. This unification enables the realization and resource estimation of various methods that require polynomial transformations, such as the quantum power iteration (QPI), its modification, the quantum power Lanczos (QPL), and the quantum inverse iteration (QII) methods. Moreover, we introduce the quantum folded spectrum method (QFSM) for excited states preparation. We analyze the success probability, query complexity, and qubit requirements of the algorithms, showing that they achieve near-optimal query complexity. Explicit circuits are constructed and numerically tested for several molecular Hamiltonians.

Chapter 4: Symmetry-Adapted State Preparation for Quantum Chemistry on Fault-Tolerant Quantum Computers

We implement symmetry projectors tailored to the continuous symmetries of molecular electronic structure problems, such as particle-number and spin symmetry. These projectors are constructed using LCU, GQSP, and GQSVT approaches, which allows them to be efficiently implemented within FTQC. The chapter analyzes the CNOT and T gates scaling behavior of these constructions, examines their numerical stability under realistic gate synthesis precision, and evaluates their performance as state filters for QPE workflows. Numerical tests for molecular systems illustrate that symmetry projection significantly improves overlap with chemically relevant states and is feasible for large strongly correlated systems, such as FeMoco.

Chapter 5: Enhancing Quantum Computations with the Synergy of Auxiliary-Field Quantum Monte Carlo and Computational Basis Tomography

We introduce a workflow that combines quantum prepared states with classical AFQMC by using CBT as an efficient readout technique. The method classically reconstructs multiconfigurational wave functions with a practical number of

shots and gate operations, enabling its use even on near-term quantum hardware. Compared with other tomography-based hybrid approaches, our QC-CBT-AFQMC method has reduced prefactor and improved scaling, making it possible to treat larger molecular systems.

Chapter 6: Conclusions

This chapter summarizes the main contributions of the dissertation and outlines promising directions for future research in quantum algorithms for quantum chemistry.

Chapter 2

Preliminaries

This chapter collects background material in quantum chemistry and quantum computing that will be used throughout the remainder of this dissertation. The aim is not to provide an exhaustive review of either field, but to fix the notation, recall standard concepts, and emphasize those aspects that are most relevant for quantum algorithms for electronic structure and for the post-processing strategies developed in later chapters.

Section 2.1 introduces the fundamentals of quantum chemistry. It begins with a review of the electronic structure problem, followed by a formulation of the theory in the language of second quantization. Then it presents mean-field methods and surveys a range of approaches for capturing electron correlation. The section concludes by summarizing key computational and conceptual challenges that motivate the search for quantum-enhanced methods.

Section 2.2 presents the fundamentals of quantum computing. It opens with a brief overview of the emergence of quantum computing and then introduces the deterministic aspects of quantum computation, including quantum states, gates, and circuits. The section then turns to measurements, the probabilistic component of quantum computation, and culminates in a discussion of quantum state tomography and classical shadow techniques that will be directly relevant for later chapters. This is followed by a survey of quantum algorithms applicable to quantum chemistry. The section closes by outlining open problems at the interface of quantum computing and quantum chemistry, providing context for the contributions developed in this dissertation.

Throughout this dissertation, symbols may be reused in different contexts. All symbols are defined locally, and when a symbol is redefined in a later chapter, the definition given in that chapter supersedes any earlier usage.

2.1 Fundamentals of Quantum Chemistry

2.1.1 Electronic Structure Problem

2.1.1.1 Quantum Mechanical Description of Molecules

Quantum mechanics emerged in the 1920s as a response to experimental phenomena that classical physics could not explain, such as discrete atomic spectra, the stability of atoms, and blackbody radiation. From the very beginning, its methods were applied not only to isolated atoms but also to molecules, since the same formalism naturally describes systems of interacting nuclei and electrons. Almost immediately after Schrödinger formulated his wave equation in 1926 [45], the attention turned to molecules.

According to the postulates of quantum mechanics, the state of a physical system is represented by a time-dependent vector $|\Psi(t)\rangle$ in a Hilbert space. A coordinate representation of this state vector is called a wave function, $\langle\vec{r}|\Psi(t)\rangle = \Psi(\vec{r},t)$, where $|\vec{r}\rangle$ is a coordinate basis vector. All observable outcomes of the experiments can be predicted using this state vector. The fundamental equation governing the dynamics of the wave function is the Schrödinger equation:

$$i\hbar\frac{\partial}{\partial t}|\Psi(t)\rangle = \hat{H}|\Psi(t)\rangle, \tag{2.1}$$

where $\hbar$ is the Dirac constant and $\hat{H}$ is the Hamiltonian operator of a quantum system.

From a mechanical point of view, a molecule is a collection of nuclei and electrons that interact through the electromagnetic field, predominantly via the Coulomb interaction. When talking about molecules, it is convenient to change from the International System of Units (SI) to atomic units. In the atomic units system, the electron rest mass m_e, the magnitude of the electronic charge e, the Dirac constant $\hbar$, and the vacuum permittivity ϵ_0 multiplied by 4π are set to unity:

$$\hbar = e = m_e = 4\pi\epsilon_0 = 1. \tag{2.2}$$

All other units are redefined using this basis. In this dissertation, we use atomic units, unless explicitly stated otherwise.

In coordinate space, the molecular Hamiltonian can be written as

$$\begin{aligned}\hat{H} = &\underbrace{-\sum_A \frac{1}{2M_A}\nabla_A^2}_{\hat{T}_\mathrm{n}} \underbrace{-\sum_i \frac{1}{2}\nabla_i^2}_{\hat{T}_\mathrm{e}} \underbrace{-\sum_{i,A}\frac{Z_A}{|\vec{r}_i-\vec{R}_A|}}_{\hat{V}_\mathrm{en}} \\ &+\underbrace{\sum_{i<j}\frac{1}{|\vec{r}_i-\vec{r}_j|}}_{\hat{V}_\mathrm{ee}} + \underbrace{\sum_{A<B}\frac{Z_A Z_B}{|\vec{R}_A-\vec{R}_B|}}_{\hat{V}_\mathrm{nn}} + \hat{H}_\mathrm{es} + \hat{H}_\mathrm{hfs},\end{aligned} \tag{2.3}$$

where $\vec{\nabla}_p = \left(\frac{\partial}{\partial X_p}, \frac{\partial}{\partial Y_p}, \frac{\partial}{\partial Z_p}\right)^T$, and indices A, B label nuclei with positions $\vec{R}_A$, charges Z_A, and masses M_A, while i, j label electrons with positions $\vec{r}_i$. The first

two terms correspond to nuclear and electronic kinetic energies, $\hat{T}_\mathrm{n}$ and $\hat{T}_\mathrm{e}$, the third to electron–nucleus attraction $\hat{V}_\mathrm{en}$, the fourth to electron–electron repulsion $\hat{V}_\mathrm{ee}$, and the fifth to nucleus–nucleus repulsion $\hat{V}_\mathrm{nn}$. The term $\hat{H}_\mathrm{es}$ contains spin-dependent electronic interactions between electron spin magnetic moments and the magnetic moments associated with the orbital motion of electrons and nuclei (spin–orbit, spin–spin couplings, etc.), while $\hat{H}_\mathrm{hfs}$ corresponds to interactions of nuclear spin magnetic moments with other magnetic moments in the molecule (hyperfine couplings). For the simplicity of further considerations we will not take into account $\hat{H}_\mathrm{es}$ and $\hat{H}_\mathrm{hfs}$ terms, however, they are important for subtle effects, such as fine and hyperfine structure in spectra, nuclear magnetic resonance and electron paramagnetic resonance observables.

Since the molecular Hamiltonian in Eq. (2.3) does not explicitly depend on time, we can seek stationary solutions of Eq. (2.1) of the form

$$\hat{H}|\Psi_k\rangle = E_k|\Psi_k\rangle, \quad |\Psi_k(t)\rangle = e^{-iE_k t}|\Psi_k\rangle, \tag{2.4}$$

where $|\Psi_k\rangle$ are the Hamiltonian eigenstates and E_k are the corresponding energy eigenvalues.

In principle, solving this eigenvalue problem exactly would provide a complete description of molecular properties, yielding reaction energies, optimal structures, and spectral data. In practice, however, the dimensionality of the Hilbert space grows exponentially with system size, making exact solutions feasible only for the smallest systems. This motivates the introduction of systematic simplifications.

2.1.1.2 Separation of Center of Mass

The first natural simplification is the removal of the translational degrees of freedom associated with motion of the molecule as a whole. Because the full molecular Hamiltonian is translationally invariant, it is convenient to introduce the center of mass (COM) coordinate

$$\vec{R}_\mathrm{COM} = \frac{1}{M_\mathrm{tot}}\left(\sum_A M_A \vec{R}_A + \sum_i \vec{r}_i\right), \qquad M_\mathrm{tot} = \sum_A M_A + N_\mathrm{elec}, \tag{2.5}$$

where M_tot is the total mass of the molecule in atomic units and N_elec is the number of electrons. The position of each particle p (nucleus or electron) can then be written as a sum of the COM coordinate and a relative coordinate,

$$\vec{R}_p = \vec{R}_\mathrm{COM} + \vec{R}'_p. \tag{2.6}$$

Upon transforming the Hamiltonian to the new coordinates, one obtains

$$\begin{aligned}\hat{H}' = &\underbrace{-\frac{1}{2M_\mathrm{tot}}\nabla^2_\mathrm{COM}}_{\hat{T}_\mathrm{COM}} \underbrace{-\sum_p \frac{1}{2M_p}\nabla'^{\,2}_p}_{\hat{T}^0} \underbrace{-\frac{1}{2M_\mathrm{tot}}\sum_{p,r}\nabla'_p\cdot\nabla'_r}_{\hat{T}_\mathrm{cross}} \\ &+ \hat{V}(\vec{R}'_2,\ldots,\vec{R}'_N) = \hat{T}_\mathrm{COM} + \hat{H}_\mathrm{rve},\end{aligned} \tag{2.7}$$

where N is the total number of particles, $\hat{V}(\vec{R}'_2, \ldots, \vec{R}'_N)$ consists of $\hat{V}_{\mathrm{en}}$, $\hat{V}_{\mathrm{ee}}$, and $\hat{V}_{\mathrm{nn}}$ terms from Eq. (2.3). $\hat{H}_{\mathrm{rve}}$ is the rotational–vibrational–electronic (rovibronic) Hamiltonian. The kinetic energy separates into the kinetic energy of the center of mass $\hat{T}_{\mathrm{COM}}$, the kinetic energy of the relative internal motion $\hat{T}^0$, and a mass-polarization or cross term $\hat{T}_{\mathrm{cross}}$. This cross term arises because the COM coordinate is defined in terms of all particle coordinates, which couples the relative motions.

Since the Hamiltonian in Eq. (2.7) is separable with respect to $\vec{R}_{\mathrm{COM}}$, the molecular state can be expressed as a product of COM translation motion part $|\Psi_{\mathrm{COM}}\rangle$ and rovibronic part $|\Phi_{\mathrm{rve}}\rangle$:

$$\begin{aligned}\Psi(\vec{R}_1, \ldots, \vec{R}_N) &= (\langle \vec{R}_{\mathrm{COM}}| \otimes \langle \vec{R}'_2, \ldots, \vec{R}'_N|)(|\Psi_{\mathrm{COM}}\rangle \otimes |\Phi_{\mathrm{rve}}\rangle) \\ &= \Psi_{\mathrm{COM}}(\vec{R}_{\mathrm{COM}})\, \Phi_{\mathrm{rve}}(\vec{R}'_2, \ldots, \vec{R}'_N),\end{aligned} \tag{2.8}$$

where $\Psi_{\mathrm{COM}}(\vec{R}_{\mathrm{COM}})$ is the COM translational wave function and $\Phi_{\mathrm{rve}}(\vec{R}'_2, \ldots, \vec{R}'_N)$ describes all internal rovibronic degrees of freedom.

The COM wave function is a plane wave and contributes only the kinetic energy of the free translation of the molecule. Because most molecular observables depend solely on internal structure and dynamics, this translational contribution can usually be neglected. From this point onward, all considerations refer exclusively to the internal Hamiltonian $\hat{H}_{\mathrm{rve}}$.

2.1.1.3 Separation of Electronic and Nuclear Motion

The internal (rovibronic) Schrödinger equation can be written as

$$\begin{aligned}\Bigg\{ -\sum_{p=2}^{N} \frac{1}{2M_p} \nabla'^2_p - \frac{1}{2M_{\mathrm{tot}}} \sum_{p,r=2}^{N} \nabla'_p \cdot \nabla'_r \\ + \hat{V}(\vec{R}'_2, \ldots, \vec{R}'_N) \Bigg\} \Phi_{\mathrm{rve}}(\vec{R}'_2, \ldots, \vec{R}'_N) \\ = E_{\mathrm{rve}}\, \Phi_{\mathrm{rve}}(\vec{R}'_2, \ldots, \vec{R}'_N),\end{aligned} \tag{2.9}$$

where the indices $p, r = 2, \ldots, N$ label the internal coordinates obtained after removing the overall center-of-mass motion.

Further simplification is achieved by separating electronic and nuclear motions. To this end, we perform one more coordinate transformation and introduce the nuclear center of mass (NCOM) coordinate

$$\vec{R}_{\mathrm{NCOM}} = \frac{1}{M_{\mathrm{nuc}}} \sum_A M_A \vec{R}'_A, \qquad M_{\mathrm{nuc}} = \sum_A M_A, \tag{2.10}$$

and coordinates relative to the NCOM,

$$\vec{R}''_p = \vec{R}'_p - \vec{R}_{\mathrm{NCOM}}, \tag{2.11}$$

for all particles p (nuclei and electrons). We denote by $\vec{R}_{\mathrm{nuc}}$ and $\vec{r}_{\mathrm{elec}}$ the sets of nuclear and electronic coordinates relative to the NCOM and drop the double primes for simplicity.

In these coordinates, the internal Hamiltonian can be written as

$$\left\{\hat{T}_{\mathrm{elec}} + \hat{T}_{\mathrm{nuc}} + \hat{V}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}})\right\}\Phi_{\mathrm{rve}}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}}) = E_{\mathrm{rve}}\,\Phi_{\mathrm{rve}}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}}), \quad (2.12)$$

with

$$\hat{T}_{\mathrm{elec}} = \underbrace{-\frac{1}{2}\sum_i \vec{\nabla}_i^2}_{\hat{T}^0_{\mathrm{elec}}} - \frac{1}{2M_{\mathrm{nuc}}}\sum_{i,j}\vec{\nabla}_i \cdot \vec{\nabla}_j, \quad (2.13)$$

$$\hat{T}_{\mathrm{nuc}} = -\sum_A \frac{1}{2M_A}\vec{\nabla}_A^2. \quad (2.14)$$

Here, the first term in $\hat{T}_{\mathrm{elec}}$ is the usual electronic kinetic energy $\hat{T}^0_{\mathrm{elec}}$, while the second term is the mass-polarization term, which couples the motion of different electrons and is proportional to $1/M_{\mathrm{nuc}}$.

At this point, we note that nuclear masses are typically much larger than the electron mass; for example, $m_p/m_e \approx 1836$ for a proton, and other nuclei are even more massive. Consequently, electrons move much faster and can adjust almost instantly to the comparatively slow motion of the nuclei. This speed separation motivates the Born–Oppenheimer (BO) approximation, introduced in 1927 [46], in which electronic and nuclear motions are approximately separated.

Shortly thereafter, the adiabatic theorem was formulated by Born and Fock in 1928 [47]. It states that a physical system remains in its instantaneous eigenstate if a perturbation acts on it sufficiently slowly and if there is a finite energy gap between this eigenvalue and the rest of the spectrum. In our context, the "physical system" is the electronic subsystem in the field of the nuclei, and the "perturbation" is the nuclear motion. In turn, BO treatment assumes not only that the electrons adiabatically follow the nuclear motion but also neglects the mass-polarization term (the second term in $\hat{T}_{\mathrm{elec}}$ in Eq. (2.13)), since it is suppressed by the small factor $1/M_{\mathrm{nuc}}$.

Under these assumptions, the rovibronic equation decouples into separate equations for electronic and nuclear motion,

$$\underbrace{\left\{\hat{T}^0_{\mathrm{elec}} + \hat{V}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}})\right\}}_{\hat{H}_{\mathrm{elec}}}\Phi_{\mathrm{elec}}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}}) = E_{\mathrm{elec}}(\vec{R}_{\mathrm{nuc}})\,\Phi_{\mathrm{elec}}(\vec{R}_{\mathrm{nuc}}, \vec{r}_{\mathrm{elec}}) \quad (2.15)$$

$$\left\{\hat{T}_{\mathrm{nuc}} + E_{\mathrm{elec}}(\vec{R}_{\mathrm{nuc}})\right\}\Phi_{\mathrm{rv}}(\vec{R}_{\mathrm{nuc}}) = E^0_{\mathrm{rve}}\,\Phi_{\mathrm{rv}}(\vec{R}_{\mathrm{nuc}}), \quad (2.16)$$

where $E_{\mathrm{elec}}(\vec{R}_{\mathrm{nuc}})$ is the electronic energy as a function of the nuclear geometry, often called the potential energy surface (PES). Usually, only a cross section

of PES is investigated along one internal coordinate. This cross section is called the potential energy curve (PEC). $E_{\text{elec}}(\vec{R}_{\text{nuc}})$ also includes the nuclear-nuclear repulsion term V_{nn}, which is a constant for a given fixed nuclear position. Determining $E_{\text{elec}}(\vec{R}_{\text{nuc}})$ is the central goal of the electronic structure problem. The function $\Phi_{\text{elec}}(\vec{R}_{\text{nuc}}, \vec{r}_{\text{elec}})$ is the electronic wave function, which depends parametrically on the nuclear coordinates. The second equation is the nuclear rotation–vibration Schrödinger equation, which governs the nuclear motion on the PES $E_{\text{elec}}(\vec{R}_{\text{nuc}})$. The solutions to this equation and the approximations used are discussed in Ref. [48]. The quantity E^0_{rve} is the rovibronic energy within the BO approximation.

The validity of the BO separation relies on a sufficiently large energy gap between the electronic state of interest and other electronic states in the region of configuration space explored by the molecular geometry. When electronic states become nearly degenerate, nonadiabatic couplings between them can become large, and the separation of electronic and nuclear motion is no longer valid. In such regimes, nuclear motion can strongly influence electronic dynamics. This occurs, for instance, in vibronic transitions and in the vicinity of conical intersections, where even small nuclear displacements can induce transitions between electronic states.

Even after BO separation, the resulting electronic Schrödinger equation (see Eq. (2.15)) remains extremely challenging to solve. In principle, one could attempt to represent the electronic wave function directly on a spatial grid, and such approaches can be numerically exact for very small systems [49]. However, the number of grid points grows exponentially with system size, and such methods do not provide a practical, systematically improvable route for larger molecules.

A key conceptual step in the development of quantum chemistry was the idea of expanding the electronic wave function in a compact set of basis functions. This basis set approximation converts the electronic Schrödinger equation into a matrix eigenvalue problem of controllable dimension, whose accuracy can be improved systematically by enlarging or refining the basis. This transition from direct numerical representations to basis set methods underlies the modern framework of electronic structure theory [50].

2.1.2 Electronic Wave Function

The electronic Hamiltonian in Eq. (2.15), for fixed nuclear positions $\{\vec{R}_A\}$ within the Born–Oppenheimer approximation, has the form

$$\hat{H}_{\text{elec}} = -\frac{1}{2}\sum_i \vec{\nabla}_i^2 + \sum_{i<j}\frac{1}{|\vec{r}_i - \vec{r}_j|} - \sum_{A,i}\frac{Z_A}{|\vec{R}_A - \vec{r}_i|} + \left(\sum_{A<B}\frac{Z_A Z_B}{|\vec{R}_A - \vec{R}_B|}\right). \quad (2.17)$$

The last term corresponding to nuclear repulsion is independent on the electronic coordinates. Thus, it is a constant value for a fixed nuclear coordinate, which we will omit for simplicity of expressions.

It is convenient to separate $\hat{H}_{\text{elec}}$ into a sum of one-electron operators and

the electron–electron interaction,

$$\hat{H}_{\text{elec}} = \hat{H}^0_{\text{elec}} + \hat{V}_{\text{ee}}, \tag{2.18}$$

where

$$\hat{H}^0_{\text{elec}} = \sum_i \hat{h}_i = \sum_i \left\{ -\frac{1}{2}\vec{\nabla}_i^2 - \sum_A \frac{Z_A}{|\vec{R}_A - \vec{r}_i|} \right\}, \quad \hat{V}_{\text{ee}} = \sum_{i<j} \frac{1}{|\vec{r}_i - \vec{r}_j|}. \tag{2.19}$$

Thus, $\hat{H}^0_{\text{elec}}$ is the sum of independent one-electron terms.

This observation suggests a zeroth-order approximation in which the spatial part of the electronic wave function, $\Phi^{\text{space}}_{\text{elec}}(\vec{r}_{\text{elec}})$, is written as a product of one-electron functions:

$$\Phi^{\text{space}}_{\text{elec}}(\vec{r}_{\text{elec}}) = \prod_i \phi_{k_i}(\vec{r}_i), \tag{2.20}$$

where $\phi_{k_i}(\vec{r}_i)$ are one-electron spatial orbitals. The orbitals are eigenfunctions of the one-electron operator $\hat{h}_i$, satisfying

$$\hat{h}_i \, \phi_k(\vec{r}_i) = \varepsilon_k \, \phi_k(\vec{r}_i), \tag{2.21}$$

with eigenvalues ε_k called molecular orbital energies and $\phi_k(\vec{r})$ the corresponding molecular orbitals (MO).

In practice, molecular orbitals are expanded in a finite basis of one-particle functions $\{\chi_\mu(\vec{r})\}$, according to

$$\phi_k(\vec{r}) = \sum_\mu c_{\mu k} \, \chi_\mu(\vec{r}), \tag{2.22}$$

where the coefficients $c_{\mu k}$ are determined in the course of the electronic structure calculation. In most applications, the molecular orbitals are chosen as linear combinations of atomic orbitals (LCAO), which retain physical intuition while providing a mathematically tractable representation of the electronic wave function. Alternative representations also exist; for example, one may expand orbitals in delocalized functions such as plane waves, which are particularly useful for periodic systems in solid-state physics.

Early quantum chemists recognized that analytic hydrogen-like functions, in the form of Slater-type orbitals (STOs), provide a realistic description of atomic electronic structure but lead to very complicated multi-center integrals. STOs are defined as

$$\chi^{\text{STO}}_{n\ell m}(\vec{r}, \zeta) = \mathcal{N} \, |\vec{r}|^{n-1} e^{-\zeta|\vec{r}|} Y_{\ell m}(\vec{r}/|\vec{r}|), \tag{2.23}$$

where $\mathcal{N}$ is a normalization constant, ζ is an effective nuclear charge, n is an integer parameter, analogous to the principal quantum number, and $Y_{\ell m}$ are spherical harmonics with angular momentum ℓ and its projection m.

A major practical breakthrough occurred with the introduction of Gaussian-type orbitals (GTOs) by Boys in 1950 [51]. GTOs uses a Gaussian of the form:

$$\chi^{\text{GTO}}_{n\ell m}(\vec{r}, \alpha) = \mathcal{N} \, |\vec{r}|^{n-1} e^{-\alpha|\vec{r}|^2} Y_{\ell m}(\vec{r}/|\vec{r}|), \tag{2.24}$$

where α is a real parameter. Although GTOs resemble exact atomic orbitals less closely than STOs, they make multi-center integrals analytically tractable. This compromise enabled large-scale quantum chemical calculations.

Subsequent developments have produced families of systematically improvable basis sets, including:

- **STO-nG (minimal)** [52]: each atomic orbital is represented as a contracted Gaussian function that fits a Slater-type orbital, using n primitive Gaussians per STO.
- **Pople split-valence basis sets** [53]: the core is described by a fixed contraction, while the valence space is split into two or more contractions (e.g. 3-21G, 6-31G), with optional polarization functions (e.g. *, **) and diffuse functions (e.g. +, ++) to increase expressivity.
- **Dunning correlation-consistent basis sets (cc-pVXZ)** [54]: hierarchical atomic basis sets for correlated methods where the cardinal numbers $X = \mathrm{D}, \mathrm{T}, \mathrm{Q}, \ldots$ indicate the angular momentum content (optionally augmented with diffuse functions, e.g. aug-cc-pVXZ).
- **Ahlrichs def2 basis sets (def2-SV(P), -TZVP, -TZVPP, -QZVP, ...)** [55]: a modern segmented-contracted split-valence family, where polarization is indicated by P/PP, designed to provide a balanced accuracy–cost trade-off across density functional theory and correlated methods.
- **Plane wave basis sets** [56]: delocalized basis functions compatible with periodic boundary conditions, defined by reciprocal-lattice vectors up to a kinetic-energy cutoff.

Despite the success of basis set approximations, they introduce subtle artifacts. One of the most important is the basis set superposition error (BSSE) [57]. Consider two molecules A and B forming a dimer AB. When computed separately, each monomer can use only the basis functions centered on its own nuclei. However, in the dimer calculation, the electrons of A can also make use of the basis functions of B, and vice versa. This artificial "borrowing" of functions lowers the dimer energy relative to the sum of the monomer energies and leads to an overestimation of the binding energies.

Ultimately, the goal of electronic structure calculations is to approach the complete basis set (CBS) limit, i.e. the energy that would be obtained with an infinite one-particle basis. Because actual basis sets are finite, various extrapolation schemes have been developed to estimate the CBS limit from calculations with basis sets of increasing size. These are discussed in the Appendix C2.

So far we have considered only the spatial part of the electronic wave function and treated electrons implicitly as distinguishable. Electrons are indistinguishable spin-1/2 particles (fermions), and the Pauli exclusion principle implies that no more than two electrons with opposite spins can occupy the same spatial orbital.

Mathematically, this is implemented by requiring that the total electronic wave function changes sign upon the exchange of any two electrons.

Let $\xi_i = (\vec{r}_i, \sigma_i)$ denote the combined spatial and spin coordinates of the electron i. The spin coordinate σ_i labels the spin degree of freedom and takes two possible values. We introduce two orthonormal spin functions, $\alpha(\sigma_i)$ and $\beta(\sigma_i)$, which form a basis for the one-electron spin space and are usually associated with spin projections $+\frac{1}{2}$ and $-\frac{1}{2}$, respectively.

A spin-orbital $\phi_k(\xi_i)$ is then a one-electron function defined on the combined spatial and spin coordinates, and can be written as a product of a spatial orbital and a spin function,

$$\phi_k(\xi_i) = \phi_p(\vec{r}_i)\,\omega_s(\sigma_i), \tag{2.25}$$

where $\phi_p(\vec{r}_i)$ is a spatial orbital and $\omega_s(\sigma_i)$ is either $\alpha(\sigma_i)$ or $\beta(\sigma_i)$. In this notation the index k labels the pair (p, s). The properly antisymmetrized N_{elec}-electron wave function built from spin-orbitals $\phi_k(\xi)$ is

$$\begin{aligned}\Phi^0_{\text{elec}}(\xi_1, \ldots, \xi_{N_{\text{elec}}}) &= \frac{1}{\sqrt{N_{\text{elec}}!}} \sum_{\pi \in S_{N_{\text{elec}}}} \operatorname{sgn}(\pi) \prod_{j=1}^{N_{\text{elec}}} \phi_{k_{\pi(j)}}(\xi_j) \\ &= \left(\phi_{k_1} \wedge \cdots \wedge \phi_{k_{N_{\text{elec}}}}\right)(\xi_1, \ldots, \xi_{N_{\text{elec}}}),\end{aligned} \tag{2.26}$$

where $S_{N_{\text{elec}}}$ is the permutation group of N_{elec} objects and $x \wedge y = -\,y \wedge x$ denotes the antisymmetric (wedge) product. This antisymmetrized product is equivalently written as a Slater determinant,

$$\begin{aligned}&\Phi^0_{\text{elec}}(\xi_1, \ldots, \xi_{N_{\text{elec}}}) \\ &\qquad = \frac{1}{\sqrt{N_{\text{elec}}!}} \begin{vmatrix} \phi_{k_1}(\xi_1) & \phi_{k_2}(\xi_1) & \cdots & \phi_{k_{N_{\text{elec}}}}(\xi_1) \\ \phi_{k_1}(\xi_2) & \phi_{k_2}(\xi_2) & \cdots & \phi_{k_{N_{\text{elec}}}}(\xi_2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{k_1}(\xi_{N_{\text{elec}}}) & \phi_{k_2}(\xi_{N_{\text{elec}}}) & \cdots & \phi_{k_{N_{\text{elec}}}}(\xi_{N_{\text{elec}}}) \end{vmatrix}.\end{aligned} \tag{2.27}$$

The antisymmetry encoded in the Slater determinant has profound energetic consequences, excluding many configurations through the Pauli principle and introducing exchange effects. A more compact and flexible framework for handling antisymmetric many-electron states is provided by second quantization, which will be introduced in the next section.

2.1.3 Second Quantization

Before moving on to specific quantum chemical algorithms for solving the electronic structure problem, it is useful to introduce the formal language in which they are most naturally expressed: the framework of second quantization. This formalism provides a compact algebraic notation for constructing and manipulating many-electron states, and it also underlies powerful diagrammatic techniques, which lie beyond the scope of this preliminary discussion.

The concept of second quantization was developed in the 1920s–1930s in the context of quantum field theory. The key idea is to shift from an explicit

description of the many-electron wave function in coordinate or momentum space to a description in terms of occupation numbers of one-particle states. For fermions, this reformulation is particularly convenient, since the Pauli principle and antisymmetry are automatically enforced by the operator algebra.

Given a set of orthonormal spin-orbitals $\phi_k(\vec{r},\sigma)$, we can represent a many-electron state by a vector of occupation numbers,

$$|\Psi\rangle = |n_{1\alpha}\, n_{1\beta}\, \ldots\, n_{k\alpha}\, n_{k\beta}\, \ldots\rangle \equiv |\vec{n}\rangle, \tag{2.28}$$

where $n_{k\alpha} \in \{0,1\}$ is the occupation number of the spin-orbital $\phi_{k\alpha}(\vec{r}) = \phi_k(\vec{r},\frac{1}{2})$, and similarly for the β spin. The set of all such states forms the Fock space. Basis vectors with at least one differing occupation number are orthogonal:

$$\langle \vec{m}|\vec{n}\rangle = \prod_k \delta_{m_k n_k}. \tag{2.29}$$

In particular, the state with all occupations zero is the vacuum state $|0\rangle$.

We now introduce fermionic creation and annihilation operators, $\hat{a}_i^\dagger$ and $\hat{a}_i$, which add or remove an electron in spin-orbital i, respectively. Their action on an occupation number state can be defined as

$$\hat{a}_{k\alpha}^\dagger|\ldots n_{k\alpha}\ldots\rangle = (-1)^{s_k}(1-n_{k\alpha})|\ldots(n_{k\alpha}+1)\ldots\rangle, \tag{2.30}$$

$$\hat{a}_{k\alpha}|\ldots n_{k\alpha}\ldots\rangle = (-1)^{s_k} n_{k\alpha}|\ldots(n_{k\alpha}-1)\ldots\rangle, \tag{2.31}$$

where s_k is the number of occupied spin-orbitals to the left of $n_{k\alpha}$ in the chosen ordering. Attempts to create an electron in an already occupied spin-orbital or to annihilate it from an empty one yield zero, as enforced by the factors $(1-n_{k\alpha})$ and $n_{k\alpha}$.

The creation and annihilation operators obey canonical anticommutation relations

$$\{\hat{a}_p, \hat{a}_q^\dagger\} = \hat{a}_p\hat{a}_q^\dagger + \hat{a}_q^\dagger\hat{a}_p = \delta_{pq}, \tag{2.32}$$

$$\{\hat{a}_p, \hat{a}_q\} = \{\hat{a}_p^\dagger, \hat{a}_q^\dagger\} = 0. \tag{2.33}$$

These algebraic rules encode fermionic antisymmetry, so one no longer needs to construct explicitly antisymmetrized wave functions.

In this formalism, the reference state can be chosen as the vacuum $|0\rangle$. An N-electron Slater determinant corresponding to occupied spin-orbitals $k_1,\ldots,k_N$ is then written simply as

$$|k_1 k_2 \ldots k_N\rangle = \hat{a}_{k_1}^\dagger \hat{a}_{k_2}^\dagger \ldots \hat{a}_{k_N}^\dagger|0\rangle. \tag{2.34}$$

Thus, many-electron wave functions become compact operator expressions.

To obtain second-quantized forms of operators, it is useful to introduce field operators,

$$\hat{\Phi}^\dagger(\vec{r},\sigma) = \sum_k \phi_{k\sigma}^*(\vec{r})\,\hat{a}_{k\sigma}^\dagger, \tag{2.35}$$

$$\hat{\Phi}(\vec{r},\sigma) = \sum_k \phi_{k\sigma}(\vec{r})\,\hat{a}_{k\sigma}, \tag{2.36}$$

which create and annihilate a fermion at position $\vec{r}$ with spin σ, respectively. These satisfy continuous anticommutation relations consistent with Eq. (2.32).

A general k-body operator $\hat{O}(\vec{r}_1,\sigma_1,\ldots,\vec{r}_k,\sigma_k)$ can be written in second-quantized form as

$$\hat{O} = \frac{1}{k!}\sum_{\sigma_1,\ldots,\sigma_k}\int d\vec{r}_1\cdots d\vec{r}_k\,\hat{\Phi}^\dagger(\vec{r}_1,\sigma_1)\cdots\hat{\Phi}^\dagger(\vec{r}_k,\sigma_k)\,\hat{O}(\vec{r}_1,\sigma_1,\ldots,\vec{r}_k,\sigma_k) \times\hat{\Phi}(\vec{r}_k,\sigma_k)\cdots\hat{\Phi}(\vec{r}_1,\sigma_1). \tag{2.37}$$

The one-body and two-body parts of the electronic Hamiltonian (Eq. (2.19)) acquire particularly simple forms in second quantization. A general one-electron operator $\hat{h}_i$ can be written without the index i, since it is the same for each electron:

$$\hat{H}^{(1)} = \sum_{\sigma}\int d\vec{r}\,\hat{\Phi}^\dagger(\vec{r},\sigma)\,\hat{h}(\vec{r})\,\hat{\Phi}(\vec{r},\sigma) = \sum_{rs\sigma} h_{rs}\,\hat{a}^\dagger_{r\sigma}\hat{a}_{s\sigma}, \tag{2.38}$$

with one-electron matrix elements

$$h_{rs} = \int d\vec{r}\,\phi^*_{r\sigma}(\vec{r})\,\hat{h}(\vec{r})\,\phi_{s\sigma}(\vec{r}), \tag{2.39}$$

which are independent of the spin labels.

Similarly, the electron–electron Coulomb interaction can be written as

$$\hat{H}^{(2)} = \frac{1}{2}\sum_{\sigma\tau}\int d\vec{r}_1 d\vec{r}_2\,\hat{\Phi}^\dagger(\vec{r}_1,\sigma)\hat{\Phi}^\dagger(\vec{r}_2,\tau)\frac{1}{|\vec{r}_1-\vec{r}_2|}\hat{\Phi}(\vec{r}_2,\tau)\hat{\Phi}(\vec{r}_1,\sigma). \tag{2.40}$$

In an orbital basis, this becomes

$$\hat{H}^{(2)} = \frac{1}{2}\sum_{rstu}\sum_{\sigma\tau}\langle r\sigma\,s\tau|t\sigma\,u\tau\rangle\,\hat{a}^\dagger_{r\sigma}\hat{a}^\dagger_{s\tau}\hat{a}_{u\tau}\hat{a}_{t\sigma}, \tag{2.41}$$

where the two-electron integrals are defined as

$$\langle r\sigma\,s\tau|t\sigma\,u\tau\rangle = \int d\vec{r}_1 d\vec{r}_2\,\phi^*_{r\sigma}(\vec{r}_1)\phi^*_{s\tau}(\vec{r}_2)\frac{1}{|\vec{r}_1-\vec{r}_2|}\phi_{t\sigma}(\vec{r}_1)\phi_{u\tau}(\vec{r}_2). \tag{2.42}$$

Absorbing the spin labels into the orbital indices (so that r,s,t,u denote spin-orbitals), the one- and two-body parts of the Hamiltonian can be written in the compact form

$$\hat{H}^{(1)} = \sum_{rs} h_{rs}\,\hat{a}^\dagger_r\hat{a}_s, \tag{2.43}$$

$$\hat{H}^{(2)} = \frac{1}{2}\sum_{rstu}\langle rs|tu\rangle\,\hat{a}^\dagger_r\hat{a}^\dagger_s\hat{a}_u\hat{a}_t. \tag{2.44}$$

A major advantage of second quantization is that strings of fermionic creation and annihilation operators can be simplified systematically using Wick's

theorem [58]. To formulate it, we first define normal ordering, denoted by $\{\cdots\}$. Normal ordering is the rearrangement in which all creation operators appear to the left of all annihilation operators, with appropriate sign changes due to the anticommutation relations in Eq. (2.32). For example,

$$\{\hat{a}_p^\dagger \hat{a}_q\} = \hat{a}_p^\dagger \hat{a}_q, \qquad \{\hat{a}_q \hat{a}_p^\dagger\} = -\hat{a}_p^\dagger \hat{a}_q. \tag{2.45}$$

A *contraction* between two operators is defined as the difference between their product and its normal-ordered form:

$$\overset{\sqcap}{\hat{x}\hat{y}} \;\equiv\; \hat{x}\hat{y} - \{\hat{x}\hat{y}\}. \tag{2.46}$$

Wick's theorem states that any product of creation and annihilation operators can be expressed as the sum of its normal-ordered form and all possible normal-ordered terms with contractions. Schematically,

$$\hat{A} = \{\hat{A}\} + \sum_{\text{all contractions}} \{\text{Contractions}(\hat{A})\}. \tag{2.47}$$

The Wick's theorem can also be formulated for two normal-ordered strings A and B:

$$\hat{A}\hat{B} = \{\hat{A}\hat{B}\} + \{\overset{\sqcap}{\hat{A}\hat{B}}\}, \tag{2.48}$$

where $\{\overset{\sqcap}{\hat{A}\hat{B}}\}$ denotes all possible contractions between operators of one string and another.

The one- and two-body Hamiltonian terms in Eqs. (2.43) are already normal ordered if the vacuum $|0\rangle$ is taken as the reference state. In quantum chemistry, however, it is often more convenient to choose a Slater determinant as a reference. Normal ordering with respect to such a reference modifies the notion of creation and annihilation, and Wick's theorem then generates terms relative to this new vacuum.

2.1.4 Mean-Field Foundations

One of the earliest practical methods of quantum chemistry is the Hartree–Fock (HF) method. Douglas Hartree (1928) introduced the idea of a self-consistent field (SCF) in which each electron moves in the average Coulomb field of all others [59], while Vladimir Fock (1930) reformulated the approach in terms of antisymmetric Slater determinants [60], ensuring compliance with the Pauli exclusion principle. Originally developed in the language of first quantization, HF theory is now most naturally expressed in second quantization, where its mean-field character emerges in a transparent way.

We introduce a Slater determinant $|\Phi_0\rangle$ as a reference state. In this state a subset of spin-orbitals is occupied (hole or core orbitals, indexed by i, j, k, l), and the remaining spin-orbitals are unoccupied (particle or virtual orbitals, indexed by a, b, c, d). General spin-orbitals will be denoted by r, s, t, u. The determinant

$|\Phi_0\rangle$ plays the role of a "Fermi vacuum": with respect to this reference, there are no holes and no particles. Acting with an annihilation operator on an occupied spin-orbital creates a hole, while acting with a creation operator on an unoccupied spin-orbital creates a particle excitation.

Normal ordering with respect to $|\Phi_0\rangle$, denoted by $\{\cdots\}$, is defined such that the expectation value of any normal-ordered operator vanishes in the reference state,

$$\langle\Phi_0|\{\hat{A}\}|\Phi_0\rangle = 0. \tag{2.49}$$

Applying Wick's theorem with $|\Phi_0\rangle$ as the reference to the one- and two-body parts of the Hamiltonian (Eqs. (2.43)) yields

$$\hat{H}^{(1)} = \sum_{rs} h_{rs}\, \hat{a}_r^\dagger \hat{a}_s = \sum_i h_{ii} + \sum_{rs} h_{rs}\, \{\hat{a}_r^\dagger \hat{a}_s\}, \tag{2.50}$$

$$\begin{aligned}\hat{H}^{(2)} &= \frac{1}{2}\sum_{rstu} \langle rs|tu\rangle\, \hat{a}_r^\dagger \hat{a}_s^\dagger \hat{a}_u \hat{a}_t \\ &= \frac{1}{2}\sum_{ij}\big(\langle ij|ij\rangle - \langle ij|ji\rangle\big) + \sum_{rsi}\big(\langle ri|si\rangle - \langle ri|is\rangle\big)\{\hat{a}_r^\dagger \hat{a}_s\} \\ &\quad + \frac{1}{2}\sum_{rstu}\langle rs|tu\rangle\, \{\hat{a}_r^\dagger \hat{a}_s^\dagger \hat{a}_t \hat{a}_u\}.\end{aligned} \tag{2.51}$$

The first terms in Eq. (2.50) and Eq. (2.51) together give the expectation value of the Hamiltonian in the reference state, which is the HF energy:

$$E_{\mathrm{HF}} = \langle\Phi_0|\hat{H}|\Phi_0\rangle = \sum_i h_{ii} + \frac{1}{2}\sum_{ij}\big(\langle ij|ij\rangle - \langle ij|ji\rangle\big). \tag{2.52}$$

The coefficients of the normal-ordered one-body operators define the matrix elements of the Fock operator,

$$f_{rs} = h_{rs} + \sum_i \big(\langle ri|si\rangle - \langle ri|is\rangle\big). \tag{2.53}$$

The term $\langle ri|si\rangle$ is the Coulomb contribution (often denoted J_i in the first quantization) and describes the classical electrostatic interaction with the charge distribution of electron i, while $\langle ri|is\rangle$ is the exchange contribution (denoted K_i in the first quantization), arising purely from antisymmetry.

As demonstrated in Eq. (2.22), molecular orbitals are expressed as LCAO. We can perform unitary rotations of the orbital basis. Since creation and annihilation operators are tied to particular spin-orbitals, a rotation in orbital space induces a corresponding unitary transformation in the space of fermionic operators:

$$\hat{\tilde{a}}_i^\dagger = \sum_j U_{ij}\hat{a}_j^\dagger, \tag{2.54}$$

$$\hat{\tilde{a}}_i = \sum_j U_{ji}^\dagger \hat{a}_j, \tag{2.55}$$

where U is a unitary matrix. This rotation can be generated by a one-body anti-Hermitian operator

$$\hat{\Lambda} = \sum_{rs} \Lambda_{rs}\, \hat{a}_r^\dagger \hat{a}_s, \qquad U = e^{i\Lambda}, \tag{2.56}$$

so that the corresponding rotated Slater determinant $|\tilde{\Phi}_0\rangle$ is given by the Thouless theorem [61]:

$$|\tilde{\Phi}_0\rangle = e^{i\hat{\Lambda}}|\Phi_0\rangle. \tag{2.57}$$

The HF method seeks the set of molecular orbitals (equivalently, the Slater determinant $|\tilde{\Phi}_0\rangle$) that minimizes the energy functional

$$E(\hat{\Lambda}) = \langle\tilde{\Phi}_0|\hat{H}|\tilde{\Phi}_0\rangle. \tag{2.58}$$

Requiring stationarity with respect to orbital rotations, $\delta E = 0$, leads to the conditions

$$f_{ia} = f_{ai} = 0, \tag{2.59}$$

meaning that all off-diagonal Fock matrix elements between occupied (hole) and virtual (particle) orbitals vanish. These are the HF equations in matrix form. In practice, they are solved iteratively in a SCF procedure. A further analysis of the Hessian of $E(\hat{\Lambda})$ shows that the converged HF solution corresponds to the local minimum.

Even after imposing Eq. (2.59), there remains some freedom in the choice of orbitals. If the orbitals diagonalize the Fock matrix, they are called *canonical* orbitals. If instead they diagonalize the one-electron density operator, they are known as *natural* orbitals. Different choices of orbitals can be advantageous for particular correlation methods.

Brillouin's theorem follows directly from the condition in Eq. (2.59) [62]. Determinants with a single excitation from the HF state have zero first-order coupling to the reference under the exact Hamiltonian. Equivalently,

$$\langle\Phi_i^a|\hat{H}|\Phi_0\rangle = 0, \tag{2.60}$$

where $|\Phi_i^a\rangle$ denotes a singly excited determinant from orbital i to orbital a. This property underpins many post-HF correlation methods, where the correlation energy arises first from double and higher excitations.

HF theory typically recovers the dominant fraction of the total electronic energy, providing a qualitatively correct mean-field picture of electron interactions. However, the remaining correlation energy, although numerically much smaller, is chemically essential. It is responsible for effects such as dispersion interactions, accurate bond breaking, and reliable barrier heights. Accurately capturing this correlation energy is the central aim of post-HF electronic structure methods, which are built systematically upon the HF reference determinant.

2.1.5 Capturing Correlation

Although HF provides a self-consistent mean-field description of electrons, it neglects most of the correlation effects between them. Electronic correlation is usually defined as the difference between the exact non-relativistic BO energy and the HF energy. It reflects the fact that the exact electronic wave function cannot, in general, be represented by a single Slater determinant. Instead, it must be written as a linear combination of determinants with different distributions of electrons among orbitals:

$$|\Phi\rangle = c_0|\Phi_0\rangle + c_1|\Phi_1\rangle + c_2|\Phi_2\rangle + \cdots , \tag{2.61}$$

where $|\Phi_0\rangle$ is a reference determinant (often the HF solution) and the coefficients c_i describe the weight of the corresponding excited configurations $|\Phi_i\rangle$.

Two qualitatively different types of correlation are usually distinguished. *Static* (or *strong*) correlation occurs when several determinants contribute with comparable, relatively large coefficients c_i. This situation arises when there are nearly degenerate electronic configurations, so that no single determinant provides an adequate qualitative description of the system. Static correlation is essential, for example, in bond dissociation, near transition states, and in systems with open shells or near-degenerate frontier orbitals.

Dynamic (or *weak*) correlation, on the other hand, arises from the collective effect of many determinants with small coefficients c_i. These excitations refine the mean-field picture by accounting for the instantaneous electron–electron repulsion, the tendency of electrons to avoid each other in space. Dynamic correlation is typically well described by single-reference methods such as Møller–Plesset perturbation theory (see Sec. 2.1.5.1) and coupled cluster expansions (see Sec. 2.1.5.3), and is crucial for achieving quantitative accuracy even when a single determinant provides a qualitatively correct starting point.

2.1.5.1 Perturbation Theory

The idea of expanding around the HF solution using perturbation theory dates back to Møller and Plesset (1934) [63]. To apply perturbation theory, the electronic Hamiltonian is partitioned into a zeroth-order part, which is solvable (preferably exactly), and a residual interaction treated as a perturbation. In Møller–Plesset (MP) perturbation theory, the Hamiltonian is written as

$$\hat{H}_{\mathrm{MP}}(\lambda) = \hat{H}^{(0)} + \lambda\hat{V} = \hat{f} + \lambda(\hat{H} - \hat{f}), \tag{2.62}$$

where $\hat{f}$ is the normal-ordered Fock operator and λ is a formal perturbation parameter that is set to $\lambda = 1$ at the end of the derivation. The residual $(\hat{H} - \hat{f})$ acts as the perturbation.

The exact state vector and energy are expanded as power series in λ:

$$|\Phi_{\mathrm{MP}}(\lambda)\rangle = |\Phi^{(0)}\rangle + \lambda|\Phi^{(1)}\rangle + \lambda^2|\Phi^{(2)}\rangle + \cdots , \tag{2.63}$$

$$E_{\mathrm{MP}}(\lambda) = E^{(0)} + \lambda E^{(1)} + \lambda^2 E^{(2)} + \cdots , \tag{2.64}$$

with the zeroth-order state taken as the HF solution, $|\Phi^{(0)}\rangle = |\Phi^{\mathrm{HF}}\rangle$. Substituting these expansions into the Schrödinger equation

$$\hat{H}_{\mathrm{MP}}(\lambda)|\Phi_{\mathrm{MP}}(\lambda)\rangle = E_{\mathrm{MP}}(\lambda)|\Phi_{\mathrm{MP}}(\lambda)\rangle \tag{2.65}$$

and collecting terms of equal order in λ yields the standard Rayleigh–Schrödinger perturbation expressions. For the MP partitioning, $E^{(0)} = E_{\mathrm{HF}}$, so that the correlation energy is given by

$$E_{\mathrm{MP}}^{\mathrm{corr}} = E(\lambda{=}1) - E_{\mathrm{HF}} = E^{(2)} + E^{(3)} + \cdots . \tag{2.66}$$

There is no first-order correction to the HF energy because of Brillouin's theorem, which ensures that single excitations do not couple to the HF reference at first order.

The second-order MP (MP2) correlation energy can be written compactly as

$$E^{(2)} = \frac{1}{4}\sum_{ijab}\frac{|\langle ij||ab\rangle|^2}{\varepsilon_i + \varepsilon_j - \varepsilon_a - \varepsilon_b}, \tag{2.67}$$

where i, j denote occupied orbitals, a, b virtual orbitals, ε_p are the HF orbital energies, and $\langle ij||ab\rangle = \langle ij|ab\rangle - \langle ij|ba\rangle$ are antisymmetrized two-electron integrals. In its canonical implementation, MP2 scales as $O(N^5)$ with the system size N, but is relatively inexpensive in practice and often yields a substantial improvement over HF for calculation of thermochemical properties.

Rayleigh–Schrödinger perturbation series are, in most cases, only asymptotic: they may approximate the exact result well up to some finite order, but higher-order terms often grow and the series ultimately diverges. In some systems, divergence or pathological behavior can already appear at low orders. These problems become particularly severe when the HF reference is qualitatively poor, for example, in systems with small orbital energy gaps or strong near-degeneracies, such as stretched bonds, transition-metal complexes, or excited and near-degenerate states where static correlation is important.

To treat such situations, multireference perturbation theories (such as CASPT2 and NEVPT2) generalize the expansion to multireference states. This improves the description of electronic correlation at the cost of increased computational complexity. These approaches will be discussed in more detail later in Sec. 2.1.5.5.

2.1.5.2 Configuration Interaction

Configuration interaction (CI) is one of the earliest systematic post-HF methods. The basic idea is to expand the electronic wave function as a linear combination of Slater determinants built from a fixed set of spin-orbitals,

$$|\Psi_{\mathrm{CI}}\rangle = c_0|\Phi_0\rangle + \sum_{ia} c_i^a\,|\Phi_i^a\rangle + \frac{1}{4}\sum_{ijab} c_{ij}^{ab}\,|\Phi_{ij}^{ab}\rangle + \cdots , \tag{2.68}$$

where $|\Phi_0\rangle$ is the HF reference determinant and $|\Phi_i^a\rangle$, $|\Phi_{ij}^{ab}\rangle$, etc. are singly, doubly, and higher excited determinants formed by exciting electrons from occupied orbitals $i, j, \ldots$ to virtual orbitals $a, b, \ldots$.

The CI coefficients $c_{ij\ldots}^{ab\cdots}$ can be obtained by projecting the Schrödinger equation onto each determinant in the expansion,

$$\langle \Phi_{ij\ldots}^{ab\cdots} | \hat{H} | \Psi_{\mathrm{CI}} \rangle = E_{\mathrm{CI}}\, c_{ij\ldots}^{ab\cdots}, \tag{2.69}$$

$$\langle \Phi_0 | \hat{H} | \Psi_{\mathrm{CI}} \rangle = E_{\mathrm{CI}}\, c_0, \tag{2.70}$$

which can be written in matrix form as

$$\mathbf{H}\,\vec{c} = E_{\mathrm{CI}}\,\vec{c}, \tag{2.71}$$

where $\mathbf{H}$ is the electronic Hamiltonian represented in the basis of Slater determinants, and $\vec{c}$ is the vector of CI coefficients. The diagonalization of $\mathbf{H}$ yields approximate energies and wave functions; the lowest eigenvalue corresponds to the CI approximation to the ground state.

If all possible determinants for the chosen basis set are included, the method is called full configuration interaction (FCI). In that case, CI provides the exact solution of the electronic Schrödinger equation within the given basis set. However, the dimension of the FCI space grows combinatorially with system size (like $\binom{M}{N}$ for N electrons in M spin-orbitals), and FCI is feasible only for very small molecules and modest basis sets.

In practice, the CI expansion is truncated. The truncation level defines the specific CI method. For example, including only single and double excitations leads to CISD, adding triple excitations gives CISDT, and so on. CISD is variational and computationally tractable (scaling as $O(N^6)$ with system size N) and often yields a significant improvement over HF. However, truncated CI suffers from two fundamental deficiencies:

- It is *not size consistent*: for two non-interacting subsystems A and B, the energy of the combined system is not exactly equal to $E(A) + E(B)$ when each is treated at the same truncated CI level.

- It is *not size extensive*: the correlation energy does not scale linearly with the number of electrons or fragments, so errors grow unphysically with system size.

These problems become particularly severe for larger systems and are a major reason why truncated CI has largely been superseded by coupled cluster methods in high-accuracy quantum chemistry.

Despite these limitations, truncated CI remains important as a conceptually simple reference method and as a convenient testing ground for new ideas. FCI continues to play a central role as the formal benchmark against which approximate electronic structure methods are assessed in small systems.

2.1.5.3 Coupled Cluster Theory

Coupled cluster (CC) theory, originally developed in nuclear physics [64] and later adopted in quantum chemistry [65], introduces the exponential wave function ansatz

$$|\Psi_{\mathrm{CC}}\rangle = e^{\hat{T}}|\Phi_0\rangle, \qquad \hat{T} = \sum_{k=1}^{N_{\mathrm{CC}}} \hat{T}_k, \tag{2.72}$$

where $|\Phi_0\rangle$ is a reference Slater determinant (typically HF) and $\hat{T}_k$ is the cluster operator that generates all k-fold excitations from $|\Phi_0\rangle$. The integer N_{CC} denotes the highest excitation rank included in the cluster operator. If $N_{\mathrm{CC}} = N_{\mathrm{elec}}$, where N_{elec} is the number of correlated electrons, the ansatz is formally exact within the given basis set and this limit is sometimes referred to as full CC. The exponential ansatz has two key advantages: it ensures the size consistency and the size extensivity.

In second quantization, for example, the double excitation cluster operator is

$$\hat{T}_2 = \frac{1}{4}\sum_{ijab} t_{ij}^{ab}\hat{a}_a^\dagger\hat{a}_b^\dagger\hat{a}_j\hat{a}_i, \tag{2.73}$$

with amplitudes t_{ij}^{ab} for exciting electrons from occupied orbitals i, j to virtual orbitals a, b. The truncation level N_{CC} defines the CC method: if $\hat{T} = \hat{T}_1 + \hat{T}_2$ the method is called CC with singles and doubles (CCSD), while $\hat{T} = \hat{T}_1 + \hat{T}_2 + \hat{T}_3$ defines CC with singles, doubles and triples (CCSDT), etc.

The CC amplitudes are determined from a projected Schrödinger equation, similarly to Eq. (2.69):

$$\langle\Phi_{ij\ldots}^{ab\cdots}|\hat{H}|\Psi_{\mathrm{CC}}\rangle = 0, \tag{2.74}$$

$$\langle\Phi_0|\hat{H}|\Psi_{\mathrm{CC}}\rangle = E_{\mathrm{CC}}, \tag{2.75}$$

where $|\Phi_{ij\ldots}^{ab\cdots}\rangle$ runs over all excited determinants included up to the truncation level N_{CC}. These equations are nonlinear in the cluster amplitudes $t_{ij\ldots}^{ab\cdots}$ and are usually solved iteratively. The nonlinearity implies that multiple solutions may exist, and convergence is not always guaranteed.

To reduce the computational cost of higher-level CC methods, various approximate schemes have been developed. For example, CC2 and CC3 are approximate CCSD and CCSDT models, respectively, derived by combining CC theory with PT [66, 67]. Linear response CC [68] and equation-of-motion CC [69] methods provide access to excited states and response properties based on the ground state.

It is well established that triple and higher excitations are important for quantitative accuracy in many molecular properties. Although CCSD is relatively tractable (formal scaling $O(N^6)$ with system size N), CCSDT ($O(N^8)$) is already prohibitively expensive. A widely used compromise is to include the effect of triple excitations perturbatively on top of CCSD, leading to the CCSD(T) method [70]. CCSD(T) often achieves high accuracy for a broad range of

molecules near equilibrium geometries and has become known as the "gold standard" of quantum chemistry.

In general, CC methods provide high accuracy for systems dominated by dynamic correlation and close to a single-reference description, particularly near equilibrium geometries. However, they struggle in the presence of static correlation, such as in bond-breaking regions or in many transition-metal complexes, where the single-determinant reference becomes qualitatively inadequate. Moreover, the steep polynomial scaling ($O(N^6)$ for CCSD, $O(N^7)$ for CCSD(T) and worse for higher excitations) limits their direct application to very large systems.

An alternative formulation is variational CC (VCC) [71], in which the cluster amplitudes are determined by minimizing the energy functional

$$E_{\mathrm{VCC}}(\hat{T}) = \frac{\langle \Phi_0 | e^{\hat{T}^\dagger} \hat{H} e^{\hat{T}} | \Phi_0 \rangle}{\langle \Phi_0 | e^{\hat{T}^\dagger} e^{\hat{T}} | \Phi_0 \rangle}. \tag{2.76}$$

This approach restores a variational upper bound to the energy, but the evaluation of $E_{\mathrm{VCC}}(\hat{T})$ for general $\hat{T}$ involves an exponentially large number of terms. The variational and unitary CC ansatz have been proposed as promising candidates for quantum algorithms, where the exponential complexity may be mitigated [72, 73]. These aspects will be discussed in the section on quantum computation.

2.1.5.4 Multiconfigurational Self-Consistent Field

In systems with strong correlation, several electronic configurations contribute with comparable weights, and any single-reference description is qualitatively inadequate. To address this, the multireference counterpart of HF was developed, the multiconfigurational self-consistent field (MCSCF) method.

In MCSCF the spin-orbital space is partitioned into three subspaces: inactive (or occupied), active, and external (or virtual). Inactive orbitals are doubly occupied in all determinants of the expansion, active orbitals may be variably occupied (0, 1, or 2 electrons) depending on the configuration, and external orbitals are unoccupied in all determinants. The most common choice of configuration space is the complete active space (CAS), in which all possible distributions of N_{act} active electrons among M_{act} active orbitals are included. The corresponding method is called complete active space self-consistent field (CASSCF) [74].

Let $\{|m\rangle\}_{m=0}^{K-1}$ denote the orthonormal set of Slater determinants generated by all possible distributions of the active electrons among the active orbitals, where $|0\rangle$ is a chosen reference determinant. Each determinant $|m\rangle$ can be obtained from $|0\rangle$ by an excitation operator,

$$\hat{s}_m^\dagger = |m\rangle\langle 0|. \tag{2.77}$$

A general unitary transformation within the configuration space can then be

written as

$$e^{i\hat{S}} = \exp\left(i \sum_{m\neq 0}^{K-1} (S^*_{0m}\hat{s}^\dagger_m + S_{0m}\hat{s}_m) \right), \tag{2.78}$$

where the matrix elements S_{0m} parametrize rotations among the CAS determinants.

As in the HF case, we also allow unitary rotations of the spin-orbital basis, generated by a one-body anti-Hermitian operator (see Eq. (2.56)) The general CASSCF ground state can then be written as

$$|\tilde{0}\rangle = e^{i\hat{\Lambda}}\, e^{i\hat{S}}\, |0\rangle. \tag{2.79}$$

It can be found by a simultaneous optimization in both the space of orbitals and the space of CI coefficients within the chosen configuration space. The CASSCF energy functional is

$$E(\hat{\Lambda}, \hat{S}) = \langle\tilde{0}|\hat{H}|\tilde{0}\rangle. \tag{2.80}$$

Requiring stationarity of the energy with respect to both orbital and configuration rotations, $\delta E(\hat{\Lambda}, \hat{S}) = 0$, leads to two coupled sets of equations, which can be expressed compactly as

$$\langle 0|[\hat{H}, \hat{\Lambda}]|0\rangle = 0, \tag{2.81}$$

$$\langle 0|[\hat{H}, \hat{S}]|0\rangle = 0. \tag{2.82}$$

The first condition is equivalent to the HF case (Eq. (2.59)). The second condition is equivalent to $\langle 0|\hat{H}|m\rangle = 0$ for all CASSCF states $m \neq 0$. In practice, these coupled equations are solved iteratively. An analysis of the Hessian of $E(\hat{\Lambda}, \hat{S})$ is required to confirm whether the resulting solution corresponds to a ground state or an excited state.

In many applications, one is interested not only in a single electronic state but in a set of low-lying states of the same symmetry. For this purpose, state-averaged CASSCF (SA-CASSCF) is used. In SA-CASSCF the functional in Eq. (2.80) is replaced by a state-averaged energy,

$$E^{\mathrm{SA}}(\hat{\Lambda}, \hat{S}) = \sum_{m=0}^{N_{\mathrm{SA}}-1} w_m \, \langle\tilde{m}|\hat{H}|\tilde{m}\rangle, \tag{2.83}$$

where $|\tilde{m}\rangle = e^{i\hat{\Lambda}}e^{i\hat{S}}|m\rangle$ are the CASSCF excited states, w_m are their weights, and N_{SA} is a number of averaged states. In the first stationarity condition (see Eq. (2.81)), the Fock operator is generalized to the weighted Fock operator for several states. The second stationarity condition (see Eq. (2.81)) then transforms to $\langle n|\hat{H}|m\rangle = 0$ for all CASSCF states $n \neq m$ from state averaging, ensuring a balanced description of several states that share the same set of optimized orbitals.

The same formalism allows one to define the complete active space configuration interaction (CASCI) method: it is exactly $\langle n|\hat{H}|m\rangle = 0$ requirement

without orbital optimization. CASCI therefore offers a compromise between flexibility and computational cost, but lacks the additional variational freedom of CASSCF.

CASSCF provides an accurate description of static correlation by allowing multiple determinants to contribute significantly within the active space. However, the number of CAS determinants grows combinatorially with the size of the active space, limiting the practical choice of N_{act} and M_{act}. Moreover, dynamic correlation outside the active space is not captured and must be added using post-CASSCF methods, such as multireference perturbation theories or multireference configuration interaction.

2.1.5.5 Multireference Perturbation, Configuration Interaction, and Coupled Cluster

Although CASSCF provides an accurate treatment of static correlation within a chosen active space, it lacks most of the dynamic correlation required for quantitative calculations. Many multireference methods can be formulated in a common language using Bloch's effective Hamiltonian formalism [75].

Let the state-averaged CASSCF reference define a N_{m}-dimensional model subspace $\mathcal{P}$ spanned by orthonormal states $\{|\tilde{m}\rangle\}_{m=0}^{N_{\text{m}}-1}$. The projector onto this model space is

$$\hat{P} = \sum_{m=0}^{N_{\text{m}}-1} |\tilde{m}\rangle\langle\tilde{m}| = \sum_{m=0}^{N_{\text{m}}-1} \hat{P}_m, \tag{2.84}$$

with the complementary projector $\hat{Q} = 1 - \hat{P}$. The goal is to construct a wave operator $\hat{\Omega}$ that maps each vector of the model space to the corresponding exact eigenstate of the full electronic Hamiltonian,

$$|\Psi_m\rangle = \hat{\Omega}\,|\tilde{m}\rangle. \tag{2.85}$$

In Bloch's formulation, one imposes the intermediate normalization condition

$$\hat{P}\hat{\Omega}\hat{P} = \hat{P}, \tag{2.86}$$

which ensures that the projection of $|\Psi_m\rangle$ onto $\mathcal{P}$ coincides with $|\tilde{m}\rangle$. The corresponding effective Hamiltonian acting within the model space is

$$\hat{H}^{\text{eff}} = \hat{P}\hat{H}\hat{\Omega}\hat{P}, \tag{2.87}$$

and satisfies the Bloch equation

$$\hat{H}\,\hat{\Omega}\hat{P} = \hat{\Omega}\hat{P}\,\hat{H}^{\text{eff}}\hat{P}. \tag{2.88}$$

Its eigenvalues reproduce the targeted (typically low-lying) eigenvalues of $\hat{H}$, while the eigenvectors are the projections of the exact states onto $\mathcal{P}$.

It is convenient to separate the "external" part of the wave operator via the *correlation operator*

$$\hat{X} \equiv \hat{Q}\hat{\Omega}\hat{P}, \qquad \Rightarrow \qquad \hat{\Omega}\hat{P} = \hat{P} + \hat{X}, \tag{2.89}$$

so that $\hat{P}\hat{X} = 0$ and $\hat{Q}\hat{X} = \hat{X}$ by construction. Different multireference methods correspond to different parameterizations/approximations of $\hat{\Omega}$ (equivalently, of $\hat{X}$).

- **Multireference perturbation theory (MRPT).** One introduces an explicit Hamiltonian partition

$$\hat{H} = \hat{H}_0 + \hat{V}, \tag{2.90}$$

typically choosing $\hat{H}_0$ to be block-diagonal with respect to $\hat{P}$ and $\hat{Q}$. The correlation operator is expanded perturbatively,

$$\hat{X} = \hat{X}^{(1)} + \hat{X}^{(2)} + \cdots, \tag{2.91}$$

and can be calculated using recursion

$$\begin{aligned} \hat{X}^{(1)} &= \sum_{m=0}^{N_\mathrm{m}-1} \frac{\hat{Q}}{E_m^{(0)} - \hat{H}_0} \hat{V}\hat{P}_m, \\ \hat{X}^{(k)} &= \sum_{m=0}^{N_\mathrm{m}-1} \frac{\hat{Q}}{E_m^{(0)} - \hat{H}_0} \left(\hat{V}\hat{X}^{(k-1)} - \sum_{l=1}^{k-1} \hat{X}^{(k-l)} \hat{V} \hat{X}^{(l-1)} \right) \hat{P}_m, \end{aligned} \tag{2.92}$$

with $E_m^{(0)} = \langle \tilde{m} | \hat{H}_0 | \tilde{m} \rangle$.

The expression for the effective Hamiltonian is as follows:

$$\hat{H}^{\mathrm{eff}} = \hat{P}\hat{H}_0\hat{P} + \hat{P}\hat{V}\hat{X}. \tag{2.93}$$

This effective Hamiltonian can then be exactly diagonalized, giving the energies and state vectors. Complete active space perturbation theory (CASPT2) [76, 77], n-electron valence state perturbation theory (NEVPT2) [78, 79], multireference MP (MRMP2) [80], and (extended) multiconfigurational quasi-degenerate perturbation theory (X)MCQDPT2 [81, 82] mainly differ in their definitions of $\hat{H}_0$, in how the first-order interacting space is represented, and in how quasi-degeneracy is handled.

- **Multireference configuration interaction (MRCI).** MRCI is obtained by diagonalizing $\hat{H}$ in a truncated space $\mathcal{P} \oplus \mathcal{Q}_\mathrm{trunc}$, where $\mathcal{Q}_\mathrm{trunc}$ is typically spanned by single and double excitations out of $\mathcal{P}$ (MRCISD). Equivalently, using a wave operator $\hat{\Omega}$, one may write a linear parameterization

$$\hat{\Omega} \approx \hat{P} + \hat{X}, \qquad \hat{X} : \mathcal{P} \to \mathcal{Q}_\mathrm{trunc}. \tag{2.94}$$

Modern implementations employ internally contracted configurations, i.e. basis functions in $\mathcal{Q}_\mathrm{trunc}$ are formed as excitation operators acting on the reference space [83]. MRCI is variational within the chosen truncated space, but due to linear truncation it is not rigorously size-extensive/size-consistent; therefore a Davidson-type correction (MRCI+Q) is often used in practice [84].

- **Multireference coupled cluster (MRCC).** MRCC employs an exponential parameterization and can recover near size-extensivity. Two widely used wave operator forms correspond to different notions of universality:

  (i) State-universal MRCC (SUMRCC (Jeziorski–Monkhorst) [85] / Hilbert space MRCC):

  $$\hat{\Omega} = \sum_{m=0}^{N_m - 1} e^{\hat{T}^{(m)}} \hat{P}_m, \tag{2.95}$$

  i.e. a *sum* of exponentials with a separate cluster operator $\hat{T}^{(m)}$ for each basis state of the model space.

  (ii) Valence-universal MRCC (VUMRCC [86]/ Fock space MRCC):

  $$\hat{\Omega} = e^{\hat{S}}, \tag{2.96}$$

  where a single cluster operator $\hat{S}$ generates correlated states from the entire model space, producing a valence-universal description across multiple target states.

  In both cases the amplitudes are obtained from projected nonlinear equations derived from the Bloch equation, with practical truncations, such as singles and doubles, in $\hat{T}^{(m)}$ or $\hat{S}$.

Viewed this way, MRPT, MRCI, and MRCC are different approximations to the same object, the Bloch wave operator $\hat{\Omega}$ built on a CASSCF model space. As one moves from MRPT to MRCI to MRCC, the description of correlation generally becomes more accurate, but also more computationally demanding. In large or strongly correlated systems, the cost of these multireference approaches still remains a significant bottleneck, motivating the exploration of stochastic methods, such as Monte Carlo approaches.

2.1.6 Quantum Monte Carlo Approaches

Deterministic post-HF methods often exhibit steep scaling with system size, which limits their applicability to large molecules. An alternative route is to evaluate the expectation values of a quantum system stochastically. Quantum Monte Carlo (QMC) is a family of methods that uses randomness to estimate the parameters of a quantum system, potentially offering high accuracy with more favorable scaling. Although limited by statistical noise and, for fermions, by the sign problem [21], QMC has developed into a versatile set of techniques for quantum chemistry and condensed matter physics.

2.1.6.1 Variational Monte Carlo

In variational Monte Carlo (VMC), one chooses a parameterized trial wave function $\Psi_T(\vec{r}; \lambda)$, where $\vec{r}$ collects all electron coordinates and λ denotes a set

of variational parameters. The energy is evaluated by stochastic integration:

$$E(\lambda) = \frac{\int d\vec{r}\,|\Psi_T(\vec{r};\lambda)|^2\, E_{\mathrm{loc}}(\vec{r};\lambda)}{\int d\vec{r}\,|\Psi_T(\vec{r};\lambda)|^2} = \int d\vec{r}\,\rho(\vec{r};\lambda)\, E_{\mathrm{loc}}(\vec{r};\lambda), \tag{2.97}$$

with local energy $E_{\mathrm{loc}}(\vec{r};\lambda) = \hat{H}\Psi_T(\vec{r};\lambda)/\Psi_T(\vec{r};\lambda)$ and $\rho(\vec{r},\lambda)$ is the probability density for electronic coordinates. The functional can be estimated as an average of local energies taken at points $\vec{r}_k$ sampled from $\rho(\vec{r};\lambda)$ using the Metropolis–Hastings algorithm [87].

One common way to organize Metropolis sampling in VMC is as follows. Starting from an initial set of electronic coordinates $\vec{r}_{\mathrm{i}}$, a small Gaussian random move is proposed:

$$\vec{r}_{\mathrm{i}} \to \vec{r}_{\mathrm{f}} = \vec{r}_{\mathrm{i}} + \frac{1}{2}\vec{F}_q(\vec{r}_{\mathrm{i}}) + G_{\delta\tau}, \tag{2.98}$$

where $\vec{F}_q(\vec{r}_{\mathrm{i}}) = \nabla \ln |\Psi_T(\vec{r}_{\mathrm{i}})|^2$ is a Bohmian (quantum) force and $G_{\delta\tau}$ is a random number drawn from a Gaussian distribution with standard deviation $\delta\tau$. The proposal probability is then computed using the Green function of the Fokker–Planck equation in the short-time approximation:

$$P_{\mathrm{prop}}(\vec{r}_f|\vec{r}_i) = \frac{1}{(2\pi\delta\tau)^{3N_{\mathrm{elec}}/2}} \exp\left(- \frac{\vec{r}_{\mathrm{f}} - \vec{r}_{\mathrm{i}} - \frac{1}{2}\vec{F}_q(\vec{r}_{\mathrm{i}})\delta\tau}{2\delta\tau}\right), \tag{2.99}$$

where N_{elec} is the number of electrons.

The acceptance probability for the proposed coordinates is then given by

$$P_{\mathrm{acc}}(\vec{r}_{\mathrm{f}}|\vec{r}_{\mathrm{i}}) = \min\left\{1, \frac{P_{\mathrm{prop}}(\vec{r}_{\mathrm{i}}|\vec{r}_{\mathrm{f}})\rho(\vec{r}_{\mathrm{f}})}{P_{\mathrm{prop}}(\vec{r}_{\mathrm{f}}|\vec{r}_{\mathrm{i}})\rho(\vec{r}_{\mathrm{i}})}\right\}. \tag{2.100}$$

In VMC calculations, an ensemble of many random walkers is usually employed. Statistical averaging should also account for the autocorrelation between successive moves. In practice, this is commonly handled by performing averages within batches to reduce these effects.

The energy functional $E(\lambda)$ is then minimized variationally for a chosen parametrization of the trial wave function $\Psi_T(\vec{r};\lambda)$. The VMC provides an upper bound to the ground state energy. Its cost depends on the chosen ansatz but typically lies between $O(N_{\mathrm{elec}}^2)$ and $O(N_{\mathrm{elec}}^4)$ per Monte Carlo sweep, which is much less than deterministic post-HF methods. The accuracy depends strongly on the expressiveness of the trial ansatz, for example, Slater–Jastrow wave functions [88].

2.1.6.2 Diffusion Monte Carlo

Highly accurate solutions can be obtained with Diffusion Monte Carlo (DMC), which performs an imaginary-time evolution (ITE) toward the ground state while imposing the fixed-node constraint derived from the trial state $|\Psi_T\rangle$ to control the fermion sign problem. DMC improves upon VMC by stochastically projecting out the ground state from an initial trial wave function via ITE:

$$|\Psi(\tau)\rangle = e^{-\tau(\hat{H}-E_T)}|\Psi_T\rangle, \tag{2.101}$$

where τ is the imaginary time and $E_T = \langle\Psi_T|\hat{H}|\Psi_T\rangle$ is the trial state energy. As $\tau \to \infty$, provided that $|\Psi_T\rangle$ has a nonzero overlap with the ground state, the propagated wave function projects onto the exact ground state.

In the DMC framework, the projection is motivated by the fact that the imaginary-time Schrödinger equation can be written analogously to a classical diffusion equation:

$$\begin{aligned}\frac{\partial}{\partial\tau}|\Psi(\tau)\rangle &= \left(\frac{1}{2}\nabla^2 - (V - E_T)\right)|\Psi(\tau)\rangle \\ &= \underbrace{\frac{1}{2}\nabla^2|\Psi(\tau)\rangle}_{\text{diffusion (random walk)}} - \underbrace{(V - E_T)|\Psi(\tau)\rangle}_{\text{branching (birth/death)}} .\end{aligned} \tag{2.102}$$

The Laplacian term generates diffusion and is simulated by a random walk: over a small time step, walkers are displaced by Gaussian moves, analogous to Metropolis updates in VMC. The remaining term multiplies $|\Psi(\tau)\rangle$ by a local factor determined by $(V - E_T)$ and acts as a branching process: walkers in regions where V is low relative to E_T tend to replicate, whereas walkers in regions where V is high tend to die out. In practice, DMC combines these two ingredients into a branching random walk: each time step proposes a diffusion walk and then applies a stochastic replication/deletion step to represent the weight accumulated from the potential term.

However, without symmetry restrictions, the propagation converges to the bosonic ground state. This is the sign problem, which appears in many QMC techniques in different guises. To control the fermion sign problem, the fixed-node approximation is imposed: the nodal surface of $|\Psi_T\rangle$ is enforced during the simulation, and walkers that attempt to cross the nodal surface are removed. The resulting energy is variational and becomes exact when the nodal surface is exact. Therefore, the achievable accuracy is limited by the nodal structure of the trial wave function.

2.1.6.3 Auxiliary-Field QMC

Auxiliary-field QMC (AFQMC) is closely related in spirit to DMC in that it also computes ground state properties by ITE (see Eq. (2.101)) [89]. The central difference is in the representation of the propagation. AFQMC applies a Hubbard–Stratonovich transformation to rewrite the two-body interaction in terms of an integral over fluctuating auxiliary field [90], such that the many-body propagator is expressed as a stochastic average of one-body propagators. As a result, the random walk is performed in the manifold of nonorthogonal Slater determinants rather than in the position space of N-particle configurations; the antisymmetry of each walker is therefore maintained by construction.

A major practical challenge is that the auxiliary-field propagation is generally complex-valued, leading to an exponential growth of the noise associated with random phases. At long projection times the phases of the walkers decorrelate, a manifestation of the fermion sign problem in this framework. The phaseless

AFQMC method controls this instability by introducing a constraint based on the overlap of each walker with a chosen trial wave function, thereby guiding and stabilizing the random walk. A more detailed discussion of AFQMC is given in Chap. 5.

AFQMC has demonstrated high accuracy for both molecules and solids, with computational scaling that is often comparable to mean-field approaches and with a systematic route to improvement through the choice of trial wave function.

More generally, QMC methods offer favorable scaling and excellent parallel efficiency, enabling the treatment of strongly correlated systems beyond the reach of standard post-HF techniques. Their limitations are the presence of statistical error bars and the need for wise choice of constraints. In practice, statistical errors are reduced by longer sampling, while systematic biases are mitigated by improving the trial wave function.

2.1.7 Computational and Conceptual Problems of Modern Quantum Chemistry

The methods described so far highlight a recurring theme: the electronic structure problem is easy to formulate but notoriously hard to solve. In modern quantum chemistry, one quickly runs into deep questions of principle that limit what can be computed or even decided in general. At the abstract level, the local Hamiltonian problem is **QMA**-complete, meaning that calculating the approximation to the ground state energy within inverse-polynomial additive error of an electronic Hamiltonian in finite basis is intractable for classical and even quantum computers in polynomial time [17, 18]. But the problem is even more severe, it was proven that there is no general algorithm to decide whether a given Hamiltonian has a spectral gap bigger than some arbitrary value, which undermines the hope of a universal criterion for distinguishing stable insulators from metals [19]. Even in less extreme cases, the basic methods of quantum chemistry rest on problems that are combinatorial or hard in the worst case. The FCI expansion grows exponentially with system size in second quantization and as $\binom{M}{N}$ for N electron distributed among M spin-orbitals. HF is **NP**-complete in worst case, meaning that it is not solvable in polynomial time on classical computer [20, 18]. Determining whether a reduced density matrix corresponds to some N-electron wavefunction (the N-representability problem), which is important for methods formulated in reduced density matrix framework, is also **QMA**-complete [91].

For fermionic systems, QMC was also proven to be **NP**-hard due to the infamous sign problem [21]. This problem reflects the antisymmetric nature of fermionic wavefunctions, and if it is not considered, then contributions of different parts cancel, variances explode, and stochastic sampling breaks down. The same negativity appears in the Wigner quasiprobability distribution [92] and in the hydrodynamic formulation (Madelung) [93], underscoring that no classical probability picture can fully capture quantum correlations. The density functional theory (DFT) proves the existence of an exact universal functional of the electron density (Hohenberg–Kohn theorems) [94], but its explicit form

is unknown. Furthermore, even if it were known, computing it for a particular system would be **QMA**-hard [95].

Despite the daunting theorems, quantum chemistry works in practice. The reason is twofold: first, in the real world we only need finite precision—energies to a few millihartree, geometries to fractions of an angstrom. Second, physical Hamiltonians have rich structure and symmetries. These make problems much easier than their worst-case complexity suggests. Hence, heuristics and approximations flourish, all exploiting structure rather than brute force. Still, there remains a gap between theoretical algorithms and efficient implementations, and robust software engineering often matters as much as formal complexity.

Thus, modern quantum chemistry sits at the intersection of impossibility results and pragmatic success. On one side lies undecidability, exponential scaling, and proved mathematical hardness; on the other side are practical heuristics, approximations, and carefully engineered codes that make useful predictions. The tension itself is instructive: it shows that chemistry is not about exact solutions to formal problems but about extracting insight from structured approximations. Looking forward, new computational paradigms, such as quantum computers, promise not to resolve undecidability or worst-case hardness, but to enlarge the space of heuristics, offering alternative representations and more computational power to push against these limits.

2.2 Fundamentals of Quantum Computing

2.2.1 Emergence of Quantum Computing

It is often useful to begin an account of a theory with the circumstances of its emergence. By the late 1970s, experimental control of small quantum systems, notably in interferometric and optical settings, and quantum communication had become sufficiently refined that information-theoretic questions about preparation, transmission, and measurement could be asked in operational terms. A landmark result in this direction is Holevo's 1973 theorem, which upper-bounds the amount of classical information that can be reliably extracted from measurements on quantum system [96].

A second and ultimately decisive motivation came from the numerical simulation of quantum systems. The state space of a quantum system grows exponentially with the system size, making it difficult to simulate classically. However, the corresponding physical evolution occurs in real time in the laboratory. This tension suggested that the apparent cost of simulating quantum mechanics is due to the classical description, and it prompted the idea of using controllable quantum systems as computational resources for physics itself.

These topics were articulated particularly clearly in the early 1980s. Yuri Manin argued that progress in the theory of computation and in the mathematical description of nature calls for a theory of quantum automata, emphasizing both the exponential growth of quantum state spaces and the need to specify a restricted, physically implementable set of operations if such devices are to be programmable in any meaningful sense [13]. At the same time, Paul Benioff developed quantum mechanical models of computation that showed how the operation of a Turing machine can be represented as unitary time evolution. This made explicit that the computation can be consistent with the underlying physical laws [97]. In 1981, Richard Feynman sharpened the simulation argument in his famous lecture and the subsequent paper [14]. He observed that generic quantum dynamics appears to resist efficient classical simulation and proposed that one should instead build computers whose elementary degrees of freedom obey quantum mechanics, so that they can efficiently emulate other quantum systems.

The modern conceptual synthesis is usually dated to David Deutsch's 1985 work, which formulated the notion of a *universal quantum computer* and connected it to a physical form of the Church–Turing thesis [98]. Shortly thereafter, the quantum gate viewpoint was developed, in which computations are built from a finite set of elementary quantum operations, a perspective that underlies the circuit model used throughout quantum information science. The terminology and standard primitives introduced in this period will be described in the following sections.

2.2.2 Deterministic Evolution

Although there exist several models of quantum computation based on different underlying physical realizations, these models are equivalent to each other in computational power. Consequently, without loss of generality, we adopt the circuit model for the remainder of this dissertation.

At this stage, we restrict our attention to closed-system deterministic quantum evolution. Measurement theory and probabilistic readout are deferred to Sec. 2.2.3.

2.2.2.1 State of a Quantum Computer and Its Evolution

The basic unit of quantum information is the quantum bit or qubit, which can be represented by any two-level quantum system. Typical realizations include trapped ions of alkaline earth metals (Ca^+, Sr^+, Ba^+) or lanthanoids (Yb^+), neutral atoms in highly excited (Rydberg) states, superconducting circuits with Josephson junctions, and photonic polarization degrees of freedom.

The state of a single qubit can be described by a normalized vector in a two-dimensional complex Hilbert space, $|\psi\rangle \in \mathbb{C}^2$,

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \tag{2.103}$$

The global phases are physically irrelevant: the states $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ represent the same physical situation. Thus, the physical state of a qubit corresponds to a ray in $\mathbb{C}^2$, which is mathematically isomorphic to the two-dimensional sphere, called the Bloch sphere. Each point on the Bloch sphere specifies a pure single-qubit state up to global phase.

An N-qubit register is described by the tensor product Hilbert space $\left(\mathbb{C}^2\right)^{\otimes N}$ of dimension 2^N, with computational basis $\{|x\rangle : x \in \{0,1\}^N\}$ induced by the basis $\{|0\rangle, |1\rangle\}$ on each qubit. Thus, the state space of a quantum computer grows exponentially with the number of qubits.

A key non-classical feature of multiqubit systems is *entanglement*. For a two-qubit state $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$, if there are no single-qubit states $|\psi\rangle_A$ and $|\psi\rangle_B$ such that

$$|\psi\rangle_{AB} = |\psi\rangle_A \otimes |\psi\rangle_B, \tag{2.104}$$

then $|\psi\rangle_{AB}$ is entangled. Informally, entanglement means that the joint system exhibits correlations and behaves as a united entity even if it is spatially separated, regardless how far.

The time evolution of a closed quantum system is governed by the Schrödinger equation and is described by a unitary operator:

$$|\psi(t)\rangle = U(t)\,|\psi(0)\rangle, \qquad U(t) = \mathcal{T}\exp\!\left(-i\int_0^t H(s)\,ds\right), \quad U(t) \in \mathrm{U}(2^N), \tag{2.105}$$

where $H(s)$ is the possibly time-dependent Hamiltonian and $\mathcal{T}\exp$ denotes the time-ordered exponential, ensuring that Hamiltonians at different times are applied in the right order.

The process of quantum computation can be explained as an application of the unitary operator U_{circ} on the initial state $|\psi_{\text{in}}\rangle$ and the obtaining of the output state $|\psi_{\text{out}}\rangle$. In the circuit model, this ideal unitary evolution is implemented as a sequence of elementary quantum gates U_i:

$$|\psi_{\text{out}}\rangle = U_{\text{circ}}|\psi_{\text{in}}\rangle, \qquad U_{\text{circ}} = U_m U_{m-1} \cdots U_1. \tag{2.106}$$

2.2.2.2 Quantum Gates

One-qubit gates are unitary operations that act on single qubits. In the Bloch sphere representation, they correspond to rotations of the state vector. The matrices $\mathrm{X}, \mathrm{Y}, \mathrm{Z}$ are the Pauli operators, corresponding to π rotations about the x, y, z axes of the Bloch sphere, respectively. The Hadamard gate Had maps $|0\rangle \mapsto (|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle \mapsto (|0\rangle - |1\rangle)/\sqrt{2}$, which geometrically is a π rotation about the axis lying at 45° between the x and z axes. The gate S (or P) is the phase gate, which corresponds to a $\pi/2$ rotation about the z axis, while the T gate corresponds to a $\pi/4$ rotation about the z axis. The operators $R_X(\theta), R_Y(\theta), R_Z(\theta)$ implement general single-qubit rotations about the x, y, z axes by an angle θ. Table 2.1 summarizes a set of elementary quantum gates used throughout this thesis.

A particularly important class of multi-qubit gates consists of controlled operations. In a controlled unitary U, one or more qubits are designated as control qubits and determine whether U is applied to the target register. For a single control qubit and a single target register, the controlled-U operation can be written as

$$\mathrm{C}U = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U, \tag{2.107}$$

so that

$$|0\rangle \otimes |\psi\rangle \mapsto |0\rangle \otimes |\psi\rangle, \tag{2.108}$$
$$|1\rangle \otimes |\psi\rangle \mapsto |1\rangle \otimes U|\psi\rangle. \tag{2.109}$$

In other words, the unitary U is applied to the target register if and only if the control qubit is in the state $|1\rangle$ in the computational basis. The canonical two-qubit controlled gate is the controlled-NOT (CNOT), sometimes denoted as CX. It applies the Pauli X gate to the target qubit when the control qubit is in state $|1\rangle$:

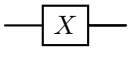
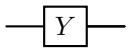
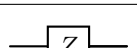
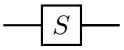
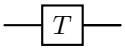
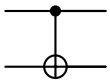
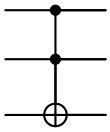
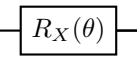
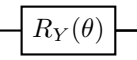
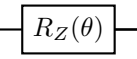
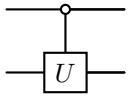
$$|0\rangle \otimes |\psi\rangle \mapsto |0\rangle \otimes |\psi\rangle, \tag{2.110}$$
$$|1\rangle \otimes |\psi\rangle \mapsto |1\rangle \otimes \mathrm{X}|\psi\rangle. \tag{2.111}$$

Thus, CNOT implements the reversible map

$$(c, t) \mapsto (c, c \oplus t), \tag{2.112}$$

where c and t are Boolean variables and $\oplus$ denotes the logical XOR (or exclusive OR).

Table 2.1: Elementary quantum logic gates and their circuit and matrix representations.

Gate name	Circuit depiction	Matrix representation
Pauli X (NOT)		$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
Pauli Y		$\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$
Pauli Z		$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$
Hadamard		$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$
Phase (S)		$\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$
$\pi/8$ (T)		$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$
Controlled-NOT (CNOT or CX)		$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$
Toffoli (CCNOT)		$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$
Rotation about x-axis		$\begin{pmatrix} \cos(\theta/2) & -i\sin(\theta/2) \\ -i\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$
Rotation about y-axis		$\begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$
Rotation about z-axis		$\begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}$
0-controlled unitary U		$\begin{pmatrix} U & 0 \\ 0 & I \end{pmatrix}$

A closely related three-qubit gate is the controlled-controlled-NOT, or Toffoli gate. It acts as the identity, unless both control qubits are in the state $|1\rangle$, in which case it flips the target qubit. In the computational basis, it implements

$$(c_1, c_2, t) \mapsto (c_1, c_2, t \oplus (c_1 \wedge c_2)), \tag{2.113}$$

where c_1, c_2 and t are Boolean variables and $\wedge$ denotes the logical AND. The Toffoli gate was introduced by Toffoli in the context of reversible classical computation [99], where it is sometimes described as an AND/NAND element. It is universal for classical reversible computation in the sense that any reversible Boolean function on a finite number of bits can be implemented using only Toffoli gates. In contrast, the Toffoli gate alone does not form a universal gate set for general quantum computation.

From the perspective of circuit synthesis, CNOT plays a central role. It is known that any controlled-U operation acting on a single target qubit can be decomposed, up to an overall global phase, into a circuit containing at most two CNOT gates and four single-qubit gates [100]. $\mathcal{B}$-controlled versions can be decomposed using ancillary qubits with $O(\mathcal{B})$ overhead in the number of control qubits, or without ancillas with $O(\mathcal{B}^2)$ overhead.

More generally, any unitary $U \in \mathrm{U}(2^N)$ can be decomposed exactly into a finite circuit consisting only of single-qubit gates and CNOT gates [100]. Such decompositions are not efficient in the worst case: a generic N-qubit unitary requires $\Theta(N^3 4^N)$ gates in any exact decomposition.[1]

In practical quantum hardware, it is often inconvenient, or even impossible, to implement arbitrary single-qubit rotations with continuous parameters at the logical level. Moreover, quantum error correcting codes, much like classical error correcting schemes, benefit from a discrete set of logical operations. These considerations motivate the search for finite, discrete universal gate sets, which will be discussed further.

2.2.2.3 Universal Gate Set

To introduce universal finite sets of quantum gates, it is convenient to recall two important subgroups of the unitary group $\mathrm{U}(2^N)$ of transformations on N-qubit states.

The N-qubit *Pauli group* consists of tensor products of single-qubit Pauli operators. Let $\sigma_1 = \mathrm{X}$, $\sigma_2 = \mathrm{Y}$, and $\sigma_3 = \mathrm{Z}$ denote the Pauli matrices, which satisfy

$$\sigma_i \sigma_j = \delta_{ij} I + i\, \epsilon_{ijk} \sigma_k, \tag{2.114}$$

where δ_{ij} is the Kronecker delta and ϵ_{ijk} is the fully antisymmetric Levi–Civita tensor. The N-qubit Pauli group can then be written as

$$\mathbf{P}_N = \langle W_0 \otimes \cdots \otimes W_{N-1} : W_i \in \{\mathrm{I}, \mathrm{X}, \mathrm{Y}, \mathrm{Z}\} \rangle. \tag{2.115}$$

[1] Here, $\Theta(\cdot)$ denotes an asymptotically tight bound: the number of gates required to decompose a generic N-qubit unitary grows proportionally to $N^3 4^N$.

The elements of $\mathbf{P}_N$ are often called Pauli strings. As a set of operations, $\mathbf{P}_N$ is quite restricted on its own, but it plays a central role in the theory of quantum error correction and in the characterization of more expressive gate sets.

The *Clifford group* is defined as the normalizer of the Pauli group in $\mathrm{U}(2^N)$:

$$\mathbf{C}_N = \left\{ C \in \mathrm{U}(2^N) : C\mathbf{P}_N C^\dagger = \mathbf{P}_N \right\}. \tag{2.116}$$

That is, $\mathbf{C}_N$ consists of all unitary operators (Clifford gates) C such that, under conjugation by C, every Pauli string is mapped to another Pauli string. In this sense, the Clifford group is the symmetry group of the set of Pauli strings. Although this definition might suggest a complicated structure, it is known that the $\mathbf{C}_N$ is generated by a remarkably small set of gates: the single-qubit Had, the phase gate S, and the two-qubit CNOT gate [101].

However, the gate set $\{\mathrm{Had}, \mathrm{S}, \mathrm{CNOT}\}$ is not universal for quantum computation. Unitaries generated by Clifford gates form only a discrete subset of $\mathrm{U}(2^N)$, and many unitaries cannot be realized using Clifford gates alone. Moreover, quantum circuits composed solely of Clifford operations, acting on computational basis inputs and followed by measurements in the computational basis, can be efficiently simulated on a classical computer (Gottesman–Knill theorem) [102].

Although the Clifford set $\{\mathrm{Had}, \mathrm{S}, \mathrm{CNOT}\}$ is not universal, a fundamental result shows that the addition of any non-Clifford gate can promote it to a universal set. A standard choice is the T gate. The gate set $\{\mathrm{Had}, \mathrm{S}, \mathrm{T}, \mathrm{CNOT}\}$ is universal for quantum computation [103]. In this picture, the non-Clifford gate is often informally said to supply the "magic" required to go beyond classically simulatable quantum dynamics. Alternatively, one can use other non-Clifford gates in place of T, such as the Toffoli gate.

However, this perspective can be misleading if taken too literally. Consider the universal gate set $\{\mathrm{Toffoli}, \mathrm{Had}\}$. The Toffoli gate alone suffices for universal classical computation and thus accounts entirely for the classical computational content; the Hadamard gate then injects intrinsically quantum behavior. In this example, the roles are reversed compared to the Clifford+T viewpoint. This illustrates that quantum features are not inherently attached to a single distinguished gate; rather, they emerge from the interplay of the chosen gate set, the structure of the computation, and the interpretation of the outputs. From a practical standpoint, the special role of T gate and other non-Clifford gates is crucial in the context of fault tolerance, which will be discussed in the Sec. 2.2.2.4.

We now know that finite universal gate sets exist. A natural question is whether such sets allow us to efficiently approximate arbitrary unitary transformations. The Solovay–Kitaev theorem provides a positive answer [104]. Generally speaking, let $\mathcal{G}$ be a finite universal gate set whose generated group is dense in $\mathrm{SU}(2)$. Then any $U \in \mathrm{SU}(2)$ can be ε-approximated by a product of elements of $\mathcal{G}$ of length

$$O\left(\log^c \frac{1}{\varepsilon}\right), \tag{2.117}$$

for some constant c (in the original proofs one can take $c \approx 4$ [105]; later refinements achieve $c = 1$ [106, 107, 108]).

This shows that a fixed finite gate set suffices to approximate any desired quantum operation to arbitrary precision with only logarithmic overhead in the target accuracy.

2.2.2.4 Circuit Cost Models and Compilation

Even though we have mathematically universal gate sets for quantum computation, in practice any algorithm must be implemented using the hardware-specific gate primitives available on a given device. This leads to two main difficulties.

First, a given architecture typically provides a fixed set of native gates from which all other operations must be synthesized. One-qubit gates are often relatively easy and high-fidelity, whereas multi-qubit entangling gates such as CNOT or Toffoli are substantially more error-prone and slower. Moreover, hardware connectivity constraints may require the insertion of additional SWAP gates to mediate interactions between distant qubits. As a consequence, when analyzing the resource requirements of an algorithm on a given architecture, it is common to use the number of entangling gates (e.g. CNOT count, or in some contexts Toffoli count) and the overall circuit depth as primary cost measures, sometimes in addition to the total number of qubits.

A second complication arises in the setting of fault-tolerant quantum computation (FTQC). It is known that no quantum error correction code can realize a full universal gate set transversally [109]. Transversal gates are desirable because they automatically prevent error proliferation. In practice, only a subset of gates can be simultaneously transversal (such as $\{\mathrm{Had}, \mathrm{S}, \mathrm{CNOT}\}$). In such error correction codes, the T gate must be realized using more elaborate protocols to achieve universality together with fault-tolerance. The standard approach is magic state distillation, which consumes many noisy copies of specially prepared ancilla states to produce a smaller number of high-fidelity magic states, enabling fault-tolerant implementation of non-Clifford gates such as T [110, 111]. The overhead of these protocols dominates the cost of large-scale fault-tolerant computation. Consequently, when estimating the resource requirements of an algorithm for FTQC, particular attention is paid to the number of T gates (or the Toffoli count).

2.2.3 Measurements

Measurements provide the interface between the quantum dynamics of a circuit and the classical information accessible to an observer. They are precisely the point at which intrinsically probabilistic behavior appears in the standard formulation of quantum mechanics: even if the evolution prior to measurement is fully deterministic and unitary, the measurement outcomes are random, with probabilities determined by the state.

2.2.3.1 Projective Measurements and the Born Rule

Let $|\psi\rangle$ denote the state of an N-qubit register, so that $|\psi\rangle$ is a unit vector in the Hilbert space $\mathcal{H} = \left(\mathbb{C}^2\right)^{\otimes N} \cong \mathbb{C}^{2^N}$. It is convenient to associate to $|\psi\rangle$ the density operator

$$\rho = |\psi\rangle\langle\psi|. \tag{2.118}$$

It is a positive semidefinite operator of unit trace acting on the Hilbert space $\mathcal{H}$.

An observable A is represented by a Hermitian operator on $\mathcal{H}$, and can be written in its spectral decomposition as

$$A = \sum_{\mathfrak{a}\in\mathcal{A}} \mathfrak{a}\,\Pi_{\mathfrak{a}}, \tag{2.119}$$

where $\mathcal{A}$ is the set of possible measurement outcomes (eigenvalues of A) and $\{\Pi_{\mathfrak{a}}\}_{\mathfrak{a}\in\mathcal{A}}$ are the corresponding orthogonal projectors onto the eigenspaces of A. These projectors satisfy

$$\Pi_{\mathfrak{a}}\Pi_{\mathfrak{a}'} = \delta_{\mathfrak{a}\mathfrak{a}'}\Pi_{\mathfrak{a}}, \qquad \sum_{\mathfrak{a}\in\mathcal{A}} \Pi_{\mathfrak{a}} = I. \tag{2.120}$$

The collection $\{\Pi_{\mathfrak{a}}\}_{\mathfrak{a}\in\mathcal{A}}$ is called a projection-valued measure (PVM).

To connect the quantum state with the statistics of measurement outcomes, one introduces the Born rule. In the standard axiomatic formulation, it is postulated that the probability of obtaining the outcome $\mathfrak{a}$ when measuring the observable A on a system in state ρ is

$$\Pr[\mathfrak{a}] = \mathrm{Tr}(\Pi_{\mathfrak{a}}\rho). \tag{2.121}$$

For pure states $\rho = |\psi\rangle\langle\psi|$, this reduces to $\Pr[\mathfrak{a}] = \|\Pi_{\mathfrak{a}}|\psi\rangle\|_2^2$. Although the Born rule can be derived from additional structural assumptions (Gleason's theorem) [112], we shall treat it here as a fundamental postulate.

Conditioned on obtaining outcome $\mathfrak{a}$, the post-measurement state of the system is given by the Lüders update rule:

$$\rho \;\longmapsto\; \rho_{\mathfrak{a}} = \frac{\Pi_{\mathfrak{a}}\rho\Pi_{\mathfrak{a}}}{\mathrm{Tr}(\Pi_{\mathfrak{a}}\rho)}. \tag{2.122}$$

2.2.3.2 What Do We Measure in Quantum Computing?

At any step of a unitary evolution, the state of an N-qubit quantum computer can be written in the computational basis as

$$|\psi\rangle = \sum_{n=0}^{2^N-1} c_n|n\rangle, \tag{2.123}$$

where $|n\rangle$ denotes the N-bit binary form of the integer n. All information about the state is encoded in the complex amplitudes c_n. However, in quantum computation, these amplitudes cannot be accessed directly.

The Holevo bound states that the amount of classical information accessible from an N-qubit system is at most N classical bits per use of the system, regardless of how the state is encoded or which measurement is performed [96]. This is a strong limitation: although specifying a generic pure state of N qubits requires 2^N complex amplitudes (hence exponentially many real parameters), a single measurement on one copy of the state yields at most N bits of classical information.

Simultaneous measurements of commuting observables are possible, but if two observables do not commute, then the information one can obtain about them jointly is constrained by uncertainty relations. Consequently, to estimate multiple expectation values or to learn more detailed properties of a state, one typically requires many identically prepared copies of that state. In general, these copies must be produced by repeating the entire state preparation and often the entire algorithm from scratch.

A further obstruction is given by the no-cloning theorem [113, 114, 115]. There cannot be a unitary operation U that, given an arbitrary and *a priori* unknown state $|\psi\rangle$, produces a second identical copy:

$$U|\psi\rangle|0\rangle = |\psi\rangle|\psi\rangle \tag{2.124}$$

for all $|\psi\rangle$. Thus, we cannot copy out a quantum state and probe it arbitrarily many times; instead, we must repeatedly prepare the state and perform fresh measurements on each copy, collecting statistics from many runs.

A particularly important and practically relevant case is the measurement of Pauli observables. Let P be an N-qubit Pauli string (a tensor product of $\{\mathrm{I}, \mathrm{X}, \mathrm{Y}, \mathrm{Z}\}$) with eigenvalues $\{\pm 1\}$ and let ρ be the prepared state. Measurement of P corresponds to a two-outcome PVM with projectors

$$\Pi_{\pm} = \frac{I \pm P}{2}. \tag{2.125}$$

A single measurement shot returns a random variable $\mathcal{R} \in \{\pm 1\}$ with

$$\Pr[\mathcal{R} = \pm 1] = \mathrm{Tr}(\Pi_{\pm}\rho), \qquad \mathbb{E}[\mathcal{R}] = \mathrm{Tr}(P\rho) = \mu. \tag{2.126}$$

Thus, the expectation value μ of P in the state ρ is encoded as the mean of the $\{\pm 1\}$-valued random variable $\mathcal{R}$.

Over D independent circuit executions we obtain independent and identically distributed samples $\mathcal{R}_1, \ldots, \mathcal{R}_D$, and estimate μ via the sample mean

$$\hat{\mu} = \frac{1}{D}\sum_{i=1}^{D} \mathcal{R}_i. \tag{2.127}$$

Unbiasedness is immediate from Eq. (2.126): $\mathbb{E}[\hat{\mu}] = \mu$. Moreover, since $|\mathcal{R}_i| \leq 1$, Hoeffding's inequality yields

$$\Pr\big(|\hat{\mu} - \mu| \geq \varepsilon\big) \leq 2\exp\big(-2D\varepsilon^2\big). \tag{2.128}$$

Therefore, it suffices to take

$$D \geq \frac{1}{2\varepsilon^2} \ln\frac{2}{\delta} \tag{2.129}$$

to ensure that $|\hat{\mu} - \mu| \leq \varepsilon$ with probability at least $1 - \delta$. This simple statistical model underlies many measurement schemes.

On current quantum hardware, measurements ultimately resolve qubits in the computational Z-basis by coupling them to a macroscopic detector. For example, in the superconducting scheme, the qubit state shifts the resonance frequency of a microwave cavity, which is then probed by microwave pulses. In trapped-ion and neutral-atom systems, measurement is typically fluorescence-based. Measurements in other single-qubit bases are performed by applying appropriate unitary rotations before the hardware-native Z-basis readout.

In summary, quantum computers do not reveal amplitudes directly, they provide samples from probability distributions defined by observables. Practical quantum algorithms are therefore designed to extract useful information from such samples, either through straightforward statistical estimation or through more advanced amplitude estimation and tomography-based techniques.

2.2.3.3 Hadamard Test

One of the standard primitives for extracting information from a quantum state is the Hadamard test. Suppose that we are given a state $|\psi\rangle$ and we wish to estimate the expectation value of a unitary operator U in that state. The Hadamard test circuit is shown in Fig. 2.1.

A short calculation shows that the probability of obtaining outcome 0 in the ancilla register is

$$\Pr[0] = \tfrac{1}{2}\big(1 + \mathrm{Re}\langle\psi|U|\psi\rangle\big), \tag{2.130}$$

so that

$$\mathrm{Re}\langle\psi|U|\psi\rangle = 2\Pr[0] - 1. \tag{2.131}$$

By inserting a S gate in the ancilla before the final Hadamard, one can similarly access the imaginary part of $\langle\psi|U|\psi\rangle$.

Repeating the Hadamard test independently $\mathcal{B}$ times produces samples of a $\{0, 1\}$-valued random variable whose expectation encodes $\mathrm{Re}\langle\psi|U|\psi\rangle$ (or $\mathrm{Im}\langle\psi|U|\psi\rangle$ in the modified version). Using the Hoeffding's inequality, one finds that $\mathcal{B} = O(1/\varepsilon^2)$ repetitions suffice to estimate the desired quantity to additive accuracy ε.

Despite its conceptual simplicity, the Hadamard test is not always the most practical tool. First, each application estimates only a single expectation value, so many repetitions and separate circuits may be required to characterize a state or an operator of interest. Second, the test requires a controlled implementation of U, which can be substantially deeper and more complex than an uncontrolled realization of U. For these reasons, in many settings one instead relies on measurement schemes that decompose observables into sums of simpler terms (such as Pauli strings) and estimate their expectation values by direct sampling, or uses more global reconstruction methods such as state tomography.

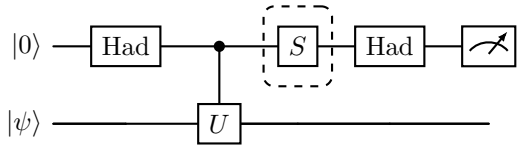


Figure 2.1: Circuit for the Hadamard test used to estimate $\mathrm{Re}\langle\psi|U|\psi\rangle$ (and $\mathrm{Im}\langle\psi|U|\psi\rangle$).

2.2.3.4 State Tomography

State tomography [116, 117] aims to learn many expectation values $\{\mathrm{Tr}(\rho O_i)\}_i$ of an unknown quantum state $\rho \in \mathbb{C}^{g\times g}$ (with $g = 2^N$ for an N-qubit system) from measurement data. Any informationally complete scheme is constrained by information-theoretic lower bounds: estimating a generic state ρ up to trace distance at most ε with constant confidence requires on the order of $\Theta(g^2/\varepsilon^2)$ copies of the state in the worst case. However, many physically relevant states are far from generic. If ρ is of low-rank, then tomography can achieve sample complexity on the order of $\tilde{O}(rg/\varepsilon^2)$, where $r \ll g$ is the rank and $\tilde{O}$ hides polylogarithmic factors. In practice, tomography focuses on choosing measurement strategies that make it easy to estimate the observables of interest with as few samples as possible, together with statistical procedures that guarantee reliability and convergence as more data are collected.

To formulate classical shadow methods in a unified framework, it is convenient to introduce the notion of a measurement channel. Consider a randomized measurement scheme: we sample a unitary U from some ensemble (e.g. Pauli strings, Clifford unitaries), apply it to the state ρ, and then measure in the computational basis, obtaining an outcome bit string b. The outcome probability is

$$\Pr[b \mid U] = \mathrm{Tr}\big(U\rho U^\dagger |b\rangle\langle b|\big). \tag{2.132}$$

Conditioned on the outcome b, the post-measurement state in the original frame can be represented by the projector

$$\Pi_{U,b} = U^\dagger |b\rangle\langle b|\, U. \tag{2.133}$$

Thus, each experimental run produces a classical record (U, b) and an associated projector $\Pi_{U,b}$.

The measurement channel $\mathcal{M}$ associated with this randomized scheme is defined as the average map

$$\mathcal{M}(\rho) = \mathbb{E}_U\Big[\sum_b \Pr[b \mid U]\,\Pi_{U,b}\Big] = \mathbb{E}_{U,b}\big[\Pi_{U,b}\big], \tag{2.134}$$

where the outer expectation is taken over the randomness in choosing U and over the measurement outcomes b.

The key idea of classical shadows is to construct an inverse map $\mathcal{M}^{-1}$ and define a single-shot estimator

$$\hat{\rho} = \mathcal{M}^{-1}\big(\Pi_{U,b}\big), \qquad \mathbb{E}[\hat{\rho}] = \rho, \tag{2.135}$$

where the expectation is taken over the randomness in the measurement procedure. The random operator $\hat{\rho}$ is called a classical-shadow estimator. The term "shadow tomography" refers to the fact that we do not reconstruct the full density matrix explicitly in all applications; instead, we use the shadows to predict many observables of interest.

In particular, classical shadows provide estimators for expectation values. Given a set of observables $\{O_i\}_{i=1}^{M}$, we define the single-shot estimate

$$\hat{o}_i = \mathrm{Tr}\big(O_i \hat{\rho}\big), \tag{2.136}$$

and approximate $\langle O_i \rangle = \mathrm{Tr}(O_i \rho)$ by aggregating $\hat{o}_i$, so that $\mathbb{E}[\hat{o}_i] = \langle O_i \rangle$. The number of shots N_{shots} needed to estimate all $\langle O_i \rangle$ within the additive error ε and the failure probability at most δ scales as

$$N_{\mathrm{shots}} = O\left(\max_i \mathrm{Var}[\hat{o}_i] \, \frac{\log(M/\delta)}{\varepsilon^2} \right), \tag{2.137}$$

where $\mathrm{Var}[\hat{o}_i]$ is the variance of estimators for $\hat{o}_i$.

The most important feature of the classical-shadow framework is the logarithmic dependence on the number of observables M (see Eq. (2.137)). This makes it possible, in principle, to predict a very large family of expectation values, provided the shadow norms remain moderate.

2.2.3.5 Pauli Shadows

The simplest classical shadows are obtained by measuring each qubit in a randomly chosen Pauli basis (i.e. X, Y, or Z), implemented by single-qubit unitaries $U_j \in \{\mathrm{I}, \mathrm{Had}, \mathrm{Had} \cdot \mathrm{S}^\dagger\}$ and computational basis readout. The corresponding N-qubit single-shot estimator has the simple inversion formula

$$\hat{\rho} = \bigotimes_{j=1}^{N} (3\Pi_j - I), \quad \Pi_j = U_j^\dagger |b_j\rangle\langle b_j| \, U_j, \tag{2.138}$$

where Π_j denotes the projector associated with the outcome on qubit j.

If the observable O acts nontrivially on k qubits, the variation satisfies $\mathrm{Var}[\hat{o}] \leq 4^k$. Therefore, estimating M such observables to additive accuracy ε and failure probability δ requires

$$N_{\mathrm{shots}} = O\left(4^k \, \frac{\log(M/\delta)}{\varepsilon^2} \right) \tag{2.139}$$

measurements. Pauli shadows are thus particularly efficient for local observables, such as short Pauli strings.

2.2.3.6 Clifford Shadows

Clifford shadows are obtained by applying a random N-qubit Clifford unitary U to the state ρ, followed by a computational-basis measurement. For such

randomized Clifford measurements, a single measurement outcome therefore yields the unbiased classical-shadow estimator

$$\hat{\rho} = (2^N + 1)\, U^\dagger |b\rangle\langle b|\, U - I. \tag{2.140}$$

For a collection of observables $\{O_i\}_{i=1}^M$, the number of measurements needed to estimate all expectation values within additive error ε and failure probability at most δ satisfies

$$N_{\text{shots}} = O\left(\max_i \operatorname{Tr}(O_i^2)\, \frac{\log(M/\delta)}{\varepsilon^2} \right). \tag{2.141}$$

This scaling is advantageous for observables with small Hilbert–Schmidt norm (e.g. low-rank projectors), but can be unfavorable for Pauli strings, since $\operatorname{Tr}(O_i^2) = 2^N$. Clifford shadows can be particularly efficient when measuring fidelities with pure states, as the number of shots is generally independent of the number of qubits.

2.2.3.7 Matchgate Shadows

Matchgate (fermionic Gaussian) shadows restrict the randomization to unitaries generated by quadratic fermionic Hamiltonians [43]. Equivalently, they are exactly the unitaries that implement a linear (Bogoliubov) transformation of the creation/annihilation operators:

$$\begin{aligned} \hat{\tilde{a}}_i^\dagger &= \sum_j \big(U_{ij}\, \hat{a}_j^\dagger + V_{ij}\, \hat{a}_j \big), \\ \hat{\tilde{a}}_i &= \sum_j \big(U_{ij}^*\, \hat{a}_j + V_{ij}^*\, \hat{a}_j^\dagger \big). \end{aligned} \tag{2.142}$$

They preserve fermionic Gaussian states; in the number-conserving case ($V = 0$), this reduces to orbital rotations and preserves Slater determinants (see Eq. (2.54)).

If O is a product of k creation and k annihilation operators in the second quantization, then [118]

$$\operatorname{Var}[\hat{o}] \le \frac{\binom{2N}{2k}}{\binom{N}{k}} \approx \frac{4^k k!}{(2k)!} N^k \quad (N \gg k), \tag{2.143}$$

where N is number of spin-orbitals (usually same as number of qubits).

This ensemble is well matched to chemistry and free-fermion phases because many target observables have small number of creation/annihilation operators, even though their qubit Pauli expansions can be highly nonlocal.

2.2.4 Quantum Algorithms

This subsection reviews core primitives and quantum algorithms relevant to electronic structure theory.

2.2.4.1 Fermion to Qubit Mapping

To simulate a second-quantized electronic Hamiltonian (see Eq. (2.43)) on a quantum device, one must represent the fermionic creation and annihilation operators $\{\hat{a}_p^\dagger, \hat{a}_p\}$ as tensor products of Pauli operators acting on qubits, while preserving the canonical anticommutation relations (see Eq. (2.32)).

The simplest and most widely used encoding is the Jordan–Wigner (JW) mapping [119]. It stores the occupation of each spin-orbital directly in the computational basis of a single qubit. Let $\mathrm{X}_p, \mathrm{Y}_p, \mathrm{Z}_p$ denote Pauli gates acting on qubit p. Defining the JW parity string $\mathrm{Z}_{<p} := \prod_{j=0}^{p-1} \mathrm{Z}_j$, the mapping is

$$\hat{a}_p^\dagger \;\mapsto\; \tfrac{1}{2}\,(\mathrm{X}_p - i\,\mathrm{Y}_p)\,\mathrm{Z}_{<p}, \tag{2.144}$$

$$\hat{a}_p \;\mapsto\; \tfrac{1}{2}\,(\mathrm{X}_p + i\,\mathrm{Y}_p)\,\mathrm{Z}_{<p}, \tag{2.145}$$

$$\hat{N}_p := \hat{a}_p^\dagger \hat{a}_p \;\mapsto\; \tfrac{1}{2}\,(\mathrm{I} - \mathrm{Z}_p), \tag{2.146}$$

where $\hat{N}_p$ is the number operator for spin-orbital p.

The JW mapping is conceptually transparent and straightforward to implement. A well-known drawback is that the parity string $\mathrm{Z}_{<p}$ can introduce long Pauli strings, which may increase circuit depth for certain operators.

Other encodings, such as the parity mapping and the Bravyi–Kitaev mapping [120], modify the distribution of parity information across qubits and can shorten typical Pauli strings. Since they are not used elsewhere in this dissertation, we do not discuss them further.

2.2.4.2 Linear Combination of Unitaries

A central primitive for many quantum algorithms in quantum chemistry is the representation of an operator as a linear combination of unitaries (LCU) [121].

Let an operator A be decomposed as:

$$A = \sum_{k=0}^{L-1} w_k\, U_k, \qquad \lambda_{\mathrm{LCU}} = \sum_k |w_k|, \tag{2.147}$$

where U_k are unitaries and L is the number of LCU terms. The map $|\psi\rangle \mapsto (A/\lambda_{\mathrm{LCU}})|\psi\rangle$ can be realized using an ancilla register by defining operators

$$\text{PREP}: \quad |0\rangle \mapsto |\omega\rangle = \sum_{k=0}^{L-1} \sqrt{|w_k|/\lambda_{\mathrm{LCU}}}\,|k\rangle, \tag{2.148}$$

$$\text{SELECT}: \quad |k\rangle|\psi\rangle \mapsto |k\rangle U_k\big|\psi\rangle, \tag{2.149}$$

where PREP prepares the state $|\omega\rangle$ encoding the coefficients w_k of Eq. (2.147), while SELECT applies the indexed unitary to a working register.

The composite circuit

$$V_A \;=\; (\text{PREP}^\dagger \otimes I)\,\text{SELECT}\,(\text{PREP} \otimes I) \tag{2.150}$$

satisfies

$$\big(\langle 0| \otimes I\big) V_A \big(|0\rangle \otimes I\big) = A/\lambda_{\mathrm{LCU}}. \tag{2.151}$$

Thus, V_A is a $(\lambda_{\mathrm{LCU}}, a, 0)$ block encoding of an operator A, since it prepares this operator scaled by λ_{LCU} in the $\{|0\rangle, |0\rangle\}$ block of the matrix V_A using $a = \lceil \log_2 L \rceil$ ancilla qubits. The third parameter equals to 0 meaning this block encoding is exact. The LCU circuit is shown in Fig. 2.2.

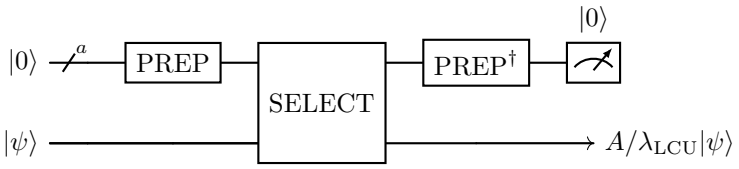


Figure 2.2: LCU circuit preparing an $(\lambda_{\mathrm{LCU}}, a, 0)$ block encoding of A. The state $A/\lambda_{\mathrm{LCU}}|\psi\rangle$ is up to corresponding normalization.

The postselection success probability on input $|\psi\rangle$ is

$$\Pr[\mathrm{succ}] = \big\|(A/\lambda_{\mathrm{LCU}})|\psi\rangle\big\|_2^2 \leq \|A\|_{\mathrm{LCU}}^2/\lambda_{\mathrm{LCU}}^2 \leq 1, \tag{2.152}$$

where $\|A\|_{\mathrm{LCU}} = \min_{\{w_k\}} \sum_k |w_k|$, is the LCU norm, corresponding to the smallest possible sum of the LCU coefficients through all possible decompositions.

The PREP operator can be synthesized via a binary-tree ladder of controlled rotations [122], applying $O(L)$ controlled rotations to $\lceil \log_2 L \rceil$ ancilla qubits. To achieve the precision of overall state preparation ε, each rotation must be accurate to $O(\varepsilon/L)$, which produces $O(\log(L/\varepsilon))$ T gates per rotation.

However, the T gate count of PREP can be improved using a quantum read-only memory (QROM) scheme [123]. The idea of QROM is that the square roots of normalized LCU coefficients $\sqrt{|w_i|/\lambda_{\mathrm{LCU}}}$ are stored as classical binary data in the table of L entries. These values are loaded in the quantum state:

$$|0\rangle|0\rangle \xrightarrow{\mathrm{QROM}} \sum_k C_k |k\rangle\, |d_k\rangle, \tag{2.153}$$

where d_k is a classical binary data, associated with index k. This allows one to prepare a state,

$$|0\rangle^{\otimes(1+2\log L+2\mu)} \to \sum_{k=0}^{L-1} \sqrt{\frac{\tilde{w}_k}{\lambda_{\mathrm{LCU}}}} |k\rangle|\mathrm{temp}_k\rangle, \tag{2.154}$$

with a μ-bit approximation $\tilde{w}_k/\lambda_{\mathrm{LCU}}$ and garbage state $|\mathrm{temp}_k\rangle$. Although this uses $1 + 2\lceil \log L \rceil + 2\mu$ ancilla qubits, it requires only $4L + O(\log(1/\varepsilon))$ T gates with $\mu = O(\log(1/\varepsilon))$, which gives linear scaling in L with a precision independent prefactor.

A basic SELECT applies indexed multi-controlled U_k and, without optimization, costs $O(L \log L)$ T gates. Using clean ancillas and gate merging [123]

reduces the multi-control overhead to $4L - 4$ T gates. These counts capture only the multi-control scaffolding; if the U_k are rotation unitaries, their own synthesis costs add and can dominate, whereas Pauli strings may be generated at lower cost.

2.2.4.3 Qubitization

The LCU block encoding represents A only probabilistically, since the desired action appears in the $\{|0\rangle, |0\rangle\}$ block after postselection. Qubitization improves this construction by producing a deterministic unitary whose eigenvalues directly encode the spectrum of A [124].

Let $|\omega\rangle = \mathrm{PREP}\,|0\rangle$ be the ancilla state prepared in the LCU construction, and define the projector $P = |\omega\rangle\langle\omega| \otimes I$. A reflection about this subspace is $R = 2P - I$. The qubitized walk operator is then defined as

$$W_A = R \cdot \mathrm{SELECT}\,. \tag{2.155}$$

The key property of W_A is its spectrum. If $A|a_j\rangle = a_j|a_j\rangle$ with $|a_i| \leq \lambda_{\mathrm{LCU}}$, then W_A acts invariantly on the two-dimensional subspace spanned by $\{|\omega\rangle|a_j\rangle, \mathrm{SELECT}\,|\omega\rangle|a_j\rangle\}$ and has eigenvalues $e^{\pm i \arccos(a_j/\lambda_{\mathrm{LCU}})}$. Thus, the eigenvalues of A are encoded in the rotation angles of W_A. In other words, each eigenspace of A is mapped to an SU(2) invariant subspace of W_A, so that it acts as a single-qubit rotation in each such subspace.

2.2.4.4 Grover Search and Amplitude Amplification

Grover's algorithm was among the first quantum algorithms to demonstrate an asymptotic advantage over existing classical algorithms [5]. It allows one to search for $\mathcal{G}$ marked ("good") elements within an unstructured database of size $\mathcal{D}$ using only $O(\sqrt{\mathcal{D}/\mathcal{G}})$ oracle queries, whereas the best classical algorithm requires $\Theta(\mathcal{D}/\mathcal{G})$ queries.

Let $|\psi\rangle$ be the initial guess state, typically a uniform superposition over all database elements. We define a projector onto the marked subspace Π_{good} and the corresponding phase oracle

$$O_{\mathrm{good}} = I - 2\Pi_{\mathrm{good}}, \tag{2.156}$$

which performs a reflection about the "good" subspace. The second reflection is taken about the guess state $|\psi\rangle$:

$$S_\psi = I - 2|\psi\rangle\langle\psi|. \tag{2.157}$$

The product

$$G = S_\psi O_{\mathrm{good}} \tag{2.158}$$

is the Grover iterate.

The action of G is confined to the two-dimensional subspace spanned by the normalized projections of $|\psi\rangle$ onto the "good" and "bad" subspaces:

$$|\psi_{\text{good}}\rangle = \frac{\Pi_{\text{good}}|\psi\rangle}{\|\Pi_{\text{good}}|\psi\rangle\|_2}, \qquad |\psi_{\text{bad}}\rangle = \frac{(I-\Pi_{\text{good}})|\psi\rangle}{\|(I-\Pi_{\text{good}})|\psi\rangle\|_2}. \tag{2.159}$$

Let

$$\Pr[\text{succ}] = \langle\psi|\Pi_{\text{good}}|\psi\rangle \tag{2.160}$$

denote the initial success probability. Then the initial state decomposes as

$$|\psi\rangle = \sqrt{1-\Pr[\text{succ}]}\,|\psi_{\text{bad}}\rangle + \sqrt{\Pr[\text{succ}]}\,|\psi_{\text{good}}\rangle. \tag{2.161}$$

Introducing an angle $\theta \in [0, \pi/2]$ such that $\sin^2\theta = \Pr[\text{succ}]$, the Grover iterate G performs a rotation by angle 2θ in the $\{|\psi_{\text{bad}}\rangle, |\psi_{\text{good}}\rangle\}$ plane. After t Grover iterations, the state evolves as

$$|\psi\rangle \mapsto \sin\big((2t+1)\theta\big)\,|\psi_{\text{good}}\rangle + \cos\big((2t+1)\theta\big)\,|\psi_{\text{bad}}\rangle. \tag{2.162}$$

Thus, Grover's search can be viewed as a sequence of SU(2) rotations by a fixed angle 2θ in this two-dimensional subspace. The optimal number of iterations that maximizes the success probability is

$$t_\star = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor. \tag{2.163}$$

In practice, however, θ (and hence $\Pr[\text{succ}]$) is usually unknown, so repeated applications of G can overshoot the optimal angle, rotating past the good subspace and decreasing the success probability. This issue is addressed in fixed-point modification [9].

This geometric picture extends naturally to a more general setting known as amplitude amplification (AA) [8]. In AA, an initial state $|\psi\rangle$ and a projector Π_{good} are given defining the desired measurement outcome. The same iteration pattern as in Grover's algorithm is then used to amplify the amplitude (and hence the probability) associated with the good subspace. The underlying theory is identical to that of Grover's search.

2.2.4.5 Quantum Signal Processing

Quantum signal processing (QSP) refers to a family of quantum circuit constructions that implement polynomial transformations of a signal encoded in the phase of a unitary operator [125]. It forms the mathematical backbone of many recent quantum algorithms.

Consider a signal unitary $S(\mathfrak{x}) = R_X(2\arccos(\mathfrak{x})) = e^{-i\arccos(\mathfrak{x})\mathrm{X}}$ denoting a rotation about the x-axis on a single qubit. A degree-d QSP sequence is defined as

$$U_\Phi(\mathfrak{x}) = \prod_{k=1}^{d} \big[R_Z(\phi_k)\,S(\mathfrak{x})\big]\,R_Z(\phi_0), \tag{2.164}$$

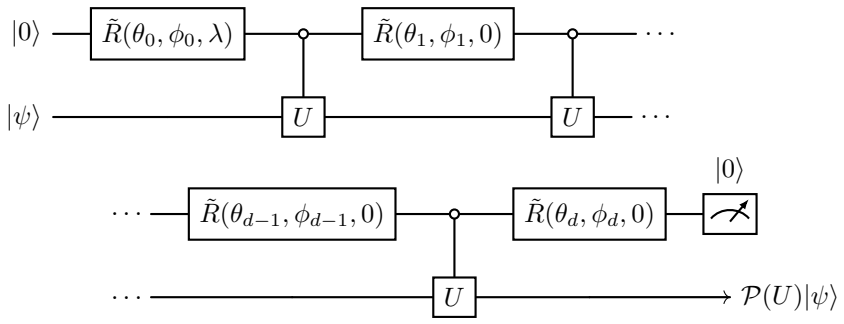


Figure 2.3: GQSP circuit. The circuit comprises one signal processing qubit and a system register, which $\mathcal{P}(U)$ acts on. The state $\mathcal{P}(U)|\psi\rangle$ is up to corresponding normalization.

where $\Phi = (\phi_0, \ldots, \phi_d)$ are programmable phase angles. The resulting unitary $U_\Phi(\mathfrak{x}) \in \mathrm{SU}(2)$ has the matrix form

$$U_\Phi(\mathfrak{x}) = \begin{pmatrix} \mathcal{P}(\mathfrak{x}) & i\mathcal{Q}(\mathfrak{x})\sqrt{1-\mathfrak{x}^2} \\ i\mathcal{Q}^*(\mathfrak{x})\sqrt{1-\mathfrak{x}^2} & \mathcal{P}^*(\mathfrak{x}) \end{pmatrix}, \tag{2.165}$$

where $\mathcal{P}(z)$ and $\mathcal{Q}(z)$ are complex polynomials of degree at most d and $d-1$, respectively, satisfying the boundedness condition

$$|\mathcal{P}(\mathfrak{x})|^2 + \left(1-\mathfrak{x}^2\right)|\mathcal{Q}(\mathfrak{x})|^2 = 1 \quad \text{for } \forall \mathfrak{x} \in [-1, 1]. \tag{2.166}$$

The parity of $\mathcal{P}(z)$ is $d \bmod 2$ and the parity of $\mathcal{Q}(z)$ is $(d-1) \bmod 2$. Any pair $(\mathcal{P}(z), \mathcal{Q}(z))$ that obeys these parity and boundedness conditions can be realized by a suitable choice of phases Φ. Thus, $\langle 0|U_\Phi(\mathfrak{x})|0\rangle$ produces the $a \mapsto \mathcal{P}(\mathfrak{x})$ map, and the less restrictive $\langle +|U_\Phi(\mathfrak{x})|+\rangle$ produces $\mathfrak{x} \mapsto \mathrm{Re}(\mathcal{P}(\mathfrak{x})) + i\,\mathrm{Re}(\mathcal{Q}(\mathfrak{x}))\sqrt{1-\mathfrak{x}^2}$.

When the signal unitary $S(\mathfrak{x})$ is obtained from the qubitized walk operator W_A associated with a block encoding of A/λ_{LCU} (see Sec. 2.2.4.3), the QSP procedure is known as quantum singular value transformation (QSVT) [126]. In this setting, each singular subspace of A forms a two-dimensional invariant subspace on which the walk operator acts as $S(\sigma_j/\lambda_{\mathrm{LCU}}) = e^{-i \arccos(\sigma_j/\lambda_{\mathrm{LCU}})\mathrm{X}}$, where σ_j is a singular value of A. Applying the QSP sequence (2.164) then implements the transformation

$$\sigma_j/\lambda_{\mathrm{LCU}} \;\mapsto\; \mathcal{P}(\sigma_j/\lambda_{\mathrm{LCU}}), \tag{2.167}$$

for each singular value σ_j, while preserving the corresponding singular vectors.

This framework unifies a wide range of quantum algorithms, from AA to Hamiltonian simulation and linear systems solvers, as special cases of polynomial transformations of the spectrum of suitably block encoded operators [127].

2.2.4.6 Generalized Quantum Signal Processing

Generalized quantum signal processing (GQSP) enables the implementation of an arbitrary complex polynomial $\mathcal{P}(z)$ acting on a unitary U [128]. Unlike conventional QSP, which is restricted to real-valued and fixed parity polynomials, GQSP allows arbitrary complex-valued and mixed parity polynomial transformations. The only limitation is that this polynomial should obey $|\mathcal{P}(z)| \leq 1$ in complex unit circle $|z| = 1$. It is proven that for any such polynomial $\mathcal{P}(z)$ the complementary polynomial $\mathcal{Q}(z)$ exists [128].

Define a controlled signal operator CU as in Eq. (2.107). The GQSP sequence alternates CU with single-qubit rotations $\tilde{R}(\theta, \phi, \lambda)$ on the ancilla[2]:

$$\begin{bmatrix} \mathcal{P}(U) & \cdot \\ \mathcal{Q}(U) & \cdot \end{bmatrix} = \prod_{k=1}^{d} \Big(\tilde{R}(\theta_k, \phi_k, 0)\, \mathrm{C}U \Big) \tilde{R}(\theta_0, \phi_0, \lambda), \tag{2.168}$$

with angles $\{\theta_k, \phi_k, \lambda\}$ chosen to realize a degree d polynomial [128, 129, 130]. Various angle finding algorithms are discussed in Appendix A1. This yields a block encoding of $\mathcal{P}(U)$ (see Fig. 2.3). The rotation is defined as

$$\tilde{R}(\theta, \phi, \lambda) = \begin{bmatrix} e^{i(\lambda+\phi)} \cos(\theta) & e^{i\phi} \sin(\theta) \\ e^{i\lambda} \sin(\theta) & -\cos(\theta) \end{bmatrix}. \tag{2.169}$$

In general, the GQSP protocol is probabilistic, as the $\mathcal{P}(U)$ transformation may not be unitary. The measurement of the ancilla in $|0\rangle$ is successful with probability

$$\Pr[\text{succ}] = \big\| \mathcal{P}(U)|\psi\rangle \big\|_2^2. \tag{2.170}$$

Employing AA, the procedure can be made nearly deterministic, thereby ensuring high-fidelity realization of the target polynomial transformation.

2.2.4.7 Quantum Phase Estimation

The unitary time evolution operator in quantum mechanics $U(t) = e^{-iHt}$, has eigenvalues $e^{-iE_j t}$ corresponding to the energies E_j of the Hamiltonian H. If one can extract these oscillatory components, the associated energies become accessible. Quantum phase estimation (QPE) [6, 7] achieves this by applying controlled time evolutions $U(t_k)$ for a sequence of times t_k in superposition, followed by an inverse quantum Fourier transform on the time register to resolve phases [131].

Assume that $H|\phi_j\rangle = E_j|\phi_j\rangle$, so that $U(t)|\phi_j\rangle = e^{-iE_j t}|\phi_j\rangle$. If we prepare an arbitrary superposition $|\psi\rangle = \sum_j c_j|\phi_j\rangle$, then the application of QPE produces an energy estimate $\tilde{E}_j$ with probability $|c_j|^2$, and collapses the system into the corresponding eigenstate $|\phi_j\rangle$. The precision of the energy estimate depends on the total evolution time and hence on the number of phase qubits.

[2]The parameter λ appearing in $\tilde{R}(\theta, \phi, \lambda)$ denotes a single-qubit rotation phase parameter and should not be confused with λ_{LCU} used elsewhere.

In practical quantum algorithms, the time evolution operator e^{-iHt} is typically not available directly as a primitive. Two principal approaches are used to construct efficient approximations: Trotterization and qubitization.

In the Trotter–Suzuki approach, the Hamiltonian is decomposed into a sum of simpler terms, each of which can be efficiently exponentiated:

$$H = \sum_{\ell=1}^{m} H_\ell. \tag{2.171}$$

The evolution operator is then approximated by a product formula,

$$e^{-iHt} \approx \left(\prod_{\ell=1}^{m} e^{-iH_\ell t/r} \right)^r, \tag{2.172}$$

where r is the number of Trotter steps.

The Trotterized approach is conceptually simple and aligns naturally with molecular Hamiltonians expressed as sums of Pauli strings. However, the gate count required to achieve a fixed target precision can grow rapidly with the system size and the simulation time.

The qubitization framework, described in Sec. 2.2.4.3, instead achieves exact unitary signal synthesis from a block encoding of H/λ_{LCU}. In this setting, a walk operator W_H is built such that, for each eigenstate $|\phi_j\rangle$ of H with eigenvalue E_j,

$$W_H|\phi_j\rangle = e^{\pm i \arccos(E_j/\lambda_{\mathrm{LCU}})}|\phi_j\rangle, \tag{2.173}$$

for some scaling parameter $\lambda_{\mathrm{LCU}} \geq \|H\|_{\mathrm{LCU}}$. Using W_H as a time evolution in QPE directly yields an estimate of $\arccos(E_j/\lambda_{\mathrm{LCU}})$, from which the energy can be recovered.

The precision and cost of QPE are determined by the number of controlled applications of W_H. Achieving the energy precision ε_E requires $\Theta(\lambda_{\mathrm{LCU}}/\varepsilon_E)$ of queries. To resolve the ground state from the first excited state, the energy precision ε_E should be proportional to the energy gap $\Delta = E_1 - E_0$. For electronic structure Hamiltonians, λ_{LCU} is determined by the sum of the absolute values of the LCU coefficients (see Eq. (2.147)) and typically scales as $O(N^4)$ in the number of orbitals. Tensor hypercontraction (THC) reduces the effective norm parameter λ_{LCU} [132]. This directly decreases the QPE runtime, since the number of controlled applications of W_H required for a target energy precision scales proportionally to λ_{LCU}.

Regardless of how accurately the underlying evolution is implemented, the overall success probability of QPE is determined by the overlap $\gamma = |\langle\phi_0|\psi\rangle|$, where $|\phi_0\rangle$ is the desired eigenstate (typically the ground state). If the overlap is small, many repetitions may be required to obtain the correct eigenvalue, with an expected number of repetitions on the order of $1/\gamma^2$. AA can be used to quadratically improve the scaling to $1/\gamma$ (see Sec. 2.2.4.4). Thus, the overall scaling of QPE is $\Theta(\lambda_{\mathrm{LCU}}/(\Delta\gamma))$. Consequently, effective state preparation is a crucial ingredient in any practical QPE-based quantum algorithm.

2.2.4.8 Quantum Selected Configuration Interaction

Quantum selected configuration interaction (QSCI) [29], also known as sample-based quantum diagonalization (SQD) [30], combines quantum sampling with classical subspace diagonalization to approximate electronic energies. The idea is to prepare a correlated state $|\psi\rangle$ on a quantum device, sample computational-basis configurations $\{|b\rangle\}$ with probabilities $\Pr[b] = |\langle b|\psi\rangle|^2$, and select the most significant ones to form a truncated CI subspace $\mathcal{S}$.

A classical computer then diagonalizes the Hamiltonian projected onto this subspace,

$$H_{\mathcal{S}\mathcal{S}} = \left\{ \langle b|H|b'\rangle \right\}_{b,b'\in\mathcal{S}}, \tag{2.174}$$

by solving the eigenproblem

$$H_{\mathcal{S}\mathcal{S}}\,\vec{c} = E\,\vec{c}. \tag{2.175}$$

The resulting eigenvectors give approximate electronic states expressed as linear combinations of sampled determinants.

The motivation behind QSCI is that quantum devices can efficiently sample from highly entangled wave functions, potentially identifying important configurations that are inaccessible to classical selected-CI heuristics. However, the quality of these samples depends entirely on the prepared state $|\psi\rangle$: it should contain all important configurations to represent the exact state and, ideally, with equal probabilities of sampling.

2.2.4.9 Quantum Krylov Subspace Methods

Quantum Krylov subspace methods construct a low-dimensional subspace

$$\mathcal{K}_m(H, |\psi_0\rangle) = \mathrm{span}\{\, f_0(H)|\psi_0\rangle, \ldots, f_{m-1}(H)|\psi_0\rangle \,\} \tag{2.176}$$

estimate the Hamiltonian matrix representation in this basis on a quantum device and then perform subspace diagonalization on a classical computer. The obtained eigenvalues approximate the low-lying energies.

Different choices of the generator set $\{f_k(H)\}$ define concrete Krylov schemes with different resource requirements and robustness:

- **Quantum filter diagonalization (QFD) [34]** Choosing $f_k(H) = e^{-it_kH}$ generates a basis from real-time snapshots $\{e^{-it_kH}|\psi_0\rangle\}_k$. QFD is a prominent example of using time-propagated states and a Rayleigh–Ritz procedure to extract multiple eigenstates.

- **Quantum Lanczos (QLanczos) [35]** Taking $f_k(H) = e^{-\tau_kH}$ produces a sequence of "cooled" states $\{e^{-\tau_kH}|\psi_0\rangle\}_k$. Combining quantum imaginary-time evolution (QITE) with a Lanczos projection provides access to low-lying energy spectra.

- **Block encoding Lanczos method [36]** One may choose $f_k(H) = T_k(H)$, where $T_k(x)$ are Chebyshev polynomials. These polynomials can be realized naturally using qubitization. Since products of Chebyshev polynomials

produce linear combinations of Chebyshev polynomials, matrix elements can be estimated easily using a quantum computer.

- **Multireference Quantum Krylov Algorithm** [31] Starting from several physically motivated reference states $\{|\psi_0^{(a)}\rangle\}_a$, one forms a block subspace span$\{f_k(H)|\psi_0^{(a)}\rangle\}_{k,a}$. The generator set consists of time evolution operators of the form $f_k(H) = e^{-it_k H}$.

These methods provide a systematic and well-defined route to ground and excited states by enlarging the Krylov space. However, their precision and convergence strongly depend on the choice and preparation of the initial state $|\psi_0\rangle$.

2.2.5 Quantum State Preparation

In previous discussion, we saw that many quantum algorithms, such as QPE, QSCI or Krylov subspace method, are critically dependent on choice of the wave function. Here, we will go through the most well-known algorithms for quantum state preparation.

2.2.5.1 Variational Quantum Eigensolver

The variational quantum eigensolver (VQE) [22] is one of the most widely used hybrid quantum–classical algorithms for ground state energy estimation and state preparation. A parameterized quantum circuit prepares a trial wave function $|\psi(\vec{\theta})\rangle$, and a classical optimizer iteratively updates the parameters $\vec{\theta}$ to minimize the energy expectation value estimated on a quantum device.

Given a molecular Hamiltonian H, VQE aims to solve the variational problem

$$E_{\mathrm{VQE}} = \min_{\vec{\theta}} \langle\psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle, \tag{2.177}$$

with the variational principle guaranteeing that $E_{\mathrm{VQE}} \geq E_0$, where E_0 is the exact ground state energy, and equality is achieved when $|\psi(\vec{\theta})\rangle$ coincides with the ground state. The potential quantum advantage of VQE arises from the efficient preparation and evaluation of highly entangled trial states that are difficult to simulate classically.

The expressivity and practicality of VQE depend crucially on the ansatz chosen for $|\psi(\vec{\theta})\rangle$. Two main ansatzes used in this dissertation are:

- **Unitary Coupled-Cluster (UCC) ansatz [72]:** Based on the classical coupled cluster expansion (see Sec. 2.1.5.3), it prepares the state

$$|\psi_{\mathrm{UCC}}(\vec{\theta})\rangle = e^{T(\vec{\theta})-T^\dagger(\vec{\theta})}|\phi_0\rangle, \tag{2.178}$$

where $|\phi_0\rangle$ is the HF reference state and $T(\vec{\theta})$ is an excitation operator. In practice, the excitation manifold is truncated, e.g., to singles and doubles (UCCSD) or to doubles only (UCCD).

- **Hardware-efficient ansatz (HEA) [133]:** Designed to match the connectivity and native gate set of a given device, typically using alternating layers of parameterized single-qubit rotations and entangling gates. The main advantage is its relatively low cost and circuit depth.

A major challenge for VQE is the occurrence of barren plateaus, regions of parameter space where the gradient of the cost function becomes negligible, making gradient-based optimization intractable for deep or unstructured ansatzes.

VQE-prepared states can also serve as initial guesses for QPE, QSCI, and Krylov subspace methods, so VQE functions not only as a standalone variational algorithm, but also as a practical state preparation primitive for a broader family of quantum chemistry algorithms.

2.2.5.2 Imaginary-Time Evolution

ITE prepares ground states by applying the nonunitary map $e^{-\tau H}$ to an initial state $|\psi(0)\rangle$:

$$|\phi(\tau)\rangle \;=\; \frac{e^{-\tau H}|\psi(0)\rangle}{\|e^{-\tau H}|\psi(0)\rangle\|}. \tag{2.179}$$

As $\tau \to \infty$, the lowest-energy component dominates, projecting onto the ground state provided that the initial overlap is nonzero. The convergence speed is governed by the spectral gap and the initial overlap with the ground state.

Several quantum implementations of ITE have been developed, including variational ITE [134], QITE [35], probabilistic nonunitary schemes [135], and QSP/QSVT realizations [136], each offering different trade-offs in circuit depth and robustness.

Recent work shows that ITE arises as a solution of Brockett's double-bracket flow for the density operator,

$$\frac{\partial \rho(\tau)}{\partial \tau} \;=\; [[\rho(\tau), H], \rho(\tau)], \qquad \rho(\tau) = |\psi(\tau)\rangle\langle\psi(\tau)|, \tag{2.180}$$

and synthesizes double-bracket QITE (DB-QITE) circuits that coherently implement small steps of this flow without postselection [137]. Each step provably lowers the energy by an amount proportional to the current energy variance and increases ground state fidelity, providing a deterministic alternative to projective schemes.

2.2.5.3 State Filters

State filtering refers to methods that amplify desired spectral components of a Hamiltonian while suppressing the rest. Let H have eigenpairs $\{(E_j, |\phi_j\rangle)\}$ and let the initial state be $|\psi\rangle = \sum_j c_j|\phi_j\rangle$. Given a filter function $f : \mathbb{R} \to \mathbb{C}$ that is large near a target energy and small elsewhere, the filtered state is

$$|\psi_{\text{filtered}}\rangle = \frac{f(H)|\psi\rangle}{||f(H)|\psi\rangle||_2} = \frac{1}{||f(H)|\psi\rangle||_2} \sum_j c_j f(E_j)\,|\phi_j\rangle. \tag{2.181}$$

In practice, $f(H)$ is implemented approximately using polynomial constructions (often via QSP/QSVT) or LCU of time evolution operators [138, 139, 140, 141, 142].

A common limitation is that $f(H)$ is generally non-unitary, so it is implemented via a block encoding and ancilla postselection. The success probability typically scales as $\Pr[\text{succ}] = \|f(H)|\psi\rangle\|_2^2$. For narrow or selective filters this can be small, but it can be boosted by AA.

2.2.6 Classical Post-Processing

It is expected that the electronic structure Hamiltonian of realistic molecular systems will be too large to be represented directly even on FTQC. Therefore, similarly to classical approaches in quantum chemistry, it is necessary to reduce the problem size using an active space approximation (see Sec. 2.1.5.4). Within this framework, only the active space problem is solved on a quantum device. The resulting wave function obtained from the quantum computation is then measured and subsequently post-processed on a classical computer to account for electronic correlation outside the active space. Several quantum–classical hybrid methods can be employed for this purpose, which will be discussed in the following.

2.2.6.1 Quantum-Classical Hybrid NEVPT

One of the quantum–classical hybrid approaches that can be used for classical post-processing employs the NEVPT2 to capture dynamic electronic correlation outside the active space [143]. NEVPT2 is one of MRPT methods (see Sec. 2.1.5.5), in which the total electronic Hamiltonian is partitioned as

$$\hat{H} = \hat{H}_{\text{Dyall}} + \hat{V} = \sum_{a,\sigma} \varepsilon_a \hat{a}^\dagger_{a\sigma} \hat{a}_{a\sigma} + \hat{H}_{\text{act}} + \hat{V}, \tag{2.182}$$

where $\hat{H}_{\text{Dyall}}$ is the Dyall Hamiltonian [144] and $\hat{H}_{\text{act}}$ is the active space Hamiltonian. ε_a correspond to the energies of the external orbitals. The Dyall Hamiltonian is chosen to be zeroth-order approximation and $\hat{V}$ is a perturbation.

In this hybrid approach, a VQE calculation is first performed for $\hat{H}_{\text{act}}$ to obtain an approximation to the ground state $|\Psi_0\rangle$ and its energy E_0. The second-order perturbative correction to the energy is then given by

$$E^{(2)} = -\sum_{m \neq 0} \frac{|\langle \Psi_m | \hat{V} | \Psi_0 \rangle|^2}{E_m - E_0}, \tag{2.183}$$

where $|\Psi_m\rangle$ denote excited states of the Dyall Hamiltonian with corresponding energies E_m.

This expression can be reformulated in terms of ionized states rather than explicit eigenstates $|\Psi_m\rangle$. To obtain these ionized states, the quantum subspace expansion (QSE) is employed. Within QSE, the effective Hamiltonian and overlap

matrices are constructed from states generated by single and double ionization of the reference state $|\Psi_0\rangle$ and are subsequently diagonalized separately within each relevant Fock-space sector.

The principal limitation of this approach is its unfavorable scaling. Specifically, the construction of QSE matrices scales as $O(N_{\text{act}}^8)$, while their diagonalization scales as $O(N_{\text{act}}^{12})$, which severely restricts the size of tractable active spaces. However, an important advantage is that the method scales only as $O(N_{\text{ext}}^2)$ with respect to the number of orbitals outside the active space.

2.2.6.2 Quantum-Classical Hybrid AFQMC

AFQMC can also be employed as a classical post-processing scheme for quantum-computed wave functions. As discussed in Sec. 2.1.6.3, AFQMC is highly efficient and accurate, having a computational scaling of $O(N^3) \sim O(N^4)$ with system size. This scaling makes AFQMC suitable for large systems and particularly for quantum–classical hybrid approaches.

In AFQMC, the ground state energy is evaluated as

$$E = \frac{\sum_l w_l e^{i\theta_l} E_l^{\text{loc}}}{\sum_l w_l e^{i\theta_l}}, \qquad E_l^{\text{loc}} = \frac{\langle \Psi_T | \hat{H} | \phi_l \rangle}{\langle \Psi_T | \phi_l \rangle}, \tag{2.184}$$

where $|\phi_l\rangle$ denote the stochastic walkers, w_l their associated weights, and θ_l the corresponding phases. The quantity E_l^{loc} is the local energy evaluated with respect to a trial wave function $|\Psi_T\rangle$, which serves as an approximation to the true ground state. The nodal structure imposed by $|\Psi_T\rangle$ constrains the AFQMC evolution and plays a crucial role in controlling the sign problem.

In the quantum–classical AFQMC (QC-AFQMC) framework, a quantum algorithm, typically VQE, is used to prepare the trial state $|\Psi_T\rangle$. In order to interface the quantum-generated state with the classical AFQMC procedure, quantum state tomography techniques are required (see Sec. 2.2.3.4). The early implementation used Clifford shadow tomography [42], demonstrating high precision and resilience to hardware noise.

However, the use of Clifford shadows leads to an exponential classical post-processing cost ($O(\exp(N_{\text{act}}))$) to evaluate the local energy $\langle \Psi_T | \hat{H} | \phi_l \rangle$ with the size of the active space N_{act}. This bottleneck can be alleviated by employing matchgate shadow techniques [44], which reduce the scaling to $O(N_{\text{act}}^{8.5} \log N_{\text{act}})$. Despite this improvement, the resulting cost remains prohibitively large for practical applications.

More recently, the matchgate shadow framework was further refined through the use of analytic differentiation to calculate the local energy and force bias, reducing the overall computational complexity to $O(N_{\text{act}}^{5.5})$ [145].

In this dissertation, we develop an approach that further improves the scaling of QC-AFQMC post-processing. The details of this method and its performance are presented in Chap. 5.

2.2.6.3 Quantum-Classical Hybrid Tailored Coupled Cluster

The tailored CC (TCC) theory [146] can also be used as a classical post-processing scheme. In classical quantum chemistry, this approach enables the incorporation of static electronic correlation into single-reference CC methods (see Sec. 2.1.5.3). The procedure begins with an active space calculation, typically performed at the CASCI level. The resulting CI vector of the ground state is then used to determine the cluster amplitudes associated with excitations within the active orbitals, denoted by $\hat{T}^{\text{act}}$. The remaining cluster amplitudes corresponding to excitations involving orbitals outside the active space, $\hat{T}^{\text{ext}}$, are obtained by solving the standard CC equations (see Eq.(2.74)).

Building on this framework, the quantum–classical TCC method based on computational basis tomography (QC-CBT-TCC) was proposed in Ref. [147]. In this approach, the active space problem is solved on a quantum device. The resulting quantum state is then measured using computational basis tomography (CBT), which is described in detail in Sec. 5.2.1. The reconstructed CBT state is subsequently used to tailor the active space cluster amplitudes $\hat{T}^{\text{act}}$.

The QC-CBT-TCC wave function is given by

$$|\Psi_{\text{QC-CBT-TCC}}\rangle = e^{\hat{T}^{\text{ext}}} e^{\hat{T}^{\text{act}}} |\Psi_0\rangle, \tag{2.185}$$

where $|\Psi_0\rangle$ denotes the HF reference state. After solving the CC amplitude equations, the total energy is evaluated as

$$E_{\text{QC-CBT-TCC}} = \langle \Psi_0 | \hat{H} | \Psi_{\text{QC-CBT-TCC}} \rangle. \tag{2.186}$$

The statistical error arising from the CBT measurement process can be mitigated using a corrected energy expression,

$$E_{\text{QC-CBT-TCC(c)}} = E^{\text{act}}_{\text{QC}} + \left(E_{\text{QC-CBT-TCC}} - E^{\text{act}}_{\text{QC-CBT-TCC}} \right), \tag{2.187}$$

where $E^{\text{act}}_{\text{QC}}$ is the active space energy obtained directly from the quantum computation, and $E^{\text{act}}_{\text{QC-CBT-TCC}}$ is the corresponding active space energy evaluated within the TCC framework.

This method has demonstrated high accuracy in capturing both static and dynamic electronic correlation. However, its computational scaling is identical to that of the underlying classical CC approach. For example, the QC-CBT-TCCSD method scales as $O(N^6_{\text{ext}})$ with respect to the number of orbitals outside the active space, which limits its applicability to systems with moderate external spaces.

2.2.7 Open Problems of Quantum Computing for Quantum Chemistry

As discussed in the previous sections, numerous quantum algorithms have been proposed for electronic structure calculations. A fundamental challenge common to all of them is the preparation of a quantum state with sufficiently large overlap

with the target electronic eigenstate. This remains the principal bottleneck in practical quantum algorithms for quantum chemistry.

If such a state could be prepared efficiently, QPE would, in principle, enable the extraction of ground and excited state energies with exponential speedup to any known classical method. Moreover, accurate state preparation would allow efficient sampling in QSCI and the construction of accurate effective Hamiltonians in Krylov subspace methods.

The quantum computer can provide an advantage in the way it prepares the state due to its ability to run an N-qubit quantum circuit with 2^N of classical parameters, requiring only $O(N)$ computational time. Although physical insight from quantum chemistry enables the construction of a well-motivated ansatz, determining its optimal parameters remains a challenging problem. Since the Hamiltonian is the only fundamental source of information about its eigenstates, it is natural to use it to guide state preparation. This idea underlies methods such as ITE and power iteration. In Chap. 3 we explore power-iteration-based approaches to prepare molecular ground and excited states on quantum hardware.

Molecular electronic states obey exact symmetries, including the particle number and spin. Quantum state preparation methods, particularly VQE using HEA, may violate these symmetries. Restoration of these symmetries helps to achieve a better overlap with the desired state. Techniques for projecting out states with correct symmetries are therefore essential and are discussed in Chap. 4.

Finally, it is unlikely that quantum computers will directly treat the full molecular Hamiltonian of realistic systems in the near term. Instead, quantum devices are expected to solve active space problems, whereas the remaining electronic correlation is recovered using classical post-processing methods. Several such schemes were reviewed in Sec. 2.2.6, and a new efficient post-processing approach is introduced in Chap. 5.

Progress in these areas is crucial for making quantum computing a practical tool for quantum chemistry.

Chapter 3

Quantum Power Iteration Unified Using Generalized Quantum Signal Processing

We propose a unifying framework for the state preparation using quantum power method algorithms based on GQSP. We apply GQSP to realize quantum analogs of classical power iteration, power Lanczos, inverse iteration, and folded spectrum methods, all within a single coherent framework. GQSP allows efficient realization of methods that require complex polynomials while avoiding the limitations of approaches based on linear combinations of time-evolution operators. Our constructions, including a Trotter-decomposition-free quantum inverse iteration, achieve near-optimal query scaling, together with reduced qubit requirements. The same formalism yields a quantum folded spectrum method for excited state preparation that avoids explicitly forming powers of the Hamiltonian or performing variational optimization. We provide a theoretical analysis of success probabilities and resource scaling, and we validate the methods numerically using molecular Hamiltonians. The results show that quantum power Lanczos lowers the computational cost and provides robust convergence compared to naive quantum power iteration. Our findings reveal that GQSP-based implementations of power methods combine scalability, flexibility, and robust convergence, paving the way for practical initial state preparations on fault-tolerant quantum devices. This chapter is based on the preprint [Khinevich, Lee, Yoshioka and Mizukami, arXiv:2507.11142 (2025)].

3.1 Introduction

Driven by advances in hardware architecture and algorithmic design, quantum computing has developed rapidly. Quantum computers are widely believed to outperform classical computers in condensed matter physics, quantum chemistry,

and materials science [148, 149]. Solving the eigenvalue problem is a key computational challenge in these domains.

The most direct method of obtaining eigenvalues on a quantum computer is the QPE algorithm [15, 150]. Despite its conceptual elegance, QPE typically requires substantial computational resources and FTQCs. Consequently, proposals to reduce overhead [151, 152] face the fundamental challenge of preparing a high-fidelity initial approximation of the target state.

The VQE [22], which uses a hybrid quantum-classical approach, is a widely used algorithm to prepare an approximate quantum state. The VQE offers the advantage of shallow circuits that are compatible with existing hardware [153, 133, 154]. However, its reliance on classical optimization makes it vulnerable to issues such as barren plateaus [155], which limit its scalability for more expressive ansatzes.

ITE offers another route for ground state preparation by approximating a non-unitary propagator [35]. Although classical ITE has long been a workhorse for quantum many-body simulations, quantum adaptations still face challenges. Many implementations inherit Trotter errors of time-evolution [134, 156, 157], high resource costs to estimate the quantum Fisher information matrix [158], or postselection overheads [159]. Furthermore, ITE is naturally tailored to find ground states, thereby limiting its versatility for broader classes of problems.

By contrast, power methods represent a more general and flexible family of iterative algorithms with a long history of use in classical numerical linear algebra. In addition to ground-state preparation, power methods include algorithms for computing excited states, inverting matrices, solving linear systems, and performing principal component analysis. Therefore, they are widely used in numerous areas such as computer-aided engineering, finance, and machine learning. In quantum chemistry, power-iteration-based schemes have also been developed for vibrational Hamiltonians, including tensor-structured power iteration [160], Chebyshev methods [161], and inverse iteration [162], demonstrating that such iterative approaches can be effective even in high-dimensional many-body settings.

Recent studies have explored various ways of implementing quantum power methods. Techniques include LCU to implement powers of cosine transformed Hamiltonians [163], finite-difference approximations of time evolution operators [164], direct encoding using rotational operators [165], HHL-based approach [166, 167] and QSVT [126, 168, 169]. Inverse iteration schemes have also been implemented [170, 171, 172, 173]. Krylov subspace methods [174, 31, 33] extend these ideas to more sophisticated approaches. However, many of these implementations face practical limitations, including Trotter errors, high qubit costs, and restrictions to only real-valued polynomials.

This study demonstrates that, for the state preparation, various power methods can be implemented in a unified way using GQSP [128, 175, 176] (or more specifically, the generalized quantum singular value transformation (GQSVT) [177]). The key advantage of GQSP over QSP [125] is that it can encode any general complex polynomial without additional overhead. This paves the way for the implementation of the power Lanczos method, which may require

complex coefficients, as well as for the future development of Green's function and complex energy shifts in the inverse iteration method. Furthermore, the GQSP gives advantage over methods that use a linear combination of time-evolution operators. For instance, our implementation of the quantum inverse iteration is free of Trotter error and requires fewer qubits in numerical tests than the existing approaches.

In addition, using the same technique, we introduced the quantum folded spectrum method (QFSM) to prepare the excited state wave functions, if approximate energy estimation is available. This method was previously used in the VQE framework, which required the explicit construction of a Hamiltonian squared. This results in numerous terms being measured [178]. GQSP allows us to avoid the explicit construction of the powers of Hamiltonian and variational optimization.

Throughout this work, we assume a FTQC setting, as the techniques we analyze require coherent implementation of polynomial transformations and precise phase operations. We provide a theoretical analysis of these quantum power methods, including estimates of success probability, query complexity, and qubit requirements. We show that the resulting algorithms achieve near-optimal scaling. We also construct explicit quantum circuits for each method and numerically validate their performance. In addition, we examine practical resource requirements, such as the total number of queries and qubits needed in realistic settings. Importantly, the proposed formulations remain mathematically equivalent to their classical counterparts; they can be simulated classically given exponential memory, and we confirm the expected convergence behavior.

The remainder of this chapter is organized as follows. Section 3.2 reviews the fundamentals of GQSP and presents our implementation of various power methods. Section 3.3 provides computational details and a link to the reference code. Section 3.4 presents a theoretical analysis of resource scalings, followed by numerical demonstrations of the algorithms and comparisons with existing techniques, with particular attention to the convergence behavior for molecular Hamiltonians and the associated computational costs. Finally, Sec. 3.5 summarizes our findings, discusses potential improvements, and presents directions for future research and development.

3.2 Theory

3.2.1 Polynomial Transformation of a Hamiltonian

GQSP is a powerful framework that can be used to efficiently implement polynomial transformations of a Hamiltonian H (see Sec. 2.2.4.6). It is first embedded in a larger unitary operator using a block encoding procedure, such as LCU (see Sec. 2.2.4.2). Within the block encoding framework, H is rescaled as $\mathcal{H} = H/\lambda_{\mathrm{LCU}}$, where λ_{LCU} is the normalization constant associated with the LCU construction. The corresponding eigenvalues are therefore scaled as $\tilde{E}_i = E_i/\lambda_{\mathrm{LCU}}$.

The LCU block encoding unitary acts, in each invariant subspace associated with an eigenvalue $\tilde{E}_i$, as

$$U_{\mathrm{LCU}} = \begin{pmatrix} \mathcal{H} & \sqrt{I-\mathcal{H}^2} \\ \sqrt{I-\mathcal{H}^2} & -\mathcal{H} \end{pmatrix}, \tag{3.1}$$

which is a reflection. For GQSP it is convenient to convert this reflection into a rotation by multiplying with a reflection about $|0\rangle$ on the ancilla:

$$U = (Z \otimes I)\, U_{\mathrm{LCU}} = \begin{pmatrix} \mathcal{H} & \sqrt{I-\mathcal{H}^2} \\ -\sqrt{I-\mathcal{H}^2} & \mathcal{H} \end{pmatrix}, \tag{3.2}$$

which acts as a SU(2) rotation in each invariant subspace associated with $\tilde{E}_i$. This is exactly what the qubitization technique does with LCU operator (see Sec. 2.2.4.3).

GQSP uses a 0-controlled block encoding operator as the signal operator, denoted by CU. This operator is defined as in Eq. (2.107).

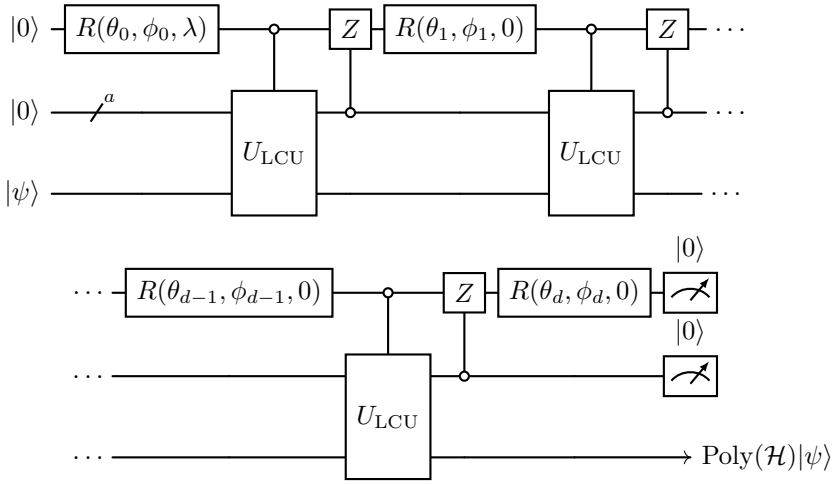


Figure 3.1: GQSP circuit. The circuit comprises three registers: the first qubit register controls the GQSP procedure, the second a qubits register performs the block encoding, and the third register holds the state on which the Hamiltonian acts on. The state $\mathrm{Poly}(\mathcal{H})|\psi\rangle$ is up to corresponding normalization.

It is essential for the GQSP method to determine the angle sequences $\vec{\theta}$, $\vec{\phi}$, and the phase parameter λ that produce the desired polynomial $\mathcal{P}(z)$. As the first step, for any target polynomial $\mathcal{P}(z)$, a complementary polynomial $\mathcal{Q}(z)$ that satisfies $|\mathcal{P}(z)|^2 + |\mathcal{Q}(z)|^2 = 1$ for all z on the complex unit circle must be constructed. It has been rigorously proven that this complementary polynomial exists for any $\mathcal{P}(z)$ with $|\mathcal{P}(z)| \leq 1$ for all z on the complex unit circle [128].

Various methods exist to determine the appropriate angles, and further details are provided in the Appendix A1.

In a practical GQSP implementation, the LCU-to-qubitization reflection (see Eq. (3.2)) can be absorbed into the GQSP sequence. This is achieved by inserting a controlled-Z gate whose target is the signal-processing ancilla and whose control is the $|0\rangle$ state of the LCU ancilla, effectively multiplying U_{LCU} by the required reflection. The resulting quantum circuit is illustrated in Fig. 3.1, and its structure is analogous to the GQSVT construction in Appendix B1.

The next task involves obtaining a polynomial function $\mathrm{Poly}(\mathcal{H})$ from a polynomial function $\mathcal{P}(U)$. This is accomplished by noting that the powers of U generate Chebyshev polynomials of $\mathcal{H}$:

$$U^k = \begin{pmatrix} T_k(\mathcal{H}) & * \\ * & * \end{pmatrix}, \tag{3.3}$$

where $T_k(\mathcal{H})$ denotes the k-th Chebyshev polynomial evaluated at $\mathcal{H}$. This identity is a standard consequence of the qubitization construction: the unitary U has eigenvalues $e^{\pm i \arccos(\tilde{E}_i)}$ for each eigenvalue $\tilde{E}_i$ of $\mathcal{H}$, so U^k implements $\cos\big(k \arccos\big(\tilde{E}_i\big)\big) = T_k(\tilde{E}_i)$. Consequently, the operator $\mathcal{P}(U)$ constructed via the GQSP block encodes a linear combination of Chebyshev polynomials of $\mathcal{H}$, which in turn can represent any polynomial function $\mathrm{Poly}(\mathcal{H})$.

3.2.2 Quantum Power Iteration Method

Power iteration is a well-known technique for determining the largest magnitude eigenvalue of a matrix. We use this approach to find a ground state of the Hamiltonian $\mathcal{H}$. The underlying idea is straightforward: one chooses an initial guess for the eigenvector $|\psi^{(0)}\rangle$, which has a nonzero overlap with the true ground state, assuming that the ground state energy $\tilde{E}_0$ has the largest absolute value among all the eigenvalues of $\mathcal{H}$. The iterative procedure is defined as follows:

$$|\psi^{(n+1)}\rangle = \frac{\mathcal{H}|\psi^{(n)}\rangle}{||\mathcal{H}|\psi^{(n)}\rangle||_2}. \tag{3.4}$$

The error of this method scales as $O(|\tilde{E}_1/\tilde{E}_0|^n)$, where $\tilde{E}_2$ is the second largest eigenvalue in magnitude. Consequently, a larger spectral gap between $\tilde{E}_0$ and $\tilde{E}_1$ generally improves convergence.

Equation (3.4) reveals that the n-th iteration can be viewed as applying a polynomial of $\mathcal{H}$ to the initial state:

$$|\psi^{(n)}\rangle = \frac{\mathrm{Poly}(\mathcal{H})|\psi^{(0)}\rangle}{||\,\mathrm{Poly}(\mathcal{H})|\psi^{(0)}\rangle||_2} = \frac{\mathcal{H}^n|\psi^{(0)}\rangle}{||\mathcal{H}^n|\psi^{(0)}\rangle||_2}. \tag{3.5}$$

This polynomial transformation can be implemented using the GQSP algorithm. We denote this method as quantum power iteration (QPI). In

particular, the powers of $\mathcal{H}$ can be expressed as linear combinations of Chebyshev polynomials. The general formula is as follows:

$$\begin{aligned}\mathcal{H}^n = {} & 2^{-n}\binom{n}{\frac{n}{2}} T_0(\mathcal{H}) \\ & + 2^{1-n} \sum_{\substack{k=1 \\ k\equiv n \,(\mathrm{mod}\, 2)}}^{n} \binom{n}{\frac{n-k}{2}} T_k(\mathcal{H}),\end{aligned} \tag{3.6}$$

where the first term contributes only when n is even, and $T_k(\mathcal{H})$ denotes the k-th Chebyshev polynomial of $\mathcal{H}$.

The main disadvantage of QPI and all other methods presented in this study is that its success probability can be very small. It equals the square of the denominator in Eq. (3.5). A higher probability of success can be achieved with a better initial guess wave function $|\psi^{(0)}\rangle$. However, AA is generally required, which deepens the circuit (see Sec. 3.4.1).

3.2.3 Quantum Power Lanczos Method

The power Lanczos (PL) method is a modification of the standard power iteration approach aimed at accelerating convergence by improving the initial guess for the eigenvector using Lanczos polynomials [179]. The Lanczos algorithm constructs a Krylov subspace and determines an approximate ground-state eigenvector by diagonalizing a tridiagonal projection of the original Hamiltonian. More rapid convergence can be achieved compared with a naive power iteration by treating this approximate ground state as the initial vector in subsequent power iterations.

The classical Lanczos algorithm is a well-known technique used to compute several of the largest absolute eigenvalues in a matrix. To accomplish this, the Hamiltonian $\mathcal{H}$ is projected onto the Krylov subspace $\mathcal{K}_d(\mathcal{H}, |\psi^{(0)}\rangle)$, which can be defined as follows:

$$\begin{aligned}\mathcal{K}_d(\mathcal{H}, |\psi^{(0)}\rangle) = \mathrm{span}\big\{ & |\psi^{(0)}\rangle,\, \mathcal{H}|\psi^{(0)}\rangle,\, \mathcal{H}^2|\psi^{(0)}\rangle, \\ & \ldots,\, \mathcal{H}^{d-1}|\psi^{(0)}\rangle\big\},\end{aligned} \tag{3.7}$$

where d is the dimension of the subspace and $|\psi^{(0)}\rangle$ is the initial state. Numerically stable implementations of the Lanczos algorithm use an orthonormal basis for this subspace. One standard method to build an orthonormal basis is through recurrence:

$$\begin{aligned}|\phi^{(0)}\rangle &= \mathcal{H}|\psi^{(0)}\rangle - a_0|\psi^{(0)}\rangle, \\ |\phi^{(n)}\rangle &= \mathcal{H}|\psi^{(n)}\rangle - a_n|\psi^{(n)}\rangle - b_n|\phi^{(n-1)}\rangle,\end{aligned} \tag{3.8}$$

where $|\phi^{(n)}\rangle = b_{n+1}|\psi^{(n+1)}\rangle$ form a new orthonormal basis, $a_n = \langle\psi^{(n)}|\mathcal{H}|\psi^{(n)}\rangle$ and $b_n = \sqrt{\langle\phi^{(n-1)}|\phi^{(n-1)}\rangle}$.

Thus, the matrix representation of the Hamiltonian becomes tridiagonal, drastically simplifying its partial diagonalization. The eigenvalues and eigenvectors of this tridiagonal subproblem approximate those of the original matrix $\mathcal{H}$. The approximation can be systematically improved by increasing the dimensions of the Krylov subspace d; however, the rate and quality of convergence also depend on the distribution of the eigenvalues and their relative gaps.

Although the standard Lanczos method provides a direct method to compute the tridiagonal representation and its eigendecomposition, it can be advantageous for treating the coefficients in the resulting Krylov vector expansion as variable parameters. In this approach, the energy $\langle\Psi_k^{(0)}|\mathcal{H}|\Psi_k^{(0)}\rangle$ is minimized in the trial state $|\Psi_k^{(0)}\rangle$ in the Krylov subspace. The key advantage of this approach lies in that it eliminates the need to precisely calculate the matrix elements of the effective Hamiltonian in the Krylov space. For the ground state problem, this minimization reproduces the usual Lanczos result; however, it can offer improved numerical stability.

For instance, consider a Krylov subspace with dimension $k+1$. A generic trial state then takes the following form:

$$\begin{aligned} |\Psi_1^{(0)}\rangle &= |\psi^{(0)}\rangle + C_1\mathcal{H}|\psi^{(0)}\rangle, \\ |\Psi_2^{(0)}\rangle &= |\psi^{(0)}\rangle + C_1\mathcal{H}|\psi^{(0)}\rangle + C_2\mathcal{H}^2|\psi^{(0)}\rangle, \end{aligned} \tag{3.9}$$

and so on. The coefficients C_i can be determined by solving the linear equations derived from the tridiagonal matrix or, as advocated here, by performing a variational minimization of the energy. This optimization involves only a small number of parameters since k is often small; therefore, it avoids most of the problems related to VQE.

Once the optimized Krylov (Lanczos) state $|\Psi_k^{(0)}\rangle$ is determined, it is used as a more refined starting point for power iteration; that is, instead of $|\psi^{(0)}\rangle$, we initialize our iterative procedure with $|\Psi_k^{(0)}\rangle$. The subsequent n-th power iteration state can then be written as follows:

$$\begin{aligned} |\Psi_k^{(n)}\rangle =& \frac{1}{N}\mathcal{H}^n|\Psi_k^{(0)}\rangle \\ =& \frac{1}{N}\mathcal{H}^n(1 + C_1\mathcal{H} + \cdots + C_k\mathcal{H}^k)|\psi^{(0)}\rangle, \end{aligned} \tag{3.10}$$

where N is a normalization factor. We will call this method quantum power Lanczos of order k (QPLk), specifying the degree of the Lanczos polynomial. The polynomial acting on the states can be naturally implemented within the GQSP framework. Here, GQSP demonstrates its advantage over conventional QSP, since the parametric C_i can become complex during optimization in order to improve convergence. Although the coefficients are ultimately real when calculated without considering the magnetic interaction, entering the complex plane in this manner helps finding a workaround to the minimum.

In practice, the QPL method proceeds in two main steps:

1. **Variational Lanczos state preparation.** Build $|\Psi_k^{(0)}\rangle$ using a polynomial in $\mathcal{H}$ (truncated at order k) using GQSP, then optimize the coefficients $\{C_i\}$ to minimize $\langle\Psi_k^{(0)}|\mathcal{H}|\Psi_k^{(0)}\rangle$. While one can solve for $\{C_i\}$ analytically using the Lanczos tridiagonalization, this requires computing expectation values of multiple powers of $\mathcal{H}$. In many quantum settings, a small variational optimization with a handful of coefficients is simpler.

2. **Quantum power iteration.** After obtaining $|\Psi_k^{(0)}\rangle$, perform the standard power iteration (see Eq. (3.10)) using this improved initial guess.

We further show that even a small Krylov subspace ($k = 1$ or $k = 2$) is often sufficient for providing a significantly better initial state. A larger k yields a more difficult optimization landscape but often offers insignificant gains in terms of accuracy. Thus, the QPL method can serve as a middle ground: it leverages a minimal Lanczos subspace to accelerate convergence in the power iteration by introducing only a few parameters. This approach can noticeably outperform pure power iteration for ground state computations, particularly when the spectrum of $\mathcal{H}$ has a large gap or when a small Krylov space already captures a significant fraction of the ground state component in $|\psi^{(0)}\rangle$.

Most other Lanczos-inspired algorithms for quantum computers rely on the explicit construction of the Lanczos tridiagonal matrix. For example, Kirby et al. [36] employ a Chebyshev-polynomial Krylov subspace, using a quantum device to evaluate matrix elements that are subsequently diagonalized on a classical computer. Baker [180] follows the standard Lanczos recursion to construct the Krylov basis and likewise performs explicit tridiagonalization.

In contrast, the QPL approach replaces explicit matrix construction and diagonalization with variational optimization of the Krylov-space coefficients. This is the central conceptual difference from other methods. By avoiding the calculation of effective Hamiltonian matrix elements in the Krylov basis, the method reduces numerical overhead and potential sources of instability. For ground state estimation, the variational minimization recovers the usual Lanczos result while providing improved numerical stability [179].

3.2.4 Quantum Inverse Iteration Method

Inverse iteration is another classical technique used to solve eigenvalue problems iteratively. Unlike the power iteration, which is naturally suited for finding the dominant eigenvalue and its corresponding eigenvector, the inverse iteration can be used to approximate any eigenpair, provided that an appropriate shift is chosen. Even though in our implementation it can only find the ground state, having a good enough energy approximation, the inverse iteration can converge much faster than the power iteration. Specifically, one starts with an initial guess of the eigenvector $|\psi^{(0)}\rangle$ and an approximate eigenvalue ϵ_0 close to the true eigenvalue $\tilde{E}_0$. The iterative step is given by

$$|\psi^{(n+1)}\rangle = \frac{(\mathcal{H} - \epsilon_0 I)^{-1}|\psi^{(n)}\rangle}{||(\mathcal{H} - \epsilon_0 I)^{-1}|\psi^{(n)}\rangle||_2}. \tag{3.11}$$

The convergence rate depends on how well ϵ_0 approximates the target eigenvalue $\tilde{E}_0$. Quantitatively, if $\tilde{E}_1$ is the eigenvalue closest to ϵ_0 other than $\tilde{E}_0$, the error drops as $O\left(\left|(\epsilon_0 - \tilde{E}_0)/(\epsilon_0 - \tilde{E}_1)\right|^n\right)$. Hence, choosing ϵ_0 closer to $\tilde{E}_0$ accelerates the convergence.

By the repeated application of $(\mathcal{H} - \epsilon_0 I)^{-1}$, we see that the n-th iterate can be written as follows:

$$|\psi^{(n)}\rangle = \frac{(\mathcal{H} - \epsilon_0 I)^{-n}|\psi^{(0)}\rangle}{||(\mathcal{H} - \epsilon_0 I)^{-n}|\psi^{(0)}\rangle||_2}. \tag{3.12}$$

Quantum inverse iteration (QII) requires a polynomial approximation of the operator $(\mathcal{H} - \epsilon_0 I)^{-n}$ to implement this step using the GQSP framework. The Taylor power series expansion for the negative powers of a shifted operator is given by

$$(\mathcal{H} - \epsilon_0 I)^{-n} = (-1)^n \sum_{k=0}^{\infty} \binom{n+k-1}{k} \frac{1}{\epsilon_0^{n+k}} \mathcal{H}^k. \tag{3.13}$$

In principle, an infinite series can be truncated to obtain a polynomial in $\mathcal{H}$. The radius of convergence for this expansion is $\|\mathcal{H}\|_2 < |\epsilon_0|$. In practice, this condition can be restrictive for excited state calculations because we require ϵ_0 to be as close as possible to $\tilde{E}_0$. We need ϵ_0 to be outside the spectral radius of $\mathcal{H}$, that is, $\|\mathcal{H}\|_2 < |\epsilon_0|$, to guarantee the convergence of the series. This is naturally satisfied if $\tilde{E}_0$ is the ground-state energy and ϵ_0 is chosen to be slightly smaller than $\tilde{E}_0$. In this scenario, one can often achieve a more rapid convergence compared to QPI, owing to the favorable shift.

Similarly to the power iteration method, a key step in QII involves mapping the truncated power series onto a polynomial expansion in terms of Chebyshev polynomials, which are naturally compatible with QSP. Once the polynomial representation is finalized, the same QSP techniques used for power iteration can be used to apply $(\mathcal{H} - \epsilon_0 I)^{-n}$ to a quantum state.

The full power of GQSP became apparent when a complex energy shift $z = \epsilon_0 + i\zeta$ was introduced. While we did not employ it in the examples below, this shift regularizes the pole at $E = \epsilon_0$ and enlarges the radius of convergence of the Taylor power series expansion by displacing the singularities off the real axis. Specifically, the shifted resolvent $(\mathcal{H} - zI)^{-1}$ remains finite at $E = \epsilon_0$ and, in the spectral representation, its squared magnitude takes the Lorentzian form $\left|(\mathcal{H} - zI)^{-1}\right|^2 = \left((\mathcal{H} - \epsilon_0 I)^2 + \zeta^2 I\right)^{-1}$. The width parameter ζ governs the trade-off between spectral selectivity (a narrow peak for a small ζ) and numerical stability or convergence speed (a larger ζ produces better conditioning). Finally, note that $(\mathcal{H} - zI)^{-1}$ is precisely the Green's function of $\mathcal{H}$ evaluated at z, establishing a direct connection to the contour integral and filter diagonalization eigensolvers and other Green's function methods.

3.2.5 Quantum Folded Spectrum Method

An iterative version of the folded spectrum method can be employed to compute the excited states within the GQSP framework. The key concept was the construction of a modified operator:

$$\mathcal{H}' \;=\; I \;-\; C\left(\mathcal{H} - \epsilon_m I\right)^2, \tag{3.14}$$

where ϵ_m is an estimate of the target eigenvalue and C is an appropriately chosen normalization constant. By focusing on $\left(\mathcal{H} - \epsilon_m I\right)^2$, one effectively "folds" the spectrum around ϵ_m. After scaling with C and shifting by identity, the eigenvalue closest to ϵ_m in the original spectrum becomes the largest eigenvalue of $\mathcal{H}'$; thus, this can be achieved via a power iteration scheme:

$$|\psi_m^{(n+1)}\rangle = \frac{(I - C(\mathcal{H} - \epsilon_m I)^2)|\psi_m^{(n)}\rangle}{||(I - C(\mathcal{H} - \epsilon_m I)^2)|\psi_m^{(n)}\rangle||_2}. \tag{3.15}$$

Repeating this step n times yields

$$|\psi_m^{(n)}\rangle = \frac{(I - C(\mathcal{H} - \epsilon_m I)^2)^n|\psi_m^{(0)}\rangle}{||(I - C(\mathcal{H} - \epsilon_m I)^2)^n|\psi_m^{(0)}\rangle||_2}. \tag{3.16}$$

Because $\left(I - C(\mathcal{H} - \epsilon_m I)^2\right)^n$ is a polynomial in $\mathcal{H}$, it is naturally suited for GQSP implementations. We call this approach the quantum folded spectrum method (QFSM). In practice, the Hamiltonian is first shifted by replacing $\mathcal{H}$ with $\mathcal{H}_{\text{shifted}} = \mathcal{H} - \epsilon_m I$. Specifically, shifting alters only the identity term in the LCU construction. After this shift, the polynomial

$$(1 - C\mathcal{H}_{\text{shifted}}^2)^n = \sum_{k=0}^{n} \binom{n}{k} (-C)^k \mathcal{H}_{\text{shifted}}^{2k} \tag{3.17}$$

can be applied exactly to the quantum state. However, the convergence in this folded spectrum scheme depends sensitively on the choice of ϵ_m, the initial guess $|\psi_m^{(0)}\rangle$, and the constant C. In typical implementations, C is chosen such that the shifted square spectrum lies between 0 and 1, which facilitates a faster convergence. However, the initial state can have an even more profound effect on the performance, as illustrated in the next section.

Table 3.1 summarizes the polynomial forms of all the methods discussed in this study.

Method	Poly($\mathcal{H}$)
QPI	$\mathcal{H}^n$
QPLk	$\mathcal{H}^n \sum_{i=0}^{k} C_i \mathcal{H}^i$
QII	$(-1)^n \sum_{k=0}^{\infty} \binom{n+k-1}{k} \frac{1}{\epsilon_0^{n+k}} \mathcal{H}^k$
QFSM	$\sum_{k=0}^{n} \binom{n}{k} (-C)^k \mathcal{H}_{\text{shifted}}^{2k}$, where $\mathcal{H}_{\text{shifted}} = \mathcal{H} - \epsilon_m I$

Table 3.1: Summary of the polynomial expansions that underlie each quantum iterative method (QPI, QPL, QII, and QFSM). Here, ϵ_m is the target eigenvalue approximation, and C is chosen to ensure the folded spectrum operator is suitably bounded.

3.3 Computational Details

Here, we provide a computational setup for proof-of-concept numerical simulations of our proposed methods. The reference CASCI calculations were performed using the PySCF package [181]. The JW transformed molecular Hamiltonians (see Sec. 2.2.4.1) used in our simulations were generated using the Quket package [182].

Except for the QFSM, all simulations employed the qubit tapering technique to reduce the number of qubits [183, 184]. For instance, the H_2 molecule in a $(2, 2)$ active space only required 1 qubit instead of 4, CH_2 in a $(6, 6)$ active space required 8 qubits instead of 12, and N_2 in a $(6, 6)$ space required 7 qubits instead of 12.

The GQSP circuit illustrated in Fig. 2.3 was implemented using the Qulacs simulator [185]. The angle estimation was performed using Prony's method [129], as described in detail in Appendix A1. The *capitalization* step required an additional qubit.

Block encoding was implemented using the LCU method only for the H_2 system, which required two additional qubits. To maintain the computational feasibility for the other systems, an explicit block encoding scheme was employed using a single additional qubit, as illustrated in Eq. (3.2). While in practical implementation on a quantum device the LCU is necessary, we used this technique for performance verification purposes. Further details are provided in Appendix A1. The complete implementation is available on our GitHub repository: https://github.com/mizukami-group/power-lanczos-data-FY2025.git.

3.4 Results and Discussion

3.4.1 Theoretical Scaling

Let $\mathcal{H}$ be a Hamiltonian with eigenpairs $\{(\tilde{E}_i, |\tilde{E}_i\rangle)\}_{i=0}^{n}$ ordered as $\tilde{E}_0 < \tilde{E}_1 < \cdots < \tilde{E}_n$. Let the spectral gap be $\tilde{\Delta} := \tilde{E}_1 - \tilde{E}_0 > 0$. We assume that $\tilde{E}_0$ and $\tilde{E}_1$ are negative and have the largest absolute values among the spectrum, and define their ratio as $\rho := \tilde{E}_1/\tilde{E}_0$, which satisfies $0 < \rho < 1$. Assume that $\tilde{\Delta} \ll |\tilde{E}_0|$ so that $\log(1/\rho) \approx \tilde{\Delta}/|\tilde{E}_0|$.

GQSP implements a polynomial filter $\mathrm{Poly}(\mathcal{H})$ on an initial state $|\psi^{(0)}\rangle$,

$$|\psi^{(0)}\rangle \longmapsto \frac{\mathrm{Poly}(\mathcal{H})|\psi^{(0)}\rangle}{\|\mathrm{Poly}(\mathcal{H})|\psi^{(0)}\rangle\|_2}, \tag{3.18}$$

with postselection success probability

$$\begin{aligned} \Pr[\mathrm{succ}] &= \left\|\mathrm{Poly}(\mathcal{H})|\psi^{(0)}\rangle\right\|_2^2 \\ &= \langle\psi^{(0)}|\,\mathrm{Poly}(\mathcal{H})^\dagger\,\mathrm{Poly}(\mathcal{H})\,|\psi^{(0)}\rangle. \end{aligned} \tag{3.19}$$

AA increases this probability to

$$\Pr^{(m)}[\mathrm{succ}] = \sin^2((2m+1)\theta) \tag{3.20}$$

with $\theta = \arcsin\sqrt{\Pr[\mathrm{succ}]}$. Choosing

$$m = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor \approx \frac{\pi}{4\sqrt{\Pr[\mathrm{succ}]}} - \frac{1}{2} \tag{3.21}$$

maximizes the success probability. Each Grover step uses the underlying block a constant number of times, so the cost scales by $(2m+1)$.

QPI. The initial guess wave function can be written as $|\psi^{(0)}\rangle = \sum_{i=0}^{n} a_i|\tilde{E}_i\rangle$, where we assume that the ground state contribution is lower bounded $|a_0| \geq \gamma$. After the k QPI steps, the ground state overlap is

$$\begin{aligned} p_k &= \frac{|a_0|^2}{|a_0|^2 + \sum_{i>0}|a_i|^2\,(\tilde{E}_i/\tilde{E}_0)^{2k}} \\ &\geq \frac{1}{1 + \frac{1-|a_0|^2}{|a_0|^2}\,\rho^{2k}}. \end{aligned} \tag{3.22}$$

The postselection success probability satisfies

$$\Pr^{(k)}[\mathrm{succ}] \;\geq\; \frac{\left(|a_0|^2 - (1-|a_0|^2)\rho^k\right)^2}{|a_0|^2 + (1-|a_0|^2)\rho^{2k}}, \tag{3.23}$$

so as $k \to \infty$ we have $p_k \to 1$ and $\Pr^{(k)}[\mathrm{succ}] \to |a_0|^2$. With AA and $|a_0| \geq \gamma$, the number of amplification steps can be chosen as $m_\star \approx \pi/(4\gamma) - 1/2$. Since degree k requires k queries to block encoding, this is multiplied by $(2m_\star + 1)$,

which gives about $\pi k/(2\gamma)$ of queries to block encoding to implement QPI with AA.

To reach an overlap with the ground state $(1-\varepsilon)$, a sufficient condition can be derived from Eq. (3.22):

$$\begin{aligned} k &\geq \frac{1}{2\log(1/\rho)}\Big[\log\frac{1}{\varepsilon}+2\log\frac{1}{\gamma}\Big] \\ &= \frac{1}{2\log(1/\rho)}\log\frac{1}{\gamma^2\varepsilon}. \end{aligned} \tag{3.24}$$

Using $\log(1/\rho)\approx\tilde{\Delta}/|\tilde{E}_0|$, the resulting query complexity is

$$\text{Queries}_{\text{QPI+AA}} = \frac{\pi}{4}\frac{1}{\tilde{\Delta}\,\gamma}\log\frac{1}{\gamma^2\varepsilon}. \tag{3.25}$$

This is a near-optimal scaling for state preparation without energy estimation [186].

The QPI prepared state can be used for QPE to obtain the ground state energy. We assume that the prepared state is almost exact, so no further AA is required for QPE. The qubitization implementation of QPE to estimate the energy up to precision $\tilde{\varepsilon}_E$ requires the following number of block encoding queries [132]:

$$\text{Queries}_{\text{QPE}} = \left\lceil\frac{\pi}{2\,\tilde{\varepsilon}_E}\right\rceil. \tag{3.26}$$

Thus, when the energy variables are returned from the rescaled Hamiltonian $\mathcal{H}=H/\lambda_{\text{LCU}}$ to the original units, the number of queries becomes

$$\begin{aligned} \text{Queries}_{\text{QPI+AA+QPE}} &= \frac{\pi}{4}\frac{\lambda_{\text{LCU}}}{\Delta\,\gamma}\log\frac{1}{\gamma^2\varepsilon} \\ &\quad+\left\lceil\frac{\pi\,\lambda_{\text{LCU}}}{2\,\varepsilon_E}\right\rceil. \end{aligned} \tag{3.27}$$

and if one sets $\varepsilon_E=\Theta(\Delta)$ to resolve the ground energy within the gap, the second term is $\Theta(\lambda_{\text{LCU}}/\Delta)$.

QPL. QPL optimizes the polynomial filter over the Krylov subspace. This modification does not influence the scaling, but changes the prefactor. We model the effect of a Lanczos step as increasing the ground state amplitude a_0 by a factor $\beta\geq 1$, so that the effective lower bound becomes $\beta\gamma$. As a result,

$$\text{Queries}_{\text{QPL+AA}} = \frac{\pi}{4}\frac{\lambda}{\Delta\,\gamma}\cdot\frac{1}{\beta}\log\frac{1}{\beta^2\gamma^2\varepsilon}. \tag{3.28}$$

In practice, this constant-factor improvement reduces iteration counts; in our numerical simulations shown in the next section, it saves up to roughly 50 iterations.

QII. When an estimate of the ground state energy is available, QII provides an efficient route for ground state preparation. Let $\mathcal{H}$ be a Hamiltonian

with eigenpairs $(\tilde{E}_i, |\tilde{E}_i\rangle)$ ordered as $\tilde{E}_0 < \tilde{E}_1 < \cdots < \tilde{E}_n$, and let ϵ_0 be an approximation to the ground state energy satisfying $|\epsilon_0 - \tilde{E}_0| \leq \tilde{\varepsilon}_E$.

Writing $|\psi^{(0)}\rangle = \sum_i a_i |\tilde{E}_i\rangle$ with $|a_0| \geq \gamma$, the overlap with the ground state after k iterations is

$$p_k = \frac{|a_0|^2}{|a_0|^2 + \sum_{i>0} |a_i|^2 \left| \frac{\tilde{E}_0 - \epsilon_0}{\tilde{E}_i - \epsilon_0} \right|^{2k}}. \tag{3.29}$$

Let

$$\eta_{\text{inv}} := \max_{i>0} \left| \frac{\tilde{E}_0 - \epsilon_0}{\tilde{E}_i - \epsilon_0} \right| \leq \frac{\tilde{\varepsilon}_E}{\Delta - \tilde{\varepsilon}_E}, \qquad \tilde{\Delta} = \tilde{E}_1 - \tilde{E}_0, \tag{3.30}$$

assuming $\tilde{\varepsilon}_E < \tilde{\Delta}/2$.

Then the lower bound of the overlap becomes

$$p_k \geq \frac{1}{1 + \frac{1-|a_0|^2}{|a_0|^2} \eta_{\text{inv}}^{2k}}. \tag{3.31}$$

The success probability for QII is also analogous to the QPI formula with replacement of ρ with η_{inv} (see Eq. (3.23)).

To achieve overlap $(1 - \varepsilon)$, the sufficient number of steps is

$$k \geq \frac{1}{2\log(1/\eta_{\text{inv}})} \log \frac{1}{\gamma^2 \varepsilon}. \tag{3.32}$$

Thus, the number of required iterations decreases rapidly as the energy estimate $\tilde{\varepsilon}_E$ approaches zero.

In the GQSP framework, the inverse expression $f_k(\mathcal{H}) = (\mathcal{H} - \epsilon_0 I)^{-k}$ is implemented by a polynomial approximation $P_d(\mathcal{H})$ of degree d. Since $\mathcal{H}$ is already the rescaled block-encoded Hamiltonian with spectrum contained in $[-1, 1]$, the only singularities of $f_k(x)$ arise at points where $x = \epsilon_0$.

Let $\tilde{g} := \tilde{\Delta} - \tilde{\varepsilon}_E$. Under the assumption $\tilde{\varepsilon}_E < \tilde{\Delta}/2$, the excited state eigenvalues $\tilde{E}_i$ satisfy $|\tilde{E}_i - \epsilon_0| \geq \tilde{g}$, so $f_k(x)$ is analytic on the subset of the spectrum obeying $|x - \epsilon_0| \geq \tilde{g}$. In other words, the nearest singularity of $f_k(x)$ lies at distance $\tilde{g}$ from the spectral interval on which the approximation must hold.

For analytic functions on $[-1, 1]$ with nearest singularity at distance $\tilde{g}$, standard Chebyshev/Bernstein–ellipse bounds imply the existence of a polynomial P_d satisfying

$$\sup_{x:\, |x-\epsilon_0| \geq \tilde{g}} \left| f_k(x) - P_d(x) \right| \leq \delta. \tag{3.33}$$

whenever

$$d = \Theta\left(\frac{1}{\tilde{g}} \left[k + \log \frac{1}{\delta} \right] \right). \tag{3.34}$$

This equation gives the polynomial degree required to achieve operator norm error δ in a single-shot implementation of $f_k(\mathcal{H})$.

A single round of AA boosts postselection, costing a factor of $(2m_\star + 1)$, where $m_\star \approx \pi/(4\gamma) - 1/2$. Choosing $\delta = \Theta(\varepsilon)$, so one application of polynomial, and returning back to the original units, gives the total query complexity

$$\text{Queries}_{\text{QII+AA}} \;=\; \Theta\!\left(\frac{1}{\gamma}\cdot\frac{\lambda_{\text{LCU}}}{\Delta-\varepsilon_E}\left[k+\log\frac{1}{\varepsilon}\right]\right). \tag{3.35}$$

Substituting (3.32) for k,

$$\begin{aligned}\text{Queries}_{\text{QII+AA}} \;=\; \tilde{O}\Bigg(&\frac{\lambda_{\text{LCU}}}{\gamma\,(\Delta-\varepsilon_E)}\\ &\times\frac{\log\!\big(1/(\gamma^2\varepsilon)\big)}{\log\!\big((\Delta-\varepsilon_E)/\varepsilon_E\big)}\Bigg).\end{aligned} \tag{3.36}$$

If $\epsilon_0 = \tilde{E}_0$ (i.e., $\varepsilon_E = 0$), one can take $k = 1$ in exact arithmetic; practically, the polynomial degree $d = \Theta((\lambda_{\text{LCU}}/(\Delta))\log(1/\varepsilon))$ then dominates.

QFSM. To prepare the excited state, QFSM can be used. Given a target excited level E_j and an energy estimate ϵ_j with $|\epsilon_j - \tilde{E}_j| \le \tilde{\varepsilon}_E$, QFSM applies the polynomial filter $F_k(\mathcal{H}) \;=\; \big(I - C\,(\mathcal{H}-\epsilon_j I)^2\big)^k$, which concentrates weight near the eigenvalue closest to ϵ_j. Choose C so that $0 \le 1 - C(\tilde{E}_i - \epsilon_j)^2 \le 1$ on the spectrum; a sufficient choice is

$$0 < C \le \frac{1}{\tilde{R}^2}, \qquad \tilde{R} \;\ge\; \max_i |\tilde{E}_i - \epsilon_j|. \tag{3.37}$$

Write $|\psi^{(0)}\rangle = \sum_{i=0}^{n} a_i |\tilde{E}_i\rangle$ and assume $|a_j| \ge \gamma$ for the target state $|\tilde{E}_j\rangle$. Let

$$\begin{aligned}\tilde{g} \;&:=\; \min_{i\neq j}|\tilde{E}_i - \epsilon_j| \;\ge\; \tilde{\Delta}_j - \tilde{\varepsilon}_E,\\ \tilde{\Delta}_j &:= \min\{\tilde{E}_j - \tilde{E}_{j-1},\, \tilde{E}_{j+1} - \tilde{E}_j\},\end{aligned} \tag{3.38}$$

so $\tilde{g} > 0$ whenever $\tilde{\varepsilon}_E < \tilde{\Delta}_j$ and $\tilde{E}_j$ is the unique closest eigenvalue to ϵ_j if $\tilde{\varepsilon}_E < \tilde{\Delta}_j/2$.

Define the contraction factor

$$\begin{aligned}\eta_{\text{fsm}} \;&:=\; \frac{\max_{i\neq j}\,|1 - C(\tilde{E}_i - \epsilon_j)^2|}{|1 - C(\tilde{E}_j - \epsilon_j)^2|}\\ &\le\; \frac{1 - C\tilde{g}^2}{|1 - C\tilde{\varepsilon}_E^2|} \;<\; 1.\end{aligned} \tag{3.39}$$

After applying and normalizing $F_k(\mathcal{H})$, the ground-state–style overlap bound carries over to the target excited state:

$$p_k \;:=\; \frac{|\langle\tilde{E}_j|F_k(\mathcal{H})|\psi^{(0)}\rangle|^2}{\|F_k(\mathcal{H})|\psi^{(0)}\rangle\|_2^2} \;\ge\; \frac{1}{1 + \frac{1-|a_j|^2}{|a_j|^2}\,\eta_{\text{fsm}}^{2k}}. \tag{3.40}$$

A lower bound on the postselection success probability mirrors the power iteration case:

$$\Pr^{(k)}[\text{succ}] \;\geq\; \frac{\left(|a_j|^2 - (1-|a_j|^2)\,\eta_{\text{fsm}}^{k}\right)^2}{|a_j|^2 + (1-|a_j|^2)\,\eta_{\text{fsm}}^{2k}} \;\xrightarrow{k\to\infty}\; |a_j|^2. \tag{3.41}$$

To reach overlap $1-\varepsilon$, it suffices that

$$k \;\geq\; \frac{1}{2\log(1/\eta_{\text{fsm}})}\;\log\frac{1}{\gamma^2\varepsilon}. \tag{3.42}$$

With the conservative choice $C = 1/\tilde{R}^2$ and in the regime $C\tilde{g}^2$, $C\tilde{\varepsilon}_E^2 \ll 1$, we have $\log(1/\eta_{\text{fsm}}) \approx C\,(g^2 - \varepsilon_E^2)$, yielding the convenient estimate

$$k \;\gtrsim\; \frac{\tilde{R}^2}{2\,(\tilde{g}^2 - \tilde{\varepsilon}_E^2)}\;\log\frac{1}{\gamma^2\varepsilon}. \tag{3.43}$$

Implement $F_k(\mathcal{H})$ as a single QSVT polynomial of degree $d = 2k$ in $\mathcal{H} = H/\lambda_{\text{LCU}}$. As in previous sections, AA with $m_\star \simeq \pi/(4\gamma) - 1/2$ boosts postselection, multiplying the polynomial cost by $(2m_\star + 1)$. The state-preparation query complexity is therefore

$$\text{Queries}_{\text{QFSM+AA}} \;=\; \Theta\!\left(\frac{\pi}{4\gamma}\cdot d\right) \;=\; \Theta\!\left(\frac{\pi}{2\gamma}\cdot k\right), \tag{3.44}$$

which, combined with (3.42), and returning back to the original units gives

$$\begin{aligned} \text{Queries}_{\text{QFSM+AA}} \;&=\; \Theta\!\left(\frac{1}{\gamma}\cdot\frac{\log\!\big(1/(\gamma^2\varepsilon)\big)}{\log\!\big(1/\eta_{\text{fsm}}\big)}\right) \\ &\approx\; \tilde{O}\!\left(\frac{R^2}{\gamma\,(g^2-\varepsilon_E^2)}\;\log\frac{1}{\gamma^2\varepsilon}\right), \end{aligned} \tag{3.45}$$

hiding constant and secondary logarithmic factors. Using $R \lesssim 2\lambda_{\text{LCU}}$ and $g \sim \Delta_j$ one can get estimation:

$$\text{Queries}_{\text{QFSM+AA}} \;\approx\; \tilde{O}\!\left(\frac{\lambda_{\text{LCU}}^2}{\gamma\,(\Delta_j^2-\varepsilon_E^2)}\;\log\frac{1}{\gamma^2\varepsilon}\right). \tag{3.46}$$

Unlike QPI, QPL, and QII, which target the ground state, QFSM prepares excited states using only an energy estimate ϵ_j. Its scaling in the spectral parameters is generally not optimal (it carries a $\lambda_{\text{LCU}}^2/(\Delta_j^2 - \epsilon_E^2)$ factor rather than the first powers of gap and norm of the Hamiltonian), but it can be practical. All query scalings and comparisons with existing algorithms are summarized in Tab. 3.2.

The query complexity table indicates that the GQSP–based filters proposed here are competitive with near-optimal ground state algorithms while remaining simple to implement and tune. QPL+AA achieves ground state preparation with

complexity $\tilde{O}(\lambda/(\Delta\,\gamma))$ without any prior energy estimate or a costly energy search over shifts. This makes QPI a practical default when E_0 is unknown and one wants a guaranteed, monotone improvement in overlap at a cost that matches the standard $1/(\gamma\Delta)$ scaling up to constants. Lanczos step reduces the iteration count by a constant factor $\beta > 1$ relative to QPI while preserving the same asymptotic. In practice, this translates into tangible depth savings (tens of iterations in our numerics) with no extra oracles beyond those already required for QPI.

When a rough energy is available, inverse filtering wins. If an a priori estimate satisfies $\epsilon_E < \Delta/2$, the single-shot inverse polynomial QII achieves $\tilde{O}\big(\frac{\lambda}{\gamma(\Delta-\epsilon_E)} \cdot \frac{\log\big(1/(\gamma^2\varepsilon)\big)}{\log\frac{\Delta-\epsilon_E}{\epsilon_E}}\big)$ queries. As ϵ_E improves, the effective gap $g = \Delta - \epsilon_E$ grows and the cost drops toward the $1/(\gamma\Delta)$ frontier.

QFSM extends the toolbox to excited states using the folded spectrum filter $(I - C(H - EI)^2)^k$. Although its dependence on spectral parameters is not optimal ($\tilde{O}(\lambda^2/((\Delta^2 - \varepsilon_E^2)\,\gamma))$), it is operationally attractive. For moderate target fidelities, good E (within the local gap), and re-use across multiple excitations, this simplicity can outweigh the non-optimal scaling.

The methods discussed above highlight distinct operating regimes. QPI and its Lanczos-enhanced variant QPL require no prior spectral information and provide reliable ground state preparation with modest depth. When an energy estimate with error $\varepsilon_E \leq \Delta/2$ is available, QII becomes the method of choice: its contraction factor improves rapidly as the estimate sharpens, and its overall query complexity approaches the near-optimal gap-inverse scaling. For preparing specific excited states, QFSM offers a complementary strategy. Provided an energy guess lies inside the local spectral gap of the target level, its polynomial filter selectively amplifies the desired component. Together, these methods form a practical toolkit for state preparation across the range of situations encountered in Hamiltonian simulation and quantum chemistry.

Method / task	**Queries to block encoding**
This work (GQSP polynomials with AA)	
QPI + AA (ground state, overlap $1-\varepsilon$)	$\frac{\pi}{4}\,\frac{\lambda}{\Delta\,\gamma}\,\log\frac{1}{\gamma^2\varepsilon}$
QPI + AA + QPE (energy to precision ε_E)	$\frac{\pi}{4}\,\frac{\lambda}{\Delta\,\gamma}\,\log\frac{1}{\gamma^2\varepsilon}\;+\;\left\lceil\frac{\pi\,\lambda}{2\,\varepsilon_E}\right\rceil$
QPL + AA (ground state, overlap $1-\varepsilon$)	$\frac{1}{\beta}\cdot\frac{\pi}{4}\,\frac{\lambda}{\Delta\,\gamma}\,\log\frac{1}{\beta^2\gamma^2\varepsilon}\quad(\beta>1)$
QII + AA (ground state, overlap $1-\varepsilon$)	$\widetilde{\mathcal{O}}\!\left(\frac{\lambda}{\gamma\,(\Delta-\varepsilon_E)}\,\frac{\log\frac{1}{\gamma^2\varepsilon}}{\log\frac{\Delta-\varepsilon_E}{\varepsilon_E}}\right)$
	$\Theta\!\left(\frac{\lambda}{\gamma\Delta}\log\frac{1}{\varepsilon}\right)\quad(\varepsilon_E=0)$
QFSM + AA (excited state, overlap $1-\varepsilon$)	$\widetilde{\mathcal{O}}\!\left(\frac{\lambda^2}{\gamma\,(\Delta^2-\varepsilon_E^2)}\,\log\frac{1}{\gamma^2\varepsilon}\right)$;
Literature baselines	
Lin–Tong (2020): ground state, E_0 bound known	$O\!\left(\frac{\lambda}{\gamma\,\Delta}\,\log\frac{1}{\varepsilon}\right)$ [186]
Lin–Tong (2020): ground energy to precision ε_E	$\tilde{O}\!\left(\frac{\lambda}{\gamma\,\varepsilon_E}\,\log\frac{1}{\vartheta}\right)$ [186]
Lin–Tong (2020): ground state, E_0 bound unknown	$\tilde{O}\!\left(\frac{\lambda}{\gamma\,\Delta}\,\log\frac{1}{\vartheta\,\varepsilon}\right)$ [186]
Ge–Tura–Cirac (2019): ground state, E_0 well known	$\tilde{O}\!\left(\frac{\lambda}{\gamma\,\Delta}\right)$ [163]
Ge–Tura–Cirac (2019): ground energy to precision ε_E	$\tilde{O}\!\left(\frac{\lambda^{3/2}}{\gamma\,\varepsilon_E^{3/2}}\right)$ [163]
Ge–Tura–Cirac (2019): ground state, E_0 bound unknown	$\tilde{O}\!\left(\frac{\lambda^{3/2}}{\gamma\,\Delta^{3/2}}\right)$ [163]

Table 3.2: Query complexities, stated in a unified notation. Here λ denotes the block encoding normalization (corresponding to α in Refs. [186, 163]), Δ the state gap, γ a lower bound on the initial overlap, ε the state preparation error, ε_E the energy estimation precision (denoted h in Ref. [186]), and ϑ a failure probability parameter. Our methods use GQSP polynomials combined with AA. For our bounds, $\widetilde{\mathcal{O}}(\cdot)$ hides only polylogarithmic factors in $1/\Delta$, $1/\varepsilon$, $1/\varepsilon_E$ and $1/\gamma$. In the literature block we retain the original authors' asymptotic notation $O(\cdot)$, $\tilde{O}(\cdot)$ as written in Refs. [186, 163], which hide different logarithmic dependencies.

3.4.2 Numerical Comparison of QPI and QPL

First, we performed numerical tests of the QPI and QPL methods. Because QPL can be viewed as a direct modification of QPI, we discuss both in this section. Following the notation introduced in Eq. (3.10), we denote the QPL method of the Lanczos polynomial degree k by QPLk. In this notation, the QPI method corresponds to QPL0.

Figure 3.2 illustrates the PECs for the hydrogen molecule using the cc-pVDZ basis set with 2 electrons in 2 active orbitals. The HF wave function served as the initial guess.

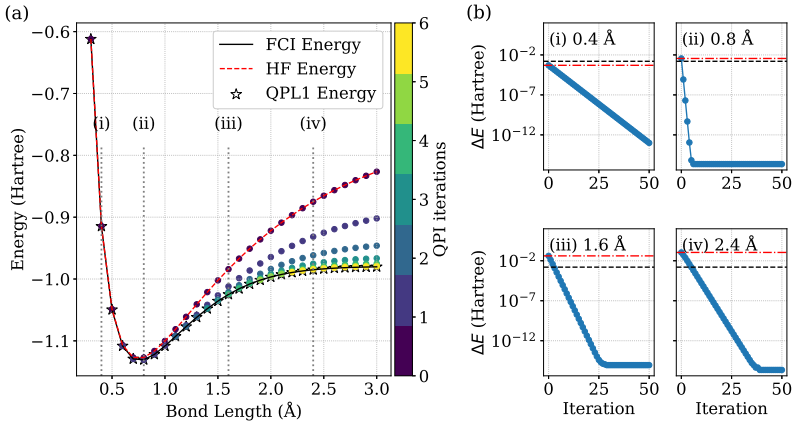


Figure 3.2: (a) PECs for the hydrogen molecule obtained with the QPI method (QPL0). The basis set is cc-pVDZ, and the active space is $(2, 2)$. Star markers indicate the corresponding QPL1 energies immediately after the variational optimization step.
(b) The right panels show the energy deviation from CASCI$(2, 2)$ at each iteration for the first 50 iterations. The black dashed line indicates chemical accuracy (1 kcal/mol), and the red dot-dashed line is the HF energy.

We observe that the QPI (QPL0) method achieves the CASCI$(2, 2)$ energy within chemical accuracy in no more than 6 iterations, even when starting from the HF wave function. For longer bond lengths, more iterations are required to achieve chemical accuracy, which reflects the increased multireference character of the wavefunction. This behavior became more pronounced as the bond was stretched, indicating that the method may converge more slowly in systems with stronger static correlation.

The convergence speed of the power method depends on the ratio E_1/E_0, where E_0 and E_1 denote the ground state and first excited state energies, respectively. The panels on the right side of Fig. 3.2 show that the convergence rates are approximately proportional to $\log_{10}(E_1/E_0)$, with a proportionality coefficient between 1.85 and 2 for different bond lengths. This behavior is

fully consistent with the classical power iteration algorithm, whose convergence is governed by the same ratio between largest eigenvalues. The quantum implementation therefore reproduces the expected classical scaling.

Next, we examine QPL1, where the final wave function after optimizing the Lanczos polynomial at degree 1 reads

$$|\mathrm{QPL1}\rangle = \frac{1}{N}(a_0 + a_1\mathcal{H})|\mathrm{HF}\rangle. \tag{3.47}$$

where N is a normalization factor. Remarkably, this ansatz can match the CASCI energy immediately after the polynomial optimization and before any additional power method steps. This means that QPL1 can reduce the number of block encoding queries. For example, at a bond length of 3 Å, QPI requires 6 queries, whereas QPL1 achieves the same result with only one.

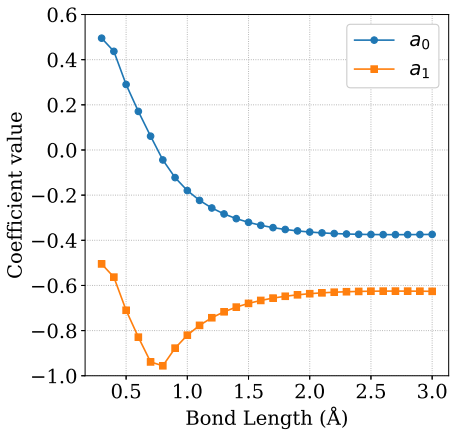


Figure 3.3: Bond length dependence of the Lanczos polynomial coefficients for the hydrogen molecule. Coefficients are obtained after the variational optimization of the Lanczos polynomial (see Eq. (3.47)).

Figure 3.3 illustrates how the Lanczos polynomial coefficients a_0 and a_1 vary with bond length. We observe that a_0 changes sign around equilibrium geometry of H_2. Because a_0 represents the HF component of the wave function, its sign change reflects a decreasing HF contribution as the bond stretches. Meanwhile, a_1, which weighs $\mathcal{H}|\mathrm{HF}\rangle$, reaches its maximum near the equilibrium geometry, and then decreases. This behavior is expected because $\mathcal{H}|\mathrm{HF}\rangle$ is not orthogonal to $|\mathrm{HF}\rangle$; therefore, the relative amplitudes of these two contributions can offset each other.

Subsequently, we computed the PECs of the singlet and triplet states of the CH_2 radical. A cc-pVDZ basis set with 6 electrons in 6 active orbitals was employed, and both C–H bonds were stretched symmetrically while keeping the H–C–H angle fixed at 101.89°. For the singlet state, we used the restricted Hartree-Fock (RHF) wave function as an initial guess, and for the triplet state,

we used the restricted open-shell Hartree-Fock (ROHF) wave function. The results are presented in Fig. 3.4.

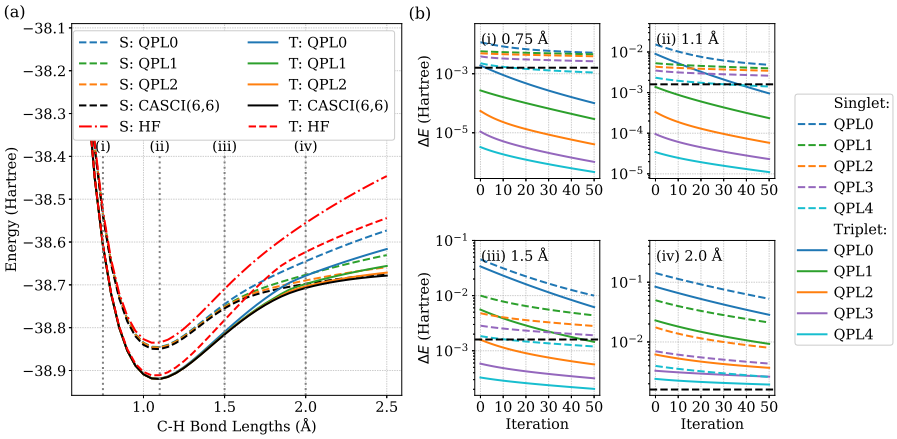


Figure 3.4: (a) Potential energy curves for singlet (S) and triplet (T) CH_2 obtained using QPL at various polynomial degrees ($k = 0, 1, 2$). All QPLk curves correspond to energies after 50 iterations. The basis set is cc-pVDZ with an active space of $(6, 6)$. (b) The right panels show deviations from CASCI$(6, 6)$ for the first 50 iterations of the power method using QPL with $k = 0, 1, 2, 3, 4$. The dashed black lines indicate chemical accuracy (1 kcal/mol). Solid lines correspond to the triplet state and dashed lines to the singlet state.

Notably, the triplet state lies below the singlet state around the equilibrium geometry of CH_2. For both spin states, QPLk yielded a marked improvement over the respective HF initial guesses. However, unlike the hydrogen molecule, QPL0 (QPI) failed to achieve chemical accuracy after 50 iterations for bond lengths beyond approximately 1.5 Å in the triplet and for all bond lengths in the singlet. Hence, we need to use Lanczos polynomials to reach chemical accuracy. As the bond stretches, this requirement becomes more pronounced, which is consistent with the growing multireference character of the wave function.

There is a consistent advantage from using higher-degree Lanczos polynomials. In our simulations, each additional unit of polynomial degree typically replaces up to 50 power iteration steps, in line with the constant-factor improvement predicted for QPL over QPI. In practice, however, optimizing polynomial coefficients becomes increasingly difficult at large k. For this reason, moderate polynomial degrees are generally preferable: pushing QPL to very high degrees risks turning it into a VQE-like procedure with hundreds of tunable parameters and diminishing practical benefit. This observation is consistent with those reported in Ref. [179].

Finally, the singlet state, with its lower spin, displays a stronger multireference character than the triplet state, generally requiring higher degree Lanczos polynomials to achieve comparable accuracy.

As a further test, we applied QPL to a nitrogen molecule, which is known to

have a highly multireference ground state character in stretched geometries. The cc-pVDZ basis set with 6 electrons in 6 active orbitals was used. Two different initial guesses were tested:

1. The RHF wave function.

2. A partially converged VQE wave function with a UCC method (see Sec. 2.2.5.1) with double excitations only (UCCD) ansatz (after only 2 iterations), which is thus multireference but not fully optimized.

The VQE optimization was halted after two optimization steps, without reaching a convergence threshold, so the resulting state is an intermediate, not fully optimized, multireference wave function. The corresponding PECs are presented in Fig. 3.5.

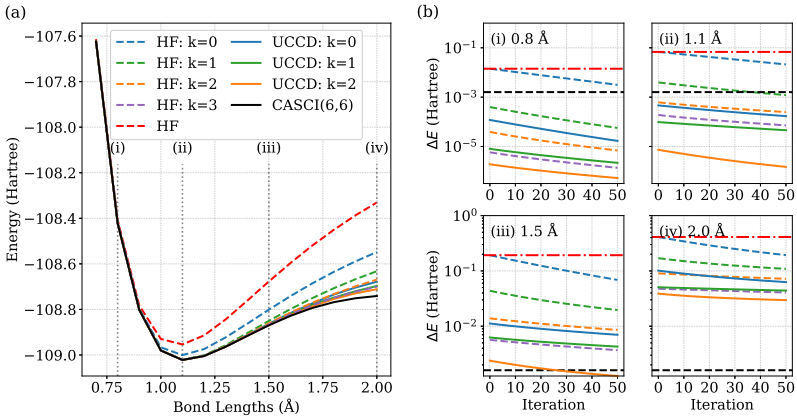


Figure 3.5: (a) Potential energy curves for the nitrogen molecule obtained using QPL with different polynomial degrees ($k = 0, 1, 2, 3$) and initial guesses (HF in dashed, and partially converged UCCD in solid). All QPLk curves correspond to energies after 50 iterations. The basis set is cc-pVDZ, and the active space is $(6, 6)$. (b) The right panels display energy differences from CASCI$(6, 6)$ over the first 50 iterations of the power method. The black dashed line marks chemical accuracy (1 kcal/mol), and the red dot-dashed line is the HF energy.

For the equilibrium geometry, QPL0/HF cannot match the CASCI$(6, 6)$ energy within 50 iterations, but QPL1/HF succeeds. However, even at longer bond lengths (1.5 Å and 2.0 Å), QPL3/HF does not achieve chemical accuracy. These results reiterate that highly multireference systems pose convergence challenges for QPI and QPL, although higher-degree polynomials are often helpful. However, each Lanczos polynomial of degree k achieves a better result after a single iteration than a polynomial of degree $k - 1$ does after 50 iterations. This reflects the constant advantage provided by QPL.

A more effective approach is to use a better initial guess than HF. Indeed, starting from the partially converged UCCD wave function significantly improves the convergence, allowing us to achieve chemical accuracy at longer bond lengths. However, for stretched geometries (for example, 2.0 Å in Fig. 3.5), although this improved initial guess is insufficient without further refinement. The constant advantage offered by QPL is evident here as well. At a bond length of 1.1 Å, where QPL2/UCCD after a single iteration yields an error nearly an order of magnitude lower than QPL1/UCCD after 50 iterations.

In addition, these results demonstrate that QPI and QPL are far more effective in refining a reasonably good reference wave function that captures most of the static correlation rather than attempting to recover both static and dynamic correlation from a simple HF state. This use case motivated for the original power Lanczos approach in Ref. [179], in which the method was used to add a dynamic correlation to Monte Carlo wave functions that already contain strong static correlation.

3.4.3 Numerical Results for QII

In this section,the performance of the GQSP-based QII algorithm is described. First, we present the PEC calculation as an example. We then compared our QII implementation with two existing methods, I-Iter and Q-Inv, which rely on linear combinations of time evolution operators [171].

We begin by examining the dissociation curve of the hydrogen molecule. The calculations used the cc-pVDZ basis set with 2 electrons in 2 active orbitals. In QII, the inverse operator is approximated by truncating the Taylor polynomial to a certain degree n. The results for $n = 2$, $n = 5$, and $n = 10$ are presented. It is important to note that the number of block encoding queries equals the polynomial degree n and does not depend on the iteration count. The number of iterations only determines the shape of the approximated function. QII also requires an initial estimate of the ground state energy, which we consider to be the HF energy in all calculations. Figure 3.6 shows the resulting PECs.

We observed that the ability to reach the CASCI$(2, 2)$ energy within chemical accuracy depends on the degree n of the truncated polynomial. For instance, at a bond length of 1.6 Å, QII$(n = 2)$ converges to a value above the chemical accuracy; similarly, QII$(n = 5)$ struggles at 2.4 Å. Conversely, for a sufficiently large n $(n = 50)$, the method achieves chemical accuracy in only one iteration for all bond lengths. This rapid convergence relative to the QPI method reflects the well-known classical advantage of inverse iteration when a good approximation for the ground state energy estimation is available. In other words, the faster convergence is not specific to the quantum adaptation but follows the traditional hierarchy between inverse and power iterations.

Next, we compare our GQSP implementation of QII with the I-Iter and Q-Inv algorithms, which evaluate H^{-1} or H^{-k} by numerically integrating the

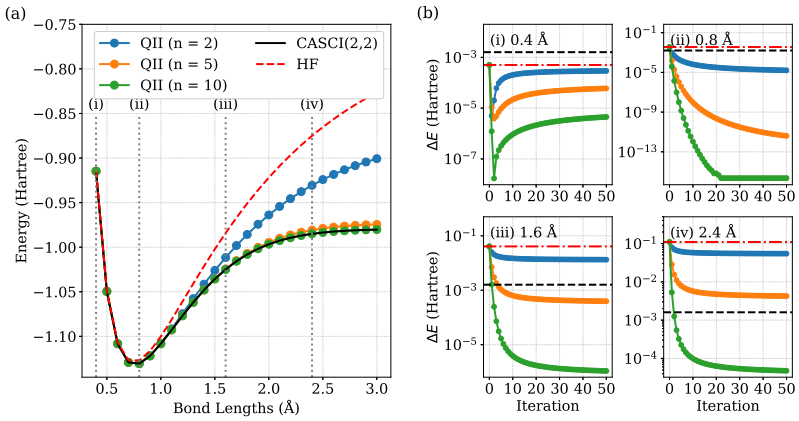


Figure 3.6: (a) Potential energy curves for the hydrogen molecule obtained using QII with different polynomial truncation degrees ($n = 2, 5, 10$). The HF energy was used as an energy shift. All QII curves correspond to energies after 50 iterations. The basis set is cc-pVDZ, and the active space is $(2, 2)$. (b) The right panels show the deviation from the CASCI$(2, 2)$ energy over the first 50 iterations. The black dashed lines mark chemical accuracy (1 kcal/mol), and the red dot-dashed lines represent the HF energy.

time evolution operators [171]. The I-Iter method approximates H^{-1} using

$$H^{-1} = \frac{i}{\sqrt{2\pi}} \int_0^\infty \mathrm{d}y \int_{-\infty}^\infty \mathrm{d}z \, z \, e^{-z^2/2} e^{-iyzH}. \tag{3.48}$$

and repeated application of H^{-1} (k times) yields H^{-k}. The Q-Inv method proposed by Kyriienko [170] directly approximates H^{-k} by

$$H^{-k} = \frac{iN_k}{\sqrt{2\pi}} \int_0^\infty \mathrm{d}y \int_{-\infty}^\infty \mathrm{d}z \, z \, y^{k-1} e^{-z^2/2} e^{-iyzH}, \tag{3.49}$$

where N_k is the normalization constant. Both I-Iter and Q-Inv use the Gauss-Legendre rule to calculate the integrals, and the precision of this procedure depends on the order of the Legendre polynomials used. In Ref. [171], comparisons were made for two orders n_y for the integration of the variable y, in particular $n_y = 1$ and $n_y = n_{\mathrm{eq}}$, where n_{eq} is the order sufficient to achieve energy convergence. The hybrid Q-Inv+I-Iter strategy was also introduced. In this approach, the Q-Inv algorithm is first applied up to the iteration order k at which its numerical error reaches a minimum, since the accuracy of the direct H^{-k} approximation deteriorates for larger k due to limitations in the integral discretization. Once this optimal point is reached, a small number of subsequent I-Iter steps are performed to further suppress the remaining error. This hybrid procedure leverages the rapid early-stage convergence of Q-Inv while avoiding

its high-k breakdown, and it reduces the total cost of inverse iterations because only a few applications of the I-Iter update are required.

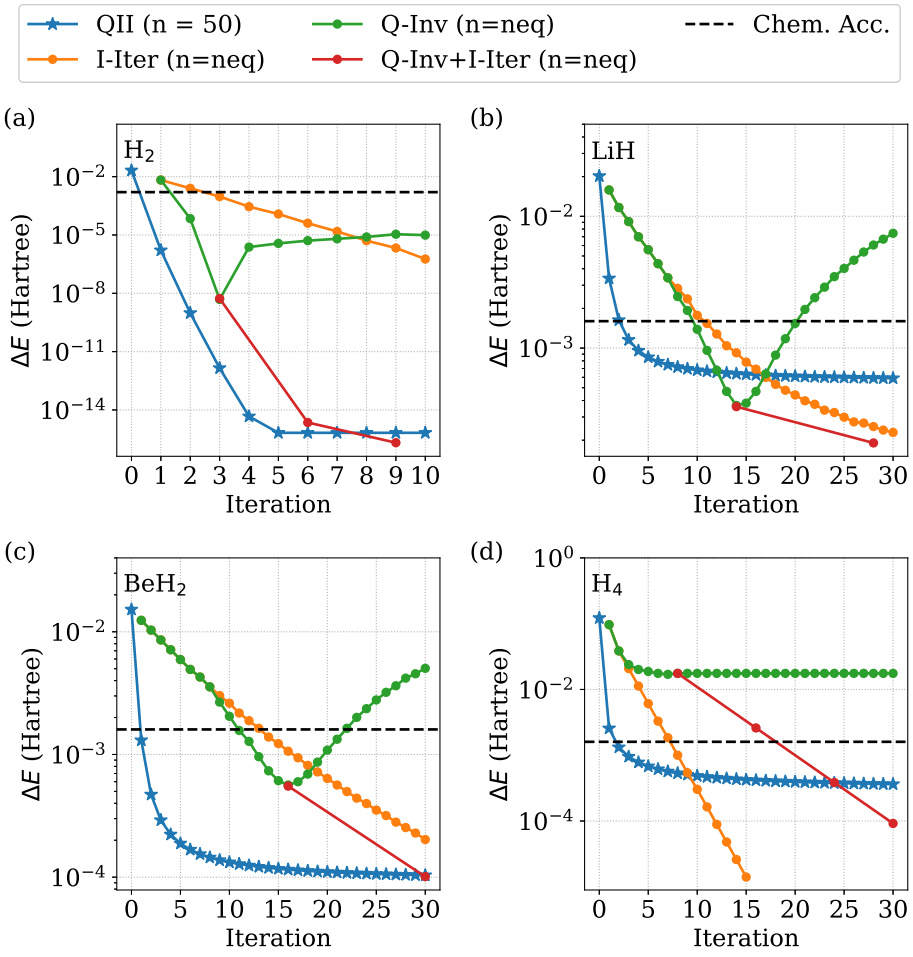


Figure 3.7: Comparison of QII, I-Iter, and Q-Inv [171] for H_2, LiH, BeH_2, and H_4. The deviation from CASCI energies as a function of the iteration number is shown. For QII, $n = 50$ refers to the Taylor polynomial truncation degree, while for I-Iter and Q-Inv, n_y is the Gauss-Legendre polynomial order. n_{eq} is the order at which energy convergence is reached.

Analogous to the reference paper, we investigated four molecules: H_2, LiH, BeH_2, and a square H_4, using a STO-6G basis set. For H_2 and H_4 full active spaces were used $(2, 2)$ and $(4, 4)$, respectively. For LiH, 2 electrons were used in 5 orbitals, while 4 electrons were used in 5 orbitals for BeH_2. All calculations were single-point with geometries of 0.75 Å for H_2, 1.23 Å for the H–H distance in H_4, 1.6 Å for the Li–H bond in LiH, and 1.326 Å for the Be–H bond in BeH_2.

Figure 3.7 compares only QII with truncation $n = 50$ with Q-Inv and I-Iter with $n_y = n_{\text{eq}}$. The results for different truncations are discussed in Appendix A2. The iteration number used for Q-Inv, I-Iter, and the hybrid Q-Inv+I-Iter refers to the power k in the approximate application of H^{-k}. For I-Iter, this corresponds to applying the numerically integrated approximation of H^{-1} repeatedly k times, whereas Q-Inv constructs a direct approximation of H^{-k} through its Fourier-integral representation. For QII, the iteration number only reflects the shape of the approximated function, while the number of queries remains constant at 50. When comparing the number of qubits between methods, we assumed LCU block encoding. We also include qubit counts using qubit tapering.

For the simplest system H_2, chemical accuracy was achieved quickly using all methods. QII and Q-Inv+I-Iter achieved machine precision. Q-Inv began to diverge after 3 iterations. For QII the number of qubits is constant for all iterations here and is equal to 7 (with qubit tapering 4). The number of qubits for Q-Inv and I-Iter is equal to at least 13. However, because we used a truncation degree 50, we required the same number of queries as the block encoded Hamiltonian, which makes the circuit deep. Q-Inv and I-Iter do not require queries for block encoding; however, they require a Trotter approximation of the time-evolution operator, which can also be relatively deep.

For LiH, all methods eventually achieved chemical accuracy. However, QII performed this task in 3 iterations, whereas Q-Inv and I-Iter required more than 10. Q-Inv also began to diverge significantly after 14 iterations. This divergent behavior was fixed in the combined implementation of Q-Inv and I-Iter. For QII the number of qubits was 11 (with qubit tapering 7). The number of qubits for Q-Inv and I-Iter was equal to at least 17. We used the same truncation degree 50; therefore, we required the same number of queries as the block encoded Hamiltonian.

For BeH_2, QII achieved chemical accuracy in one iteration, whereas Q-Inv and I-Iter required more than 10. Generally, Q-inv and I-Iter demonstrated a similar behavior to that of LiH. The resource requirements were also the same as those for LiH. QII requires 11 qubits (with qubit tapering 7), whereas Q-Inv and I-Iter require at least 17. The number of queries for block encoding was the same as before.

The final tested system was the square H_4. It was a challenging case for Q-Inv since it converged before chemical accuracy was reached. QII and I-Iter were more successful. Although QII with $n = 50$ demonstrates a very fast reach to the chemical accuracy with smaller truncation, it has the problem of early convergence, analogous to Q-Inv (see Appendix A2). QII requires 9 qubits (with qubit tapering 7), whereas Q-Inv and I-Iter require at least 15.

In summary, the proposed GQSP-based QII implementation offers several advantages. First, implementing the time-evolution operator in I-Iter and Q-Inv typically involves Suzuki–Trotter decompositions, which introduce additional errors and increase the circuit depth. Second, I-Iter and Q-Inv require LCU approaches to approximate the integral, necessitating $O(\log(n_y n_z))$ auxiliary qubits, where n_y and n_z correspond to the number of nodes in the numerical integration of Eq. (3.48) and Eq. (3.49). The authors noted that n_z depends on

the power of the inverse Hamiltonian k. Moreover, the application of H^{-1} k times using I-Iter is not explicitly explained. It may be necessary to perform state tomography after each step or the use of QSP. The same problem applies to Q-Inv+I-Iter. Conversely, GQSP requires only one signal processing qubit, one qubit for capitalization (if needed), and additionally $O(\log L)$ qubits for LCU block encoding, where L is the number of Pauli terms in the LCU expansion (see Eq. (2.147)). In the worst case, L scales as $O(N_{\text{orb}}^4)$, leading to $O(\log N_{\text{orb}})$ ancilla qubits, where N_{orb} denotes the number of orbitals in active space. Third, QII reaches chemical accuracy quickly, having provided a sufficiently large truncation value. However, QII requires a good initial estimate of the ground state energy, whereas I-Iter and Q-Inv do not. Nevertheless, if such an estimate is available, QII can converge to a high accuracy in very few iterations, particularly for larger polynomial degrees.

3.4.4 Excited States Calculations Using QFSM

Finally, we performed numerical tests of the QFSM implemented using GQSP. Unlike the other algorithms discussed in this study, QFSM enables access to excited states.

In the first test, we computed the PECs of the ground and excited states in the hydrogen molecule using the cc-pVDZ basis set with an active space of 2 electrons in 2 orbitals. The convergence of QFSM critically depends on the choice of the initial guess wave function, which should have a large overlap with the target excited state. For instance, starting from an RHF wave function does not allow convergence to the $1^3\Sigma_u^+$ triplet state or $1^1\Sigma_u^+$ singlet excited state.

To ensure sufficient overlap with specific target states, we employed the following initial guess states:

$$\begin{aligned} |\Psi_{\Sigma_g}\rangle &= |0011\rangle, \\ |\Psi_{\Sigma_u}\rangle &= |1001\rangle, \end{aligned} \tag{3.50}$$

where the ordering corresponds to $|\sigma_\beta^* \sigma_\alpha^* \sigma_\beta \sigma_\alpha\rangle$. As an energy shift in the folded Hamiltonian, we used the corresponding CASCI(2, 2) energies for testing. The results are shown in Fig. 3.8.

The singlet states $X^1\Sigma_g^+$ and $2^1\Sigma_g^+$ converged to chemical accuracy across all bond lengths when starting from $|\Psi_{\Sigma_g}\rangle$. However, at 0.4 Å and 0.8 Å, divergence was observed after 25 iterations owing to the loss of orthogonality in floating point arithmetic. This is a typical drawback of power iteration methods. Convergence to the $1^1\Sigma_u^+$ and $1^3\Sigma_u^+$ states was more challenging when starting from $|\Psi_{\Sigma_u}\rangle$. At short bond lengths (0.4 Å and 0.8 Å), these states were unable to achieve chemical accuracy, and their convergence was very slow. The right-hand plots show overlapping data points for these two states, indicating near-degeneracy. This behavior can be attributed to the presence of quasi-degenerate eigenstates in the folded Hamiltonian, which causes the wave function to be simultaneously attracted toward both, thereby hindering the convergence. Improved initial guesses could resolve this issue.

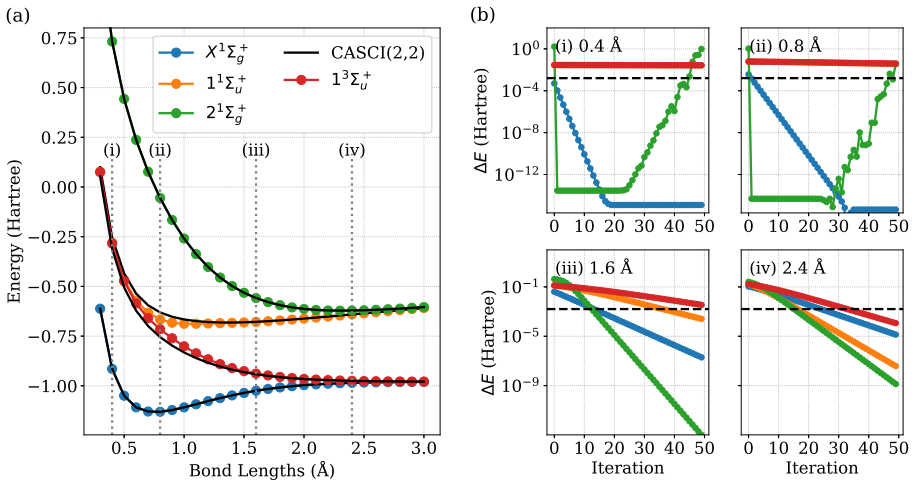


Figure 3.8: (a) Potential energy curves for the hydrogen molecule computed using QFSM. The basis set is cc-pVDZ and the active space is $(2, 2)$. All QFSM curves correspond to converged energies or energies after 50 iterations, if convergence was not achieved. (b) The right plots show the energy deviation from CASCI$(2, 2)$ over the first 50 iterations. The black dashed line denotes chemical accuracy (1 kcal/mol).

We also calculated the PECs for the four low-lying electronic states of ethylene along the internal rotation around the double bond. The cc-pVDZ basis set was used with 2 electrons in 2 active π orbitals. As in the case of hydrogen, the convergence strongly depends on the initial guess. The following estimate states were used:

$$\begin{aligned} |\Psi_{A_g}\rangle &= |0011\rangle, \\ |\Psi_{B_{1u}}\rangle &= |1001\rangle, \end{aligned} \tag{3.51}$$

where orbital ordering is $|\pi^*_\beta \pi^*_\alpha \pi_\beta \pi_\alpha\rangle$. As before, CASCI$(2, 2)$ energies were used as the energy shifts. The results are presented in Fig. 3.9.

At 0°, all states, except for the triplet $^3B_{1u}$, reached chemical accuracy within 50 iterations. The triplet state likely converges with additional iterations. The convergence was fastest for the singlet ^1A_g states. At 90°, none of the states reached chemical accuracy within 50 iterations, although the convergence trends were still visible. Up to 60°, both singlet ^1A_g states achieved chemical accuracy.

A key advantage of our GQSP-based QFSM implementation over variational quantum eigensolver VQE-based folded spectrum approaches is that we do not need to explicitly construct or measure the squared Hamiltonian. For instance, Ref. [178] showed that even for H_2 in the STO-3G basis, constructing a squared Hamiltonian requires the measurement of 24 Pauli strings, and for H_2O, over 111000 Pauli strings are required. Conversely, the proposed approach does not necessitate additional Pauli measurements and completely avoids variational

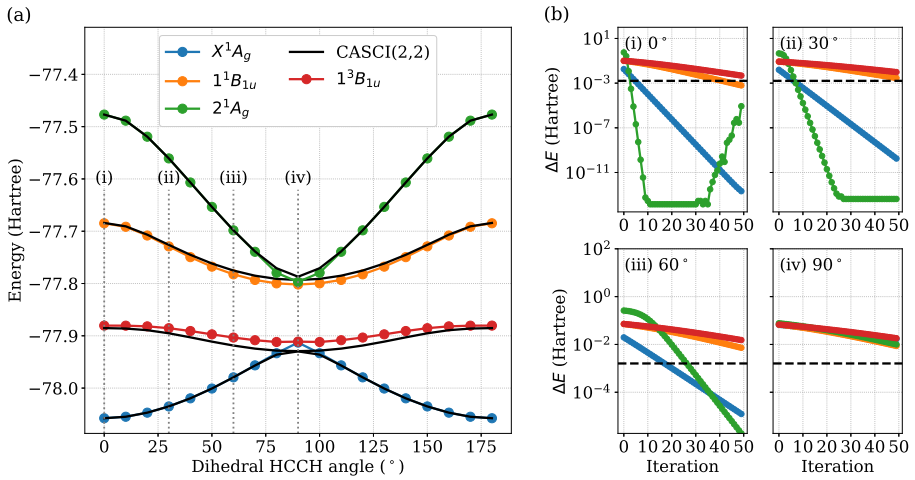


Figure 3.9: (a) Potential energy curves for ethylene along the torsional rotation of the double bond, computed using QFSM. The basis set is cc-pVDZ and the active space is $(2, 2)$. (b) The left plots show energy deviation from CASCI$(2, 2)$ over the first 50 iterations. The black dashed line indicates chemical accuracy (1 kcal/mol).

optimization. The main limitation of our method is that convergence is generally slower and is highly dependent on the quality of both the initial guess wave function and the target energy. Nevertheless, QFSM is a promising tool for accessing excited states, especially when a good initial state is available.

3.5 Conclusion

We developed a family of GQSP-based polynomial filters for quantum state preparation, covering ground and excited state tasks through the QPI, QPL, QII, and QFSM schemes. A query complexity analysis shows that QPI and QPL achieve ground state preparation with near-optimal $\tilde{O}(1/(\gamma\Delta))$ scaling without requiring prior energy estimates, while QII becomes the most efficient option once a coarse energy approximation is available. QFSM extends the framework to excited states and remains practical when an energy guess lies within the local gap.

The GQSP construction enables these filters to be implemented with predictable resource requirements, avoiding Trotterization and relying only on a single block encoding of the Hamiltonian and straightforward AA. Numerical proof-of-concept simulations confirm the theoretical trends: QPL achieves consistent gains over QPI, QII improves over earlier inverse iteration methods, such as I-Iter and Q-Inv, and QFSM reaches excited states without variational optimization. The main limitation is the low success probability inherent to

non-unitary filters, which necessitates postselection or AA and thus increases circuit depth. In our analysis, these resource costs are evaluated in a FTQC setting, where coherent polynomial transformations and AA can be implemented reliably.

Overall, these results show that classical power method intuition transfers cleanly to the quantum setting through GQSP, yielding simple, modular, and competitive state preparation algorithms suitable for both ground and excited states.

However, many aspects can be improved before the practical implementation of these methods. First, eliminating the need for variational optimization in QPL by deriving polynomial coefficients from perturbation theory can improve efficiency. The QII method could benefit from replacing the current polynomial approximation of the inverse function with one that has a larger radius of convergence, potentially enabling accurate computation of excited states. In addition, imaginary energy shifts can be beneficial, and, more generally, Green's function methods can be implemented using this technique. In the case of QFSM, it would be advantageous to identify a function that better distinguishes energy levels, because it is currently common for multiple eigenvalues in the folded spectrum to cluster near 1, greatly worsening convergence. Moreover, it is in principle possible to extend the present GQSP-based methods to block variants aimed at preparing several excited states simultaneously. Within the Bloch effective Hamiltonian framework (see Sec. 2.1.5.5), the corresponding wave operator can be written as an function of the Hamiltonian or of the perturbation, allowing its polynomial approximation through GQSP [187]. The main challenges involve constructing the required model space projectors and approximating resolvent-like functions whose spectral structure demands high degree polynomials [188, 189], thereby increasing the circuit depth and sensitivity to noise. Finally, the integration of GQSP-based methods with other quantum-classical hybrid approaches may open the door for tackling larger and more complex molecular systems on near-term fault-tolerant quantum devices.

Chapter 4

Symmetry-Adapted State Preparation for Quantum Chemistry on Fault-Tolerant Quantum Computers

We present systematic and resource-efficient constructions of continuous symmetry projectors, particularly U(1) particle number and SU(2) total spin, tailored for FTQCs. Our approach leverages LCU, GQSP and GQSVT to implement projectors that can be coherently applied as state filters prior to QPE. We analyze their asymptotic gate complexities and provide explicit circuit realizations. For the particle number and S_z symmetries, GQSP offers a particularly favorable resource usage features owing to its low ancilla qubit requirements and robustness to finite precision rotation gate synthesis. For the total spin projection, the structured decomposition of $\hat{P}_{S,M_S}$ reduces the projector T gate count. Numerical simulations show that symmetry filtering substantially increases the QPE success probability, leading to a lower overall cost compared to that of unfiltered approaches across representative molecular systems. Resource estimates further indicate that the cost of symmetry filtering is 3 to 4 orders of magnitude lower than that of the subsequent phase estimation step, which is especially relevant in large, strongly correlated systems, such as FeMoco, a standard strongly correlated open-shell benchmark. These results establish continuous-symmetry projectors as practical and scalable tools for state preparation in quantum chemistry and provide a pathway toward realizing more efficient fault-tolerant quantum simulations. This chapter is based on the preprint [Khinevich and Mizukami, arXiv:2601.08533 (2026)].

4.1 Introduction

Quantum computing has been significantly developed in recent years, supported by progress in hardware platforms and the design of quantum algorithms that are capable of addressing classically intractable problems [190, 191, 192, 193]. Among the most promising applications is the simulation of quantum many-body systems, particularly in quantum chemistry, where the ability to describe electronic structure accurately is directly related to the understanding of chemical reaction mechanisms, material properties, and molecular spectroscopy [194, 148, 195, 196, 197]. Although many quantum algorithms have been proposed to compute ground- and excited-state energies [7, 22, 32, 198, 29], a recurring challenge is the preparation of high-fidelity initial states suitable for such computations.

State preparation plays a central role in algorithms such as QPE [6, 7], where the probability of successful projection onto an eigenstate is determined by the overlap between the prepared input state and the desired eigenstate of the target Hamiltonian. In practice, the overlap is often limited when generic or heuristic initial states are employed. Various methods have been explored to improve the state preparation, including variational optimization [22, 133, 72, 199], adiabatic state preparation [150, 200], tensor networks [201, 202, 203], and state filters [140, 204, 141, 138, 142]. However, these strategies can suffer from limitations such as barren plateau behavior, large circuit depths, or the need for classical preoptimization.

An alternative approach leverages the inherent symmetries of physical Hamiltonians. In particular, molecular Hamiltonians possess well-defined symmetry sectors and physical eigenstates reside within specific irreducible representations of the corresponding symmetry group. Projecting an initial state onto a chosen symmetry sector can significantly enhance the overlap with physically relevant eigenstates and reduce the effective Hilbert space explored by the algorithm.

The traditional use of symmetries may require non-scalable resources or rely on repetitive measurements, limiting their applicability in fault-tolerant settings [205, 206]. One method to optimize the state of a particular symmetry sector is to use penalty terms [207]. They usually introduce many additional terms into the Hamiltonian to measure, and require parameter fitting. The QPE can also be used for symmetry-state filtering [208]. Another approach uses Löwdin spin projectors to modify the Hamiltonian itself [209]. This leads to an exponential number of terms, which is impractical for such a realization. However, in this chapter, we demonstrate how to implement this type of projector efficiently. There is also a recent general approach for constructing any symmetry projector, but only for finite-symmetry groups, such as point groups [210].

In this chapter, we constructed projectors for continuous symmetries relevant to the electronic structure, including the particle number and spin projection (U(1)), and total spin (SU(2)). We implemented them using an LCU technique [121], GQSP [128], and GQSVT [177] frameworks. We analyze the theoretical scaling of these constructions under FTQC assumptions and evaluated their numerical accuracy and sensitivity to the finite precision of rotational gate

synthesis. We further provide detailed resource estimates demonstrating their practicality as state filters for QPE in molecular Hamiltonian simulations, showing that symmetry projectors enhance state overlap in realistic quantum chemistry settings. Finally, we estimate the required cost for the FeMoco complex, a standard benchmark for strongly correlated open-shell systems [211, 212].

The remainder of the chapter is organized as follows. Section 4.2 presents the theoretical background and formal construction of the symmetry projectors, including their properties. Section 4.3 describes the computational stack and provides the source code for reproducibility. Section 4.4 reports theoretical scaling and identifies the optimal parameter regimes, followed by numerical calculations and resource estimates, with representative molecular case studies. This section also outlines the open challenges and prospective improvements of this methodology. Finally, Sec. 4.5 summarizes the main contributions and discusses the potential directions for future research.

4.2 Theory

4.2.1 Symmetry Operators

We aimed to prepare the quantum states in the fixed-symmetry sector of the electronic Hamiltonian. We denote $\hat{N}$, $\hat{S}_z$, and $\hat{S}^2$ as the particle number, spin projection onto the selected z-axis, and total spin-squared operators, respectively. Each operator is Hermitian and we aim to prepare quantum states belonging to the common subspace defined by

$$\begin{aligned} \hat{N}|\psi\rangle &= N_{\text{elec}}|\psi\rangle, \\ \hat{S}_z|\psi\rangle &= M_S|\psi\rangle, \\ \hat{S}^2|\psi\rangle &= S(S+1)|\psi\rangle, \end{aligned} \tag{4.1}$$

where N_{elec} is the number of particles (electrons in this case), M_S is the spin projection, and S is the total spin quantum number.

To express these operators in the second quantized form, we adhere to the convention that even spin-orbital indices correspond to the spin state $|\alpha\rangle$ (or $|\uparrow\rangle$), and odd indices correspond to spin $|\beta\rangle$ (or $|\downarrow\rangle$). Indexing starts at 0. The particle number operator is expressed as

$$\hat{N} = \sum_{i=0}^{N_{\text{SO}}-1} \hat{a}_i^\dagger \hat{a}_i, \tag{4.2}$$

where $\hat{a}_i^\dagger$ and $\hat{a}_i$ are fermionic creation and annihilation operators acting on spin-orbital i, N_{SO} is the number of spin-orbitals.

Similarly, the spin-projection operator is

$$\hat{S}_z = \frac{1}{2} \sum_{i=0}^{N_{\text{SO}}-1} (-1)^i \hat{a}_i^\dagger \hat{a}_i. \tag{4.3}$$

The second quantized form of the total spin-squared operator $\hat{S}^2$ is more complicated and requires the introduction of spin ladder operators:

$$\hat{S}_+ = \sum_{i=0}^{N_{\mathrm{SO}}/2-1} \hat{a}_{2i}^\dagger \hat{a}_{2i+1}, \quad \hat{S}_- = \sum_{i=0}^{N_{\mathrm{SO}}/2-1} \hat{a}_{2i+1}^\dagger \hat{a}_{2i}. \tag{4.4}$$

The total spin operator can then be written as:

$$\hat{S}^2 = \hat{S}_z^2 + \hat{S}_z + \hat{S}_-\hat{S}_+ = \hat{S}_z^2 - \hat{S}_z + \hat{S}_+\hat{S}_-. \tag{4.5}$$

To use these operators on a quantum computer, we must map the fermionic operators to qubit operators via a fermion-to-qubit transformation. Let the fermionic system be encoded onto N_{SO} qubits using the JW transformation (see Sec. 2.2.4.1). In this mapping, the number operator becomes

$$\hat{N} \xrightarrow{\mathrm{JW}} \sum_{j=0}^{N_{\mathrm{SO}}-1} \frac{I - Z_j}{2} = \frac{N_{\mathrm{SO}}}{2} I - \frac{1}{2} \sum_{j=0}^{N_{\mathrm{SO}}-1} Z_j. \tag{4.6}$$

Thus, the number operator is a linear combination of Pauli strings with a total number of terms $L = N_{\mathrm{SO}} + 1$.

Similarly, the following simple expression holds for $\hat{S}_z$:

$$\hat{S}_z \xrightarrow{\mathrm{JW}} \frac{1}{4} \sum_{p=0}^{N_{\mathrm{SO}}/2-1} (Z_{2p+1} - Z_{2p}). \tag{4.7}$$

For $\hat{S}_z$, the number of distinct Pauli terms was $L = N_{\mathrm{SO}}$.

The JW image of $\hat{S}^2$ is more complex. By introducing the local spin-projection operator $D_p := Z_{2p+1} - Z_{2p}$ and the ladder operators $\sigma^\pm = (X \pm iY)/2$, we obtain the following:

$$\begin{aligned} \hat{S}^2 \xrightarrow{\mathrm{JW}} & \sum_{p=0}^{N_{\mathrm{SO}}/2-1} \tfrac{3}{8}(1 - Z_{2p}Z_{2p+1}) \\ & + \sum_{p<q}^{N_{\mathrm{SO}}/2-1} \Big[\tfrac{1}{8} D_p D_q + \big(\sigma_{2p}^+ \sigma_{2p+1}^- \sigma_{2q}^- \sigma_{2q+1}^+ \\ & \qquad + \sigma_{2p}^- \sigma_{2p+1}^+ \sigma_{2q}^+ \sigma_{2q+1}^-\big) \Big]. \end{aligned} \tag{4.8}$$

The number of terms in this expansion scales as $L = \frac{1}{2}(3N_{\mathrm{SO}}^2 - 5N_{\mathrm{SO}}) + 1$, which is quadratic with N_{SO}.

Although the exact eigenstates of the electronic Hamiltonian respect these symmetries, variational ansatzes and state preparation techniques in quantum computing often break them down. Therefore, explicit projection onto the desired symmetry sectors can therefore substantially improve the quality of the approximate states and the accuracy of the measurements.

4.2.2 Symmetry Projectors

By definition, orthogonal projectors are operators satisfying $\hat{P}^2 = \hat{P} = \hat{P}^\dagger$. In the present context, we seek operators that project an arbitrary state $|\psi\rangle$ onto the eigenspace of a symmetry operator $\hat{O}$ corresponding to a specific eigenvalue o_i:

$$\hat{O}\,\hat{P}_{o_i}|\psi\rangle = o_i\,\hat{P}_{o_i}|\psi\rangle. \tag{4.9}$$

We assume that the initial state $|\psi\rangle$ exhibits a nonzero overlap with the target eigenspace.

Such a projector can be constructed using several methods [209]. In the following, we outline the approaches used in this chapter, which are compatible with modern quantum computing frameworks, such as LCU (see Sec. 2.2.4.2), GQSP (see Sec. 2.2.4.6), and GQSVT (see Appendix B1).

4.2.2.1 Integral Projectors

For operators $\hat{N}$ and $\hat{S}_z$, the exact projector onto eigenvalue o_i can be expressed as a Fourier-type U(1) group integral over a unit circle:

$$\hat{P}_{o_i} = \frac{1}{2\pi}\int_0^{2\pi} e^{i\phi(o_i-\hat{O})}\,d\phi, \tag{4.10}$$

where $\hat{O}$ is the symmetry operator, and o_i is the eigenvalue corresponding to the desired subspace.

In practice, this continuous integral is replaced by a discrete Fourier sum, resulting in a finite linear combination of unitary operators as explained in Sec. 4.2.3.

The projector onto a definite total-spin subspace S with projection M_S can be written as a SU(2) group integral:

$$\begin{aligned}\hat{P}_{S,M_S} = \frac{2S+1}{8\pi^2}\int_0^{2\pi} d\alpha \int_0^{\pi} d\beta\,\sin\beta \int_0^{2\pi} d\gamma \\ \times D^S_{M_S M_S}(\alpha,\beta,\gamma)^* R(\alpha,\beta,\gamma).\end{aligned} \tag{4.11}$$

where $R(\alpha,\beta,\gamma) = e^{-i\alpha\hat{S}_z}e^{-i\beta\hat{S}_y}e^{-i\gamma\hat{S}_z}$ is a spin-rotation operator expressed in terms of the Euler angles. $D^S_{M_S M_S}(\alpha,\beta,\gamma)$ is the corresponding Wigner D-matrix element.

This expression can be rewritten as follows:

$$\begin{aligned}&\hat{P}_{S,M_S} = \frac{1}{2\pi}\int_0^{2\pi} e^{i(M_S-\hat{S}_z)\alpha}d\alpha \\ &\times \frac{2S+1}{2}\int_0^{\pi} e^{-i\beta\hat{S}_y} d^S_{M_S M_S}(\beta)^*\sin\beta d\beta \\ &\times \frac{1}{2\pi}\int_0^{2\pi} e^{i(M_S-\hat{S}_z)\gamma}d\gamma = \\ &\hat{P}_{M_S}\left[\frac{2S+1}{2}\int_0^{\pi} e^{-i\beta\hat{S}_y} d^S_{M_S M_S}(\beta)^*\sin\beta d\beta\right]\hat{P}_{M_S},\end{aligned} \tag{4.12}$$

where $d^S_{M_S M_S}(\beta) = e^{iM_S\alpha} D^S_{M_S M_S}(\alpha,\beta,\gamma) e^{iM_S\gamma}$ denotes the reduced Wigner d-matrix. The integral in square brackets we will call a $\hat{P}_S$ projector.

If the initial wave function is an eigenfunction of $\hat{S}_z$ with eigenvalue M_S, which is the case for some ansatzes, such as UCC (see Sec. 2.2.5.1), then the first $\hat{P}_{M_S}$ projector is a trivial identity. Thus, we only need to apply a reduced $\hat{P}_S$ projector followed by a $\hat{P}_{M_S}$ projector. This construction provides a computational advantage to $\hat{S}_z$-conserving ansatzes.

As in the case of Eq. (4.10), these integrals must be discretized for implementation. In the case of $\hat{P}_{S,M_S}$ and $\hat{P}_S$ the quadrature nodes and weights are generally nonuniform, limiting the direct applicability of GQSP. Nevertheless, the resulting discretization yields an LCU-compatible form (see Sec. 4.2.3).

4.2.2.2 Lagrange Interpolation Projectors

An alternative and fully algebraic approach for constructing projectors is based on the Lagrange interpolation formula. The objective is to construct a polynomial function of the symmetry operator $\hat{O}$ that nullifies all the undesired eigenvalues o_j $(j \neq i)$ while retaining the target eigenvalue o_i. The general expression is as follows:

$$\hat{P}_{o_i} = \prod_{o_i \neq o_j} \frac{\hat{O} - o_j}{o_i - o_j}. \tag{4.13}$$

This polynomial depends solely on the Hermitian operator $\hat{O}$ and therefore can be efficiently implemented using GQSVT. This provides an exact construction of an arbitrary symmetry operator. In particular, by using this approach, the Löwdin projector can be constructed for the operator $\hat{S}^2$ [213], which will be used in future sections.

Note that because we work on a finite spin-orbital basis, all the symmetry operators $\hat{O}$ are finite Hermitian matrices. Consequently, the integral projector (4.10) and Lagrange interpolation projector (4.13) represent the same polynomial function as in the Cayley–Hamilton theorem. This equivalence holds only when a single eigenvalue o_i is targeted and the number of quadrature nodes in the integral is sufficiently large. Despite their mathematical equivalence, these two expressions allow distinct and potentially advantageous implementations on quantum computers.

4.2.3 Projectors Implementation

4.2.3.1 Linear Combination of Unitaries

$\hat{N}$ and $\hat{S}_z$ operators.

The projectors onto the selected symmetry subspaces for $\hat{N}$ and $\hat{S}_z$ can be constructed using the integral representation introduced in Sec. 4.2.2. In practice, this integral is replaced by a finite Fourier sum over discrete phase points.

$$\hat{P}_{o_i} = \frac{1}{N_\phi} \sum_{k=0}^{N_\phi - 1} e^{i\phi_k (o_i - \hat{O})}, \qquad \phi_k = \frac{2\pi k}{N_\phi}, \tag{4.14}$$

where N_ϕ is selected to resolve all the distinct eigenvalues of $\hat{O}$. For the number operator $\hat{N}$ acting on N_{SO} qubits, $N_{\mathrm{SO}}+1$ distinct eigenvalues exist, and thus, $N_\phi = N_{\mathrm{SO}}+1$ integration nodes are sufficient for an exact representation.

The exponential terms in Eq. (4.14) is a unitary operator of the form $e^{-i\phi_k\hat{O}}$ with the corresponding phase $e^{i\phi_k o_i}$. The resulting linear combination of unitaries can be implemented using a standard LCU framework with an ancillary register comprising $\lceil \log_2 N_\phi \rceil$ qubits (see Sec. 2.2.4.2). The controlled unitaries corresponding to different values ϕ_k are conditionally applied, based on the ancillary register state.

The unitaries $e^{i\phi_k(N_{\mathrm{elec}}-\hat{N})}$ for $\hat{P}_N$ projector can be constructed using Eq. (4.6). These unitaries are diagonal on a computational basis and can be implemented efficiently using single-qubit R_Z rotations (see Fig. 4.1).

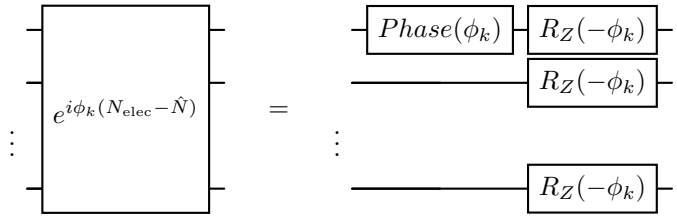


Figure 4.1: Decomposition of the unitary operator $e^{i\phi_k(N_{\mathrm{elec}}-\hat{N})}$ into single-qubit R_Z rotations and a phase gate $Ie^{i\phi_k(N_{\mathrm{elec}}-N_{\mathrm{SO}}/2)}$ on the first qubit.

The same construction applies to $\hat{P}_{M_S}$ spin projectors, differing only in the form of unitary exponentials $e^{-i\phi_k\hat{S}_z}$. This expression can be obtained from Eq. (4.7), which corresponds to the relative phase rotations between α and β spin-orbitals, as shown in Fig. 4.2.

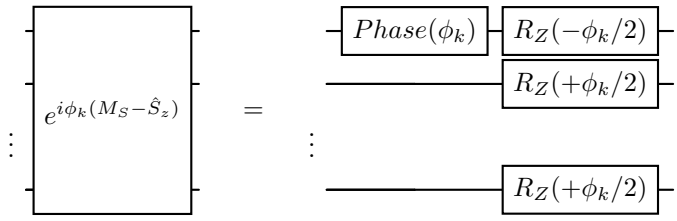


Figure 4.2: Decomposition of the unitary operator $e^{i\phi_k(M_S-\hat{S}_z)}$ into single-qubit R_Z rotations and a phase gate $Ie^{i\phi_k M_S}$ on the first qubit.

All weights in Eq. (4.14) are equal. Consequently, the LCU ancilla preparation operator (PREP) reduces to Hadamards when the number of terms is a power of two. The SELECT operator consists of controlled unitaries $e^{i\phi_k(o_i-\hat{O})}$ for different angles ϕ_k. The complete LCU circuits shown in Eq. (4.14) with optimal PREP is shown in Fig. 4.3. The measurement and postselection of the ancilla

register in the $|0\rangle$ state allows the implementation of projectors $\hat{P}_N$ and $\hat{P}_{M_S}$.

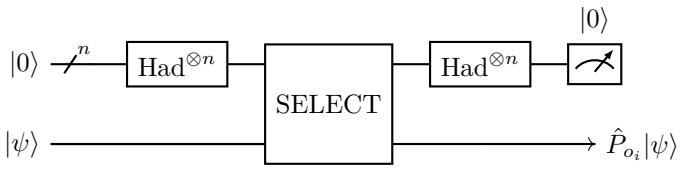


Figure 4.3: The LCU scheme for $\hat{P}_N$ and $\hat{P}_{M_S}$ projectors, if number of LCU terms is power of two. The state $\hat{P}_{o_i}|\psi\rangle$ is up to corresponding normalization.

$\hat{S}^2$ operator.

The total-spin projector $\hat{P}_{S,M_S}$ can be implemented using numerical discretization of the group integral in Eq. (4.11). We discretize the SU(2) projector for a fixed (S, M_S) sector using Gauss–Legendre nodes at the polar angle and uniform trapezoid grids at the azimuthal angles as follows:

$$\hat{P}_{S,M_S} \approx \sum_{b=1}^{N_\beta} \sum_{a=0}^{N_\alpha-1} \sum_{g=0}^{N_\gamma-1} w_{bag}\, R(\alpha_a, \beta_b, \gamma_g), \tag{4.15}$$

where the weights are defined as:

$$\begin{aligned} w_{bag} &= \frac{2S+1}{2N_\alpha N_\gamma}\, w_b^{\mathrm{GL}}\, D^S_{M_S M_S}(\alpha_a, \beta_b, \gamma_g)^*, \\ \beta_b &= \arccos(x_b). \end{aligned} \tag{4.16}$$

Here, x_b and w_b^{GL} are the nodes and weights for the Gauss–Legendre quadrature with N_β nodes. In the implementation, the LCU uses $|w_{bag}|$ as the amplitudes, whereas the phase is absorbed into $R(\alpha_a, \beta_b, \gamma_g)$ using an additional phase gate.

Each $R(\alpha, \beta, \gamma) = e^{-i\alpha\hat{S}_z} e^{-i\beta\hat{S}_y} e^{-i\gamma\hat{S}_z}$ consists of two $\hat{S}_z$ rotations, whose circuit is illustrated in Fig. 4.2. The $\hat{S}_y$ rotation can also be constructed efficiently but requires CNOT gates [206]. Figure 4.4 shows a single two-spin-orbital fragment is shown, which is repeated for each spin-orbital pair. The complete circuit for $R(\alpha, \beta, \gamma)$ is the consecutive application of the $\hat{S}_z$, $\hat{S}_y$ and $\hat{S}_z$ rotations.

The overall projectors in Eq. (4.15) is then realized as a weighted LCU over these rotation unitaries, requiring the ancilla register, consisting of $\lceil \log_2(N_\alpha N_\beta N_\gamma) \rceil$ qubits. The PREP operator corresponding to β angles is nontrivial and cannot be replaced by Hadamard gates.

As mentioned in Sec. 4.2.2, the qubit cost can be reduced if the initial guess is of definite $\hat{S}_z$ symmetry. In this case, only a sequential application of $\hat{P}_S$ and $\hat{P}_{M_S}$ projectors is required. Then, the number of ancillary qubits is $\lceil \log_2(N_\alpha N_\beta) \rceil$. The gate cost is estimated in Sec. 4.4.

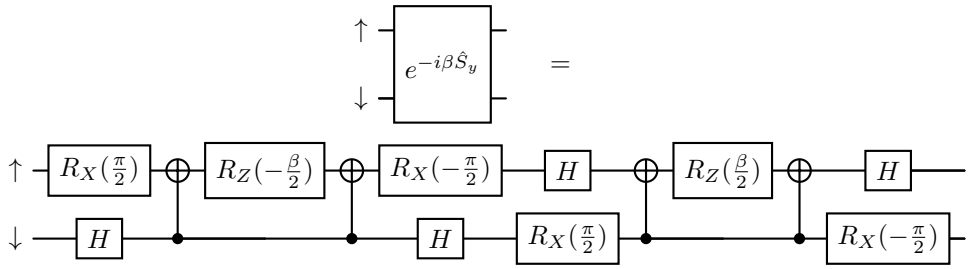


Figure 4.4: Circuit representation of $e^{-i\beta S_y}$ for a single spin-orbital pair $(\uparrow,\downarrow)$. This fragment is repeated for all spatial orbitals.

4.2.3.2 Generalized Quantum Signal Processing

Both the $\hat{N}$ and $\hat{S}_z$ operators have uniformly spaced eigenvalues that allow their projectors to be expressed as polynomials of the corresponding unitary operator. Formally, the projector in Eq. (4.14) can be rewritten as a polynomial

$$\hat{P}_{o_i} = \mathrm{Poly}(\hat{U}) = \sum_{k=0}^{d} c_k \hat{U}^k, \tag{4.17}$$

where $\hat{U} = e^{i\frac{2\pi}{N_\phi}\hat{O}}$ is the signal unitary and $c_k = \frac{1}{N_\phi} e^{i\frac{2\pi k}{N_\phi} o_i}$ are complex polynomial coefficients.

The resulting polynomial transformation can be implemented using GQSP with only a single ancilla qubit. GQSP is essential because the polynomial $\mathrm{Poly}(\hat{U})$ is of mixed parity and has complex coefficients. The degree of the polynomial d required for exact projection is equal to the number of spin-orbitals N_{SO}. The general structure of a GQSP circuit that implements such a projector is explained in Sec. 2.2.4.6.

A detailed comparison of qubit requirements and circuit depths for LCU and GQSP implementations is presented in Sec. 4.4.

4.2.3.3 Generalized Quantum Singular Value Transformation

Another approach to constructing symmetry projectors relies on the Lagrange interpolation formula introduced in Sec. 4.2.2. In this method, the projector is expressed as a polynomial with the corresponding symmetry operator.

$$\hat{P}_{o_i} = \mathrm{Poly}(\hat{O}) = \sum_{k=0}^{d} p_k \hat{O}^k, \tag{4.18}$$

where p_k denotes the polynomial coefficients, defined in Eq. (4.13). The major limitation of the Lagrange interpolation approach, when implemented directly, is the need to compute the high powers of $\hat{O}$, which results in an exponentially growing number of Pauli terms, and therefore, quickly becomes impractical.

The GQSVT solves this issue by allowing for the efficient implementation of polynomial transformations of block-encoded operators.

The difference from the previous approach is that we use LCU not to encode the projector operators ($\hat{P}_N$, $\hat{P}_{M_S}$, $\hat{P}_S$), but the symmetry operators ($\hat{N}$, $\hat{S}_z$, and $\hat{S}^2$) using Eqs. (4.6), (4.7) and (4.8). Subsequently, for each operator, we construct a corresponding polynomial that implements a Löwdin-like projector, which selectively amplifies the desired eigenvalue sector.

Considering that the symmetry operators are Hermitian and therefore diagonalizable, the singular value transformation is equivalent to a polynomial transformation for the symmetry operators. The GQSVT protocol is then applied to the block-encoding operators to realize the projector operator exactly. The degree of the polynomial for $\hat{P}_N$ and $\hat{P}_{M_S}$ projectors is equal to N_{SO}, whereas for $\hat{P}_{S^2}$ it is shorter and equal to the number of distinct eigenvalues for the $\hat{S}^2$ operator $N_{\mathrm{SO}}/2$. The degree of the implemented polynomial determined the circuit depth. This construction uses 1 signal ancilla and $\lceil \log_2 L \rceil$ ancilla qubits, one for the signal register and additional ancillas for LCU block encoding. The general architecture of the GQSVT projector circuit is shown in the Appendix B1.

The primary advantage of this approach is its generality: the same framework can be used to construct projectors for arbitrary symmetry operators without modification. Furthermore, multiple eigenvalue sectors can be targeted simultaneously by appropriately designing a polynomial $\mathrm{Poly}(\hat{O})$. Unlike the numerically discretized integral constructions used for $\hat{P}_{S,M_S}$, the GQSVT-based approach provides an exact polynomial realization and therefore offers higher theoretical accuracy. However, the Löwdin projector for $\hat{S}^2$ does not project onto a particular M_S sector.

Nevertheless, the practical implementation of GQSVT-based projectors remains challenging owing to the circuit depth and required gate precision. These aspects are analyzed in detail in Sec. 4.4, where we compare the resource requirements and discuss their potential applications as state-preparation filters for QPE.

4.3 Computational Details

The software stack used in this study is summarized below.

Resource estimation and fault–tolerant circuit compilation were performed using the QURI Parts framework (QSub subpackage) [214]. Numerical circuit simulations were performed using Qulacs [185]. The second-quantized Hamiltonian and JW mappings were generated using OpenFermion [215], and the molecular electronic integrals were obtained from PySCF [181].

All the scripts and notebooks used to produce the figures and tables in this chapter have been released at: https://github.com/mizukami-group/gqsp_symmetry_projectors.

4.4 Results and Discussion

4.4.1 Theoretical Scaling

We assessed the projector implementations in terms of the asymptotic gate counts and ancilla requirements. T gate counts are crucial in fault-tolerant settings because each T gate requires magic-state distillation. CNOT counts are also informative because they correlate with the circuit depth and routing overhead in limited-connectivity architectures, where SWAP operations can introduce significant errors. Temporary work ancilla qubits that are uncomputed deterministically and can be reused freely are not included.

We begin our analysis with LCU based projectors for the $\hat{N}$ and $\hat{S}_z$ operators. For this construction, the number of ancilla qubits required for indexing scales as $O(\log N_{\mathrm{SO}})$. The costs are determined by the PREP (ancilla state preparation) and SELECT (controlled operator selection) components.

When the number of LCU terms is a power of two, PREP reduces to a layer of Hadamard gates. In the general case, PREP prepares a superposition of the $O(N_{\mathrm{SO}})$ terms. A binary-tree state preparation scheme yields $O(N_{\mathrm{SO}} \log(N_{\mathrm{SO}}/\epsilon))$ T and $O(N_{\mathrm{SO}})$ CNOT gates. Using QROM reduces the T count to $O(N_{\mathrm{SO}})$ at the expense of $N_{\mathrm{SO}} \log(N_{\mathrm{SO}}/\epsilon)$ CNOT gates and a modest increase in ancilla qubits by $O(\log(1/\epsilon))$ qubits, depending on the required PREP precision ϵ.

SELECT uses indexed single-qubit unitaries. Considering the unitary constructions shown in Fig. 4.1 and Fig. 4.2, it is necessary to apply N_{SO} single qubit rotations per LCU term. The number of LCU terms is $O(N_{\mathrm{SO}})$. Therefore, any implementation of SELECT that uses indexed operations, requires $O(N_{\mathrm{SO}}^2)$ CNOT gates. The T gate cost includes both the indexing overhead and the synthesis of many small-angle rotations, $O(N_{\mathrm{SO}}^2)+N_{\mathrm{SO}}^2 T(R_z)$, where $T(R_z)$ is the T gate scaling for a single rotational gate. Corresponding to the Solovay–Kitaev theorem (see Sec. 2.2.2.3), the rotation gates require $O(\log(N_{\mathrm{SO}}/\epsilon))$ T gates for a global precision budget ϵ, which yields a total T gate count of $O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$.

Thus, for both $\hat{P}_N$ and $\hat{P}_{M_S}$, the cost is dominated by $O(N_{\mathrm{SO}}^2)$ controlled rotations and does not depend heavily on the specific PREP/SELECT implementation.

For the $\hat{S}^2$ projector, we begin with the general implementation expressed in Eq. (4.15), which requires $\lceil \log(N_\alpha N_\beta N_\gamma) \rceil$ ancilla qubits. N_α and N_γ scale as N_{SO}, as in the case of $\hat{P}_N$ and $\hat{P}_{M_S}$. N_β can be considered to be independent of N_{SO} and dependent only on the chosen spin S, which will be explained in Sec. 4.4.3. Therefore, the ancillary qubit count is $O(\log N_{\mathrm{SO}})$, which is the same as that for $\hat{N}$ and $\hat{S}_z$ but with a prefactor 2.

For a general $\hat{P}_{S,M_S}$ projector, $O(N_{\mathrm{SO}}^2)$ LCU terms exist, each of which consists of $O(N_{\mathrm{SO}})$ of unitaries shown in Fig. 4.4 and are surrounded by the single qubit rotations shown in Fig. 4.2. Thus, analogous to the previous cases of $\hat{P}_N$ and $\hat{P}_{M_S}$ projectors, the LCU is dominated by the number of unitary operators to be applied, which scales as $O(N_{\mathrm{SO}}^3)$. The resulting gate costs are $O(N_{\mathrm{SO}}^3)$ CNOT gates and $O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$ T gates. However, by expressing $\hat{P}_{S,M_S}$ as the product of $\hat{P}_{M_S}\hat{P}_S\hat{P}_{M_S}$ (see Eq. (4.12)) significantly reduced the

scaling to that of a single $\hat{P}_{M_S}$ projector. If the input is already an $\hat{S}_z$ eigenstate, one of the $\hat{P}_{M_S}$ projectors can be omitted without affecting the asymptotic scaling.

Consider now GQSP implementation for $\hat{N}$ and $\hat{S}_z$. The polynomial degree d in GQSP scales linearly with $O(N_{\mathrm{SO}})$, as shown in Sec. 4.2.3.2. The circuit requires d of SU(2) rotation gates for signal processing and $(d-1)N_{\mathrm{SO}}$ controlled single qubit rotations to build the signal operator, yielding $O(N_{\mathrm{SO}}^2)$ CNOT gates and $O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$ T gates. This matches the asymptotic cost of LCU, but requires only one ancilla qubit. The prefactors and robustness to angle precision differed and were evaluated separately.

GQSVT uses qubitization, which requires $O(\log L)$ ancilla qubits for L LCU terms. For $\hat{N}$ and $\hat{S}_z$ L scales linearly $O(N_{\mathrm{SO}})$; however, for $\hat{S}^2$, $O(N_{\mathrm{SO}}^2)$. This implies the same scaling of the ancilla qubits but a larger prefactor for P_{S^2}. The degree of the interpolation polynomial is equal to the number of distinct eigenvalues minus one: $d = N_{\mathrm{SO}}$ for $\hat{N}$ and $\hat{S}_z$ and $d = N_{\mathrm{SO}}/2$ for $\hat{S}^2$. The degree defines the number of times a qubitization operator will be accessed.

For $\hat{N}$ and $\hat{S}_z$, the operators are decomposed into single-Pauli strings (see Eq. (4.6) and Eq. (4.7)). Using the binary-tree PREP and indexed SELECT, we obtain $O(N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}})$ CNOT gates and $O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$ T gates. QROM trades these counts to $O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$ CNOT gates and $O(N_{\mathrm{SO}}^2)$ T gates with ancilla $O(\log(N_{\mathrm{SO}}/\epsilon))$.

For $\hat{S}^2$, the longest Pauli strings have constant length (up to 4); however, the number of LCU terms is $O(N_{\mathrm{SO}}^2)$ (see Eq. (4.8)). This increases the cost by one power of N_{SO}, giving $O(N_{\mathrm{SO}}^3 \log N_{\mathrm{SO}})$ CNOT gates and $O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$T gates for binary-tree implementation, and $O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$ CNOT gates and $O(N_{\mathrm{SO}}^3)$ T gates for QROM. All scalings are summarized in Tab. 4.1.

4.4.2 Success Probability

The success probabilities for all these schemes are essentially the same in the best case and are equal to

$$\Pr[\mathrm{succ}] = ||\hat{P}_{o_i}|\psi\rangle||_2^2 = |\langle\psi|\hat{P}_{o_i}|\psi\rangle|, \tag{4.19}$$

where $\hat{P}_{o_i}$ is the projector onto the subspace, corresponding to o_i and $|\psi\rangle$ is the initial wave function.

This probability can be increased by using AA. After m rounds of amplification,

$$\Pr^{(\mathrm{m})}[\mathrm{succ}] = \sin^2((2m+1)\theta), \quad \theta = \arcsin\left(\sqrt{\Pr[\mathrm{succ}]}\right). \tag{4.20}$$

Choosing m to maximize $\Pr^{(\mathrm{m})}[\mathrm{succ}]$ yields:

$$m = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor \approx \frac{\pi}{4} \frac{1}{\sqrt{\Pr[\mathrm{succ}]}}, \tag{4.21}$$

Table 4.1: Asymptotic scaling of T and CNOT gate counts and number of ancilla qubits for different projector implementations. N_{SO} is the number of spin-orbitals and ϵ is the target projector precision.

Method	**Projector**	**Implementation**	**T gates**	**CNOT gates**	**Ancillas**
LCU	$\hat{N}, \hat{S}_z$	Any LCU	$O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^2)$	$O(\log N_{\mathrm{SO}})$
LCU	$(\hat{S}^2, \hat{S}_z)$	Any LCU/$\hat{P}_{S,M_S}$ (Eq. (4.11))	$O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^3)$	$O(\log N_{\mathrm{SO}})$
LCU	$(\hat{S}^2, \hat{S}_z)$	Any LCU/$\hat{P}_{M_S}\hat{P}_S\hat{P}_{M_S}$(Eq. (4.12))	$O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^2)$	$O(\log N_{\mathrm{SO}})$
GQSP	$\hat{N}, \hat{S}_z$	—	$O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^2)$	1
GQSVT	$\hat{N}, \hat{S}_z$	Naïve LCU	$O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}})$	$O(\log N_{\mathrm{SO}})$
GQSVT	$\hat{N}, \hat{S}_z$	QROM LCU	$O(N_{\mathrm{SO}}^2)$	$O(N_{\mathrm{SO}}^2 \log(N_{\mathrm{SO}}/\epsilon))$	$O(\log(N_{\mathrm{SO}}/\epsilon))$
GQSVT	$\hat{S}^2$	Naïve LCU	$O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$	$O(N_{\mathrm{SO}}^3 \log N_{\mathrm{SO}})$	$O(\log N_{\mathrm{SO}})$
GQSVT	$\hat{S}^2$	QROM LCU	$O(N_{\mathrm{SO}}^3)$	$O(N_{\mathrm{SO}}^3 \log(N_{\mathrm{SO}}/\epsilon))$	$O(\log(N_{\mathrm{SO}}/\epsilon))$

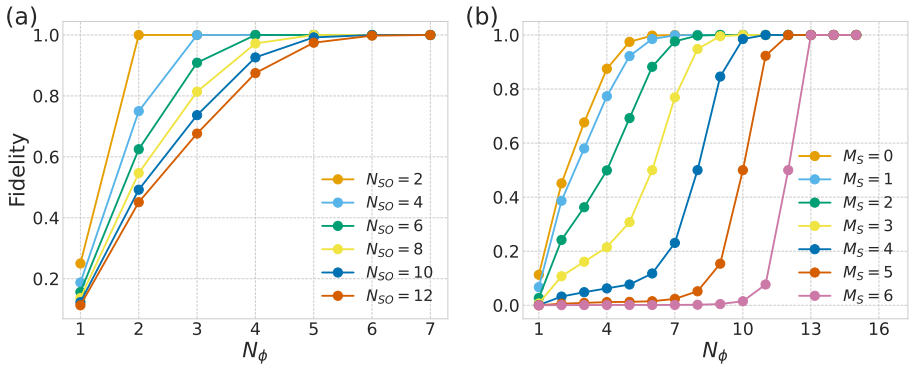


Figure 4.5: Fidelity of the $\hat{P}_{M_S}$ projector. (a) Dependence of the fidelity on the number of integration nodes for different numbers of spin-orbitals. (b) Dependence of the fidelity on the number of integration nodes for different target values of M_S.

where Pr[succ] is assumed small. Each round of amplification uses the underlying block a constant number of times; therefore, the gate count and depth grow linearly in m (approximately by a factor of $2m + 1$, plus the cost of reflections).

The number of iterations required to start from a uniform superposition can be estimated. The uniform superposition state of all possible occupations of N_{SO} spin-orbitals was projected onto a subspace with $N_{\mathrm{elec}} = N_{\mathrm{SO}}/2$ electrons. Then

$$|\langle\psi|\hat{P}_{N_{\mathrm{elec}}}|\psi\rangle| = \frac{\binom{N_{\mathrm{SO}}}{N_{\mathrm{elec}}}}{2^{N_{\mathrm{SO}}}} \sim \frac{1}{\sqrt{N_{SO}\pi/2}}. \tag{4.22}$$

By substituting this expression into Eq. (4.19), and then into Eq. (4.21), we observe that the number of amplification steps increases only with the fourth root of N_{SO}, which is typically modest in practice.

4.4.3 Numerical Accuracy of Projectors

To enable meaningful resource analysis, we quantified the numerical accuracy of the projectors and their dependence on discretization parameters. In the integral construction of Eqs. (4.10) and (4.12), the only tunable parameter is the number of quadrature nodes. This choice sets both the number of LCU terms and, in GQSP/QSVT-style realizations, the degree of the resulting polynomial.

When the eigenvalues of a symmetry operator $\hat{O}$ are integer-spaced, the integral form of the corresponding projector can be evaluated exactly using uniform discretization as shown in Eq. (4.14). In the eigenbasis $\hat{O}|o_j\rangle = o_j|o_j\rangle$ the integrand in Eq. (4.10) is reduced to a single Fourier mode $e^{i\phi(o_i - o_j)}$ with integer frequency $q = o_i - o_j \in \mathbb{Z}$, so the continuous projector simply enforces $\frac{1}{2\pi}\int_0^{2\pi} e^{iq\phi}d\phi = \delta_{q,0}$. Replacing the integral with a uniform grid $\phi_k = 2\pi k/N_\phi$

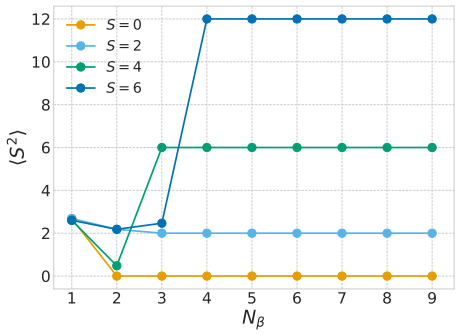


Figure 4.6: Expectation value of $\hat{S}^2$ for states produced by $\hat{P}_{S,M_S}$, as a function of the number of β-nodes N_β, for different target spin sectors S at fixed $M_S = 0$.

yields the corresponding discrete Fourier sum $\frac{1}{N_\phi}\sum_{k=0}^{N_\phi-1} e^{i2\pi kq/N_\phi}$, which is exactly 1 if $q \equiv 0 \pmod{N_\phi}$ and 0 otherwise. Therefore, the discrete projector is equal to the continuous projector if and only if no two distinct eigenvalues of $(o_i - \hat{O})$ coincide with the modulo N_ϕ; this is called the aliasing condition.

For $\hat{S}_z$, the spectrum in full Fock space is $\{-N_{\mathrm{SO}}/2, \ldots, N_{\mathrm{SO}}/2\}$ in unit steps. The accuracy of the discrete projector onto M_S is guaranteed when

$$N_\phi \geq \frac{N_{SO}}{2} + |M_S| + 1. \tag{4.23}$$

Therefore, in the common case $M_S = 0$ the minimum exact choice is $N_\phi = N_{\mathrm{SO}}/2+1$. This reduces the prefactor for applying $\hat{P}_{M_S}$ relative to $N_\phi = N_{\mathrm{SO}}+1$.

Figure 4.5 illustrates the fidelity of the $\hat{S}_z$ projector. We begin with uniform superposition. In panel (a), we fixed $M_S = 0$ and varied N_{SO}; in panel (b), we fixed $N_{\mathrm{SO}} = 12$ and varied the target M_S.

For the particle-number operator $\hat{N}$, the analysis is analogous. With the eigenvalues $\{0, \ldots, N_{\mathrm{SO}}\}$, the aliasing condition for the projector onto N_{elec} is

$$N_\phi \geq \frac{N_{\mathrm{SO}}}{2} + \left| N_{\mathrm{elec}} - \frac{N_{\mathrm{SO}}}{2} \right| + 1. \tag{4.24}$$

In particular, at half filling $N_{\mathrm{elec}} = N_{\mathrm{SO}}/2$, one can again take $N_\phi = N_{\mathrm{SO}}/2 + 1$.

The numerical integration of the total-spin projector $\hat{P}_{S,M_S}$ behaves differently. In our benchmarks, we target singlets with $M_S = 0$, starting with a uniform superposition restricted to $M_S = 0$. In this setting, it is convenient to apply $\hat{P}_S$ (see Fig. 4.4), followed by $\hat{P}_{M_S}$. For the $\hat{S}_z$ step, we use $N_\phi = N_{\mathrm{SO}} + 1$ nodes to make the discretization exact for all M_S, and for $M_S = 0$, we use the minimal exact value $N_\phi = N_{\mathrm{SO}}/2 + 1$.

From Fig. 4.6 we observe that the number of nodes N_β for $\hat{P}_S$ required to saturate the correct projection grows with S. In the range tested, we find $N_\beta = 2$ for $S = 0$, $N_\beta = 3$ for $S = 2$ and $S = 4$, and $N_\beta = 4$ for $S = 6$.

This behavior can be understood by expanding the β-kernel in Eq. (4.12) in a basis with fixed M_S and S',

$$\begin{aligned}
&\langle S' M_S | e^{-i\beta S_y} | S' M_S \rangle \, d^S_{M_S M_S}(\beta)^* \\
&= d^{S'}_{M_S M_S}(\beta) \, d^S_{M_S M_S}(\beta)^* \\
&= \sum_{L=|S'-S|}^{S'+S} B_L(S, S', M_S) \, P_L(\cos\beta),
\end{aligned} \tag{4.25}$$

where the coefficients $B_L(S, S', M_S)$ are independent of explicit N_{SO}. In particular, for the target component $S' = S$ the expansion contains only degrees up to $L_{\max} = 2S$.

If one aims for a sufficient condition for an exact evaluation of the $S' = S$ contribution, the Gauss–Legendre quadrature with N_β nodes integrates polynomials up to the degree $2N_\beta - 1$, implying a conservative requirement $N_\beta \geq S + 1$. However, in our benchmarks, the expectation values were saturated for a significantly smaller N_β, as reported above.

In contrast, GQSVT implementations based on Lagrange interpolation do not allow for this freedom. Their polynomial degree must be equal to the number of possible eigenvalues of the symmetry operator to filter out all the undesired eigenvalues. Otherwise, if some eigenvalues are not filtered, the Lagrange polynomial can take large values in these sectors. If the initial wave function is not perfectly zero in those sectors, it will be multiplied uncontrollably in those sectors, resulting in an even worse state than before the projection. Another problem is that if the polynomial grows significantly in these sectors, it requires normalization proportional to the largest polynomial value. This reduces the probability of success by the factor of the normalization constant. Thus, the polynomial degree for GQSVT should always be $d = N_{\mathrm{SO}} + 1$ for $\hat{N}$, $\hat{S}_z$ and $d = N_{\mathrm{SO}}/2 + 1$ for $\hat{S}^2$.

4.4.4 Gate Precision Sensitivity

The second practical issue for implementation and resource estimation is the robustness of the projectors to finite precision in single-qubit rotation gates. This directly indicates the T gate counts required to synthesize these rotations and the resulting fidelity. We modeled the finite precision by quantizing each continuous rotation angle ϕ to the nearest point on a uniform grid with a spacing $2\epsilon_r$. Concretely, we replace

$$\phi \;\mapsto\; \phi_d \equiv 2\epsilon_r \left\lfloor \frac{\phi}{2\epsilon_r} + \frac{1}{2} \right\rfloor \tag{4.26}$$

that guarantees that the bound

$$|\phi - \phi_d| \leq \epsilon_r. \tag{4.27}$$

A smaller ϵ_r corresponds to a finer angular resolution, whereas a larger ϵ_r yields a coarser discretization.

We benchmarked the LCU implementation of $\hat{P}_{M_S}$ for singlet targets starting from a uniform superposition without fixing the particle number. Figure 4.7 shows that the fidelity improves monotonically with angle precision and depends only weakly on the number of spin-orbitals N_{SO}. We consider two models: (I) SELECT-only, where only the indexed unitaries are rounded and PREP is exact (model for QROM or trivial PREP at a power-of-two register size) and (II) SELECT+PREP, where both operators are rounded. PREP errors further reduced the fidelity, but not significantly. In all tested cases, $\epsilon_r = 10^{-1}$ is sufficient to achieve high fidelity.

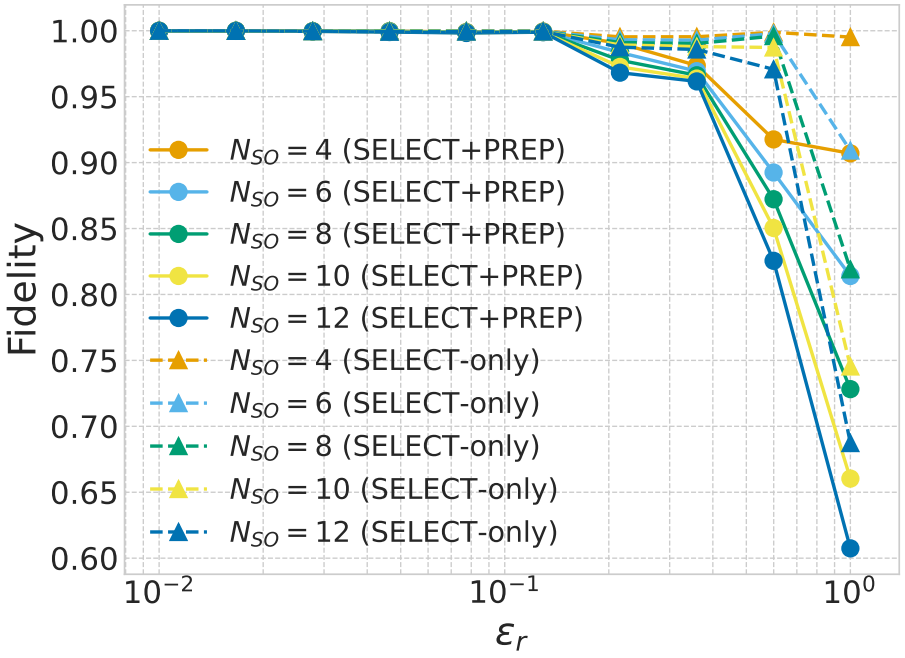


Figure 4.7: Fidelity of the LCU implementation of $\hat{P}_{M_S}$ versus single-qubit rotation precision ϵ_r for fixed $M_S = 0$ and varying N_{SO}. Solid lines with circular markers: both PREP and the indexed unitaries in SELECT are rounded (SELECT+PREP). Dashed lines with triangular markers: only the indexed unitaries in SELECT are rounded, while PREP is exact (SELECT-only).

Next, we tested the GQSP implementation for the same $\hat{P}_{M_S}$; the results are shown in Fig. 4.8. The GQSP realization is more sensitive to the angle precision than LCU, although it is still relatively robust. The signal operator angles were the most critical, and a precision of approximately $10^{-0.8}$ was generally considered sufficient. The GQSP rotation sequence itself should be slightly more precise, better than $10^{-1.5}$, to avoid noticeable fidelity drops. The dependence on N_{SO} is small, and increasing the number of spin-orbitals does not significantly affect the required precision. In all tests with $\epsilon_r = 10^{-2}$ the fidelities were close

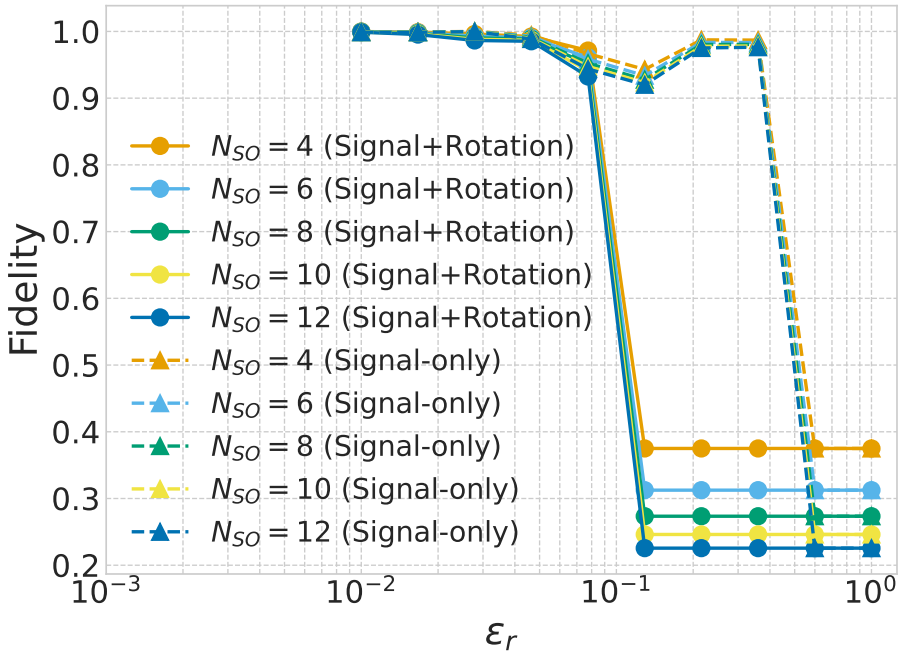


Figure 4.8: Fidelity of the GQSP implementation of $\hat{P}_{M_S}$ versus single-qubit rotation precision ϵ_r for fixed $M_S = 0$ and varying N_{SO}. Solid lines with circular markers: both the GQSP signal operator and the GQSP rotation sequence are rounded. Dashed lines with triangular markers: only the signal operator is rounded; the GQSP rotations are exact.

to one.

For the GQSVT implementation of $\hat{P}_{M_S}$ (see Fig. 4.9), the robustness was considerably lower. We examined two precision models: one in which PREP was exact and GQSVT rotations were rounded, and the other in which the rotations were exact but PREP was rounded. The results show that PREP can tolerate somewhat coarse precision, approximately 10^{-2}, without a major loss of fidelity, but GQSVT rotations require much finer control. A precision of at least $\epsilon_r \lesssim 10^{-3.7}$ is required to maintain a high fidelity. Furthermore, accurate phase choices are critical for correctly realizing the intended polynomial transformation, and the sensitivity increases with N_{SO}, owing to the accumulation of small phase errors and numerical conditioning in the polynomial construction.

Because the standard Clifford+T synthesis has a cost that scales linearly with the number of target bits, reducing ϵ_r by one decimal order (adding $\log_2 10 \approx 3.3$ bits of precision) increases the per-rotation T gate count by a moderately constant factor. Thus, the qualitative thresholds above (10^{-1} for LCU, 10^{-2} for GQSP, 10^{-4} for GQSVT) translate into predictable logarithmic increases in the T budget per rotation.

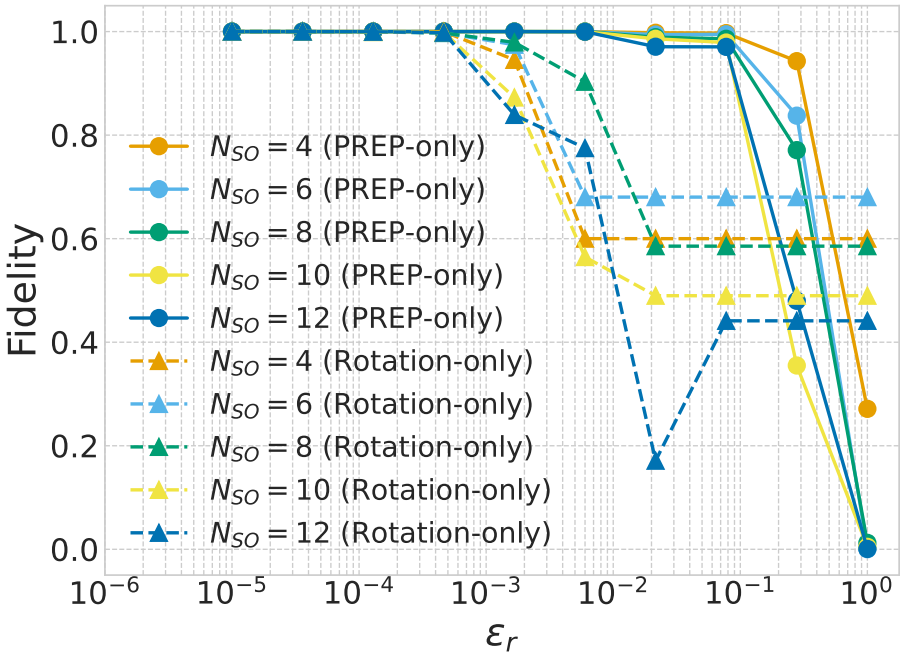


Figure 4.9: Fidelity of the GQSVT implementation of $\hat{P}_{M_S}$ versus single-qubit rotation precision ϵ_r for fixed $M_S = 0$ and varying N_{SO}. Solid lines with circular markers: GQSVT rotations are exact; PREP is rounded. Dashed lines with triangular markers: only GQSVT rotations are rounded; PREP is exact.

4.4.5 Resource Estimates

We estimated the gate counts for the concrete implementations of the projectors. This allowed us to assess practicality. We counted CNOTs, Toffolis, and single-qubit rotations, then convert to T gate counts using the angle precisions established in the previous section.

We counted the Toffoli as 6 CNOTs and 7 T gates, which is the standard decomposition [216]. A single-axis rotation to precision ϵ_r costs

$$T_{\mathrm{rot}}(\epsilon_r) \;=\; 3\log_2(1/\epsilon_r) \;+\; o(\log(1/\epsilon_r)) \tag{4.28}$$

T gates using Clifford+T synthesis [108]. When N_{rot} independent rotations appear, we allocate a per-rotation budget $\epsilon_r/N_{\mathrm{rot}}$ such that the total error is at most ϵ_r. This provides a conservative aggregate estimate, as follows:

$$T_{\mathrm{total\ rot}} \;\approx\; 3\,N_{\mathrm{rot}}\,\log_2\!\big(N_{\mathrm{rot}}/\epsilon_r\big). \tag{4.29}$$

First, we compared the T and CNOT gate counts for the $\hat{P}_{M_S}$ projector (see Fig. 4.10). The corresponding resource estimations for the particle number

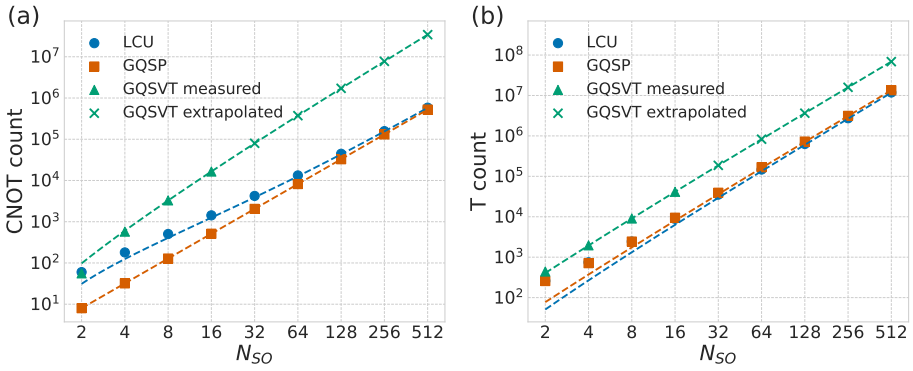


Figure 4.10: Gate counts for $\hat{P}_{M_S}$ as a function of N_{SO}: (a) CNOTs and (b) T gates, comparing LCU, GQSP, and GQSVT implementations under the precision budgets summarized in the text.

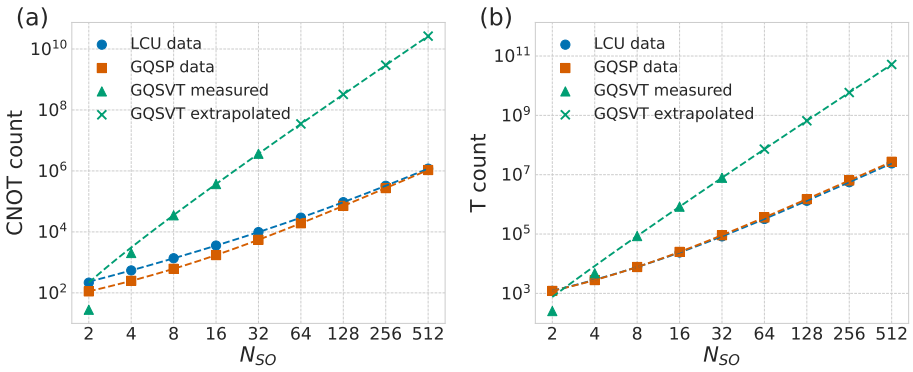


Figure 4.11: Gate counts for total-spin projectors versus N_{SO}: (a) CNOTs and (b) T gates for $\hat{P}_{S,M_S}$ via LCU and GQSP, and for $\hat{P}_{S^2}$ via GQSVT.

projector $\hat{P}_N$ are essentially the same because both rely on the same discretization structure. In all the cases, we set the number of quadrature nodes to $N_\phi = N_{\mathrm{SO}} + 1$, which guarantees the exactness of any desired value of M_S. For the LCU construction, we choose the number of LCU terms to be a power of two; thus, the PREP blocks consist only of Hadamard gates.

In Fig. 4.10(a), we observe that both the LCU and GQSP realizations of $\hat{P}_{M_S}$ exhibit quadratic scaling in the number of CNOT gates. In both cases, the dominant contribution increases as $2N_{\mathrm{SO}}$, with an additional overhead for LCU which increases approximately as $N_{\mathrm{SO}} \log N_{\mathrm{SO}}$ owing to the use of indexed unitaries in the SELECT step. Therefore, the LCU consistently requires more CNOT gates than GQSP for the same system size. However, the difference between the two becomes less pronounced as the number of spin-orbitals increases because the shared quadratic term dominates the scaling.

In contrast, the GQSVT construction exhibited worser scaling. For GQSVT, we used the standard binary-tree version of LCU. The number of CNOT gates increases roughly as $N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}}$. We were able to compute the exact counts up to $N_{\mathrm{SO}} = 16$, and the same trend was easily extrapolated to larger systems up to at least $N_{\mathrm{SO}} = 512$. This growth makes GQSVT more expensive than LCU or GQSP in terms of entangling operations.

To estimate the number of T gates, we used the values from the precision study in Sec. 4.4.4. This behavior was noticeably more uniform among the three methods (see Fig. 4.10(b)). In all cases, we observe a scaling proportional to $N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}}$, reflecting the combined effect of the number of required single-qubit rotations and the synthesis cost required to reach the target precision. Both LCU and GQSP implementations share essentially the same leading prefactor in this scaling; the numerical data indicate a coefficient of approximately 4. The two approaches differ only in additive lower-order terms; therefore, so while LCU consistently yields a slightly lower T gate count than GQSP, the difference remains small across the entire range of problem sizes investigated.

Although the T gate scaling of GQSVT with binary-tree LCU implementation has the same asymptotic form, the corresponding prefactor is significantly larger, approximately 20 in our tests. This means that at the same target fidelity, GQSVT is approximately 5 times more expensive in terms of T gates than LCU or GQSP. Thus, although all three methods share similar asymptotic scaling, GQSVT has a substantially higher constant cost and is the least practical option.

For the LCU and GQSP implementation of $\hat{P}_{S,M_S}$ projector, we use standard decomposition as $\hat{P}_{M_S}\hat{P}_S\hat{P}_{M_S}$; thus, the resource cost is determined by the implementation of $\hat{P}_{M_S}$. Operators $\hat{P}_S$ were implemented using the LCU construction in both cases. Therefore, the difference between LCU and GQSP therefore arises only in the implementation of $\hat{P}_{M_S}$. As before, we chose the number of quadrature nodes for $\hat{P}_{M_S}$ as $N_\phi = N_{\mathrm{SO}} + 1$ to ensure the exactness of any spin projection. For $\hat{P}_S$ because we focus on the singlet states ($S = 0$), only two nodes are required in the integration; therefore, the cost of this step is independent of N_{SO}.

The GQSVT approach is slightly different in that it constructs the projector

onto the desired total-spin subspace via $\hat{P}_{S^2}$ rather than by selecting a specific M_S. However, in practice, this still produces states with a well-defined total spin suitable for the same applications; therefore, we include it in the comparison.

As shown in Fig. 4.11(a), for both the LCU-based and GQSP-based implementations, the CNOT count scales quadratically with the number of spin-orbitals. The leading contribution grows as $4N_{\mathrm{SO}}$, with a slightly larger additive term for LCU because of the additional SELECT overhead inherited from $\hat{P}_{M_S}$. However, this difference becomes small in comparison with total gate count as N_{SO} increases; hence, for moderate or large system sizes, the two methods are essentially comparable in terms of CNOT cost.

By contrast, the GQSVT implementation with the binary-tree LCU exhibited a much steeper scaling. In this case, the CNOT count increases approximately as $21N_{\mathrm{SO}}^3 \log N_{\mathrm{SO}}$, reflecting both a higher-degree sequence length and a more complex signal processing structure. Consequently, GQSVT becomes significantly more expensive than LCU or GQSP when N_{SO} has more than a few orbitals. Nevertheless, for very small system sizes the situation can be reversed; for example, at $N_{\mathrm{SO}} = 2$, the GQSVT implementation requires fewer CNOTs than LCU or GQSP.

The trends in the T gate counts are shown in Fig. 4.11(b) exhibit a similar patterns. For both LCU and GQSP realizations, the T gate cost increases as $8N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}}$. In direct comparison, GQSP grows slightly faster because of its larger rotation overhead; thus, the LCU construction retains a small advantage across all the studied system sizes.

For GQSVT with the binary-tree LCU, the T gate count grows substantially more quickly, as $36.5N_{\mathrm{SO}}^3 \log N_{\mathrm{SO}}$, in line with the deeper signal-processing sequence and the increased synthesis precision required for stable polynomial filtering. Consequently, GQSVT becomes significantly more expensive than both LCU and GQSP as the system size increases. Similar to the CNOT comparison, a narrow regime exists at a very small N_{SO}, in which GQSVT can be competitive or even favorable; however, beyond this range, the cubic scaling dominates.

4.4.6 Projectors as a QPE Filter

The success probability of QPE depends directly on the overlap between the prepared input state $|\psi\rangle$ and the target eigenstate $|\phi_0\rangle$ of the Hamiltonian. In particular, the probability of obtaining the correct eigenvalue is $|\langle\phi_0|\psi\rangle|^2$. In the case of strongly correlated electronic states or when several eigenstates lie within a narrow energy window, preparing an input state with a sufficiently high overlap can be challenging. This issue is particularly pronounced for open-shell molecules or other systems with near-degeneracies.

To address this limitation, we apply symmetry projectors $\hat{P}$ prior to QPE. The operation $\hat{P}|\psi\rangle$ increases the weight of the components of $|\psi\rangle$ that are in the symmetry sector of the target eigenstate, while suppressing the components belonging to other sectors. Because $\hat{P}|\phi_0\rangle = |\phi_0\rangle$ for the eigenstates, this filtering step does not perturb the target state.

We tested the effectiveness of projector filtering in systems in which the electronic spectrum exhibited near-degeneracy.

O_2 Singlet–Triplet Gap. Molecular oxygen (O_2) is a canonical example in which the triplet ground state $X^3\Sigma_g^-$ is close to the low-lying singlets $a^1\Delta_g$ and $b^1\Sigma_g^+$. We consider an active space calculation in the CAS(8,6) space within an aug-cc-pVTZ basis, where the orbitals are optimized in a triplet CASSCF calculation. Active orbitals consisted of bonding and antibonding σ pair and bonding and antibonding 2 π pairs; the O–O bond length was set to 1.208 Å.

Within this active space, the initial state prepared using HEA (see Sec. 2.2.5.1) does not respect spin symmetry and produces a substantial overlap with both the triplet ground state and at least one singlet excited state. The numerical results are presented in Fig. 4.12. In particular, we observed comparable amplitudes in the $^3\Sigma_g^-$ and $a^1\Delta_g$ states, resulting in a QPE success probability of only $\mathrm{Pr}_{\mathrm{QPE}} \approx 0.47$.

Applying a total spin projector $\hat{P}_{S,M_S}$ with $S = 1$ and $M_S = 0$ substantially suppresses singlet contamination and increases the overlap with the triplet manifold. After symmetry filtering, the QPE success probability increases to $\mathrm{Pr}_{\mathrm{QPE}}^{\mathrm{filtered}} \approx 0.94$

We also estimated the QPE cost for the active space Hamiltonian. The qubitization-based QPE requires approximately 8 ancilla qubits to handle the energy with chemical precision (~ 1 mHa). Then, the total non-Clifford is no the order of 10^7 T gates. By contrast, the symmetry projector requires only $\sim 1.6 \times 10^4$ T gates, which is almost 3 orders of magnitude less. The intrinsic probability of success of $\hat{P}_{S,M_S}$ is ≈ 0.30. If one wishes to deterministically prepare the filtered state, one round of AA suffices to increase the success probability to ≈ 0.97, thereby the total T cost of the projection to $\sim 4.8 \times 10^4$, which remains negligible relative to the QPE cost.

Trimethylenemethane (TMM). TMM exhibits a well-studied singlet–triplet gap in the order of tens of kcal/mol. The ground state is a triplet $^3A_2'/^3B_2$, and the first excited state is a degenerate pair of singlets 1^1B_2 and 1^1A_1, located approximately 1.17 eV above the ground state [217]. We constructed a CAS(4,4) Hamiltonian in a cc-pVTZ basis, where the active orbitals spanned the full carbon π-system. The molecular orbitals were optimized using a triplet CASSCF calculation.

Using a hardware-efficient VQE ansatz to prepare an initial state nominally targeting the triplet ground state $^3A_2'/^3B_2$ results in noticeable contamination from the singlet states 1^1B_2 and 1^1A_1 (see Fig. 4.13). The resulting QPE success probability was $\mathrm{Pr}_{\mathrm{QPE}} \approx 0.46$. Applying the spin projector $\hat{P}_{S,M_S}$ with $S = 1$ and $M_S = 0$ suppresses unwanted singlet components, increases the overlap with the triplet sector, and improves the QPE success probability to $\mathrm{Pr}_{\mathrm{QPE}}^{\mathrm{filtered}} \approx 0.94$.

We also estimated the QPE cost for the active space Hamiltonian. For the CAS(4,4) Hamiltonian, qubitization-based QPE with chemically accurate energy resolution requires approximately 8 ancilla qubits on the order of 10^7 T gates. In contrast, the spin projector requires approximately 7.7×10^3 T gates, which is more than 3 orders of magnitude lower than QPE. The intrinsic success

probability of the projection is ≈ 0.29; therefore, a single step of the AA is sufficient to increase the probability above 0.98. This increases the total T cost of the projector to $\sim 2.3 \times 10^4$, which is negligible compared with the cost of the subsequent QPE.

Implications for FeMoco. The FeMoco cluster presents an extreme case with strong static correlation. Even in reduced active space models (54 [211] or 76 [212] spatial orbitals), multiple states can lie within chemically relevant energy differences. In such cases, increasing the initial overlap with the correct irreducible subspace prior to the phase estimation can significantly reduce the number of QPE repetitions required.

Recent large-scale classical calculations on the 76-orbital FeMoco model indicate that a single broken-symmetry Slater determinant can already carry a substantial amplitude in an accurate ground state approximation. In particular, overlaps between the dominant broken-symmetry determinant $|\psi_{\mathrm{BS}}\rangle$ and a highly converged matrix product state (used as an approximate ground state) reach $|\langle\psi_{\mathrm{BS}}|\phi_0\rangle| \approx 0.45$–$0.47$ for representative low-energy spin isomers, while the next-largest determinant coefficient is below 0.1 [218]. If one uses $|\psi\rangle = |\psi_{\mathrm{BS}}\rangle$ as an input guess for QPE, this corresponds to an estimated success probability $p_{\mathrm{QPE}} \approx |\langle\psi_{\mathrm{BS}}|\phi_0\rangle|^2 \approx 0.2$.

However, broken-symmetry determinants are generally not eigenstates of $\hat{S}^2$ and therefore contain components in multiple total-spin sectors. Consequently, a spin projector can serve as a natural filter that removes contributions outside the target (S, M_S) sector and concentrates the remaining weight into the desired irreducible subspace.

The QPE cost for the FeMoco system can be estimated as $\sim 10^{10}$ T gates by using tensor hypercontraction [132]. At the same time, P_{S,M_S} projector realized using GQSP would require $\sim 10^6$ T gates, which is 4 orders of magnitude less than QPE cost.

In reality, several orders will be taken by the AA, which can be roughly estimated if we assume a uniform superposition of the states of N_{elec} electrons, and total spin S and fixed M_S. The number of multiplets is

$$n(S) = \binom{N_{\mathrm{elec}}}{\frac{N_{\mathrm{elec}}}{2} + S} - \binom{N_{\mathrm{elec}}}{\frac{N_{\mathrm{elec}}}{2} + S + 1}. \tag{4.30}$$

Thus, the expected squared overlap of a random state in the $M_S = S$ sector is the ratio:

$$p_{\min} = \frac{n(S)}{\binom{N_{\mathrm{elec}}}{\frac{N_{\mathrm{elec}}}{2}}}. \tag{4.31}$$

Therefore, the estimated probability is approximately 0.03–0.07, depending on the FeMoco model (see Tab. 4.2). This probability resulted in 5-7 steps of AA. Then, the GQSP implementation of P_{S,M_S} projector together with AA costs approximately $\sim 10^7$ of the T gates, which is still a 3 order advantage over QPE itself.

Such symmetry filtering has been shown to significantly increase overlap in smaller systems, such as O_2 and TMM. By analogy, one may therefore expect

that the same strategy can substantially enhance the initial overlap for FeMoco, thereby reducing the number of required QPE repetitions. Taken together, these considerations suggest that the state preparation problem for FeMoco, while still nontrivial, may no longer be the dominant bottleneck when physically motivated broken-symmetry references are combined with symmetry projection and AA.

4.4.7 Limitations and Outlook

Although the projector-based filtering approach improves the effective overlap with the target eigenstate prior to QPE, several limitations remain. First, the probability of success of the projection step may be low when the initial state carries only a modest component in the desired symmetry sector. Although AA can boost this probability, the number of amplification rounds scales as inverse overlap. In cases in which the trial state is poorly chosen, the overhead associated with AA may become significant. This highlights the importance of constructing a good initial ansatz, either from classical multireference methods or using adaptive quantum state preparation strategies.

Although we focused on continuous symmetries ($\hat{N}$, $\hat{S}_z$, $\hat{S}^2$), the same GQSP/GQSVT frameworks can be applied to finite molecular point groups when the symmetry actions $U(g)$ are available. For any irreducible representation Γ of finite group G, the projector

$$P_\Gamma = \frac{d_\Gamma}{|G|} \sum_{g \in G} \chi_\Gamma(g)^* \, U(g) \tag{4.32}$$

selected the Γ-subspace. For abelian groups, all irreducible representations are one-dimensional and (4.32) reduces to a weighted average of symmetry operations, which is directly compatible with LCU, GQSP, and GQSVT implementations. If G decomposes as a direct or semidirect product (such as $\mathrm{D}_{3h} \cong \mathrm{C}_{3v} \times \mathrm{C}_s$ and $\mathrm{C}_{3v} \cong \mathrm{C}_3 \rtimes \mathrm{C}_s$), the corresponding projectors factorize as $P_{\mathrm{D}_{3h}} = P_{\mathrm{C}_{3v}} P_{\mathrm{C}_s}$, allowing symmetry filtering to be applied sequentially. Another method involves constructing each symmetry operator and projecting on its specific eigenvalue using the Lagrange interpolation and GQSVT.

The projector and AA can be merged into a single deterministic procedure using a double-bracket symmetry refinement. Given Hermitian generator $\hat{O}$, the evolution

$$|\psi(\tau)\rangle \propto \exp\big(-\tau[\hat{O}, [\hat{O}, \cdot]]\big) \, |\psi(0)\rangle \tag{4.33}$$

suppresses the components outside the target symmetry sector, while leaving the desired subspace invariant. A discretized implementation in terms of the controlled applications of $\hat{O}$ yields deterministic symmetry purification at the cost of an additional depth.

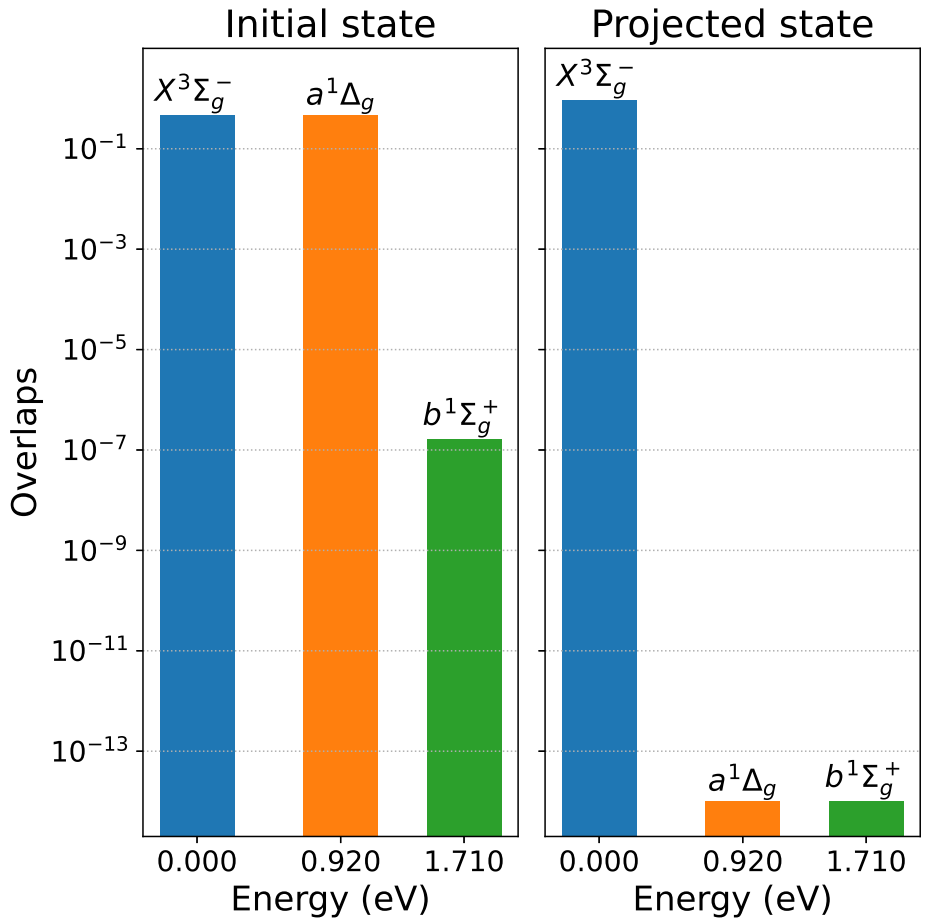


Figure 4.12: Overlap of the initial guess with the singlet ($a^1\Delta_g$, $b^1\Sigma_g^+$) and triplet ($X^3\Sigma_g^-$) states of O_2 before and after applying the $\hat{P}_{S,M_S}$ spin-symmetry projection.

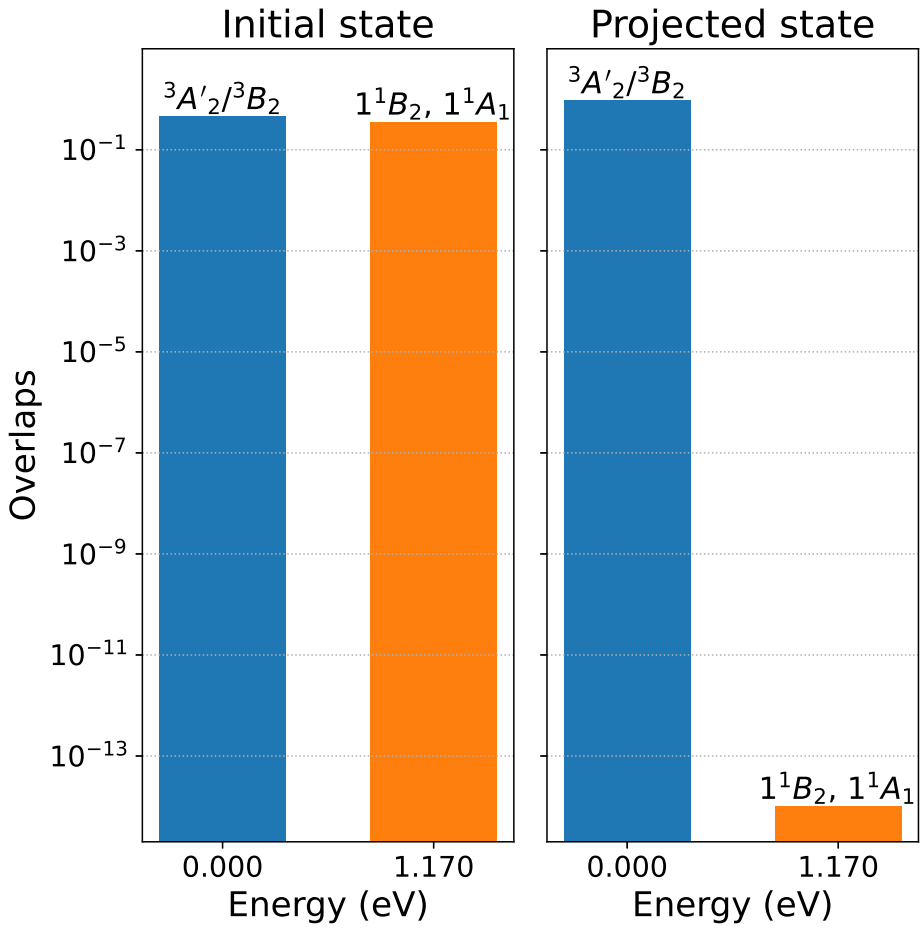


Figure 4.13: Overlap of the initial guess with the singlet (1^1B_2, 1^1A_1) and triplet ($^3A'_2/^3B_2$) states of TMM before and after applying the $\hat{P}_{S,M_S}$ spin-symmetry projection.

Table 4.2: Estimated spin-projection baseline overlaps and T gate costs for FeMoco active space models using AA. The quantity p denotes the initial overlap with the desired spin sector. Queries refers to the number of projector applications in the AA sequence. T gate counts include all projector applications in the amplification schedule.

Model	**Electrons** N_{elec}	**Target Spin** S	M_S	**Overlap** p	**Queries**	**T Gates**
FeMoco (54 orbitals, Reiher)	54	0	0	0.036	7	1.5×10^7
FeMoco (76 orbitals, Li), $S = 3/2$	113	3/2	1/2	0.068	5	1.1×10^7
FeMoco (76 orbitals, Li), $S = 1/2$	113	1/2	1/2	0.035	7	1.5×10^7

4.5 Conclusion

In this study, we constructed symmetry projectors for quantum state preparation based on LCU, GQSP, and GQSVT, together with circuit realizations and parameter choices that achieve a prescribed projection accuracy. We analyzed the numerical accuracy, gate-precision sensitivity, and asymptotic resource scaling, and translated them into concrete CNOT and T gate estimates. The comparison isolated the main trade-offs among the three approaches: LCU was the most rotation-robust and simplest to parameterize; GQSP achieved similar leading-order T gate scaling with minimal ancilla overhead (one qubit) but requires slightly tighter rotations; GQSVT attained exact polynomial filtering on discrete spectra but was more phase-sensitive and carried larger constant factors in both entangling and non-Clifford resources.

For $\hat{N}$ and $\hat{S}_z$, both LCU and GQSP exhibited quasi-quadratic scaling in the system size, with T gate counts that grew as $N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}}$ and differed only in small additive terms; LCU maintained a slight advantage in the total T cost. GQSVT showed a steeper entangling cost (CNOT scaling roughly as $N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}}$) and a larger T gate prefactor, making it several times more expensive at a fixed accuracy, although it can be competitive at very small N_{SO}. For total-spin filtering, the composition $\hat{P}_{S,M_S} = \hat{P}_{M_S}\hat{P}_S\hat{P}_{M_S}$ introduces the cost of preparing fixed states S into the same quasi-quadratic T gate scaling class as $\hat{N}$ and $\hat{S}_z$. A key result is that the β integration for $\hat{P}_S$ depends weakly on the target spin S and is independent of N_{SO}.

Our precision analysis provided practical synthesis targets. The LCU remained accurate with a single-qubit rotation precision of approximately 10^{-1}. For GQSP, signal-operator angles of approximately $10^{-0.8}$ and QSP phase rotations better than 10^{-1} were typically sufficient, with 10^{-2} being a safe choice. GQSVT required substantially tighter phase control, with a rotation precision near $10^{-3.7}$ whereas PREP-type data paths tolerate $\sim 10^{-2}$. These thresholds were mapped directly to the T gate counts via standard Clifford+T synthesis, implying only a logarithmic overhead per additional decade of precision.

We demonstrated that symmetry projectors act as effective QPE filters in representative molecular systems. In O_2 and TMM, symmetry filtering increased the QPE success probability from ~ 0.46 to ~ 0.94 while incurring projector costs is 3 orders of magnitude lower that of the subsequent QPE. Even for FeMoco-scale active spaces, the total cost of projection plus AA remained $\sim 10^7$ T gates versus $\sim 10^{10}$ for a single qubitization-based QPE with chemical accuracy, yielding a substantial net reduction in the number of repetitions required to resolve the desired eigenvalue.

Overall, symmetry filtering is a low-cost, accuracy-preserving preprocessing step that scales favorably with the system size and materially improves QPE success probability. Given its modest ancilla footprint, straightforward parameterization, and clear synthesis targets, we expect projector-based state preparation to become a standard component of fault-tolerant quantum simulation workflows, especially for strongly correlated molecules where resolving the correct symmetry sector is essential.

Chapter 5

Enhancing Quantum Computations with the Synergy of Auxiliary-Field Quantum Monte Carlo and Computational Basis Tomography

We introduce QC-CBT-AFQMC, a hybrid algorithm that incorporates CBT into the QC-AFQMC method proposed by Huggins *et al.* [*Nature* **603**, 416-420 (2022)], replacing the use of classical shadows. While the original QC-AFQMC showed high accuracy for quantum chemistry calculations, it required exponentially costly post-processing. Subsequent work using matchgate shadows [Wan *et al., Commun. Math. Phys.* **404**, 629 (2023)] improved scalability, but still suffers from prohibitive computational requirements that limit practical applications. Our QC-CBT-AFQMC approach uses shallow Clifford circuits with a quadratic reduction of two-qubit gates over the original algorithm, significantly reducing computational requirements and enabling accurate calculations under limited measurement budgets. We demonstrate its effectiveness on the hydroxyl radical, ethylene, and nitrogen molecule, producing potential energy curves that closely match established benchmarks. We also examine the influence of CBT measurement counts on accuracy, showing that subtracting the active space AFQMC energy mitigates measurement-induced errors. Furthermore, we apply QC-CBT-AFQMC to estimate reaction barriers in [3+2] cycloaddition reactions, achieving agreement with high-level references and successfully incorporating complete basis set extrapolation techniques. These results highlight QC-CBT-

AFQMC as a practical quantum-classical hybrid method that bridges the capabilities of quantum devices and accurate chemical simulations. This chapter is based on the published article [Khinevich and Mizukami, *Phys. Rev. Res.* **7**, 043344 (2025)].

5.1 Introduction

Over the past decade, quantum computing has significantly advanced, particularly in quantum chemistry [148, 149]. Many quantum algorithms for quantum chemistry have been proposed, among which the VQE [22] and QPE [15, 150] are widely recognized as representative methods. Nonetheless, existing algorithms, including both VQE and QPE, are resource intensive, making their application to full electronic structure Hamiltonians impractical.

To address this, the active space approach from classical computational chemistry is essential, which involves selecting critical molecular orbitals and electrons, thus focusing on static electron correlation. Quantum computers efficiently handle rapidly increasing configurations within the active space. However, capturing the dynamical correlation outside this space remains challenging. There are various attempts to incorporate dynamical correlation after quantum computing. One of such methods is the multireference perturbation theory [219, 143, 220, 221]. However, this requires heavy measurements of 3,4-reduced density matrices, or other quantities [222]. Other approaches involve explicitly correlated methods, such as F12 and transcorrelation methods [223, 224, 225, 226]. For F12 methods, the inclusion of the Coulomb cusp condition alone is insufficient to adequately capture dynamical correlation, while the transcorrelation methods are limited by the emergence of three-body terms and loss of Hermicity, creating practical implementation bottlenecks.

Huggins *et al.* proposed a quantum-classical hybrid algorithm, named QC-AFQMC [42], where a quantum-prepared trial wave function is used to reduce bias in phaseless AFQMC (ph-AFQMC) [227]. Although QC-AFQMC was originally introduced to improve an AFQMC calculation with quantum computing, it can also be viewed as a post-processing scheme that describes residual dynamical correlation beyond the active space [228, 229]. In the framework of QC-AFQMC, a quantum computer is used to prepare a trial wave function, employed to handle the infamous sign problem, for the AFQMC method. A significant advantage of AFQMC is its scaling $O(N^3) - O(N^4)$ with respect to the system size N [230, 231], which is close to HF. The favorable scaling ensures that AFQMC remains computationally tractable for larger systems. These advantages have led AFQMC methods to become increasingly popular for capturing dynamic correlation effects [232, 233, 234]. In the original QC-AFQMC paper, the translation of the quantum wave function into its classical counterpart was achieved using Clifford shadow tomography [117]. It was found that QC-AFQMC is inherently noise resilient, since the Monte Carlo step effectively corrected the errors introduced by the noisy quantum step. However, this method requires

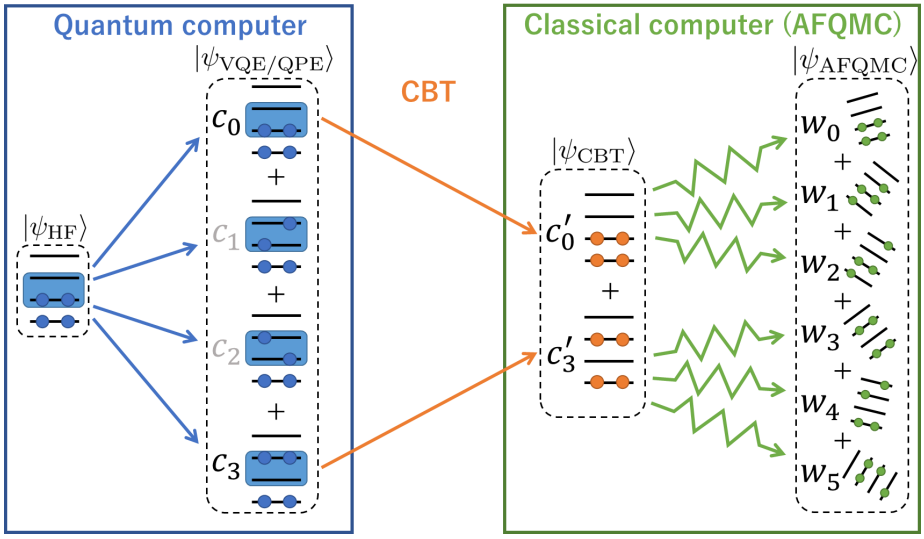


Figure 5.1: Schematic of the QC-CBT-AFQMC methodology. Blue frames highlight the selected active space for the quantum computer part. The CBT algorithm captures the most important configurations from the quantum computer state and translates them to the classical computer. AFQMC generates walkers from the CBT state, where the orbitals also rotate. This is depicted as rotated orbital diagrams.

one to use the post-processing with exponential scaling on classical computers.

In order to overcome this bottleneck, the matchgate shadow tomography technique was used [43]. Although this approach allows to reduce the scaling to polynomial, the degree of this polynomial is much higher than that of AFQMC itself [44]. The most expensive operation is to compute the local energy for each walker at each time step, which scales as $O(N_q^{8.5}\log^2(N_q))$ with the number of qubits N_q. Another problem of this approach is the large constant prefactor, due to the high sampling requirements. When specific physical quantities or information are targeted, this method may not be the optimal choice in terms of measurement count due to shot budget constraints [235]. Also, it was previously noted that all implementations of QC-AFQMC may require an exponential number of shots due to the exponential decay of the overlap of the trial state with the walkers [236, 237].

To overcome these limitations of QC-AFQMC for practical calculations, we introduce the QC-CBT-AFQMC method (see Fig. 5.1), which applies CBT [238] to QC-AFQMC as a technique to extract wave function information from quantum computers in an efficient way to use on classical computers. The practicality of our approach comes from a lower requirement for the number of shots. A smaller constant prefactor and better scaling $O(N_q^4)$ than for matchgates allowed us to do calculations for much larger systems than previously. Also,

CBT utilizes shallow circuits, built from only Clifford gates; our method can be implemented even on near-term quantum devices. CBT circuits require $O(N_q)$ number of CNOT gates. The depth of the CBT circuit can also be reduced and even made constant due to the properties of the fan-out gates [239]. Using CBT, we have already succeeded in combining coupled cluster and quantum computing through a method called tailored coupled cluster [147].

The structure of this chapter is as follows. Section 5.2 summarizes the CBT and outlines the integration of CBT into the QC-AFQMC workflow. Section 5.3 details our numerical experiments, examining potential energy curves and the effect of the shot budget on accuracy. Finally, Sec. 5.4 concludes with an outlook on further improvements and possible expansions of this methodology, including the incorporation of alternative approaches.

5.2 Methods

5.2.1 Computational Basis Tomography

In quantum computing, a wave function is typically represented as a linear combination of basis functions. An arbitrary wave function can be decomposed into a series of bit strings of length N as follows:

$$|\psi\rangle = \sum_{n=0}^{2^N-1} \langle n|\psi\rangle |n\rangle. \tag{5.1}$$

However, extracting the coefficients $\langle n|\psi\rangle$ from a quantum device poses a challenge, as they are encoded as quantum information. The retrieval of such data for use in hybrid quantum-classical algorithms requires measurement. To do this, we utilized an efficient measurement strategy known as CBT.

The CBT technique was designed to extract classical information about a wave function that is encoded on a quantum device. This extraction is achieved through projection measurements executed within the computational basis. The foundation of CBT lies in the computational basis sampling method, which was originally used for estimating expectation values, as described previously [238]. CBT can extract CI coefficients using shallow Clifford circuits, making it appealing for systems where shot budgets are limited. Although CBT may particularly benefit states where only a limited number of configurations dominate, our primary motivation is its practical feasibility: it requires fewer two-qubit gates and less circuit depth and allows efficient coefficient extraction even within the constraints of near-term quantum devices.

The initial phase of CBT involves determining the absolute magnitudes of the coefficients $\langle n|\psi\rangle$ shown in Eq. (5.1). To facilitate this, it is necessary to generate N_f identical instances of the wave function $|\psi\rangle$ and perform measurements within the computational basis, which encompasses all possible bit strings $|n\rangle_{n=0}^{2^N-1}$ of length N. Consequently, the absolute values of the coefficients are estimated

using the formula

$$|\langle n|\psi\rangle| \approx \sqrt{\frac{N_n}{N_f}}, \tag{5.2}$$

where N_n denotes the frequency of observing the specific bit string n.

After these are measured, we select several of the most significant bit strings in Eq. (5.1). We assume that, if we choose an appropriate single-particle basis, there is a relatively small set of configurations in chemistry-related systems that carry non-negligible weight, specifically those that do not exhibit exponential or high-polynomial decay with respect to system size. Given this assumption, we choose to keep only the R most probable outcomes; our method does not work when this assumption breaks down, say, for quantum states with uniformly exponentially decaying CI coefficients. This step reduces the number of necessary measurements while preserving the accuracy of the wave function.

At this point, the absolute magnitudes of the R chosen coefficients $\langle n|\psi\rangle$ have been determined. Nevertheless, this information alone is insufficient for reconstructing the wave function due to the loss of phase information. To recover the phase details, we employed the following equation:

$$e^{i(\phi_n-\phi_m)} = \frac{\langle n|\psi\rangle\langle\psi|m\rangle}{|\langle n|\psi\rangle||\langle\psi|m\rangle|}. \tag{5.3}$$

It is important to note that phase determination is achievable only up to an arbitrary global phase. Thus, the phase ϕ_m was arbitrarily set to zero, which allows the calculation of all other phases relative to ϕ_m.

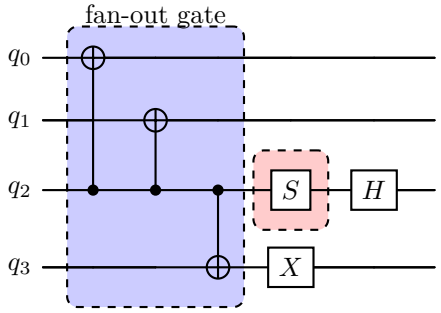


Figure 5.2: Quantum circuit implementing the operators $\langle 0000|U_{1,14} = (\langle 0001| + \langle 1110|)/\sqrt{2}$ and $\langle 0000|V_{1,14} = (\langle 0001| + i\langle 1110|)/\sqrt{2}$. The blue dashed box indicates a fan-out gate from the flag qubit q_2 to the other qubits, corresponding to different bits in binary representations of 1 and 14. The additional S gate transforms $U_{1,14}$ into $V_{1,14}$.

To utilize Eq. (5.3), it is necessary to compute the value of $\langle n|\psi\rangle\langle\psi|m\rangle$, referred to as the *interference factor*, with the denominator already determined from the previous step. The auxiliary relationship was determined as follows:

$$\langle n|\psi\rangle\langle\psi|m\rangle = \mathcal{A}_{n,m} + i\mathcal{B}_{n,m} - \frac{1+i}{2}(|\langle n|\psi\rangle|^2 + |\langle m|\psi\rangle|^2), \tag{5.4}$$

where the terms $\mathcal{A}_{n,m}$ and $\mathcal{B}_{n,m}$ are defined by the equations

$$\mathcal{A}_{n,m} \equiv |\frac{\langle n| + \langle m|}{\sqrt{2}}|\psi\rangle|^2,\ \mathcal{B}_{n,m} \equiv |\frac{\langle n| + i\langle m|}{\sqrt{2}}|\psi\rangle|^2. \tag{5.5}$$

These quantities represent the only unknowns in Eq. (5.4). Ancillary unitary operators $U_{n,m}$ and $V_{n,m}$ were constructed as described in Ref. [238]: $U^\dagger_{m,n}|0\rangle = (|m\rangle + |n\rangle)/\sqrt{2}$ and $V^\dagger_{m,n}|0\rangle = (|m\rangle - i|n\rangle)/\sqrt{2}$. These operators can be written in Clifford circuits, each of which requires only Hamming(m, n) of CNOT gates. The construction proceeds by identifying the qubit position at which the binary representations of m and n differ to use it as a flag and applying a Hadamard gate there. Then, for each remaining bit that differs, a CNOT gate is applied with the control on the flag qubit and the target on the differing bit. To construct $V_{m,n}$ instead of $U_{m,n}$, an additional S gate is applied to introduce the required relative phase. An example of this construction for $m = 1$ and $n = 14$ is shown in Fig. 5.2. This approach guarantees that each such unitary requires at most $O(N_q)$ CNOT gates and CNOT layers, whereas arbitrary matchgate circuits typically require $O(N_q^2)$ gates. Also, these CNOT gates used in $U_{n,m}$ and $V_{n,m}$ comprise fan-out gates, allowing to implement it with constant CNOT layers using midcircuit measurements [239]. The operators $U_{n,m}$ and $V_{n,m}$ were incorporated to redefine $\mathcal{A}_{n,m}$ and $\mathcal{B}_{n,m}$ as

$$\mathcal{A}_{n,m} = |\langle 0|U_{n,m}|\psi\rangle|^2, \quad \mathcal{B}_{n,m} = |\langle 0|V_{n,m}|\psi\rangle|^2. \tag{5.6}$$

From Eq. (5.6), it is evident that the values of $\mathcal{A}_{n,m}$ and $\mathcal{B}_{n,m}$ can be derived from measurements on the zero state $|0\rangle$. Consequently, to evaluate $\mathcal{A}_{n,m}$ and $\mathcal{B}_{n,m}$ for each distinct bit string pair $n \neq m$, N_a and N_b measurements must be performed, respectively, yielding:

$$\mathcal{A}_{n,m} \approx \sqrt{\frac{N_0^{(a)}}{N_a}}, \quad \mathcal{B}_{n,m} \approx \sqrt{\frac{N_0^{(b)}}{N_b}}, \tag{5.7}$$

where $N_0^{(a)}$ and $N_0^{(b)}$ indicate the counts of zero outcomes for the respective measurement sets. With these estimates, Eq. (5.4) is applied to determine the phases for all coefficients $\langle n|\psi\rangle$. Thus, the complete CBT protocol necessitates a total of $N_f + (R-1)(N_a + N_b)$ measurements of the state $|\psi\rangle$, where R is the number of selected bit strings for the correct representation of the wave function. The CBT wave function has the following form:

$$|\psi_{\text{CBT}}\rangle = \sum_{n=0}^{R} c'_n|n\rangle, \tag{5.8}$$

where $c'_n = |\langle n|\psi\rangle|e^{i(\phi_n - \phi_0)}$ are normalized coefficients, obtained from Eqs. (5.2) and (5.3). Both the absolute values and phases are estimated from sampling results, and therefore carry statistical errors.

5.2.2 Auxiliary-Field QMC

The AFQMC method is a state-of-the-art approach for obtaining high-precision solutions to the time-independent Schrödinger equation [89]. It offers computational advantages over both deterministic and stochastic high-accuracy methods. AFQMC retains the low, HF–like scaling of $O(N^3-N^4)$, yet delivers precision comparable to high-precision methods, such as FCI, MRPT, and CC. It also has advantages over DMC and full configuration interaction QMC, since it uses a Slater determinant trial wave function that controls the sign problem and can be efficiently processed on GPU or multicore CPU architectures [240].

In AFQMC, the ground-state wave function is obtained by evolving an initial state $|\psi(0)\rangle$ in imaginary time:

$$-\frac{d}{d\tau}|\psi(\tau)\rangle = (\hat{H} - E_0)|\psi(\tau)\rangle, \tag{5.9}$$

where $\tau = it$ is the imaginary time and E_0 is the energy of the ground state. The solution is as follows:

$$|\psi(\tau)\rangle = e^{-\tau(\hat{H}-E_0)}|\psi(0)\rangle. \tag{5.10}$$

As $\tau \to \infty$, $|\psi(\tau)\rangle$ is projected onto the ground state, provided the initial state $|\psi(0)\rangle$ has nonzero overlap with it. The parameter E_0 needs to be reasonably close to the true ground-state energy to ensure stable propagation; a closer estimate further improves numerical stability and reduces population fluctuations. Similarly, an initial state closer to the ground state improves convergence. In practice, E_0 is iteratively refined alongside the wave function updates.

To compute the propagation in Eq. (5.10), time discretization is used:

$$e^{-\tau(\hat{H}-E_0)} = \left(e^{-\Delta\tau(\hat{H}-E_0)}\right)^n, \tag{5.11}$$

with $\Delta\tau = \tau/n$.

Direct computation is as challenging as solving the full configuration interaction problem. AFQMC circumvents this by using a Hubbard–Stratonovich transformation to decouple the two-body interactions into one-body operators coupled to stochastic auxiliary fields. It converts the interacting system into noninteracting particles under a fluctuating external potential.

To apply this transformation, one first expresses the Hamiltonian in quadratic form via the Cholesky decomposition:

$$\hat{H} = \hat{\nu}_0 - \frac{1}{2}\sum_{i=1}^{N_{\max}} \hat{\nu}_i^2, \tag{5.12}$$

where

$$\hat{\nu}_0 = \sum_{pq} h_{pq} \sum_{\sigma} \hat{a}_{p\sigma}^\dagger \hat{a}_{q\sigma}, \hat{\nu}_i = \mathrm{i}\sum_{pq} L_{pq}^i \sum_{\sigma} \hat{a}_{p\sigma}^\dagger \hat{a}_{q\sigma}.$$

Here, $\hat{a}^\dagger_{p\sigma}$ and $\hat{a}_{q\sigma}$ are the fermionic creation and annihilation operators for an electron in spatial orbital p or q with spin projection σ. The indices p and q label a set of orthonormal single-particle spatial orbitals. The coefficients h_{pq} are the matrix elements of the one-electron part of the Hamiltonian. The quantities L^i_{pq} are obtained from the Cholesky decomposition of the tensor of the two-electron repulsion integrals V_{pqrs}, approximated as $V_{pqrs} \approx \sum_i L^i_{pr} L^i_{qs}$.

Since $\hat{\nu}_0$ and $\hat{\nu}_i$ generally do not commute, the first-order Trotter decomposition is used:

$$e^{-\Delta\tau(\hat{H}-E_0)} \approx e^{-\frac{\Delta\tau}{2}\hat{\nu}_0} \prod_i e^{\frac{\Delta\tau}{2}\hat{\nu}_i^2} e^{-\frac{\Delta\tau}{2}\hat{\nu}_0}. \tag{5.13}$$

Now, the propagator is ready to apply the Hubbard-Stratonovich transformation, yielding the following:

$$e^{-\Delta\tau(\hat{H}-E_0)} \approx \int d\vec{x}\, p(\vec{x}) \hat{B}(\vec{x}), \tag{5.14}$$

where $p(\vec{x})$ is a normal distribution over the auxiliary field $\vec{x}$ and $\hat{B}(\vec{x}) = \exp\left(-\Delta\tau(\hat{\nu}_0 - E_0) + \sqrt{\Delta\tau}\,\vec{x}\cdot\hat{\vec{\nu}}\right)$ is the evolution of a single particle. Here, $\vec{x}$ is an auxiliary field that simulates particle interactions by coupling to this field.

Equation (5.14) allows propagation of the initial wave function, with the integral over $\vec{x}$ evaluated via Monte Carlo techniques. However, direct use faces the sign problem, which can be mitigated by the ph-AFQMC [227].

In ph-AFQMC, an importance function $I(\vec{x};\psi(\tau))$ is incorporated into the propagator:

$$e^{-\Delta\tau(\hat{H}-E_0)} \approx \int d\vec{x}\, p(\vec{x}) I(\vec{x};\psi(\tau)) \hat{B}(\vec{x} - \langle\vec{x}\rangle). \tag{5.15}$$

The sign problem arises from the fact that the coefficients $\hat{\nu}_i$ may carry arbitrary complex phases, including negative real values. In the multireference case, the propagator alters the phases of Slater determinants and makes orbitals complex, potentially leading to cancellations in the Monte Carlo averaging and thus causing divergence of the wave function and energy. The phaseless approximation uses an importance function to prevent abrupt phase changes, ensuring that the coefficients remain sufficiently large. A commonly used importance function is

$$I(\vec{x};\psi(\tau)) = e^{-\Delta\tau(\mathrm{Re}(E_L)-E_0)} \times \max\left(0, \cos(\Delta\theta)\right), \tag{5.16}$$

where $\Delta\theta = \mathrm{Arg}\left(\frac{\langle\psi_T|\psi_{k+1,w}\rangle}{\langle\psi_T|\psi_{k,w}\rangle}\right)$ represents the change in the argument of the overlap between the trial wave function $|\psi_T\rangle$ and the propagated wave function. The local energy is $E_L(\tau) = \langle\psi_T|\hat{H}|\psi(\tau)\rangle / \langle\psi_T|\psi(\tau)\rangle$. The force bias $\langle\vec{x}\rangle = -\sqrt{\Delta\tau}\langle\psi_T|\hat{\vec{\nu}}|\psi(\tau)\rangle / \langle\psi_T|\psi(\tau)\rangle$ minimizes the fluctuations in $I(\vec{x};\psi(\tau))$. The trial wave function $|\psi_T\rangle$ is typically the same as the initial guess but can differ.

The evolution of Slater determinants and their weights is described by

$$\begin{aligned} |\psi_{k+1,w}\rangle &= \hat{B}(\vec{x}_{k,w} - \langle\vec{x}_{k,w}\rangle)|\psi_{k,w}\rangle, \\ W_{k+1,w} &= I(\vec{x}_{k,w};\psi_{k,w}) W_{k,w}, \end{aligned} \tag{5.17}$$

where k is the time step and w indexes the Slater determinants.

The energy is calculated using

$$E(\tau) = \frac{\sum_w W_w(\tau) E_L(\tau)}{\sum_w W_w(\tau)}, \tag{5.18}$$

with $W_w(\tau) = W_{w,k}$ at $\tau = k\Delta\tau$.

5.2.3 QC-CBT-AFQMC

The essential features of QC-CBT-AFQMC are depicted in Fig. 5.1. The initial phase of the QC-CBT-AFQMC method involves executing a quantum algorithm, such as VQE or QPE, within a predefined active space. This step is crucial for capturing the electron correlation within the selected active space, which predominantly consists of static correlation in systems that require multireference approaches. This phase yields estimates of the energy $E^{\text{act}}_{\text{VQE/QPE}}$ and the initial wave function $|\psi_{\text{VQE/QPE}}\rangle$. However, this wave function, encapsulated as quantum information, is not directly applicable to subsequent AFQMC calculations.

CBT is used to convert the acquired wave function for use in the classical part of the algorithm. By preparing the state $|\psi_{\text{VQE/QPE}}\rangle$ and conducting measurements as outlined in the CBT framework, the CI coefficients are reconstructed, and the $|\psi_{\text{CBT}}\rangle$ wave function is formulated. CBT works under the assumption that a relatively small set of configurations have reasonably large coefficients. If, despite orbital optimization, all coefficients were uniformly exponentially decaying, CBT would break down; however, in such a pathological regime, any tomography-based QC-AFQMC scheme would likewise fail, as no meaningful multideterminant trial state could be constructed. The CBT wave function serves as the initial guess for the AFQMC computation. It is important to note that the $|\psi_{\text{CBT}}\rangle$ wave function now incorporates errors originating from both the VQE/QPE and CBT processes because, in practice, the number of measurements and described configurations of $|\psi_{\text{VQE/QPE}}\rangle$ is limited.

An important advantage of QC-CBT-AFQMC over the matchgate shadows approach [43] lies in the calculation of the local energy $E_L(\tau)$ and the force bias $\langle\vec{x}\rangle$ (Eqs. (5.16) and (5.18)). In our method, these quantities are evaluated using the classically reconstructed multideterminant trial wave function $|\psi_{\text{CBT}}\rangle$, in exactly the same way as conventional ph-AFQMC with multi-Slater determinants as a trial wave function. The necessary overlaps and Hamiltonian matrix elements between $|\psi_{\text{CBT}}\rangle$ and the Slater determinant walkers can be efficiently computed, scaling polynomially with the number of qubits (active space size) N_q and linearly with the number of selected determinants R. As a result, the classical computational cost within the AFQMC loop remains dominated by standard steps that scale as $O(N_q^4)$, consistent with conventional AFQMC methods. In contrast, for matchgate shadows, the local energy must be repeatedly estimated for each walker at each time step, and this procedure is the most computationally demanding part of the approach, scaling as $O(N_q^{8.5}\log^2(N_q))$ with active space

size [44]. Moreover, CBT allows for accurate computation and storage of amplitudes and their ratios, even for truncated trial wave functions, which is particularly important for mitigating the problem of exponential decay in the overlap between the trial wave function and AFQMC walkers. This issue otherwise leads to an exponential number of shots required to reconstruct the trial wave function [236, 237].

Once the $|\psi_{\mathrm{CBT}}\rangle$ wave function is reconstructed, it is used as a trial wave function for the AFQMC calculation. No further CBT is performed, and all subsequent calculations are carried out on a classical computer. AFQMC is used to address the dynamical electron correlation missing from the active space. We assume that the errors from the CBT measurements are larger than those from the VQE or QPE steps. To isolate the correlation energy external to the active space, the energy computed using the full electronic Hamiltonian via AFQMC, $E^{\mathrm{Full}}_{\mathrm{AFQMC}}$, is adjusted by subtracting the energy determined using the active space Hamiltonian, $E^{\mathrm{act}}_{\mathrm{AFQMC}}$. Therefore, for each trial wave function $|\psi_{\mathrm{CBT}}\rangle$, two distinct AFQMC evaluations are conducted.

Upon gathering all necessary data, the overall energy within the QC-CBT-AFQMC framework, $E_{\text{QC-CBT-AFQMC}}$, is calculated as follows:

$$E_{\text{QC-CBT-AFQMC}} = E^{\mathrm{act}}_{\mathrm{VQE/QPE}} + (E^{\mathrm{Full}}_{\mathrm{AFQMC}} - E^{\mathrm{act}}_{\mathrm{AFQMC}}). \tag{5.19}$$

This equation assumes that errors from CBT measurements will cancel out, while the correlation energy external to the active space is cumulatively added to $E^{\mathrm{act}}_{\mathrm{VQE/QPE}}$. Subsequent analysis results demonstrate that this assumption is true, especially when CBT-induced errors are greater than those from the VQE/QPE stages. Virtually the same correction was used in our tailored coupled cluster study with CBT [147].

5.3 Results and Discussion

This section demonstrates the efficacy of the proposed method for constructing PECs of selected molecular systems, including the hydroxyl radical, ethylene, and nitrogen molecules. We compare the PECs generated via QC-CBT-AFQMC with those obtained from established computational techniques such as multireference configuration interaction, multireference perturbation theory, and coupled cluster methods. Furthermore, the impact of the number of CBT measurements on the accuracy of the results is evaluated, both with and without the application of Eq. (5.19). Additionally, the reaction barriers for two [3+2] cycloaddition reactions are evaluated, and the application of CBS extrapolation techniques to enhance QC-CBT-AFQMC calculations is discussed. Although a detailed resource analysis of matchgate-shadow-based QC-AFQMC has recently been published [44], the corresponding implementation has not yet been released in an open source repository. Since developing one would require considerable additional effort, we did not include the computational results of this method as a direct reference in the present study.

The initiation of the QC-CBT-AFQMC computational workflow involved executing the VQE using the UCCSD ansatz, facilitated by the Chemqulacs package [241]. We assumed that the VQE stage errors are much smaller than those introduced by the CBT measurements. As a result, we performed all VQE calculations in a noiseless setting. Subsequent steps leveraging CBT were developed using the Qulacs framework [185]. For the AFQMC step, the Ipie package was used to ensure efficient quantum Monte Carlo simulations [240]. Benchmark PECs were generated using various computational approaches, including HF, CASCI (i.e., exact diagonalization within the active space), strongly contracted NEVPT2 (SC-NEVPT2), and CCSD(T), all implemented within the PySCF package [181]. The selection of active spaces and basis sets was tailored specifically to each molecule, as described later in detail. Given the unsatisfactory performance of CCSD(T) PEC for N_2, MRCISD+Q was adopted as a more accurate reference method. Due to the well-documented size-consistency issues of MRCISD, the Davidson correction was applied to mitigate this shortcoming [84]. The MRCISD+Q method was conducted using the *ab initio* Orca package [242].

5.3.1 Potential Energy Curves

Initially, the PECs of the hydroxyl radical were mapped as a function of the O-H bond length. Utilizing both VQE/UCCSD and CASCI methodologies, we delineated an active space comprising three spatial orbitals to encapsulate the 1s atomic orbital of the hydrogen and 2s and $2p_x$ atomic orbitals of the oxygen that form the σ bond, with three electrons shared among them. The selected basis set was the augmented aug-cc-pVDZ Dunning basis, enriched with diffuse functions to accurately capture the formation of radicals during the dissociation process. For the AFQMC calculations, simulations were performed with 10000 blocks, where each block corresponds to a fixed number of propagation steps used for statistical sampling and averaging, 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}. In the CBT phase, we aimed for a maximum of $R = 5$ determinants, setting the measurement counts for N_f, N_a, and N_b at 10^6, resulting in a total of approximately 9×10^6 measurements. The maximum number of CNOT gates in the CBT circuits is 4 for fully connected qubits.

Figure 5.3 presents the constructed PECs for the hydroxyl radical. Notably, the restricted open-shell HF (ROHF) method diverges as the bond length increases indefinitely, highlighting the multireferential nature of this extensively stretched system. A similar behavior is observed with AFQMC based on the ROHF trial wave function. Conversely, both CASCI and its quantum computing counterpart, VQE, accurately capture the energy plateau at long bond lengths, with almost indistinguishable PEC results. However, neither CASCI(3,3) nor VQE(3,3) accurately depicted the PEC shape in the intermediate region (around 1.5 Å), which is due to the absence of dynamic electron correlation. In contrast, the QC-CBT-AFQMC approach yields a PEC similar to that derived from CCSD(T) and MRCISD+Q and matches the SC-NEVPT2/CASCI(3,3) PEC with minimal deviation. Only small discrepancies observed in the intermediate bond length range affirm the reliability of our method, positioning it on par with

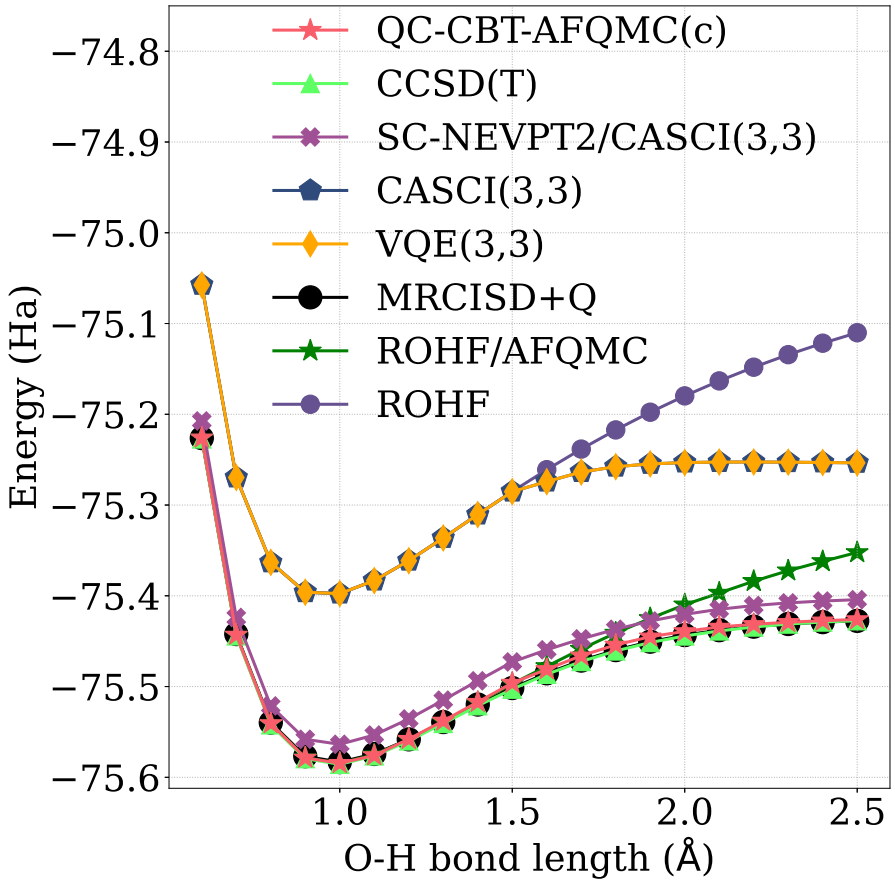


Figure 5.3: Potential energy curves for the OH radical obtained using a (3,3) active space and the aug-cc-pVDZ basis set. These results highlight how QC-CBT-AFQMC recovers dynamical correlation effects missing from the active space alone, closely matching benchmark methods and accurately describing the bond dissociation.

CCSD(T), MRCISD+Q, and NEVPT2/CASCI(3,3) for analyzing this system.

Next, PECs were plotted for the ethylene molecule, which was analyzed in terms of the changes in the dihedral HC-CH angle, specifically the rotation around the double bond. In the applications of both VQE/UCCSD and CASCI, an active space of 2 orbitals was defined to encompass the $2p_z$ carbon atomic orbitals that constitute the π bond, with 2 electrons shared between them. The cc-pVDZ Dunning basis set was chosen for this study. For the AFQMC step, the configuration included 10000 blocks, 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}. In the CBT approach, we aimed for a maximum of $R = 2$

determinants, with N_f, N_a, and N_b set to 10^6 measurements each, resulting in a total of 3×10^6 measurements. The maximum number of CNOT gates in the CBT circuits is 4 for fully connected qubits.

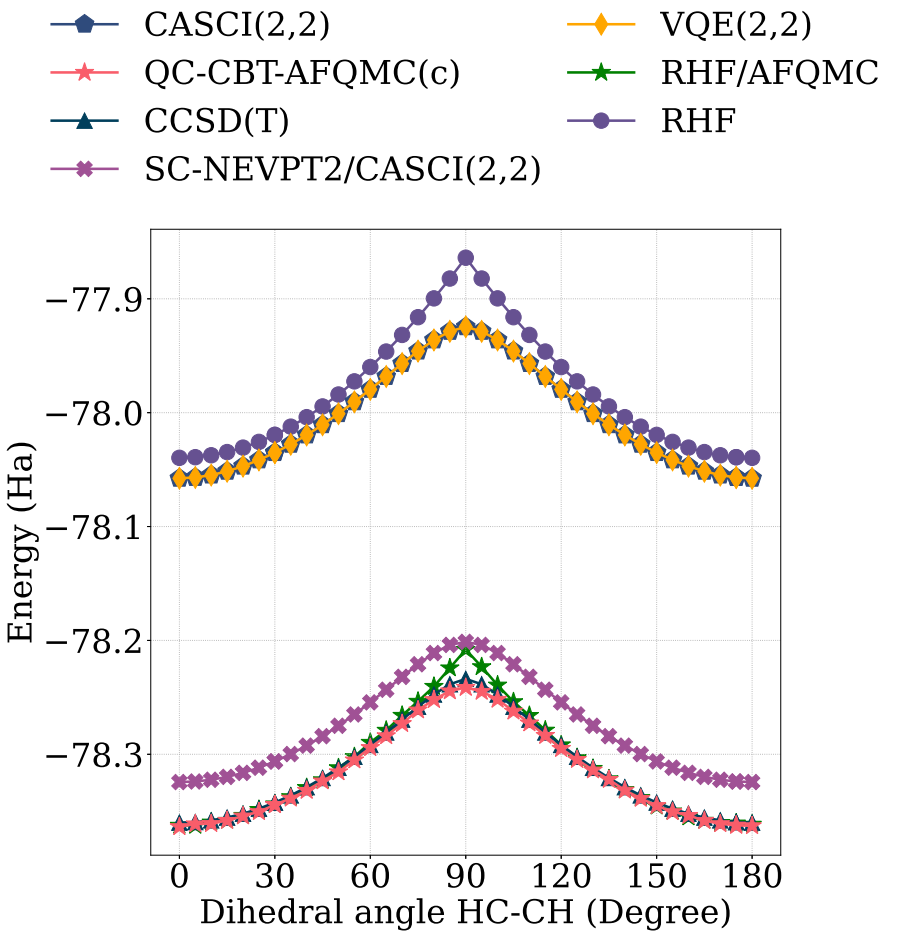


Figure 5.4: Potential energy curves for ethylene as the HC-CH dihedral angle varies, computed with a (2,2) active space and a cc-pVDZ basis set. The QC-CBT-AFQMC(c) approach improves upon CASCI and VQE.

Figure 5.4 presents the PECs for the ethylene molecule. The restricted HF (RHF) resulted in a cusp at a dihedral angle of 90°, which is a manifestation of the degeneracy of the energy states. The same incorrect behavior is visible for the AFQMC method, which uses RHF as a trial wave function. In contrast, the CASCI(2,2) and VQE(2,2) methods effectively smooth the energy barrier. However, accurately depicting the overall shape of the PEC necessitates accounting for dynamical electron correlation. MRCISD+Q was not included

as a reference in this case, since in a minimal (2,2) active space, it does not adequately describe the dynamical correlation, performing worse than CCSD(T), AFQMC, and even NEVPT2. The PEC derived using QC-CBT-AFQMC closely aligns with the CCSD(T) results, except for the region near 90°, where the QC-CBT-AFQMC PEC exhibits a gentler incline. In contrast, the PECs determined using CBT-AFQMC and a downwardly adjusted SC-NEVPT2/CASCI(2,2) show almost perfect congruence. The small discrepancy between the CCSD(T) and CBT-AFQMC methods may be attributable to the inherent single-reference nature of the CCSD(T) approach.

As a third case study, we explored the dissociation of the N_2 molecule, a scenario characterized by its complex, highly multireferential nature due to the triple bond. For both VQE/UCCSD and CASCI approaches, we selected an active space encompassing 6 orbitals to include all 2p atomic orbitals forming the N-N bond, with 6 electrons distributed amongst them. The aug-cc-pVDZ Dunning basis set was used because it contains diffuse functions crucial for accurately modeling the dissociation process. The parameters for the AFQMC calculations were set to 10000 blocks, 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}. In the CBT phase, we aimed for a maximum of $R = 102$ determinants, with the measurement counts N_f, N_a, and N_b set at 10^6 each, culminating in a total of 2.03×10^8 measurements. The maximum number of CNOT gates in the CBT circuits is 10 for fully connected qubits.

Figure 5.5 displays the PECs constructed for the N_2 molecule. As anticipated, the RHF method cannot accurately predict the PEC for the significantly stretched N_2 molecule. AFQMC also fails when using RHF as a trial wave function. Notably, even the CCSD(T) method exhibits poor performance due to its inherent single-reference limitation and perturbative nature. Consequently, MRCISD+Q was employed as the reference method for this challenging case. Similarly to the OH radical, both CASCI(6,6) and VQE(6,6) accurately predict the energy plateau at extended bond lengths but cannot depict the correct PEC shape in the intermediate region (approximately $1.5 - 2.0$ Å). Again, QC-CBT-AFQMC demonstrates commendable agreement with MRCISD+Q across both equilibrium and intermediate regions. Nevertheless, deviations from the expected plateau become apparent at longer distances. This discrepancy is attributed to the fact that while MRCI employs CASSCF(6,6)-optimized orbitals, CBT-AFQMC relies on the same orbitals as VQE(6,6), specifically the RHF canonical orbitals. Achieving closer alignment with MRCI necessitates either expanding the active space or optimizing orbitals in a CASSCF-like fashion. Furthermore, the curvature of the QC-CBT-AFQMC PEC in the intermediate region aligns more closely with MRCI than with SC-NEVPT2/CASCI(6,6), highlighting the precision of the proposed method.

Another point worth mentioning is the computation time. As stated in Ref. [44], classical post-processing for the calculation of benzene on a single CPU core would take approximately 1000 hours per time step using a (6,6) active space. In our case, the full QC-CBT-AFQMC calculation for a single bond length of N_2 within the same active space took at most around 5.5 hours on 48 CPU cores. The CBT part was not parallelized and took only several minutes

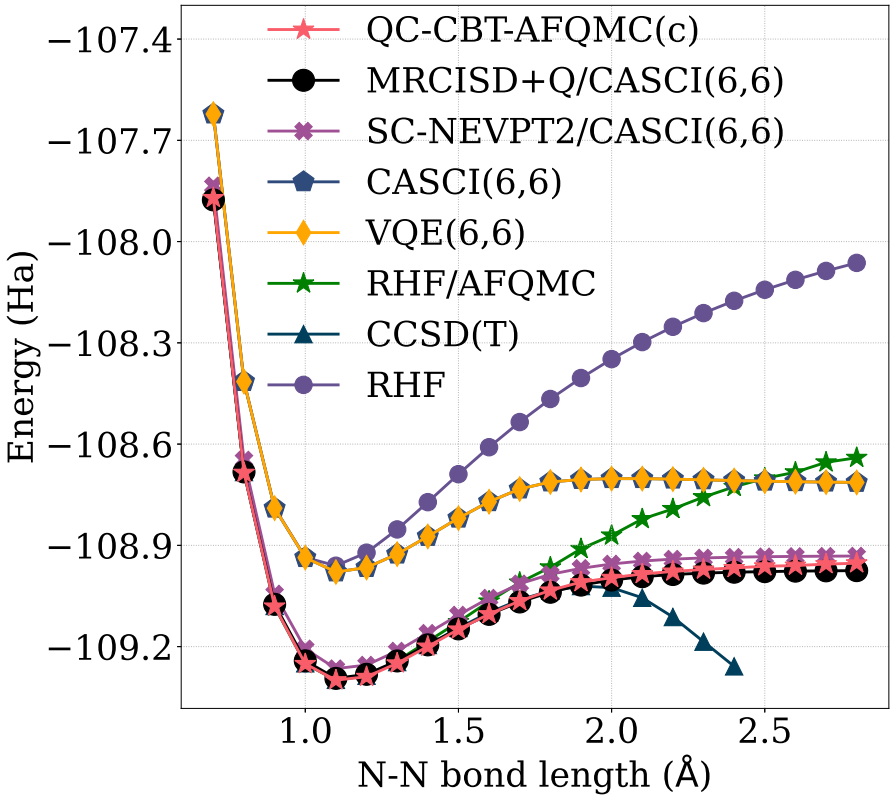


Figure 5.5: Potential energy curves for N_2, obtained with a (6,6) active space and the aug-cc-pVDZ basis set. QC-CBT-AFQMC(c) accurately captures the multireference character of the stretched N-N bond, showing improved agreement with MRCISD+Q references over other methods.

to complete. This demonstrates the much higher efficiency and practicality of our CBT approach compared to the existing matchgate shadows method.

5.3.2 Influence of the Trial Wave Function on AFQMC Convergence

To assess the impact of the trial wave function quality and demonstrate the potential advantage of quantum-prepared initial states, we performed a series of AFQMC calculations with different trial states for the N_2 molecule. All calculations employed the aug-cc-pVDZ basis set, containing diffuse functions crucial for accurately modeling the dissociation process. The (6,6) active space was used, including all 2p orbitals of the nitrogen atoms, which participate in N–N bond formation. The simulation parameters of AFQMC were fixed to 10000

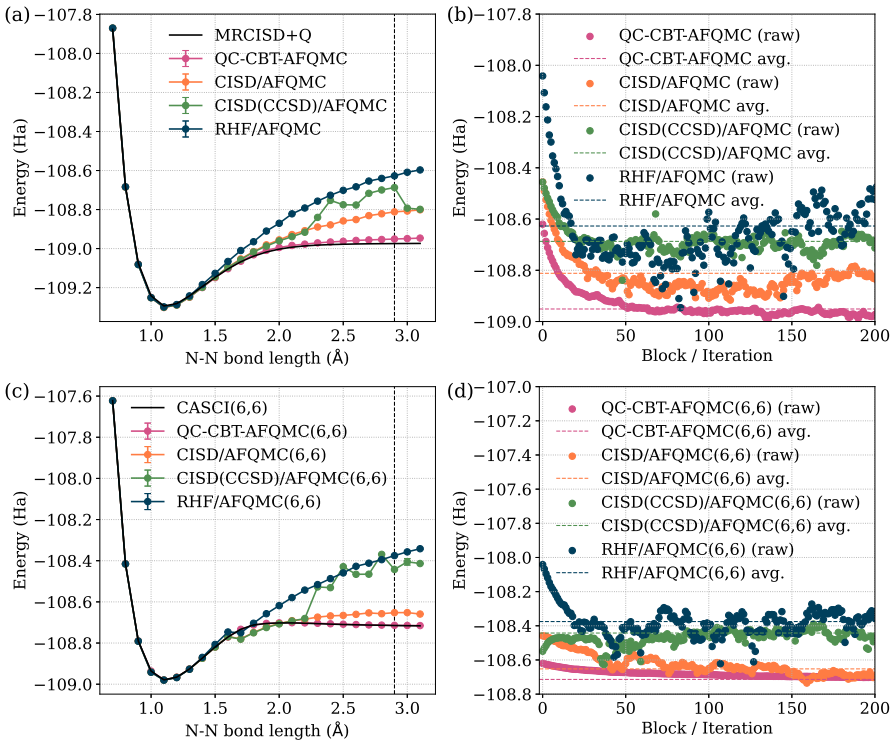


Figure 5.6: AFQMC calculations for the N_2 molecule with a (6,6) active space in the aug-cc-pVDZ basis set. (a) Potential energy curves obtained using AFQMC with different trial wave functions. The MRCISD+Q curve is shown as a reference. (b) Block-averaged AFQMC energy as a function of block index during the equilibration stage at an N–N bond length of 2.9 Å. In all calculations, we used 10000 blocks with 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}. The first 100 blocks were discarded for equilibration and were not included in the final energy estimates. (c) Potential energy curves obtained using active-space-only AFQMC with different trial wave functions. The CASCI(6,6) curve is shown as a reference. (d) Block-averaged active-space-only AFQMC energy as a function of block index during the equilibration stage at an N–N bond length of 2.9 Å. In all calculations, we used 10000 blocks with 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}. The first 100 blocks were discarded for equilibration and were not included in the final energy estimates.

blocks, 10 steps per block, 480 walkers, and a time step of 0.005 Ha^{-1}.

We considered the following trial wave functions: (i) a single-reference RHF determinant (RHF/AFQMC); (ii) a CISD wave function including single and double excitations within the active space (CISD/AFQMC); (iii) a CISD state constructed using CCSD-derived excitation amplitudes (CISD(CCSD)/AFQMC); and (iv) a multireference state obtained from a VQE calculation using the UCCSD ansatz, the CBT technique (QC-CBT-AFQMC). In the CBT scheme, we set $R = 102$ determinants, with the measurement counts N_f, N_a, and N_b set at 10^6 each, resulting in a total of 2.03×10^8 measurements. The correction formula from Eq. (5.19) was not used here.

The resulting potential energy curves are shown in Fig. 5.6(a). The QC-CBT-AFQMC energies closely follow the reference MRCISD+Q curve over the wide range of N–N bond lengths, separating in the well-stretched, strongly correlated region (approximately $2.5 - 3.0$ Å). In contrast, AFQMC initialized from RHF, CISD, and CISD(CCSD) trial states shows increasing deviations at shorter bond distances (approximately $1.7 - 2.0$ Å). Consistent with well-known limitations of single-reference methods in strongly correlated regions, the underlying CCSD description deteriorates rapidly as the multireference character grows and can perform worse than even CISD in this regime, whereas QC-CBT-AFQMC remains more accurate.

To analyze the convergence behavior, Fig. 5.6(b) reports the block-by-block AFQMC energies at the N–N bond length of 2.9 Å. These raw energies refer to the instantaneous energy computed at each block according the Eq. (5.18), before block averaging is performed. Out of the 10000 blocks, the first 100 were treated as an equilibration phase and excluded from the final energy averaging. The RHF trial leads to large fluctuations in the raw block energies. For multireference trials, the fluctuations are reduced and the energy trace becomes more compact. The CBT trial gives the most stable behavior, with narrow fluctuations and a rapid onset of a stationary regime.

Moreover, we performed active-space-only AFMQC calculations using those trial states, where all the electron correlation outside the active space is ignored and CASCI energies become exact reference values. For the active-space-only results shown in Fig. 5.6(c), the same trend is observed, except that CISD and CISD(CCSD) deviate earlier, already at 1.6 Å. In this case, the UCCSD(CBT) PEC coincides with the reference CASCI(6,6) curve. The active-space-only AFQMC calculations in Fig. 5.6(d) show the same behavior: QC-CBT-AFQMC energies converge smoothly to the average with minimal fluctuations, whereas the other methods exhibit significantly large fluctuations.

The results in Fig. 5.6 demonstrate a clear improvement in the strongly correlated region as one moves from RHF $\rightarrow$ CISD/CISD(CCSD) $\rightarrow$ UCCSD(CBT). These observations indicate that trial states that possess a sufficient multireference character, such as the UCCSD(CBT) state, provide better overlap with the exact ground state. This improves both the statistical efficiency and the robustness of the AFQMC projection, leading to smoother equilibration and higher overall accuracy in the strongly correlated region of the N_2 dissociation curve.

5.3.3 Effect of Measurement Quantity on CBT

The accuracy of CBT is inherently linked to the number of performed measurements, where more measurements typically yield more precise CI coefficients. However, conducting millions of CBT measurements in practical scenarios is time consuming. Importantly, the primary objective within the AFQMC step is to capture the dynamical correlation outside of the active space. In this context, it is assumed that errors in CI coefficients mostly affect the correlation outside the active space. By employing Eq. (5.19), we isolated this energy correction by subtracting the AFQMC energy obtained within the active space approximation from that obtained within the full space. This section examines the precision of our method in practical applications.

Focusing on the OH radical, the left side of Fig. 5.7 clearly shows the effect of the number of CBT measurements on the AFQMC energy based on the CBT-VQE initial guess. Notably, the error increases with the extension of the O-H bond, correlating with the system's increasingly multireferential character. This multireferential aspect should ideally be accounted for during the VQE stage. Consequently, applying Eq. (5.19) effectively mitigates these errors, as shown in the left segment of Fig. 5.7. It is crucial to acknowledge that the PECs obtained using a large number of measurements, whether or not Eq. (5.19) is applied, converge closely to the reference because $E^{\text{act}}_{\text{AFQMC}}$ converges to $E^{\text{act}}_{\text{VQE}}$.

Other molecular systems were also explored and yielded comparable outcomes. The detailed results of these investigations are available in Appendix C1.

5.3.4 Analysis of Reaction Barriers

To illustrate the practical applicability of our method, we explored the reaction barriers of [3+2] cycloaddition reactions as a representative example. These reactions are important in biochemistry, particularly for the *in vivo* study of biomolecules [243]. Cycloaddition enables the precise tagging of biomolecules in a process known as bio-orthogonal click reactions. In addition to real-time imaging, bio-orthogonal reactions are indispensable for pioneering applications in medicine and healthcare, such as targeted cancer therapies and *in situ* drug synthesis [244]. The search for mutually orthogonal reactions is a priority research area.

Among the extensive array of reactions, we focused on the relatively simple yet crucial reactions between ethylene and nitrous oxide N_2O or hydrogen azide HN_3 (see Fig. 5.9). The manageable size of these systems permits highly accurate calculations of reaction barriers, enabling the construction of a complete basis set extrapolation. Our findings were benchmarked against the CCSD(T)/CBS [245] and CBS-QB3 [246] methods, which have been established as reliable approaches for calculating reaction barriers.

To determine the reaction barriers, we conducted single-point calculations for reactants at their equilibrium geometries and for the transition state at the peak of the barrier. All geometric data were derived using DFT using the B3LYP functional with the def2-SVP basis set, as reported in the work by Stuyver *et al.* [247]. Given the high efficiency of DFT for providing optimal geometries

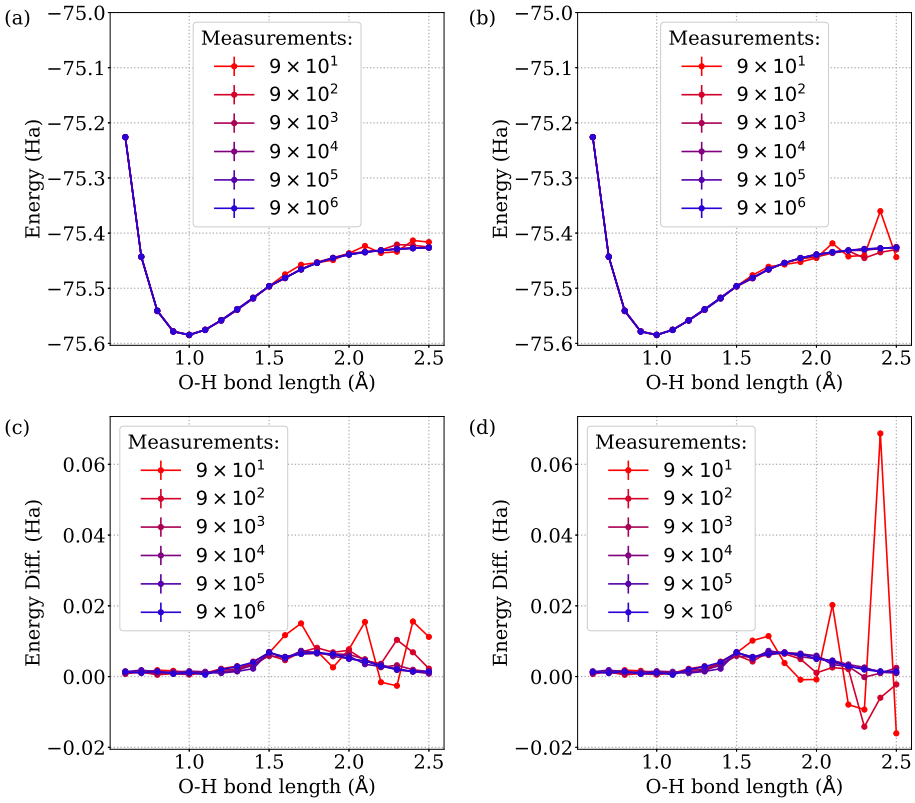


Figure 5.7: Analysis of CBT measurement effects (i.e., statistical errors) on QC-CBT-AFQMC results for the OH radical with a (3,3) active space and aug-cc-pVDZ basis set. Panels (a) and (b) compare the potential energy curves reconstructed without and with the energy correction from Eq. (5.19), respectively, as the number of CBT measurements changes. Panels (c) and (d) show the corresponding energy deviations from CCSD(T). These results demonstrate that applying the correction significantly reduces measurement-induced errors, allowing accurate computations even with limited shot budgets.

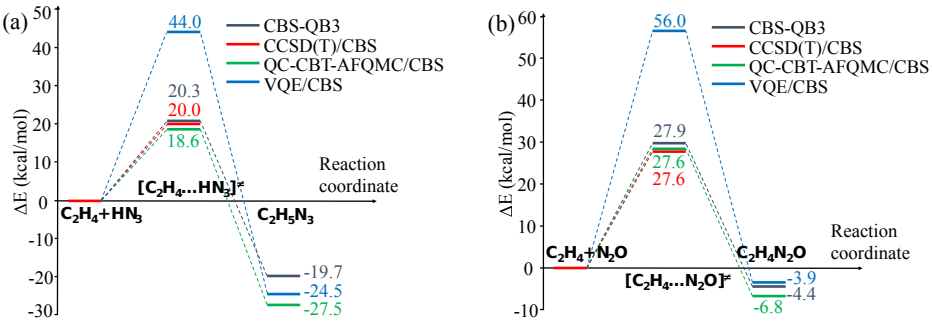


Figure 5.8: Computed energy barriers for the [3+2] cycloaddition reactions illustrated in Fig. 5.9, comparing QC-CBT-AFQMC results against CCSD(T)/CBS and CBS-QB3 reference values. Panel (a) shows the results for HN_3, while panel (b) shows those for N_2O. The close agreement underscores QC-CBT-AFQMC's ability to accurately estimate reaction barriers and demonstrates the effectiveness of incorporating complete basis set extrapolations.

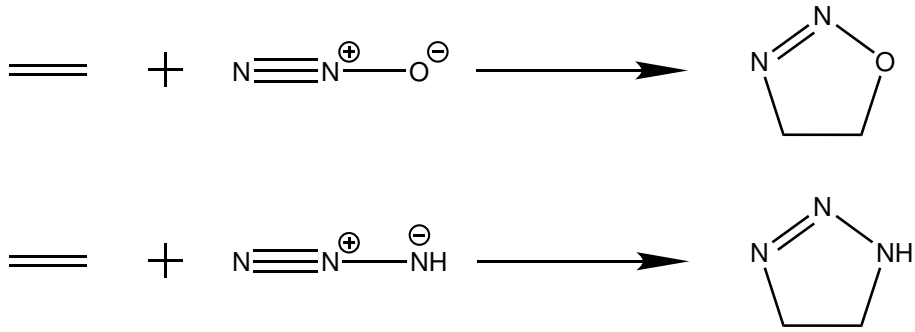


Figure 5.9: Illustration of the [3+2] cycloaddition reactions studied: ethylene reacting with N_2O and HN_3. These reactions serve as test cases for evaluating the accuracy of QC-CBT-AFQMC in computing reaction barriers.

for various high-level theoretical approaches, and considering that the DFT geometries were employed in notable studies [245, 246], we adopted a similar strategy.

The initial phase of our investigation involved executing the VQE for each geometric configuration. The VQE wave function was translated exactly, which corresponds to the CBT with an infinite number of measurements. This approximation effectively eliminates any potential errors arising from the measurement process. For the ethylene molecule, we defined an active space with 2 electrons across 2 orbitals, encompassing both bonding and antibonding π orbitals. For N_2O, the active space included 4 electrons on 4 π orbitals, with one involved in N-N bonding and simultaneously in antibonding of O-H bonds. This molecule exhibits axial symmetry, leading to sets of 2 degenerate orbitals, thus

requiring the selection of a pair of such orbitals. In the case of HN_3, an active space with 4 electrons over 3 orbitals was chosen, incorporating 2 π bonding and one nonbonding orbital. Given that both the final products and the transition states in these reactions share similar structures, featuring a single double bond, we selected active spaces of 2 electrons across 2 orbitals for all entities. These orbitals represent the π bonding and antibonding interactions associated with the double bond.

In the final phase, we conducted AFQMC calculations, employing VQE-CBT wave functions as the initial guesses. The AFQMC simulations were configured with 10000 blocks, 10 steps per block, 960 walkers, and a time step of 0.005 Ha^{-1}. Moreover, the initial 1000 blocks were omitted to allow system equilibration.

To construct a complete basis set extrapolation, we performed all calculations using a sequential series of Dunning basis sets, aug-cc-pVnZ (where n = D, T, Q), enriched with diffuse functions. This series of basis sets is known for its rapid convergence compared to those without diffuse functions. Subsequently, we applied two distinct CBS extrapolation techniques to both the VQE energies and AFQMC correlation energies. Comprehensive details are provided in Appendix C2. Eventually, we selected an exponential three-basis-set scheme for the VQE energy and the Riemann zeta function method for the AFQMC correlation energy. These results are presented in Fig. 5.8.

Figure 5.8 shows that the VQE/CBS method generally predicts the thermodynamics of reactions with accuracy approaching that of the more sophisticated perturbation theory-based CBS-QB3 method, particularly noticeable in the reaction of C_2H_4 with N_2O. However, VQE struggles to accurately predict the reaction barriers for both reactions. Furthermore, the use of AFQMC gives a more accurate prediction of the reaction barrier energies, even for the smallest basis set. This behavior is expected as reactants and products are single-reference closed-shell systems, whereas the transitional state is a highly stretched multireferential system. The Riemann zeta function method yields results comparable to the reference data, similar to the CCSD(T)/CBS for the barrier energy of the C_2H_4 + N_2O reaction, albeit with a slight overestimation of the heat effect of the reaction. For the C_2H_4 + HN_3 reaction, the same approach estimates the barrier energy within 1.5 kcal/mol of that given by CCSD(T) and the reaction's heat effect within 8 kcal/mol of that given by CBS-QB3. Remarkably, our findings closely align with those obtained using the much more costly and precise CCSDT(Q)/CBS method [248]. In our analysis of the C_2H_4 and HN_3 reaction, we determined an energy barrier of 18.6 kcal/mol, while the referenced study reported a barrier of 18.29 kcal/mol for the reaction between C_2H_4 and CH_3N_3. This similarity underscores the chemical resemblance between the two systems, with the only possible variance arising from the presence of the CH_3 group, which likely introduces a small geometrical obstacle. Similar findings were observed for the heat effect; our method gave a value of -27.5 kcal/mol, which is very close to the -28.86 kcal/mol reported in the referenced study, underscoring the high accuracy of our method.

5.4 Conclusions

In this work, we presented QC-CBT-AFQMC, which integrates the QC-AFQMC approach with CBT. Our main motivation stems from the practical limitations of current schemes of quantum tomography under realistic shot budgets. Although QC-AFQMC with matchgate shadow has polynomial asymptotic scaling, its high polynomial degree ($O(N_q^{8.5}\log^2(N_q))$), circuit requirements, and measurement overhead limit its practicality, particularly for near-term quantum devices and could remain nontrivial even in the FTQC era.

In contrast, CBT uses sparse Clifford circuits with $O(N_q)$ CNOT gates and has a lower constant prefactor, allowing us to do calculations for bigger active spaces. CBT circuits can be implemented in constant depth using midcircuit measurements. The classical post-processing cost of QC-CBT-AFQMC $O(N_q^4)$ is lower than that of the matchgate shadows. This makes CBT applicable to near-term quantum devices. In addition, the extracted wave function data via CBT is stored in double precision format, avoiding exponential scaling of the measurements. Within the QC-AFQMC framework, the wave function data measured by CBT are then used to mitigate the sign problem in AFQMC, enabling efficient incorporation of dynamical correlation.

Tests on molecular benchmarks under realistic shot constraints showed that accurate trial wave functions were obtained using CBT. We further applied QC-CBT-AFQMC to [3+2] cycloaddition reactions, which are important for biochemistry and medicine. The predicted barriers of the method matched those from established methods like CCSD(T)/CBS and CBS-QB3. We also examined the performance of various CBS extrapolation techniques for AFQMC, finding that both the inverse cube and Riemann approaches gave similar and reliable results.

Looking ahead, there are multiple directions for further improvements. Techniques such as QSCI/ADAPT-QSCI [29, 249, 250, 251, 252, 253, 254, 255], which can potentially reduce the circuit depth and the measurement demands, could be integrated to achieve even higher shot efficiency. Moreover, additional research is needed to explore the applicability of the QC-CBT-AFQMC method to larger and more complex molecular systems, as well as to investigate the combination of CBT with other quantum-classical hybrid quantum Monte Carlo approaches [256, 257, 234].

Chapter 6

Conclusions

This dissertation demonstrates how quantum computations for quantum chemistry can be made practically viable under realistic FTQC constraints. Rather than attempting to brute-force the fundamental complexity barriers of electronic structure theory, these limitations are treated explicitly as design constraints that shape feasible quantum-classical workflows. From this perspective, two factors emerge as decisive for practical performance: the preparation of chemically relevant quantum states with sufficiently large overlap with target eigenstates, and the availability of efficient classical post-processing methods capable of recovering correlation effects beyond the active space.

On the quantum side, this work shows that state preparation can be systematically improved by exploiting the Hamiltonian itself as a filtering resource. A family of power-iteration-based quantum methods is developed to refine an initial state toward ground and excited eigenstates with near-optimal query complexity across a broad range of regimes, including scenarios with and without prior energy estimates. A key conclusion is that algorithmic efficiency alone is insufficient unless the exact symmetries of molecular Hamiltonians are explicitly enforced. To this end, symmetry projectors for particle number and spin are constructed that are robust to finite gate precision and exhibit favorable scaling, enabling the preparation of symmetry-adapted states with substantially enhanced overlap at a cost negligible compared to that of quantum phase estimation. These results highlight that improving input state quality is often a more effective route to practical gains than further optimizing downstream quantum algorithms.

On the classical side, we address the inevitable limitations of quantum hardware by developing an efficient post-processing pipeline that recovers dynamical correlation outside the active space. By combining CBT with AFQMC, the proposed QC–CBT–AFQMC workflow enables the use of quantum-prepared multiconfigurational states as high-quality trial wavefunctions under realistic measurement budgets. Together with the state preparation techniques developed in this work, this approach establishes an end-to-end hybrid quantum–classical strategy in which quantum resources are focused on the hardest correlated subproblems, while classical methods extend these results to chemically relevant

accuracy.

Chapter 3 establishes GQSP as a unifying and scalable framework for initial state preparation using quantum power-based methods. The central motivation is that the electronic Hamiltonian itself constitutes the fundamental source of information about its eigenstates, naturally leading to Hamiltonian-driven refinement schemes that iteratively improve an initial guess toward states of interest, in close analogy to ITE [35]. Within this framework, QPI enables convergence to the ground state without requiring prior knowledge of the target energy, while QPL further accelerates convergence through variational subspace optimization. A rigorous query complexity analysis shows that QPI and QPL achieve near-optimal scaling $O\big(\frac{1}{\beta}\frac{\lambda}{\Delta\gamma}\log\frac{1}{\beta^2\gamma^2\varepsilon}\big)$ for ground-state preparation, where γ denotes the initial state overlap, Δ the spectral gap, ε the target preparation error, and β the Lanczos overlap improvement factor. While quantum analogues of inverse iteration are traditionally implemented using Trotterized time evolution [164, 171], the GQSP-based construction developed here enables inverse iteration via polynomial approximation, yielding the QII method. When a reliable energy estimate is available, QII achieves optimal $\Theta\big(\frac{\lambda}{\gamma\Delta}\log\frac{1}{\varepsilon}\big)$ scaling. Excited states can also be accessed using the QFSM, albeit with non-optimal scaling; however, unlike existing implementations based on variational optimization [178], the GQSP formulation avoids measurement-intensive variational procedures. Benchmarking on representative molecular systems demonstrates consistent performance across different correlation regimes, supporting the broader conclusion that GQSP-based quantum power methods provide a flexible and analytically controlled route to preparing high-quality input states for quantum chemical algorithms on FTQC.

Chapter 4 establishes symmetry filtering as a practical and resource-efficient component of state preparation by constructing and quantitatively benchmarking explicit symmetry projectors for particle number and spin. In contrast to previous approaches that are restricted to finite groups [210], modify the Hamiltonian itself [207, 206], or rely on energy penalty terms [209], and to earlier projector proposals lacking detailed resource analysis [205], we systematically compare projector implementations based on LCU, GQSP and GQSVT with explicit gate counts and precision requirements. We find that LCU and GQSP projectors for $\hat{N}$ and $\hat{S}_z$ achieve quasi-quadratic scaling in the number of spin-orbitals, $O(N_{\mathrm{SO}}^2 \log N_{\mathrm{SO}})$ T gates, while maintaining robustness to single-qubit rotation errors at the 10^{-1}–10^{-2} level, whereas GQSVT incurs significantly larger constant factors and stricter phase-precision demands. An efficient realization of total-spin filtering via the composite projector $\hat{P}_{M_S}\hat{P}_S\hat{P}_{M_S}$ brings its cost into the same scaling class as $\hat{N}$ and $\hat{S}_z$ projection. Numerical demonstrations on representative molecular systems show that symmetry filtering increases the success probability of quantum phase estimation from approximately 0.5 to above 0.9, while incurring projector T gate costs three orders of magnitude below those of QPE itself, even for FeMoco-scale active spaces. Together, these results demonstrate that explicit symmetry projection is a low-cost preprocessing step that materially improves the state quality for practical quantum chemistry simulations.

Chapter 5 addresses the classical post-processing of quantum-prepared active

space states. Although existing QC-AFQMC approaches are robust to noise and can achieve high accuracy, they suffer from prohibitive measurement and computational costs when implemented using Clifford or matchgate shadows [42, 43]. The proposed QC-CBT-AFQMC workflow overcomes these limitations by reconstructing multiconfigurational wave functions within the active space using CBT. This reconstruction requires only shallow Clifford circuits with $O(N_{\mathrm{SO}})$ CNOT gates, making it feasible under realistic time budgets. The resulting state is then employed as a trial wave function in AFQMC, where it mitigates the sign problem and enables the recovery of dynamical correlation beyond the active space. As a consequence, the overall cost of post-processing is reduced to $O(N_{\mathrm{SO}}^4)$ scaling, representing a substantial improvement over existing QC-AFQMC schemes while maintaining high accuracy. Importantly, the ability to perform reliable CBS extrapolation within the QC-CBT-AFQMC scheme enables systematic convergence toward chemically accurate, experimentally testable results rather than isolated finite-basis benchmarks. Taken together, these results demonstrate that quantum-prepared active space states offer a qualitative advantage over classically prepared counterparts by providing access to multireference trial states that would be exponentially costly to construct or manipulate on classical hardware.

Viewed as a whole, the contributions of this dissertation support a practical strategy for quantum chemistry on FTQCs: targeted, symmetry-adapted state preparation in an active space on quantum hardware, combined with efficient state tomography and classical post-processing to recover missing electronic correlation. Together, these elements can form a hybrid quantum–classical workflow, in which quantum resources are devoted to the generation of high-quality multiconfigurational states, while classical algorithms systematically extend these states to recover residual correlation effects beyond the active space. This division of labor provides a realistic pathway for leveraging quantum hardware in chemically relevant simulations, including challenging strongly correlated systems such as the FeMoco active site, where multiple electronic states with different total spin lie in close energetic proximity. In such cases, accurate state preparation critically depends on reliable spin filtering, as developed in this dissertation. The subsequent QC–CBT–AFQMC framework then incorporates correlation effects outside the active space, providing a route toward quantitatively accurate descriptions of complex multireference systems that remain difficult to treat using purely classical methods.

The main conceptual limitation of the present work is its focus on the preparation and post-processing of a single dominant electronic state. In many chemically relevant situations, however, several low-lying, quasi-degenerate states must be treated on an equal footing to maintain balanced accuracy across the manifold. Addressing this multi-state regime will likely require extensions beyond the present power-iteration framework. For example, state preparation schemes based on effective Hamiltonian formalisms can simultaneously approximate multiple low-lying eigenstates. On the classical side, analogous challenges arise because standard AFQMC targets a single state. Thus, consistent treatments of near-degenerate manifolds may require explicitly multi-state formulations.

Developing such multi-state quantum state preparation methods and integrating them with compatible classical strategies represents a natural direction for future work, particularly for systems exhibiting quasi-degeneracies.

Additional limitations arise from practical considerations. Quantum power methods are most effective when the initial guess state already captures the essential multireference character, mirroring the classical use of the PL method [179] as refinement techniques rather than as a black-box solver. The QFSM further relies on sufficiently large energy gaps, which tend to decrease for excited states in chemically relevant systems. Both quantum power methods and symmetry projectors assume the availability of AA to achieve practical success probabilities. Finally, the QC-CBT-AFQMC workflow inherits the active space selection problem common to many multireference methods and requires that all qualitatively important configurations be included in the CBT reconstruction. Although their coefficients are only approximate and subsequently refined during AFQMC, phases should be captured precisely.

At this point, a philosophical question may arise: can a truly universal method exist that would allow one to solve any molecular problem to an arbitrary precision more efficiently than brute-force FCI or direct spatial discretization? This question reflects a tension between the abstract universality of formal laws and the concrete individuality of real systems. Computational complexity theory provides strong reasons to doubt the existence of such a universal method, indicating that, when the Hamiltonian is taken as an abstract object in isolation, extracting its physical content requires exponential resources.

One may nevertheless imagine the discovery of a universal algorithm capable of efficiently solving all electronic structure problems. Yet such a discovery would have a paradoxical consequence. By abolishing the need for approximations, heuristics, and chemical intuition, it would dissolve the very conditions under which chemistry exists as a science. Chemistry persists not despite the simplicity of its fundamental Hamiltonians, but because concrete molecular systems exhibit individuality that cannot be exhausted by abstract universal rules. In this sense, such a universal algorithm would signify the negation of chemistry as a living process of knowledge, reducing it to a closed formal system.

Scientific understanding does not proceed by the mere application of abstract laws to particular cases. Rather, it unfolds through the continuous mediation between reasoning and reality. Knowledge emerges through this process of interaction, in which experiments, engineering constraints, and practical limitations play an essential role. Purely formal reasoning, detached from material engagement, cannot "outrun the locomotive" of nature itself.

Formal limitations, expressed, for instance, in the language of computational complexity theory, should therefore not be interpreted as absolute prohibitions. Instead, they reflect the current stage of our conceptual and practical interaction with the physical world. They mark points of tension where abstract universality encounters concrete specificity and where new forms of approximation, representation, and computation may emerge through their resolution.

Thus, the perspective that emerges from this work suggests that progress in quantum chemistry relies on the systematic incorporation of chemical intuition and physically motivated heuristics into computational methods, rather than on the prospect of discovering a universal algorithm. Even in the presence of highly structured molecular Hamiltonians, one must ultimately acknowledge the irreducible emergent complexity of each system. Knowledge in quantum chemistry should not be understood as automatic computation, but as a dialog between human reasoning and nature, in which computation constitutes only one moment within a broader process of scientific becoming.

Appendix

Appendix A: Supporting Information for Chapter 3

A1 Angle Finding

A crucial component of any QSP-based algorithm is the determination of the rotation angles sequence. The GQSP algorithm is no exception; it uses a procedure very similar to that of conventional QSP and QSVT. An angle finding protocol typically involves four steps: *truncation*, *partition*, *completion*, and *carving*. Each step is described in detail below.

In GQSP, our goal is to apply an arbitrary function $f(z)$ to a unitary block encoding matrix U. As the first step, we approximate the target function using a complex-valued polynomial $F(z)$, for instance, using the Remez exchange algorithm. This procedure is called truncation: restricting the function to a finite-degree polynomial that approximates $f(z)$ sufficiently well for our purposes.

Once we obtain $F(z)$, the next step is partition, in which we rescale $F(z)$ to ensure that its absolute value does not exceed 1 on a unit circle in the complex plane. Typically, this is achieved by dividing $F(z)$ by a sufficiently large constant α, and defining $\mathcal{P}(z) = F(z)/\alpha$. This results in $|\mathcal{P}(z)| \leq 1$ on the complex unit circle.

Having obtained $\mathcal{P}(z)$, we must then complete the polynomial to form the unitary that GQSP applies to U. Specifically, we seek a polynomial $\mathcal{Q}(z)$ such that

$$\mathrm{GQSP}(U) = \begin{pmatrix} \mathcal{P}(U) & * \\ \mathcal{Q}(U) & * \end{pmatrix} \tag{A1}$$

with the constraint $|\mathcal{P}(z)|^2 + |\mathcal{Q}(z)|^2 = 1$ on a unit circle. Several methods can be used to find $\mathcal{Q}(z)$; four of the most common are via optimization, via root-finding, via Prony's method, and the recently proposed exact expression using contour integration [258].

The optimization-based technique used by the original GQSP authors [128] begins by applying a Fourier transform to the condition $|\mathcal{P}(z)|^2 + |\mathcal{Q}(z)|^2 = 1$. Denoting by $\vec{a}$ and $\vec{b}$ the coefficient vectors of $\mathcal{P}(z)$ and $\mathcal{Q}(z)$ respectively, one

obtains

$$\vec{a} \star \text{reverse}(\vec{a})^* + \vec{b} \star \text{reverse}(\vec{b})^* = \vec{\delta}, \tag{A2}$$

where $\star$ denotes discrete convolution and reverse$(\cdot)$ refers to the reversal order of elements of an vector.

One can then solve

$$\arg\min_{b} \left\| \vec{a} \star \text{reverse}\big(\vec{a}\big)^* + \vec{b} \star \text{reverse}\big(\vec{b}\big)^* - \vec{\delta} \right\|^2, \tag{A3}$$

determining a suitable $\vec{b}$. Although straightforward to implement, this method can be sensitive to initial guesses and may not be precise enough.

A more direct approach uses polynomial properties in a complex plane [259]. One first defines

$$g(z) = 1 - |\mathcal{P}(z)|^2. \tag{A4}$$

Half of the roots of $g(z)$ lie inside the unit circle, whereas the other half lie outside. Let $\{\xi_j\}$ be the roots within the unit circle. Subsequently, a polynomial $\mathcal{Q}(z)$ satisfying $|\mathcal{Q}(z)|^2 = g(z)$ on $|z| = 1$ can be written up to scaling as follows:

$$\mathcal{Q}(z) \propto \prod_{|\xi_j|<1} (z - \xi_j). \tag{A5}$$

Although conceptually simpler, this requires finding all roots of a potentially high-degree polynomial, making it computationally expensive.

Prony's method was proposed to address the scaling challenges of root finding [260, 129]. Define

$$h(z) = \frac{1}{g(z)} = \frac{1}{1 - |\mathcal{P}(z)|^2}. \tag{A6}$$

In this formulation, the roots of $g(z)$ become poles of $h(z)$. Hence $h(z)$ can be expressed as follows:

$$h(z) = \sum_{\xi_j} \frac{r_j}{z - \xi_j} + \text{constant}, \tag{A7}$$

where r_j are the residues of $h(z)$ at the poles ξ_j. Next, one computes the Fourier coefficients $\hat{h}_k$ of $h(z)$:

$$\hat{h}_k = \frac{1}{2\pi i} \int_{\mathrm{G}} \frac{h(z)}{z^k} \frac{dz}{z} = - \sum_{|\xi_j|<1} r_j \xi_j^{-(k+1)}, \tag{A8}$$

where G denotes a unit circle contour. Thus, each $\hat{h}_k$ is a linear combination of the inverses of the roots inside G. Following Prony's approach, the polynomial for $\mathcal{Q}(z)$ is obtained by solving

$$H_{\text{Prony}} \begin{pmatrix} m_0 \\ m_1 \\ \vdots \\ m_d \end{pmatrix} = 0, \tag{A9}$$

where H_{Prony} is the $l \times (d+1)$ Hankel matrix

$$H_{\text{Prony}} = \begin{pmatrix} \hat{h}_{-1} & \hat{h}_{-2} & \cdots & \hat{h}_{-(d+1)} \\ \hat{h}_{-2} & \hat{h}_{-3} & \cdots & \hat{h}_{-(d+2)} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{h}_{-l} & \hat{h}_{-(l+1)} & \cdots & \hat{h}_{-(d+l)} \end{pmatrix}. \tag{A10}$$

The vector $\vec{m} = (m_0, m_1, \ldots, m_d)^T$ contains the coefficients of the $\mathcal{Q}(z)$ up to a scaling factor and can be obtained using the singular value decomposition (SVD). In particular, $H_{\text{Prony}} = U_{\text{SVD}} \Sigma V_{\text{SVD}}^{\dagger}$, where U_{SVD} and V_{SVD} are unitary matrices, and Σ is a diagonal matrix of singular values of H_{Prony}. Because $\text{rank}(H_{\text{Prony}}) < d+1$, Σ has at least one zero and the corresponding column of V_{SVD} supplies $\vec{m}$.

The most advanced technique for the completion step, which relies on complex analysis, was introduced in [258]. In this approach, one begins with the function $g(z)$ defined in Eq. (A4); however, this time, the logarithm is applied:

$$w(z) = \log(g(z)). \tag{A11}$$

Next, the Fourier series of the boundary function $w(e^{i\theta})$ is calculated, denoted by $\hat{w}(\theta)$. Then the projection operator Π is applied, which leaves only non-negative frequency components from $\hat{w}(\theta)$ and halves the constant term:

$$\Pi\left[e^{in\theta}\right] := \begin{cases} e^{in\theta}, & n \in \mathbb{Z}_{>0}, \\ \frac{1}{2}, & n = 0, \\ 0, & n \in \mathbb{Z}_{<0}. \end{cases} \tag{A12}$$

This projection ensures that $\Pi w(z)$ is holomorphic inside the unit disc; thus, the resulting polynomial $\mathcal{Q}(z)$ will contain only non-negative powers of z.

Finally, the polynomial $\mathcal{Q}(z)$ is recovered by exponentiating the projected function:

$$\mathcal{Q}(z) = \exp[\Pi w(z)]. \tag{A13}$$

After completing $\mathcal{P}(U)$ and $\mathcal{Q}(U)$, the final step is called carving, in which the GQSP sequence (see Eq. (2.168)) was unwrapped from the end:

$$\begin{aligned} &\begin{pmatrix} \hat{\mathcal{P}}(U) & \cdot \\ \hat{\mathcal{Q}}(U) & \cdot \end{pmatrix} = \text{CU}_0^{\dagger}\, \tilde{R}(\theta_j, \phi_j, 0)^{\dagger} \begin{pmatrix} \mathcal{P}(U) & \cdot \\ \mathcal{Q}(U) & \cdot \end{pmatrix} \\ =&\begin{pmatrix} e^{-i\phi_j}\cos(\theta_j)U^{\dagger} & \sin(\theta_j)U^{\dagger} \\ e^{-i\phi_j}\sin(\theta_j)I & -\cos(\theta_j)I \end{pmatrix} \begin{pmatrix} \mathcal{P}(U) & \cdot \\ \mathcal{Q}(U) & \cdot \end{pmatrix}, \end{aligned} \tag{A14}$$

yielding new polynomials $\hat{\mathcal{P}}(U)$ and $\hat{\mathcal{Q}}(U)$ of one degree lower than $\mathcal{P}(U)$ and

$\mathcal{Q}(U)$. At each iteration, the rotation angles are extracted using

$$\begin{aligned} \theta_j &= \tan^{-1}\left(\frac{|b_j|}{|a_j|}\right), \\ \phi_j &= \arg\left(\frac{a_j}{b_j}\right), \\ \lambda_0 &= \arg(b_0). \end{aligned} \tag{A15}$$

This process is repeated until the sequence is fully carved out, and any remaining global phase is stored in the final angle λ_0.

In practice, numerical problems may arise when $\mathcal{P}(z)$ has nearly zero coefficients at high or low powers. This reduces the numerical rank of H_{Prony} in Eq. (A9). To mitigate this, a *capitalization* technique is used, in which an auxiliary polynomial $\mathcal{P}'(z)$ with large coefficients at the highest and lowest powers is added and then subtracted:

$$\mathcal{P}''(z) = \frac{1}{\beta}\mathcal{P}(z) - \mathcal{P}'(z) = \frac{1}{\beta}\mathcal{P}(z) - \Upsilon(z^d + 1), \tag{A16}$$

where β and Υ are suitably chosen large constants. Subsequently, we work with $\mathcal{P}''(z)$, which is numerically more stable.

Once $\mathcal{P}''(U)$ is computed, the $\mathcal{P}(U)$ must be recovered. This requires the addition of $\mathcal{P}'(U)$ using an extra qubit, as illustrated in Fig. A1:

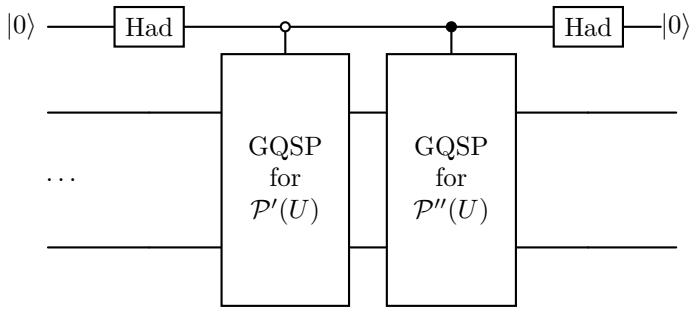


Figure A1: Quantum circuit to restore the original polynomial $\mathcal{P}(U)$ after the capitalization step, using LCU with one additional qubit. Had denotes Hadamard gates.

Although this method doubled the circuit depth and consumed an additional qubit, it stabilized the numerical procedure.

A2 Comparison of QII with Q-Inv and I-Iter for Different Truncation

Here, we analyze how the convergence behavior of QII changes with different truncations of the Taylor polynomial (see Fig. A2). As a general rule, the

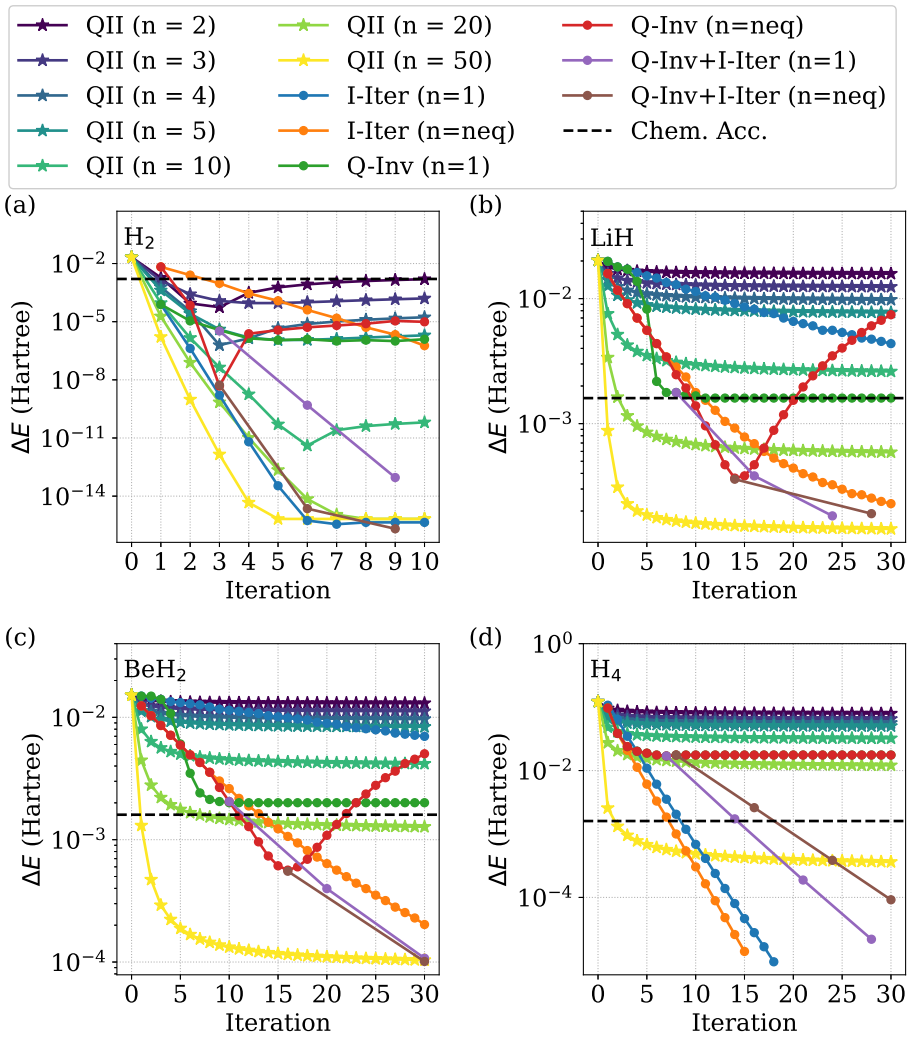


Figure A2: Comparison of QII, I-Iter, and Q-Inv [171] for H_2, LiH, BeH_2, and H_4. Shown is the deviation from CASCI energies as a function of the iteration number. For QII, n refers to the polynomial truncation degree, while for I-Iter and Q-Inv, n is the Gauss-Legendre polynomial order. n_{eq} is the order at which energy convergence is reached.

convergence limit depends on the truncation: the higher the polynomial degree, the lower the energy.

For H_2, QII achieved chemical accuracy for all tested values n. Q-Inv($n = 1$) is very close to QII($n = 5$). The curves for QII($n = 20$), QII($n = 50$), I-Iter($n = 1$) and Q-Inv+I-Iter($n = n_{\text{eq}}$) are grouped near 4×10^{-16}, which is the

most likely consequence of machine precision.

For LiH, polynomial truncation in QII becomes a more significant constraint. Chemical accuracy was achieved only for $n \geq 20$, whereas lower n values converged to energies above the chemical accuracy. However, once n is sufficiently large, QII converges faster than Q-Inv and I-Iter; QII($n = 20$) requires only two iterations to reach chemical accuracy, whereas QII($n = 50$) does so in a single iteration.

A similar trend is observed for BeH_2. QII with $n \geq 20$ achieved chemical accuracy with seven iterations at $n = 20$ and in only one iteration at $n = 50$. Again, QII($n = 50$) initially converges faster than I-Iter or Q-Inv. However, after approximately 30 iterations, the combined Q-Inv+I-Iter approach slightly exceeded QII($n = 50$). Extending QII to $n > 50$ further improves the final accuracy.

The ground state of the square H_4 system exhibits a strong multireference character, making it the most challenging case. Only QII($n = 50$) achieved chemical accuracy in only two iterations. However, at longer iteration counts (beyond 10), I-Iter eventually reached lower energies. Similarly, increasing the degree of polynomial n in QII can further enhance its performance. Conversely, the Q-Inv method alone did not achieve chemical accuracy of this system within the tested iterations.

A3 Pseudocodes

Algorithm 1 Quantum power Lanczos (QPLk)

Require: Hamiltonian $\mathcal{H}$, initial state $\left|\psi^{(0)}\right\rangle$, Krylov order d, iterations N

Initialize coefficients C_i.

Define the Krylov polynomial

$$P_d(x) = \sum_{i=0}^{d} C_i x^i.$$

Variational optimization:

repeat

  Use GQSP to prepare

$$\left|\Psi_k^{(0)}\right\rangle \propto P_d(\mathcal{H}) \left|\psi^{(0)}\right\rangle.$$

  Estimate the energy

$$E_k^{(0)} = \left\langle \Psi_k^{(0)} \right| \mathcal{H} \left| \Psi_k^{(0)} \right\rangle.$$

  Update $\{C_i\}$ to minimize $E_k^{(0)}$.

until convergence of $E_k^{(0)}$

Power iteration:

$n = 1$

repeat

  Use GQSP to prepare

$$\left|\Psi_k^{(n)}\right\rangle \propto \mathcal{H}^n P_d(\mathcal{H}) \left|\psi^{(0)}\right\rangle.$$

  Apply AA circuit

  Estimate the energy

$$E_k^{(n)} = \left\langle \Psi_k^{(n)} \right| \mathcal{H} \left| \Psi_k^{(n)} \right\rangle.$$

  $n = n + 1$

until $n = N + 1$ or convergence of $E_k^{(n)}$

return $\left|\Psi_k^{(n)}\right\rangle$ (and $E_k^{(n)}$)

Algorithm 2 Quantum inverse iteration (QII)

Require: Hamiltonian $\mathcal{H}$, shift ϵ_0, truncation order M, iterations N

1: $n = 1$
2: **repeat**
3: Compute polynomial coefficients $\{a_k\}_{k=0}^{M}$ approximating

$$(\mathcal{H} - \epsilon_0 I)^{-n} \approx \sum_{k=0}^{M} a_k \mathcal{H}^k.$$

4: Use GQSP to prepare

$$\left|\Psi^{(n)}\right\rangle \propto \sum_{k=0}^{M} a_k \mathcal{H}^k \left|\psi^{(0)}\right\rangle.$$

5: Apply AA circuit
6: Estimate the energy

$$E^{(n)} = \left\langle \Psi^{(n)} \right| \mathcal{H} \left| \Psi^{(n)} \right\rangle.$$

7: $n = n + 1$
8: **until** $n = N + 1$ or convergence of $E^{(n)}$
9: **return** $\left|\Psi^{(n)}\right\rangle$ (and $E^{(n)}$)

Algorithm 3 Quantum folded spectrum method (QFSM)

Require: Hamiltonian $\mathcal{H}$, target shift ϵ_m, constant $C > 0$, iterations N

$n = 1$

repeat

Define the shifted Hamiltonian

$$\mathcal{H}_{\text{shifted}} = \mathcal{H} - \epsilon_m I.$$

Define the folded-spectrum polynomial

$$F_n(x) = (1 - Cx^2)^n.$$

Use GQSP to prepare

$$\left|\Psi^{(n)}\right\rangle \propto F_n(\mathcal{H}) \left|\psi^{(0)}\right\rangle.$$

Apply AA circuit

Estimate the energy

$$E^{(n)} = \left\langle \Psi^{(n)} \right| \mathcal{H} \left| \Psi^{(n)} \right\rangle.$$

$n = n + 1$

until $n = N + 1$ or convergence of $E^{(n)}$

return $\left|\Psi^{(n)}\right\rangle$ (and $E^{(n)}$)

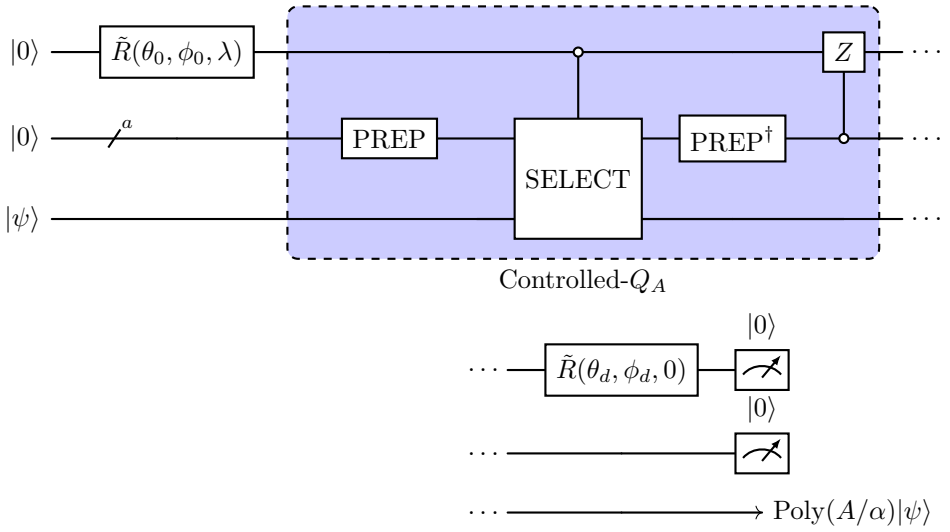


Figure B1: GQSVT circuit. The circuit comprises three registers: the first one qubit register performs signal processing; the second a qubits register performs the block encoding; and the third qubits register that holds the state, which the operator $\text{Poly}(A/\alpha)$ acts on. The state $\text{Poly}(A/\alpha)|\psi\rangle$ is up to corresponding normalization. The blue frame highlights the controlled-Q_A qubitization operator.

Appendix B: Supporting Information for Chapter 4

B1 Generalized Quantum Singular-Value Transformation

Given a $(\lambda_{\text{LCU}}, a, 0)$ block encoding V_A of A/λ_{LCU}, qubitization (see Sec. 2.2.4.3) embeds this operator into a larger unitary, whose action is reduced to two–dimensional invariant subspaces. Concretely, define

$$Q_A \;=\; (2\,|0\rangle\langle 0| \otimes I - I)\,V_A. \tag{B1}$$

Within each invariant subspace, Q_A acts as an SU(2) rotation. If σ_i are the singular values of A, then the eigenphases $\pm\theta_i$ of Q_A satisfy

$$\cos\theta_i = \sigma_i/\lambda_{\text{LCU}}, \tag{B2}$$

such that the spectrum of A is encoded at the rotation angles of Q_A.

Using Q_A as the signal operator in the GQSP framework yields the GQSVT [177], in which polynomial transformations act directly on the singular values of A. The resulting sequence has the same form as Eq. (2.168) with controlled-Q_A instead of controlled-U. If A is diagonalizable, then the procedure implements $\text{Poly}(A/\lambda_{\text{LCU}})$. The GQSVT circuit is illustrated in Fig. B1.

Therefore, the GQSVT framework provides a unified method for constructing projectors, filters, and other spectral transformations for any block encoded operator, without assuming normality. The probability of success for an input state $|\psi\rangle$ is

$$\Pr[\text{succ}] \;=\; \left\| \operatorname{Poly}(A/\lambda_{\text{LCU}})\,|\psi\rangle \right\|_2^2, \tag{B3}$$

and can be boosted using AA, similar to GQSP.

Appendix C: Supporting Information for Chapter 5

C1 Dependency Analysis of QC-CBT-AFQMC on the Number of Measurements for Ethylene and Nitrogen

For ethylene, we benchmarked our CBT-AFQMC results against SC-NEVPT2 and CCSD(T) methods (refer to the lower part of Fig. C1). In the SC-NEVPT2 analysis, the energy does not vary significantly at the 90° dihedral angle compared to CCSD(T), albeit with an offset in the energy scale. Nevertheless, both analyses reveal minimal energy discrepancies. This was expected since our system is predominantly a single reference across most geometries. However, the presence of two anomalies was corrected using Eq. (5.19).

As previously highlighted, the nitrogen molecule represents a challenging multireference scenario at extended N-N bond lengths. Figure C2 illustrates the significant escalation of errors as the bond length expands. Nonetheless, by implementing Eq. (5.19), we successfully mitigated these errors, even with a small number of measurements. This reduction is feasible because the VQE component effectively captures most of the static correlation.

C2 Complete Basis Set Extrapolation

The incompleteness of the basis set significantly contributes to discrepancies between experimental results and computational calculations. Therefore, CBS extrapolation is crucial for achieving accurate and practically viable results. Additionally, reference calculations were conducted with CBS extrapolation, prompting us to adopt a similar strategy.

Typically, for post-HF calculations, it is advisable to differentiate between the SCF energy (E_X^{SCF}) and the correlation energy (E_X^{corr}):

$$E_X^{\text{tot}} = E_X^{\text{SCF}} + E_X^{\text{corr}}, \tag{C1}$$

where E_X^{tot} denotes the total energy of the method under consideration, and X indicates the cardinal number of the basis set. Subsequently, distinct CBS schemes are applied to each component of the total energy because the convergence rate of E_X^{SCF} is typically much faster than that of E_X^{corr}. In our study, E_X^{SCF} corresponds to the VQE energy, analogous to CASCI, while E_X^{corr} represents the energy difference between the energies derived using AFQMC and

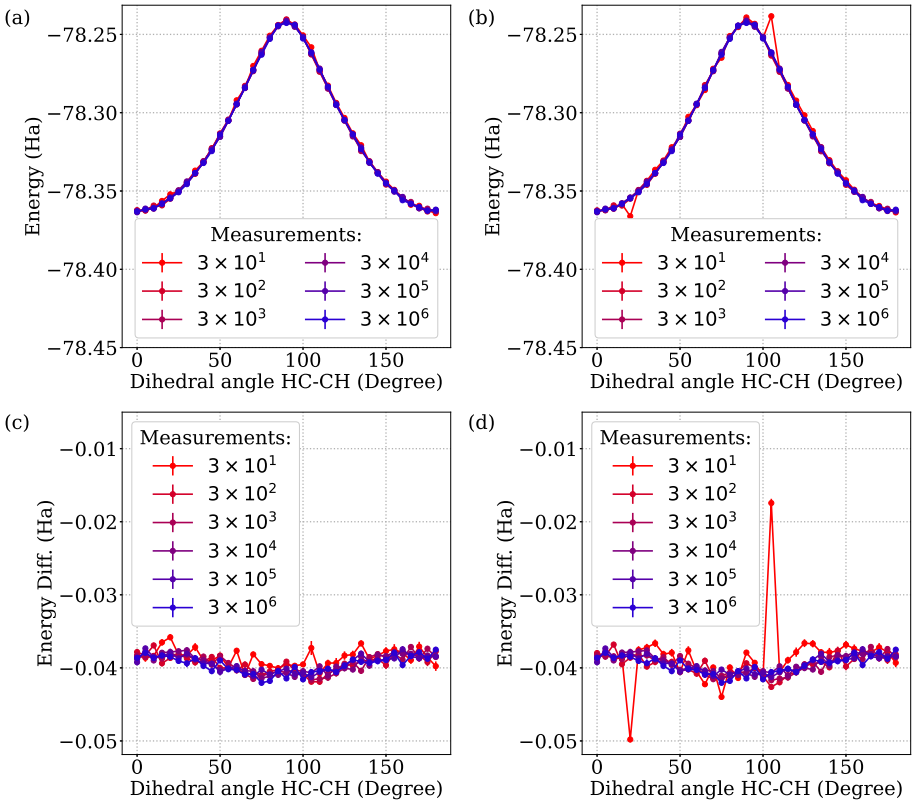


Figure C1: Potential energy curves for the ethylene molecule in the (2,2) active space with the cc-pVDZ basis set as a function of the number of CBT measurements (a) without and (b) with applying Eq. (5.19). Energy differences between QC-CBT-AFQMC and SC-NEVPT2/CASCI(2,2) as a function of the number of CBT measurements (c) without and (d) with applying Eq. (5.19).

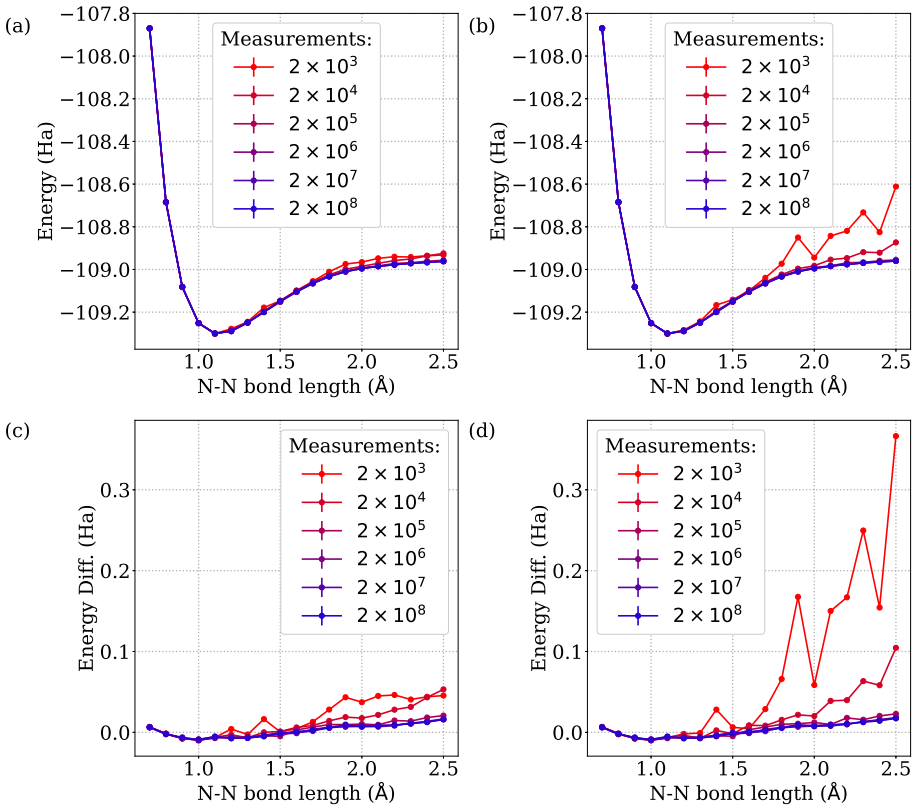


Figure C2: Potential energy curves for the N_2 molecule in the (6,6) active space with the aug-cc-pVDZ basis set as a function of the number of CBT measurements (a) without and (b) with implementing Eq. (5.19). Energy differences between QC-CBT-AFQMC and MRCISD+Q as a function of the number of CBT measurements (c) without and (d) with using Eq. (5.19).

VQE, signifying the AFQMC correlation energy. For the convergence analysis of VQE energy, we employed SCF extrapolation schemes. Table C2 illustrates the rapid convergence of VQE energies with increasing basis set size, showing a pattern similar to that observed with SCF convergence.

Several existing techniques for CBS extrapolation are applicable to SCF and CASSCF. Initially, we employed the exponential scheme [261] as described by

$$E_X^{\mathrm{SCF}} = E_\infty^{\mathrm{SCF}} + A\exp(-\beta X), \tag{C2}$$

where β, A, and E_∞^{SCF} are parameters determined through fitting. This method utilizes data from all three basis sets. The second approach adapts the exponential formula, incorporating pre-optimized parameters tailored for the CASSCF method based on molecular datasets [262]. The formula for this technique is given by

$$E_\infty^{\mathrm{CAS}} = \frac{E_{X_i}e^{\beta X_i} - E_{X_j}e^{\beta X_j}}{e^{\beta X_i} - e^{\beta X_j}}, \tag{C3}$$

where X_i represents the optimized cardinal numbers of the basis sets, which slightly deviate from their integer values, and β is another parameter obtained through fitting. For our extrapolations, we chose the aug-cc-pVTZ and aug-cc-pVQZ basis sets. Despite the different approaches, both methods yielded similar results, with a marginal discrepancy of only 1 mHa, as detailed in Tab. C1.

Furthermore, various methods are available for extrapolating the correlation energy to the CBS limit. Among these, we explored the use of a recently introduced CBS scheme that utilizes the Riemann zeta function, a method known for its generality and solid theoretical foundation [263]. The principal formula for extrapolating across three consecutive basis sets using this method is expressed as

$$E_\infty = E_L + a\left(\zeta(4) - \sum_{l=1}^{L} l^{-4}\right) + b\left(\zeta(6) - \sum_{l=1}^{L} l^{-6}\right), \tag{C4}$$

where $\zeta(x)$ denotes the Riemann zeta function (notably, $\zeta(4) = \frac{\pi^4}{90}$ and $\zeta(6) = \frac{\pi^6}{945}$). The coefficients a and b are determined as follows:

$$\begin{aligned} a &= \frac{L^6(E_L - E_{L-1}) - (L-1)^6(E_{L-1} - E_{L-2})}{2L-1}, \\ b &= L^6(E_L - E_{L-1}) - aL^2. \end{aligned} \tag{C5}$$

Here, L represents the maximum angular momentum achievable with the chosen basis set. For Dunning's basis sets, L corresponds to the cardinal number of the set. For example, in the aug-cc-pVnZ or cc-pVnZ series, $L = n$.

In addition to the Riemann zeta function approach, we applied the well-established CBS extrapolation method proposed by Helgaker *et al.* [264], which uses an inverse cubic function with two fitting parameters:

$$E_X = E_\infty + \frac{a}{X^3}, \tag{C6}$$

where X is the cardinal number of the basis set. Since this extrapolation formula necessitates the use of two basis sets, aug-cc-pVTZ and aug-cc-pVQZ were used. This scheme has been previously employed in studies focusing on the AFQMC method and has yielded plausible results [265].

Table C1 shows that the CASCI energies exhibit rapid convergence as the basis set size increases. The CBS methods, initially designed for CASSCF but applied to CASCI in our study, yield remarkably consistent energies, diverging by no more than 1.5 mHa. Such consistency underscores the reliability of the obtained results.

Conversely, the convergence of AFQMC correlation energies is notably slower. The CBS methods are tailored for correlation energy calculations, resulting in more substantial discrepancies in energies across different molecular systems. Nonetheless, when these methods are applied to computing energy barriers, the differences tend to offset each other, resulting in energy discrepancies within 1.5 kcal/mol, as detailed in Tab. C2.

Table C1: Calculated energies for all systems used in reaction barrier calculations (Hartree). For AFQMC calculations, only correlation energies are provided (Hartree). exp(2,3,4) is method (C2), SCF-E(3,4) is method (C3), ζ(2,3,4) is method (C4), and X^{-3}(3,4) is method (C6). nZ refers to the aug-cc-pVnZ basis set.

Method/basis	C_2H_4	N_2O	HN_3	TS-C_2H_4-N_2O	TS-C_2H_4-HN_3	$C_2H_4N_2O$	$C_2H_5N_3$
CASCI/DZ	−78.0431	−183.7117	−163.8674	−261.6730	−241.8448	−261.7693	−241.9556
CASCI/TZ	−78.0637	−183.7558	−163.9041	−261.7311	−241.8981	−261.8267	−242.0066
CASCI/QZ	−78.0683	−183.7679	−163.9143	−261.7470	−241.9124	−261.8425	−242.0211
CASCI/ exp(2,3,4)	−78.0697	−183.7725	−163.9181	−261.7529	−241.9176	−261.8485	−242.0268
CASCI/ SCF-E(3,4)	−78.0697	−183.7715	−163.9173	−261.7517	−241.9166	−261.8472	−242.0254
AFQMC/DZ	−0.3333	−0.5649	−0.5589	−0.9485	−0.9322	−0.9062	−0.8900
AFQMC/TZ	−0.4186	−0.7156	−0.7012	−1.1861	−1.1630	−1.1459	−1.1258
AFQMC/QZ	−0.4562	−0.7955	−0.7740	−1.3011	−1.2726	−1.2609	−1.2363
AFQMC/ ζ(2,3,4)	−0.4974	−0.8889	−0.8582	−1.4315	−1.3961	−1.3908	−1.3603
AFQMC/ X^{-3}(3,4)	−0.4837	−0.8537	−0.8271	−1.3849	−1.3526	−1.3447	−1.3170

Table C2: Reaction barrier energies in different basis sets and under different CBS extrapolations (kcal/mol). exp(2,3,4) is method (C2), SCF-E(3,4) is method (C3), ζ(2,3,4) is method (C4) and X^{-3}(3,4) is method (C6). nZ refers to the aug-cc-pVnZ basis set.

Method/basis	C_2H_4+N_2O barrier energy	C_2H_4+N_2O reaction energy	C_2H_4+HN_3 barrier energy	C_2H_4+HN_3 reaction energy
CASCI/DZ	51.4	−9.1	41.3	−28.2
CASCI/TZ	55.4	−4.6	43.8	−24.3
CASCI/QZ	56.0	−3.9	44.1	−24.1
CASCI/exp(2,3,4)	56.0	−3.9	44.0	−24.5
CASCI/SCF-E(3,4)	56.2	−3.7	44.2	−24.1
AFQMC/DZ	19.8	−14.1	16.2	−26.8
AFQMC/TZ	22.9	−11.9	16.7	−28.0
AFQMC/QZ	25.0	−9.7	17.5	−28.0
CASCI/exp(2,3,4) + AFQMC/ζ(2,3,4)	27.6	−6.8	18.6	−27.5
CASCI/SCF-E(3,4) + AFQMC/X^{-3}(3,4)	26.3	−8.4	17.9	−28.0
CCSD(T)/CBS [245]	27.6	No data	20.0	No data
CBS-QB3 [246]	27.9	−4.4	20.3	−19.7

Bibliography

[1] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, SIAM J. Comput. **26**, 1510 (1997).

[2] E. Bernstein and U. Vazirani, SIAM J. Comput. **26**, 1411 (1997).

[3] P. W. Shor, in *Proc. FOCS 1994* (1994) pp. 124–134.

[4] P. W. Shor, SIAM J. Comput. **26**, 1484 (1997).

[5] L. K. Grover, in *Proc. STOC 1996* (1996) p. 212–219.

[6] A. Y. Kitaev, arXiv preprint arXiv:quant-ph/9511026 (1995).

[7] D. S. Abrams and S. Lloyd, Phys. Rev. Lett. **83**, 5162 (1999).

[8] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, Contemp. Math. **305**, 53 (2002).

[9] L. K. Grover, Phys. Rev. Lett. **95**, 150501 (2005).

[10] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Proc. STOC 2014* (2014) p. 283–292.

[11] T. J. Yoder, G. H. Low, and I. L. Chuang, Phys. Rev. Lett. **113**, 210501 (2014).

[12] B. Yan, S. Wei, H. Jiang, H. Wang, Q. Duan, Z. Ma, and G.-L. Long, Sci. Rep. **12**, 14339 (2022).

[13] Y. I. Manin, *Vychislimoe i nevychislimoe (Computable and Noncomputable)* (Sovetskoye Radio, Moscow, 1980) in Russian.

[14] R. P. Feynman, Int. J. Theor. Phys. **21** (1982).

[15] S. Lloyd, Science **273**, 1073 (1996).

[16] H. Gao, S. Imamura, A. Kasagi, and E. Yoshida, J. Chem. Theory Comput. **20**, 1185 (2024).

[17] J. Kempe, A. Kitaev, and O. Regev, SIAM J. Comput. **35**, 1070 (2006).

[18] B. O'Gorman, S. Irani, J. Whitfield, and B. Fefferman, PRX Quantum **3**, 020322 (2022).

[19] T. S. Cubitt, D. Pérez-García, and M. M. Wolf, Nature **528**, 207 (2015).

[20] J. D. Whitfield and Z. Zimborás, J. Chem. Phys. **141**, 234103 (2014).

[21] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).

[22] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nat. Commun. **5**, 4213 (2014).

[23] O. Oumarou, P. J. Ollitrault, C. L. Cortes, M. Scheurer, R. M. Parrish, and C. Gogolin, J. Chem. Theory Comput. **21**, 4543 (2025).

[24] E. Farhi, J. Goldstone, and S. Gutmann, arXiv preprint arXiv:1411.4028 (2014).

[25] J. Copenhaver, A. Wasserman, and B. Wehefritz-Kaufmann, J. Chem. Phys. **154**, 034105 (2021).

[26] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, Proc. Natl. Acad. Sci. U.S.A. **105**, 18681 (2008).

[27] I. Gaidai, D. Babikov, A. Teplukhin, B. K. Kendrick, S. M. Mniszewski, Y. Zhang, S. Tretiak, and P. A. Dub, Sci. Rep. **12**, 16824 (2022).

[28] T. Navickas, R. J. MacDonell, C. H. Valahu, V. C. Olaya-Agudelo, F. Scuccimarra, M. J. Millican, V. G. Matsos, H. L. Nourse, A. D. Rao, M. J. Biercuk, C. Hempel, I. Kassal, and T. R. Tan, J. Am. Chem. Soc. **147**, 23566 (2025).

[29] K. Kanno, M. Kohda, R. Imai, S. Koh, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, arXiv preprint arXiv:2302.11320 (2023).

[30] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki, and A. Mezzacapo, Sci. Adv. **11**, eadu9991 (2025).

[31] N. H. Stair, R. Huang, and F. A. Evangelista, J. Chem. Theory Comput. **16**, 2236 (2020).

[32] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Phys. Rev. A **95**, 042308 (2017).

[33] N. Yoshioka, M. Amico, W. Kirby, P. Jurcevic, A. Dutt, B. Fuller, S. Garion, H. Haas, I. Hamamura, A. Ivrii, R. Majumdar, Z. Minev, M. Motta, B. Pokharel, P. Rivero, K. Sharma, C. J. Wood, A. Javadi-Abhari, and A. Mezzacapo, Nat. Commun. **16**, 5014 (2025).

[34] R. M. Parrish and P. L. McMahon, arXiv preprint arXiv:1909.08925 (2019).

[35] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Nat. Phys. **16**, 205 (2020).

[36] W. Kirby, M. Motta, and A. Mezzacapo, Quantum **7**, 1018 (2023).

[37] R. Zhang, G. Wang, and P. Johnson, Quantum **6**, 761 (2022).

[38] E. T. Campbell, Quantum Sci. Technol. **7**, 015007 (2021).

[39] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, PRX Quantum **5**, 020101 (2024).

[40] O. Kiss, U. Azad, B. Requena, A. Roggero, D. Wakeham, and J. M. Arrazola, Quantum **9**, 1682 (2025).

[41] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Nat. Commun. **1**, 149 (2010).

[42] W. J. Huggins, B. A. O'Gorman, N. C. Rubin, D. R. Reichman, R. Babbush, and J. Lee, Nature **603**, 416 (2022).

[43] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Commun. Math. Phys. **404**, 629 (2023).

[44] B. Huang, Y.-T. Chen, B. Gupt, M. Suchara, A. Tran, S. McArdle, and G. Galli, Phys. Rev. Res. **6**, 043063 (2024).

[45] E. Schrödinger, Ann. Phys. **384**, 361 (1926).

[46] M. Born and R. Oppenheimer, Ann. Phys. **389**, 457 (1927).

[47] M. Born and V. Fock, Z. Phys. **51**, 165 (1928).

[48] P. R. Bunker and P. Jensen, *Molecular Symmetry and Spectroscopy*, 2nd ed. (NRC Research Press, 2006).

[49] L. Frediani and D. Sundholm, Phys. Chem. Chem. Phys. **17**, 31357 (2015).

[50] T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, 2000).

[51] S. F. Boys, Proc. R. Soc. Lond. A **200**, 542 (1950).

[52] W. J. Hehre, R. F. Stewart, and J. A. Pople, J. Chem. Phys. **51**, 2657 (1969).

[53] R. Ditchfield, W. J. Hehre, and J. A. Pople, J. Chem. Phys. **54**, 724 (1971).

[54] T. H. Dunning Jr., J. Chem. Phys. **90**, 1007 (1989).

[55] F. Weigend and R. Ahlrichs, Phys. Chem. Chem. Phys. **7**, 3297 (2005).

[56] G. Kresse and J. Furthmüller, Comput. Mater. Sci. **6**, 15 (1996).

[57] S. Boys and F. Bernardi, Mol. Phys. **19**, 553 (1970).

[58] G. C. Wick, Phys. Rev. **80**, 268 (1950).

[59] D. R. Hartree, Math. Proc. Cambridge Philos. Soc. **24**, 89–110 (1928).

[60] V. Fock, Z. Phys. **61**, 126 (1930).

[61] D. Thouless, Nucl. Phys. **21**, 225 (1960).

[62] L. Brillouin, J. Phys. Radium **5**, 413 (1934).

[63] C. Møller and M. S. Plesset, Phys. Rev. **46**, 618 (1934).

[64] F. Coester, Nucl. Phys. **7**, 421 (1958).

[65] J. Čížek, J. Chem. Phys. **45**, 4256 (1966).

[66] O. Christiansen, H. Koch, and P. Jørgensen, Chem. Phys. Lett. **243**, 409 (1995).

[67] H. Koch, O. Christiansen, P. Jørgensen, A. M. Sanchez de Merás, and T. Helgaker, J. Chem. Phys. **106**, 1808 (1997).

[68] H. Koch and P. Jørgensen, J. Chem. Phys. **93**, 3333 (1990).

[69] J. F. Stanton and R. J. Bartlett, J. Chem. Phys. **98**, 7029 (1993).

[70] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, Chem. Phys. Lett. **157**, 479 (1989).

[71] J. Arponen, Ann. Phys. **151**, 311 (1983).

[72] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. Love, and A. Aspuru-Guzik, arXiv preprint arXiv:1701.02691 (2018).

[73] L. Erhart, Y. Yoshida, V. Khinevich, and W. Mizukami, arXiv preprint arXiv:2406.07364 (2024).

[74] B. O. Roos, P. R. Taylor, and P. E. Sigbahn, Chem. Phys. **48**, 157 (1980).

[75] C. Bloch, Nucl. Phys. **6**, 329 (1958).

[76] K. Andersson, P. Malmqvist, and B. O. Roos, J. Chem. Phys. **96**, 1218 (1992).

[77] T. Shiozaki, W. Győrffy, P. Celani, and H.-J. Werner, J. Chem. Phys. **135**, 081106 (2011).

[78] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J.-P. Malrieu, J. Chem. Phys. **114**, 10252 (2001).

[79] C. Angeli, S. Borini, A. Cavallini, M. Cestari, R. Cimiraglia, L. Ferrighi, and M. Sparta, Int. J. Quantum Chem. **106**, 686 (2006).

[80] K. Hirao, Chem. Phys. Lett. **190**, 374 (1992).

[81] H. Nakano, J. Chem. Phys. **99**, 7983 (1993).

[82] A. A. Granovsky, J. Chem. Phys. **134**, 214113 (2011).

[83] H. Werner and P. J. Knowles, J. Chem. Phys. **89**, 5803 (1988).

[84] S. R. Langhoff and E. R. Davidson, Int. J. Quantum Chem. **8**, 61 (1974).

[85] B. Jeziorski and H. J. Monkhorst, Phys. Rev. A **24**, 1668 (1981).

[86] L. Z. Stolarczyk and H. J. Monkhorst, Phys. Rev. A **32**, 725 (1985).

[87] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).

[88] R. Jastrow, Phys. Rev. **98**, 1479 (1955).

[89] G. Sugiyama and S. Koonin, Ann. Phys. **168**, 1 (1986).

[90] R. Stratonovich, Sov. Phys. Dokl. **2**, 416 (1957).

[91] Y.-K. Liu, M. Christandl, and F. Verstraete, Phys. Rev. Lett. **98**, 110503 (2007).

[92] E. Wigner, Phys. Rev. **40**, 749 (1932).

[93] E. Madelung, Z. Phys. **40**, 322 (1927).

[94] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).

[95] N. Schuch and F. Verstraete, Nat. Phys. **5**, 732 (2009).

[96] A. S. Holevo, Probl. Inf. Transm. **9**, 177 (1973).

[97] P. Benioff, J. Stat. Phys. **22**, 563 (1980).

[98] D. Deutsch, Proc. R. Soc. Lond. A **400**, 97 (1985).

[99] T. Toffoli, in *Proc. ICALP 1980*, Vol. 85 (1980) pp. 632–644.

[100] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Phys. Rev. A **52**, 3457 (1995).

[101] D. Gottesman, Phys. Rev. A **57**, 127 (1998).

[102] D. Gottesman, arXiv preprint arXiv:quant-ph/9807006 (1998).

[103] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, arXiv preprint arXiv:quant-ph/9906054 (1999).

[104] A. Y. Kitaev, Russ. Math. Surv. **52**, 1191 (1997).

[105] C. M. Dawson and M. A. Nielsen, arXiv preprint arXiv:quant-ph/0505030 (2005).

[106] V. Kliuchnikov, D. Maslov, and M. Mosca, Phys. Rev. Lett. **110**, 190502 (2013).

[107] N. J. Ross and P. Selinger, Quantum Info. Comput. **16**, 901–953 (2016).

[108] H. Morisaki, K. Sano, and S. Akibue, arXiv preprint arXiv:2510.05816 (2025).

[109] B. Eastin and E. Knill, Phys. Rev. Lett. **102**, 110502 (2009).

[110] E. Knill, arXiv preprint arXiv:quant-ph/0402171 (2004).

[111] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).

[112] A. M. Gleason, Indiana Univ. Math. J. **6**, 885 (1957).

[113] J. L. Park, Found. Phys. **1**, 23 (1970).

[114] W. K. Wootters and W. H. Zurek, Nature **299**, 802 (1982).

[115] D. Dieks, Phys. Lett. A **92**, 271 (1982).

[116] S. Aaronson, arXiv preprint arXiv:1711.01053 (2018).

[117] H.-Y. Huang, R. Kueng, and J. Preskill, Nat. Phys. **16**, 1050 (2020).

[118] V. Heyraud, H. Chomet, and J. Tilly, npj Quantum Inf. **11**, 65 (2025).

[119] P. Jordan and E. Wigner, Z. Phys. **47**, 631 (1928).

[120] S. B. Bravyi and A. Y. Kitaev, Ann. Phys. **298**, 210 (2002).

[121] A. M. Childs and N. Wiebe, Quantum Inf. Comput. **12** (2012).

[122] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Quantum Info. Comput. **5**, 467–473 (2005).

[123] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Phys. Rev. X **8**, 041015 (2018).

[124] G. H. Low and I. L. Chuang, Quantum **3**, 163 (2019).

[125] G. H. Low and I. L. Chuang, Phys. Rev. Lett. **118**, 010501 (2017).

[126] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in *Proc. STOC 2019* (2019) p. 193–204.

[127] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, PRX Quantum **2**, 040203 (2021).

[128] D. Motlagh and N. Wiebe, PRX Quantum **5**, 020368 (2024).

[129] S. Yamamoto and N. Yoshioka, arXiv preprint arXiv:2402.03016 (2024).

[130] B. K. Berntson and C. Sünderhauf, Commun. Math. Phys. **406**, 161 (2025).

[131] P. Shor, in *Proc. FOCS 1994* (1994) pp. 124–134.

[132] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, PRX Quantum **2**, 030305 (2021).

[133] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Nature **549**, 242 (2017).

[134] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, npj Quantum Inf. **5**, 75 (2019).

[135] H.-N. Xie, S.-J. Wei, F. Yang, Z.-A. Wang, C.-T. Chen, H. Fan, and G.-L. Long, Phys. Rev. A **109**, 052414 (2024).

[136] T. L. Silva, M. M. Taddei, S. Carrazza, and L. Aolita, Sci. Rep. **13**, 18258 (2023).

[137] M. Gluza, J. Son, B. H. Tiang, R. Zander, R. Seidel, Y. Suzuki, Z. Holmes, and N. H. Y. Ng, arXiv preprint arXiv:2412.04554 (2025).

[138] L. Lin and Y. Tong, Quantum **4**, 361 (2020).

[139] M.-Q. He, D.-B. Zhang, and Z. D. Wang, Phys. Rev. A **106**, 032420 (2022).

[140] G. Lee, M. Kang, J. Hong, S. Fomichev, and J. Huh, arXiv preprint arXiv:2510.04294 (2025).

[141] R. Sakuma, K. Wada, S. Kanno, K. Keithley, K. Sugisaki, T. Abe, H. Nakamura, and N. Yamamoto, arXiv preprint arXiv:2507.01361 (2025).

[142] R. Irmejs, M. C. Bañuls, and J. I. Cirac, Quantum **8**, 1389 (2024).

[143] A. Tammaro, D. E. Galli, J. E. Rice, and M. Motta, J. Phys. Chem. A **127**, 817 (2023).

[144] K. G. Dyall, J. Chem. Phys. **102**, 4909 (1995).

[145] L. Zhao, J. J. Goings, W. Aboumrad, A. Arrasmith, L. Calderin, S. Churchill, D. Gabay, T. Harvey-Brown, M. Hiles, M. Kaja, M. Keesan, K. Kulesz, A. Maksymov, M. Maruo, M. Muñoz, B. Nijholt, R. Schiller, Y. de Sereville, A. Smidutz, F. Tripier, G. Yao, T. Zaveri, C. Collins, M. Roetteler, E. Epifanovsky, A. Kovyrshin, L. Tornberg, A. Broo, J. R. Hammond, Z. Chandani, P. Khalate, E. Kyoseva, Y.-T. Chen, E. M. Kessler, C. Y.-Y. Lin, G. Ramu, R. Shaffer, M. Brett, B. Huang, M. R. Hugues, and T. Y. Takeshita, arXiv preprint arXiv:2506.22408 (2025).

[146] T. Kinoshita, O. Hino, and R. J. Bartlett, J. Chem. Phys. **123**, 074106 (2005).

[147] L. Erhart, Y. Yoshida, V. Khinevich, and W. Mizukami, Phys. Rev. Res. **6**, 023230 (2024).

[148] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, Chem. Rev. **120**, 12685 (2020).

[149] M. Motta and J. E. Rice, WIREs Comput. Mol. Sci. **12**, e1580 (2022).

[150] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309**, 1704 (2005).

[151] Z. Ding and L. Lin, PRX Quantum **4**, 020331 (2023).

[152] L. Lin and Y. Tong, PRX Quantum **3**, 010318 (2022).

[153] A. G. Taube and R. J. Bartlett, Int. J. Quantum Chem. **106**, 3393 (2006).

[154] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, and R. C. Pooser, npj Quantum Inf. **5**, 99 (2019).

[155] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Nat. Commun. **12**, 6961 (2021).

[156] H.-E. Li, X. Li, J.-C. Huang, G.-Z. Zhang, Z.-P. Shen, C. Zhao, J. Li, and H.-S. Hu, J. Chem. Phys. **161**, 144104 (2024).

[157] X. Yi, J. Huo, G. Liu, L. Fan, R. Zhang, and C. Cao, EPJ Quantum Technol. **12**, 43 (2025).

[158] I. Kolotouros, D. Joseph, and A. K. Narayanan, Phys. Rev. A **111**, 012424 (2025).

[159] H. H. S. Chan, D. M. Ramo, and N. Fitzpatrick, arXiv preprint arXiv:2303.06161 (2023).

[160] A. Leclerc and T. Carrington, J. Chem. Phys. **140**, 174111 (2014).

[161] A. Leclerc, P. S. Thomas, and T. Carrington, Mol. Phys. **115**, 1740 (2017).

[162] S. D. Kallullathil and J. Carrington, Tucker, J. Chem. Phys. **155**, 234105 (2021).

[163] Y. Ge, J. Tura, and J. I. Cirac, J. Math. Phys. **60**, 022202 (2019).

[164] K. Seki and S. Yunoki, PRX Quantum **2**, 010333 (2021).

[165] B. C. Britt, Quantum Inf. Process. **23**, 36 (2024).

[166] A. W. Harrow, A. Hassidim, and S. Lloyd, Phys. Rev. Lett. **103**, 150502 (2009).

[167] N. A. Nghiem and T.-C. Wei, Phys. Lett. A **488**, 129138 (2023).

[168] T. Ko, H. Park, and S. Choi, arXiv preprint arXiv:2408.08556 (2024).

[169] N. A. Nghiem, H. Sukeno, S. Zhang, and T.-C. Wei, Phys. Rev. A **111**, 012434 (2025).

[170] O. Kyriienko, npj Quantum Inf. **6**, 7 (2020).

[171] M. Cainelli, R. Baba, and Y. Kurashige, J. Chem. Theory Comput. **20**, 7855 (2024).

[172] A. M. Childs, R. Kothari, and R. D. Somma, SIAM J. Comput. **46**, 1920 (2017).

[173] T. Yoshikura, S. L. Ten-no, and T. Tsuchimochi, J. Phys. Chem. A **127**, 6577 (2023).

[174] M. Motta, W. Kirby, I. Liepuoniute, K. J. Sung, J. Cohn, A. Mezzacapo, K. Klymko, N. Nguyen, N. Yoshioka, and J. E. Rice, Electron. Struct. **6**, 013001 (2024).

[175] Z. Yu, H. Yao, M. Li, and X. Wang, arXiv preprint arXiv:2205.07848 (2022).

[176] Y. Wang, L. Zhang, Z. Yu, and X. Wang, Phys. Rev. A **108**, 062413 (2023).

[177] C. Sünderhauf, arXiv preprint arXiv:2312.00723 (2023).

[178] L. Cadi Tazi and A. J. W. Thom, J. Chem. Theory Comput. **20**, 2491 (2024).

[179] Y. C. Chen and T. K. Lee, Phys. Rev. B **51**, 6723 (1995).

[180] T. E. Baker, Phys. Rev. A **103**, 032404 (2021).

[181] Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui, J. J. Eriksen, Y. Gao, S. Guo, J. Hermann, M. R. Hermes, K. Koh, P. Koval, S. Lehtola, Z. Li, J. Liu, N. Mardirossian, J. D. McClain, M. Motta, B. Mussard,

H. Q. Pham, A. Pulkin, W. Purwanto, P. J. Robinson, E. Ronca, E. R. Sayfutyarova, M. Scheurer, H. F. Schurkus, J. E. T. Smith, C. Sun, S.-N. Sun, S. Upadhyay, L. K. Wagner, X. Wang, A. White, J. D. Whitfield, M. J. Williamson, S. Wouters, J. Yang, J. M. Yu, T. Zhu, T. C. Berkelbach, S. Sharma, A. Y. Sokolov, and G. K.-L. Chan, J. Chem. Phys. **153**, 024109 (2020).

[182] T. Tsuchimochi, Y. Mori, Y. Shimomoto, T. Nishimaki, Y. Ryo, M. Taii, T. Yoshikura, S. C. Tsang, K. Sasasako, and K. Yoshimura, Quket: Quantum unified kernel for emulator toolbox, `https://github.com/quket/quket` (2022).

[183] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, arXiv preprint arXiv:1701.08213 (2017).

[184] K. Setia, R. Chen, J. E. Rice, A. Mezzacapo, M. Pistoia, and J. D. Whitfield, J. Chem. Theory Comput. **16**, 6091 (2020).

[185] Y. Suzuki, Y. Kawase, Y. Masumura, Y. Hiraga, M. Nakadai, J. Chen, K. M. Nakanishi, K. Mitarai, R. Imai, S. Tamiya, T. Yamamoto, T. Yan, T. Kawakubo, Y. O. Nakagawa, Y. Ibe, Y. Zhang, H. Yamashita, H. Yoshimura, A. Hayashi, and K. Fujii, Quantum **5**, 559 (2021).

[186] L. Lin and Y. Tong, Quantum **4**, 372 (2020).

[187] Y. Tong, D. An, N. Wiebe, and L. Lin, Phys. Rev. A **104**, 032422 (2021).

[188] K. Mitarai, K. Toyoizumi, and W. Mizukami, Quantum **7**, 1000 (2023).

[189] C.-T. Li, T. Ong, C.-Y. Lin, Y.-C. Chen, H. Lin, and M.-H. Hsieh, arXiv preprint arXiv:2510.08088 (2025).

[190] M. AbuGhanem, J. Supercomput. **81**, 687 (2025).

[191] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar, and R. Babbush, Nature **646**, 831 (2025).

[192] G. Q. AI and Collaborators, Nature **638**, 920 (2025).

[193] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. Campora, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen, A. Chernoguzov, E. Chertkov, J. Colina, J. P. Curtis, R. Daniel, M. DeCross, D. Deen, C. Delaney, J. M. Dreiling, C. T. Ertsgaard, J. Esposito, B. Estey, M. Fabrikant, C. Figgatt, C. Foltz, M. Foss-Feig, D. Francois, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, E. Glynn, A. Hall, A. M. Hankin, A. Hansen, D. Hayes, B. Higashi, I. M. Hoffman, B. Horning, J. J. Hout, R. Jacobs, J. Johansen, L. Jones, J. Karcz, T. Klein, P. Lauria,

P. Lee, D. Liefer, S. T. Lu, D. Lucchetti, C. Lytle, A. Malm, M. Matheny, B. Mathewson, K. Mayer, D. B. Miller, M. Mills, B. Neyenhuis, L. Nugent, S. Olson, J. Parks, G. N. Price, Z. Price, M. Pugh, A. Ransford, A. P. Reed, C. Roman, M. Rowe, C. Ryan-Anderson, S. Sanders, J. Sedlacek, P. Shevchuk, P. Siegfried, T. Skripka, B. Spaun, R. T. Sprenkle, R. P. Stutz, M. Swallows, R. I. Tobey, A. Tran, T. Tran, E. Vogt, C. Volin, J. Walker, A. M. Zolot, and J. M. Pino, Phys. Rev. X **13**, 041052 (2023).

[194] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Rev. Mod. Phys. **92**, 015003 (2020).

[195] J. Huh, G. G. Guerreschi, B. Peropadre, J. R. McClean, and A. Aspuru-Guzik, Nat. Photonics **9**, 615 (2015).

[196] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Phys. Rev. X **8**, 011021 (2018).

[197] A. Kunitsa, D. Dhawan, S. Fomichev, J. M. Arrazola, M. Zhang, and T. F. Stetina, arXiv preprint arXiv:2508.15935 (2025).

[198] K. M. Nakanishi, K. Mitarai, and K. Fujii, Phys. Rev. Res. **1**, 033062 (2019).

[199] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, Nat. Commun. **10**, 3007 (2019).

[200] L. Veis and J. Pittner, J. Chem. Phys. **140**, 214111 (2014).

[201] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, Phys. Rev. Lett. **95**, 110503 (2005).

[202] D. Malz, G. Styliaris, Z.-Y. Wei, and J. I. Cirac, Phys. Rev. Lett. **132**, 040404 (2024).

[203] D. W. Berry, Y. Tong, T. Khattar, A. White, T. I. Kim, G. H. Low, S. Boixo, Z. Ding, L. Lin, S. Lee, G. K.-L. Chan, R. Babbush, and N. C. Rubin, PRX Quantum **6**, 020327 (2025).

[204] Z. Ding, L. Lin, Y. Yang, and R. Zhang, arXiv preprint arXiv:2510.07439 (2025).

[205] D. Lacroix, E. A. Ruiz Guzman, and P. Siwach, Eur. Phys. J. A **59**, 3 (2023).

[206] T. Tsuchimochi, Y. Mori, and S. L. Ten-no, Phys. Rev. Res. **2**, 043142 (2020).

[207] I. G. Ryabinkin, S. N. Genin, and A. F. Izmaylov, J. Chem. Theory Comput. **15**, 249 (2019).

[208] P. Siwach and D. Lacroix, Phys. Rev. A **104**, 062435 (2021).

[209] T.-C. Yen, R. A. Lang, and A. F. Izmaylov, J. Chem. Phys. **151**, 164111 (2019).

[210] V. M. Bastidas, N. Fitzpatrick, K. J. Joven, Z. M. Rossi, S. Islam, T. Van Voorhis, I. L. Chuang, and Y. Liu, Phys. Rev. A **111**, 052433 (2025).

[211] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Proc. Natl. Acad. Sci. U.S.A. **114**, 7555 (2017).

[212] Z. Li, J. Li, N. S. Dattani, C. J. Umrigar, and G. K.-L. Chan, J. Chem. Phys. **150**, 024302 (2019).

[213] P.-O. Löwdin, Phys. Rev. **97**, 1509 (1955).

[214] QunaSys Inc., QURI Parts: Modular quantum computing library, `https://github.com/QunaSys/quri-sdk` (2022).

[215] J. R. McClean, N. C. Rubin, K. J. Sung, I. D. Kivlichan, X. Bonet-Monroig, Y. Cao, C. Dai, E. S. Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, O. Higgott, C. Huang, J. Izaac, Z. Jiang, X. Liu, S. McArdle, M. Neeley, T. O'Brien, B. O'Gorman, I. Ozfidan, M. D. Radin, J. Romero, N. P. D. Sawaya, B. Senjean, K. Setia, S. Sim, D. S. Steiger, M. Steudtner, Q. Sun, W. Sun, D. Wang, F. Zhang, and R. Babbush, Quantum Sci. Technol. **5**, 034014 (2020).

[216] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[217] L. V. Slipchenko and A. I. Krylov, J. Chem. Phys. **118**, 6874 (2003).

[218] H. Zhai, C. Li, X. Zhang, Z. Li, S. Lee, and G. K.-L. Chan, arXiv preprint arXiv:2601.04621 (2026).

[219] M. Krompiec and D. M. Ramo, arXiv preprint arXiv:2210.05702 (2022).

[220] K. Mitarai, K. Toyoizumi, and W. Mizukami, Quantum **7**, 1000 (2023).

[221] C. L. Cortes, M. Loipersberger, R. M. Parrish, S. Morley-Short, W. Pol, S. Sim, M. Steudtner, C. S. Tautermann, M. Degroote, N. Moll, R. Santagati, and M. Streif, PRX Quantum **5**, 010336 (2024).

[222] S. Nishio, Y. Oba, and Y. Kurashige, Phys. Chem. Chem. Phys. **25**, 30525 (2023).

[223] P. Schleich, J. S. Kottmann, and A. Aspuru-Guzik, Phys. Chem. Chem. Phys. **24**, 13550 (2022).

[224] S. McArdle and D. P. Tew, arXiv preprint arXiv:2006.11181 (2020).

[225] A. Kumar, A. Asthana, C. Masteran, E. F. Valeev, Y. Zhang, L. Cincio, S. Tretiak, and P. A. Dub, J. Chem. Theory Comput. **18**, 5312 (2022).

[226] I. O. Sokolov, W. Dobrautz, H. Luo, A. Alavi, and I. Tavernelli, Phys. Rev. Res. **5**, 023174 (2023).

[227] S. Zhang and H. Krakauer, Phys. Rev. Lett. **90**, 136401 (2003).

[228] L. Zhao, J. J. Goings, W. Aboumrad, A. Arrasmith, L. Calderin, S. Churchill, D. Gabay, T. Harvey-Brown, M. Hiles, M. Kaja, *et al.*, arXiv preprint arXiv:2506.22408 (2025).

[229] J. J. Goings, K. Shin, S. Noh, W. Kyoung, D. Kim, J. Baek, M. Roetteler, E. Epifanovsky, and L. Zhao, J. Chem. Theory Comput. **21**, 12573 (2025).

[230] M. Motta and S. Zhang, WIREs Comput. Mol. Sci. **8**, e1364 (2018).

[231] J. Lee, H. Q. Pham, and D. R. Reichman, J. Chem. Theory Comput. **18**, 7024 (2022).

[232] Y. Chen, L. Zhang, W. E, and R. Car, J. Chem. Theory Comput. **19**, 4484 (2023).

[233] Y. Yoshida, L. Erhart, T. Murokoshi, R. Nakagawa, C. Mori, H. Tagami, and W. Mizukami, J. Chem. Theory Comput. **21**, 9404 (2025).

[234] M. Kiser, M. Beuerle, and F. Šimkovic, J. Chem. Theory Comput. **21**, 2256 (2025).

[235] N. Takemori, Y. Teranishi, W. Mizukami, and N. Yoshioka, arXiv preprint arXiv:2312.17452 (2023).

[236] G. Mazzola and G. Carleo, arXiv preprint arXiv:2205.09203 (2022).

[237] J. Lee, D. R. Reichman, R. Babbush, N. C. Rubin, F. D. Malone, B. O'Gorman, and W. J. Huggins, arXiv preprint arXiv:2207.13776 (2022).

[238] M. Kohda, R. Imai, K. Kanno, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, Phys. Rev. Res. **4**, 033173 (2022).

[239] E. Bäumer and S. Woerner, arXiv preprint arXiv:2408.03064 (2024).

[240] F. D. Malone, A. Mahajan, J. S. Spencer, and J. Lee, J. Chem. Theory Comput. **19**, 109 (2023).

[241] Chemqulacs, `https://github.com/wmizukami/chemqulacs` (2023).

[242] F. Neese, WIREs Comput. Mol. Sci. **2**, 73 (2012).

[243] J. C. Jewett and C. R. Bertozzi, Chem. Soc. Rev. **39**, 1272 (2010).

[244] N. K. Devaraj, ACS Cent. Sci. **4**, 952 (2018).

[245] D. H. Ess and K. N. Houk, J. Phys. Chem. A **109**, 9542 (2005).

[246] A. Karton and L. Goerigk, J. Comput. Chem. **36**, 622 (2015).

[247] T. Stuyver, K. Jorner, and C. W. Coley, Sci. Data **10**, 66 (2023).

[248] P. Vermeeren, M. Dalla Tiezza, M. E. Wolf, M. E. Lahm, W. D. Allen, H. F. Schaefer, T. A. Hamlin, and F. M. Bickelhaupt, Phys. Chem. Chem. Phys. **24**, 18028 (2022).

[249] Y. O. Nakagawa, M. Kamoshita, W. Mizukami, S. Sudo, and Y.-y. Ohnishi, J. Chem. Theory Comput. **20**, 10817 (2024).

[250] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, *et al.*, arXiv preprint arXiv:2405.05068 (2024).

[251] S. Barison, J. Robledo Moreno, and M. Motta, Quantum Sci. Technol. **10**, 025034 (2025).

[252] D. Kaliakin, A. Shajan, F. Liang, J. Robledo Moreno, Z. Li, A. Mitra, M. Motta, C. Johnson, A. A. Saki, S. Das, I. Sitdikov, A. Mezzacapo, and K. M. Merz Jr, Commun. Phys. **8**, 396 (2025).

[253] K. Sugisaki, S. Kanno, T. Itoko, R. Sakuma, and N. Yamamoto, arXiv preprint arXiv:2412.07218 (2024).

[254] M. Mikkelsen and Y. O. Nakagawa, arXiv preprint arXiv:2412.13839 (2024).

[255] J. Yu, J. R. Moreno, J. T. Iosue, L. Bertels, D. Claudino, B. Fuller, P. Groszkowski, T. S. Humble, P. Jurcevic, W. Kirby, *et al.*, arXiv preprint arXiv:2501.09702 (2025).

[256] S. Kanno, H. Nakamura, T. Kobayashi, S. Gocho, M. Hatanaka, N. Yamamoto, and Q. Gao, npj Quantum Inf. **10**, 56 (2024).

[257] N. S. Blunt, L. Caune, and J. Quiroz-Fernandez, arXiv preprint arXiv:2410.18901 (2024).

[258] B. K. Berntson and C. Sünderhauf, arXiv preprint arXiv:2406.04246 (2024).

[259] R. Chao, D. Ding, A. Gilyen, C. Huang, and M. Szegedy, arXiv preprint arXiv:2003.02831 (2020).

[260] L. Ying, Quantum **6**, 842 (2022).

[261] A. Halkier, T. Helgaker, P. Jørgensen, W. Klopper, and J. Olsen, Chem. Phys. Lett. **302**, 437 (1999).

[262] F. N. N. Pansini, A. C. Neto, and A. J. C. Varandas, Theor. Chem. Acc. **135**, 261 (2016).

[263] M. Lesiuk and B. Jeziorski, J. Chem. Theory Comput. **15**, 5398 (2019).

[264] T. Helgaker, W. Klopper, H. Koch, and J. Noga, J. Chem. Phys. **106**, 9639 (1997).

[265] M. Amsler, P. Deglmann, M. Degroote, M. P. Kaicher, M. Kiser, M. Kühn, C. Kumar, A. Maier, G. Samsonidze, A. Schroeder, M. Streif, D. Vodola, C. Wever, and Q. M. S. W. Group, J. Chem. Phys. **159**, 044119 (2023).

Acknowledgments

I would first like to express my deep gratitude to my academic advisor, Professor Keisuke Fujii of the Graduate School of Engineering Science at the University of Osaka, and to Professor Wataru Mizukami of the Center for Quantum Information and Quantum Biology at the University of Osaka. Through their guidance, encouragement, and patience I was able to find my way through the academic world, even when my own stubbornness made that a harder job than it needed to be. I am also sincerely grateful to my Master's advisor, Professor Ilya Ioffe of Moscow State University, who continued to support and encourage me even after I moved to another institution.

My thanks go as well to my lab colleagues, who became my close friends here: Dr. Luca Erhart, Leonardo Placidi, and Dr. Tomoya Shiota. The good humor in our daily interactions and their willingness to discuss ideas made this period of my life far more enjoyable and creative than I could have hoped for. I am equally grateful to Dr. Yuichiro Yoshida and Dr. Tuan Minh Do for their constructive criticism and thoughtful advice throughout my PhD. I also appreciate the support of the administrative and institutional staff who handled the practical matters that kept my research moving.

After all the formalities, I want to express my heartfelt gratitude to my family for supporting me in every possible way throughout this long journey. Their care has been a constant source of stability and confidence.

Finally, I want to thank my friends, who often believed in me more than I believed in myself. Their support played a major role in helping me choose and stay on this path.

List of Publications

Main Publications

1. V. Khinevich and W. Mizukami, "Symmetry-Adapted State Preparation for Quantum Chemistry on Fault-Tolerant Quantum Computers", arXiv:2601.08533 (2026).

2. V. Khinevich, Y. Lee and N. Yoshioka and W. Mizukami, "Quantum Power Iteration Unified Using Generalized Quantum Signal Processing", arXiv:2507.11142 (2025).

3. V. Khinevich and W. Mizukami, "Enhancing quantum computations with the synergy of auxiliary field quantum Monte Carlo and computational basis tomography", Physical Review Research, **7**, 4, 043344 (2025).

Other Publications

4. L. Erhart, Y. Yoshida, V. Khinevich and W. Mizukami, "Chebyshev Approximated Variational Coupled Cluster for Quantum Computing", arXiv:2406.07364 (2024).

5. L. Erhart, Y. Yoshida, V. Khinevich and W. Mizukami, "Coupled cluster method tailored with quantum computing", Physical Review Research, **6**, 2, 023230 (2024).

6. V. E. Khinevich, S. M. Sudarkova, and I. N. Ioffe, "Effect of exohedral functionalization on the magnetic properties in the dysprosium-containing endohedral fullerene $DySc_2N@C_{80}(CF_2)$", Physical Chemistry Chemical Physics, **26**, 42, 26765-26771 (2024).

7. S. M. Sudarkova, V. E. Khinevich and I. N. Ioffe, "Substitution pattern dependent behavior of the singlet excited states in symmetrically fluorinated biphenyls", Physical Chemistry Chemical Physics, **23**, 38, 22067-22077 (2021).

8. A. V. Belyakov, R. Yu. Kulishenko, R. D. Johnson, I. F. Shishkov, A. N. Rykov, V. Yu. Markov, V. E. Khinevich and A. A. Goryunkov, "Structure of $C_{60}F_{36}$: A Gas-Phase Electron Diffraction and Quantum Chemical Computational Study of a Remarkably Distorted Fluorofullerene", The Journal of Physical Chemistry A, **124**, 49, 10216-10224 (2020).