Finanziato dall'Unione europea
NextGenerationEU

MUR
Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

UNIVERSITÀ
DEGLI STUDI
DELL'AQUILA

DISIM
Dipartimento di Ingegneria
e Scienze dell'Informazione
e Matematica

Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica

Dottorato di Ricerca in Matematica e Modelli
XXXVIII ciclo

Quantum Computing-Based Methods for Electronic Structure in Industrial Chemistry Applications

SSD 03/CHEM-02

Supervisor e Co-Supervisor
Prof. Leonardo Guidoni
Ph.D Martina Stella

PhD candidate
Fabio Tarocco

PhD course Coordinator
Prof. Debora Amadori

year 2024/2025

Preface

The rapid development of quantum technologies has opened the possibility of rethinking the traditional computational chemistry workflow in a hybrid quantum–classical perspective. While fully fault-tolerant quantum computing remains out of reach, at least practically, hybrid strategies already offer a realistic opportunity to integrate quantum subroutines into established pipelines, improving or accelerating well-defined step of the electronic structure problem. Rather than focusing on isolated quantum algorithms, this thesis adopts a workflow-driven viewpoint: *given a realistic quantum chemistry pipeline, which specific steps can benefit from quantum computing, and how should these quantum components interact with the "classical*[1]*" quantum chemistry fundaments?*

Given a molecular system to treat, a standard electronic structure workflow, generally, consists of four major steps: (1) optimization of the target system molecular geometry on the corresponding Potential Energy Surface (PES), (2) the identification of a chemically meaningful active space, (3) the computation of accurate single-point energies on that space, and finally, (4) the recovery and refinement thanks to the inclusion of dynamical correlation from the complementary space. In this thesis, we investigate each of these stages and develop *Hybrid Quantum–Classical (HQC)* methods that target them individually, with the long-term goal of enabling quantum-enhanced end-to-end pipelines. All of these pieces, have been analyzed, independently, both methodologically and temporally, but with in mind always the idea of producing a fully interacting hybrid quantum–classical framework for managing a quantum chemistry pipeline.

We begin by addressing the selection of the active space, a step that is both chemically crucial and often subjective. To this end, we introduce the *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS)*, which uses atomic orbital information and entanglement measures to semi–automate and standardize the active space identification. The method is then applied in the field of Photodynamic Therapy (PDT), a light–based cancer treatment, which identification of efficient molecules as Photosensitizer (PS) is still a challenge for wavefunction based methods. Once the active space is fixed, the next challenge lies in preparing electronic states efficiently that approximate the ground state with reasonable accuracy. Having a efficient and accurate single–point solver or warm–start procedure is essential, especially when dealing with complex systems. We tackle this through the *Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA)*, a warm-start strategy that improves Variational Quantum Eigensolver (VQE) convergence by exploiting quantum information-theoretic criteria. Multi-QIDA is therefore tested on general molecular systems and also, not chemistry related, strongly–correlated system, i.e. 2D Heisenberg spin–lattices. However, most of the time and for complex systems, static correlation alone is not sufficient to achieve the desired chemical accuracy, as dynamical correlation outside the active space must also be captured. To address this, we extend the previously defined *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR)* algorithm, into *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)*, which incorporates portion of the total dynamical correlation through a self-consistent treatment of orbital– rotations. Additionally, given the already modular framework, we also tried to include adaptivity to the procedure, defining the *ADAPT–WAHTOR–SCF* procedure. To test the applicability of the method,

---

[1] Used wrongly here, and also in the rest of manuscript, for referring to quantum chemistry problem solved on classical machines.

we benchmarked our WAHTOR–SCF approach against general molecular systems bond dissociation, and, more specific, for industrial chemistry. In particular, we studied a free–radical model polymerization reaction that involves three different structure. Finally, having established tools for treating both static and dynamical correlation, we move to the nuclear degrees of freedom. We propose a proof-of-concept strategy for geometry refinement, the *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)*, which integrates quantum energies within a hybrid Potential Energy Surface navigation scheme, integrated with different quantum chemistry software and quantum computing simulation backed, both on CPU and on GPU. The hybrid routine is then finalized as structure optimizer for DompeKey (DK), recently formalized molecular descriptors, as well as for standard benchmark molecules.

Looking at these results as a whole, the central message can be state as follows: *Quantum computing should not, and will never, be viewed as a monolithic replacement for classical electronic structure theory, but rather as a set of specialized tools to be integrated step–by–step into already well–established pipelines.* An Hybrid Quantum–Classical pipeline is necessary for allowing "classical" to exploits quantum algorithms as sub–routines, while allowing quantum–computers to have a whole classical framework to manipulate efficiently the quantum data obtained.

The remainder of this thesis is organised according to this workflow logic. After introducing the theoretical background in Part I, split between Quantum computing theory, Chapter 1, and Quantum Chemistry fundamentals, Chapter 2, while in Chapter 3, we present the main linking tools for integrating "classical" theory into a quantum framework. These sections do not present any advancement and they are exclusively used to present, in a personal way, the theory between the two worlds, by reinterpretation of existing textbooks and theoretical manuscripts. With this section, we conclude the first part of the thesis, completely oriented to give a brief introduction of the theory required to understand the development part of the manuscript.

In the second part of the thesis, Part II, we start from Chapter 4, in which presents AEGISS, our active space selection method, and the specific application for Photodynamic Therapy. Once the importance to have a robust method for selecting the active region on which focus the analysis, single–point calculations can start. In fact, Chapter 5 introduces Multi-QIDA, our state–preparation or warm–start method based on quantum information theory measures. Chapter 6 develops WAHTOR–SCF, in order to include the missing correlation that a method focused only on the active space orbital might miss. Finally, Chapter 7 introduces our structure optimization or geometry refinement method, the QA–q–GeomOpt protocol.

As the last developmental chapter of the thesis in Part III, Chapter 8, we wrap up the functionalities and implementations available in our code base, Quantum @ L'Aquila (QuAQ), while final remarks and perspectives are given in concluding final chapter.

Contents

List of Figures

List of Tables

Acronyms

1o-RDM 1-orbital reduced density matrix

ADAPT–WAHTOR–SCF Adaptive Hardware–Efficient Ansatz Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field

ADAPT-VQE Adaptive Derivative-Assembled Pseudo-Trotter ansatz Variational Quantum Eigensolver

ADAPT-VQE–SCF ADAPT–VQE–SCF

AEGISS Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection

AVAS Atomic Valence Active Space

BOA Born–Oppenheimer Approximation

CAS Complete Active Space

CASCI Complete Active Space Configuration-Interaction

CASPT2 Complete Active Space Second–Order Perturbation Theory

CASPTn Complete Active Space Perturbation Theory (to the n-th order)

CASSCF Complete Active Space Self–Consistent Field

CC Coupled-Cluster

CCSD Coupled–Cluster Singles–Doubles

CCSD(T) Coupled–Cluster Singles–Doubles Pertubative Triples

CI Configuration Interaction

CID Configuration Interaction Doubles

CIS Configuration Interaction Singles

CISD Configuration Interaction Singles–Doubles

DFT Density Functional Theory

DK DompeKey

DMRG Density Matrix Renormalization Group

DMRG-SCF Density Matrix Renormalization Group Self–Consistent Field

ECP Effective Core Potential

F2Q Fermion–to–Qubit

FCI Full Configuration Interaction

GGA Generalized Gradient Approximations

HEA Hardware–Efficient Ansatz

HEA–WAHTOR–SCF Hardware–Efficient Ansatz Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field

HF Hartree–Fock

HPC High–Performance Computing

HQC Hybrid Quantum–Classical

IL Inter–Ligand

InChI International Chemical Identifier

INOs Iterative Natural Orbitals

ISC Inter–system Crossing

JW Jordan–Wigner

KS Kohn–Sham

LDA Local Density Approximation

MC–PDFT Multi–Configuration Paired Density Functional Theory

MCSCF Multi-Configuration Self-Consistent Field

MLCT Metal–to–Ligand Charge Transfer

MP2 Second—order Møller-–Plesset Perturbation Theory

MPS Matrix Product Ansatz

MS-CASPT2 Multi–State Complete Active Space Second–Order Perturbation Theory

Multi-QIDA Multi–threshold Quantum Information Driven Ansatz

NEVPT N-Electron Valence State Perturbation Theory

NEVPT2 2-Electron Valence State Perturbation Theory

NISQ Noisy Intermediate–Scale Quantum

NOONs Natural Orbital Occupation Numbers

OO–VQE Orbital–Optimized VQE

PDT Photodynamic Therapy

PEC Potential Energy Curve

PES Potential Energy Surface

PQC Parametrized Quantum Circuit

PS Photosensitizer

QA–q–GeomOpt Quasi-Adiabatic Quantum Geometry Optimization

QDK Quantum Developement Kit

QIDA Quantum Information Driven Ansatz

QM Quantum Mechanics

QMI Quantum Mutual Information

QPE Quantum Phase–Estimation

QPU Quantum Processing Unit

QuAQ Quantum @ L'Aquila

RDM Reduce Density Matrix

ROS Reactive Oxygen Species

SCF Self–Consistent Field

SD Slater Determinant

SMART SMILES arbitrary target specification

SMILE Simplified Molecular Input Line Entry System

SOE Single Orbital Entropy

SparQ Sparse Quantum–State Analysis

TD-DFT Time–Dependednt Density Functional Theory

UCC Unitary Coupled–Cluster

UCCSD Unitary Coupled Cluster Singles and Doubles

VQE Variational Quantum Eigensolver

WAHTOR Wavefunction–Adapted Hamiltonian Through Orbital–Rotation

WAHTOR–SCF Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field

XC Exchange–Correlation

Abstract

The advent of Noisy Intermediate–Scale Quantum (NISQ) devices has opened new perspectives for quantum chemistry, where quantum processors can act as specialized solvers embedded within classical workflows. Rather than replacing established quantum chemistry methodologies, Hybrid Quantum–Classical (HQC) approaches can be used to integrate quantum subroutines where they can most effectively accelerate or enhance the description of molecular systems. This thesis adopts a workflow-oriented perspective, asking how a realistic electronic structure pipeline, from active space definition, state preparation, correlation recovery, to geometry optimization, all of which can be systematically hybridized with quantum computing.

Four original contributions are developed to address these stages. The *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS)* provides a pseudo-automated strategy for chemically sound active space definition. The *Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA)* introduces a state-preparation, or warm-start procedure, that exploits quantum information theory measures to improve *Variational Quantum Eigensolver (VQE)* convergence and precision, linking both chemistry inspired and *Hardware–Efficient Ansatz (HEA)*. The *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)* extends the existing *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR)* procedure to capture dynamical correlation beyond the active space through orbital relaxation, exploting once again cheap *Hardware–Efficient Ansatz (HEA)*. Finally, *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)* demonstrates a proof-of-concept structure–refinement procedure oriented even in this case to Hardware–Efficient Ansatz.

Together, all these developments define a consisten vision of quantum computing-aided computational chemistry, where quantum processors are not simply isolated tools but modular components embedded into an efficient classical backbone. The results highlight the feasibility of end-to-end hybrid workflows, shrinking the gap between theoretical innovation and practical applicability in chemical and pharmaceutical discovery.

Part I

Theory

Chapter 1

Quantum Computing in a nutshell

Before we can talk about simulating molecules on quantum computers, we have to face a quite unfair truth: classical computers can be outmatched when chemistry gets serious. The complexity of treating exactly[1], even a modest molecule[2], scales quickly and becomes not manageable. This is not because classical algorithms are badly implemented or inefficient, but because nature plays by the rules of quantum mechanics, and so far, quantum chemistry algorithms have been trying to catch up with a rulebook written in another language. Quantum computers, *in theory*, speak that language natively. They approximate the quantum world, they are quantum systems. And that opens the possibility to simulate molecular interactions with a precision that is beyond the actual limit of classical methods. Quantum computers, as we already said, play by the same rules as the chemistry they are simulating. They can, at least in principle, represent molecular states exactly, scaling to systems that would send even the fastest supercomputers into a nervous breakdown. As first step, to know how this works in practice, it is worth understanding the machinery itself, i.e. how quantum computation operates, and why it might finally let us solve problems that have been a dream for classical computational quantum chemistry for decades.

What is necessary to state from the start is that, nowadays, a quantum computer is not a general–purpose system on which everything can be run. Instead, we have to see it as a component of a standard computer in which a part of its computation relies on a *Quantum Processing Unit (QPU)*. Most of the algorithms are subdivided into classical computation, in which we prepare the instance of the problem and the part of the evaluation of the solution, and quantum computation for the calculus of the solution employing quantum approaches.

Additionally, we need to point out the fact that on quantum computers is possible to reproduce any classical computation. Quantum computers can do anything that a classical computer can do with at least the same computational complexity. Nevertheless, it is clear that, there is no point in building a quantum computer if we want to use it in the same way we use a classical one. We must rewrite algorithms and develop new techniques targeting the regime in which, classical algorithms can not compete. For example, simulating complex interacting systems of electrons, like for example in quantum chemistry applications. After this, we can start to introduce the basis behind quantum computation and how we can take advantage from quantum mechanics to speed-up our algorithms.

This chapter is made by taking inspiration from standard quantum computing textbooks [1–

---

[1] We specify exactly, and not by introducing any sort of approximation, i.e. treating the whole system at exact level of theory.

[2] Extended basis set for small molecules or more than four non-hydrogen atoms.

3] and it is structured in the following way: first, in Section 1.1, we will introduce briefly the postulates of Quantum Mechanics (QM) and how they relate to the fundamental building block of quantum computers, the qubits. In Section 1.2, we will discuss about how our wavefunction can be abstracted, visualized, and manipulated using the quancum circuit paradigm. In Section 1.3, the available *Noisy Intermediate–Scale Quantum (NISQ)* devices will be presented very quickly, as well as the type of errors that affect them and some mitigation and correction techniques.

1.1 Brief dive into Quantum Mechanics

In order to give a complete overview on the whole pipeline we are going to follow inside the manuscript, we start from the building blocks of the quantum computations and the fundamentals rules that allow us to operate in the field of quantum mechanics. We will do this taking inspiration for different textbooks [1–3], as already pointed out in the introduction.

1.1.1 Single qubit representation

Differently from a classical computer, in which the smallest computational unit is the *bit*, the building block of a quantum computer, or more general a quantum algorithm, is a single *qubit* [3]. With a simple example, a bit is a unit that, according to the physical implementation, can be either 0 OR 1. A qubit instead is defined as bi-dimensional complex object that can described as being *simultaneously* (not it the time-related meaning) in the 0 and 1 state. Exploiting the *first postulate* of quantum mechanics, a single qubit can be described as

Postulate 1. *Each isolated physical system has an associated Hilbert space complex, called the state space of the system. The system is completely described by its state vector, which is a unit vector in the state space.*

and we are going to stick to the mathematical representation of it, or group of qubits, for the rest of the thesis. Therefore, mathematically, a qubit is described by a *state-vector* $\mathbf{v}$, with $\mathbf{v} \in \mathbb{C}^2$. The bi-dimensional complex vectorial space, in which $\mathbf{v}$ lives, is denoted as *Hilbert space* $\mathcal{H}$, where $\mathcal{H} = \mathbb{C}^2$. The dimension of a Hilbert space associated with a single-qubit is $dim(\mathcal{H}) = 2$. The state-vector of single-qubit, $\mathbf{v}$ being the description of a two-level quantum system, is defined as

$$\mathbf{v} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \text{ with } \alpha, \beta \in \mathbb{C}, \tag{1.1}$$

with the additional condition that the norm of $\mathbf{v}$ must be equal to 1, i.e.

$$|\alpha|^2 + |\beta|^2 = 1. \tag{1.2}$$

Being in two *simultaneously* (still used improperly) states is defined as *superposition.* The numbers of quantum superpositions are *possibly infinitely many*, but we cannot determine the quantum state of a qubit with the same precision we would obtain a classical bit. In fact, when we measure a qubit, we can only obtain 0 with a probability of $|\alpha|^2$ or the state 1 with a probability of $|\beta|^2$. The coefficients α and β are called *amplitudes of probability*, while the squared module of them is simply the *probability* of obtaining the corresponding state. This is the reason behind the fact that every qubit must be normalized and $|\alpha|^2 + |\beta|^2 = 1$, because the sum of the probability to be in each state must be 1.

The most common way to represent the component of a state-vector is by using two *orthonormal* vectors,

$$\text{State } 0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ State } 1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \tag{1.3}$$

to compose the associated *computational basis.*

Dealing with only one qubit makes the notation light-weighted, but as the system start to grow, it might become cumberstone very quickly. The *de-facto* standard is by representing the complex vector using the more compact *Dirac notation*, also referred as *bra-ket notation.* The vector $\mathbf{v}$, a column complex vector, is denoted with a *ket*, $|v\rangle$. The *hermitian conjugate* of $\mathbf{v}$, a complex row vector, is denoted using a *bra*, $\langle v|$. If $\langle u|$ is the hermitian conjugate of $|v\rangle$, then they are related by

$$|v\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \ \langle u| = |v\rangle^\dagger = (|v\rangle^T)^* = (\alpha^*, \beta^*), \tag{1.4}$$

where $(\cdot)^\dagger$ is the conjugate transpose operation or hermitian conjugate operation. Following the new notation, the computational basis reads as

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{and} \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \tag{1.5}$$

where both the states, $|0\rangle$ and $|1\rangle$, are orthonormal vectors, i.e. $\langle 0|1\rangle = \langle 1|0\rangle = 0$ and $\langle 0|0\rangle = \langle 1|1\rangle = 1$. Therefore, combining the two orthonormal states we define a *orthonormal basis* $\{|j\rangle\}$ with $|j\rangle \in \{|0\rangle, |1\rangle\}$.

A *ket* $|v\rangle$ represents a generic element of the vector space, while $|j\rangle$ can be used to refer to the j-th element of the computational basis. Using the previous definition, we can expand $|v\rangle$ in a more general version using the elements of the basis

$$|v\rangle = \sum_j \alpha_j |j\rangle,$$

where $\alpha_j \in \mathbb{C}$ and $\sum_j |\alpha_j|^2 = 1$.

Inner product

Useful operation fundamental for measuring the overlap between two wavefunction or in general for measuring observables, is the *inner product*, also know as *scalar product* or *dot product.* The inner product between two state-vectors returns a scalar. In particular, the inner product between $|\psi\rangle$ and $|\phi\rangle$, where

$$|\psi\rangle = \sum_j \alpha_j |a_j\rangle, \ |\phi\rangle = \sum_i \beta_i |a_i\rangle, \tag{1.6}$$

is defined as

$$\langle\psi|\phi\rangle = \sum_j \alpha_j^* \langle a_j| \sum_i \beta_i |a_i\rangle = \sum_j \sum_i \alpha_j^* \beta_i \langle a_j|a_i\rangle = \sum_j \sum_i \alpha_j^* \beta_i \delta_{j,i} = \sum_j \alpha_j^* \beta_j, \tag{1.7}$$

in which $\delta_{i,j}$ is the *Kronecker delta* , defined as

$$\delta_{ij} = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{if } i = j. \end{cases}$$

Obviously, the inner product between a state-vector and its hermitian conjugate yields one (if the normalization condition is respected).

Outer product

If we reverse the order of the contraction, we perform an *outer product.* Using the same previous vector $|\psi\rangle$ and $|\phi\rangle$, the outer product is defined as

$$|\psi\rangle\langle\phi| = \sum_j \alpha_j |a_j\rangle \sum_i \beta_i^* \langle a_i| = \sum_j \sum_i \alpha_j \beta_i^* |a_j\rangle\langle a_i| \tag{1.8}$$

which is equivalent of building the $N \times N$ matrix:

$$|\psi\rangle\langle\phi| = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_N \end{pmatrix} \begin{pmatrix} \beta_1^* & \dots & \beta_N^* \end{pmatrix} = \begin{pmatrix} \alpha_1\beta_1^* & \alpha_1\beta_2^* & \dots & \alpha_1\beta_N^* \\ \alpha_2\beta_1^* & \alpha_2\beta_2^* & \dots & \alpha_2\beta_N^* \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_N\beta_1^* & \alpha_N\beta_2^* & \dots & \alpha_N\beta_N^* \end{pmatrix},$$

where N is the size of the computational basis.

Projectors

When the outer-product is built using the same element of the computational basis, i.e. $|a_i\rangle$, what we obtain is a *Projector*, defined as $|a_i\rangle\langle a_i|$. A projector $\hat{P}$ is a *hermitian* operator i.e. $P^\dagger = P$ which satisfy the *idempotency* relation i.e. $P^2 = P$. If we have a set of projectors $\{P_i\}_{i\in N}$, defined as $|a_i\rangle \langle a_i|$, we can use the completeness relation $\sum_j |a_j\rangle \langle a_j| = I$ on a generic state $|\Psi\rangle$ to obtain how it is expanded in the $\{|a_i\rangle\}_{i\in N}$ basis as

$$|\Psi\rangle = I\,|\Psi\rangle = \sum_{i\in N} |a_i\rangle \langle a_i|\Psi\rangle = \sum_{i\in N} (\langle a_i|\Psi\rangle)\, |a_i\rangle\,, \tag{1.9}$$

where $(\langle a_i|\Psi\rangle) \in \mathbb{C}$ and represents the coefficient of the in the expansion.

1.1.2 Multi-Qubits representation

Tensor product

Postulate 2. *The state space of a composite physical system is the tensor product of the state spaces of the physical subsystem. Moreover, if we have system composed of n subsystems, and the j-th system is prepared in the state $|\psi_j\rangle$, then the total state of the composite system will be $|\psi_1\rangle \otimes |\psi_2\rangle \otimes \dots |\psi_j\rangle \otimes \dots \otimes |\psi_n\rangle$.*

This postulate describes how to build the state space of a quantum system composed of two or more distinct physical system from the individual qubits subspaces. For doing this, we have to introduce an operator that allows us to define composite systems and it is the *tensor product*, $\otimes$. If we have two qubits $|v\rangle$ and $|u\rangle$, the composite system is $|v, u\rangle$ defined as

$$|v, u\rangle = |v\rangle \otimes |u\rangle = \begin{pmatrix} v_0 \\ v_1 \end{pmatrix} \otimes \begin{pmatrix} u_0 \\ u_1 \end{pmatrix} = \begin{pmatrix} v_0 u_0 \\ v_0 u_1 \\ v_1 u_0 \\ v_1 u_1 \end{pmatrix}, \tag{1.10}$$

associated in order with computational basis states $|00\rangle, |01\rangle, |10\rangle$, and $|11\rangle$, and with $v_i, u_i \in \mathbb{C}$. Extending to n qubits is possible to obtain 2^n distinct states. A quantum register of n-qubits is an element of the 2^n-dimensional Hilbert space, which corresponds to the total $\mathbb{C}^{2^n}$ complex space. If we take 2^n quantum register of n-qubits we ca build the related computational basis, which we can write in the form

$$|x_1\rangle \otimes |x_2\rangle \otimes \cdots \otimes |x_j\rangle \otimes \cdots \otimes |x_n\rangle$$

where $x_j \in \{0, 1\}$ and $1 \leq j \leq n$. As explained for the single qubit case, the number of superposition of 2^n computational basis states is infinite, but after a measurement , the state will collapse to one of the relative classical state. The result of the measurement will be one of the basis state $|v_i\rangle$, where $v_i \in \{0, 1\}^n$ with probability equals to $|\alpha_j|^2$. A generic wavefunction in this basis will appear as

$$|\psi\rangle = \sum_{i \in \{0,1\}^n} c_i |v_i\rangle, \tag{1.11}$$

where v_i is a vector of the n- qubits computational basis i.e. a string of 0s and 1s.

Density Matrices

Given a state $|\Psi\rangle$ defined over the orthonormal computational basis $\{|a_i\rangle\}_{i \in N}$, the *Density matrix* associated to it is defined as

$$\rho = |\Psi\rangle \langle\Psi| = \sum_{i \in N} \sum_{j \in N} a_j^* a_i |a_i\rangle \langle a_j| \tag{1.12}$$

and it generalizes a classical probability distribution. A density matrix must be *hermitian*, i.e. $\rho = \rho^\dagger$, semi-definite positive, i.e. $\langle v| \rho |v\rangle \geq 0$ for each $|v\rangle \in \mathcal{H}$, and *unitary trace*, i.e. $Tr(\rho) = 1$. If the density matrix is defined as just showed, it is associated with a *pure* state i.e. $Tr(\rho^2) = 1$, and this value is called *purity*. A state can also be defined as a mixture of state, namely *mixed state*, as

$$\rho = \sum_i p_i \rho_i, \tag{1.13}$$

where p_i is the classical probability to obtain the state represented by ρ_i. For a mixed state, the purity is $Tr(\rho) < 1$, reaching the lowest value $\frac{1}{N}$ in the case of a *maximally-mixed* state, i.e. all the quantum states in the ensemble have equal probability.

1.1.3 Measurements procedure

In the previous paragraph, we used multiple times the concept of measurement, now we introduce this concept with the proper formalism.

Postulate 3. *The physical quantities that can be measured on a system are called Observables and they are represented by Hermitian operators. Every observable* Π *is represented by a hermitian operator on the state space of the system being observed. Quantum measurements are described by a set of measurement operators* $\{\Pi_m\}$ *which act on the state space of the measured system. The possible outcomes of a measure of the observable* Π *are denoted by the index* m*, while the eigenvalues associated are* ϵ_m*.*

An operator Π has the form

$$\Pi = \sum_m \epsilon_m \Pi_m = \sum_m \epsilon_m |\phi_m\rangle\langle\phi_m| \tag{1.14}$$

where Π_m is the projector, before denoted with P, on the subspace of the eigenvectors $|\phi_m\rangle$ of Π related to the eigenvalue m. After a measure of Π_m in the state $|\psi\rangle$, the probability of obtaining m is

$$p(m) = \langle\psi| \Pi_m^\dagger \Pi_m |\psi\rangle = |\Pi_m |\psi\rangle|^2. \tag{1.15}$$

The state of the system post-measurement with result m is

$$\frac{\Pi_m |\psi\rangle}{\sqrt{p(m)}}. \tag{1.16}$$

The Π_m operators satisfy the *completeness equation*

$$\sum_m \Pi_m = \mathbf{I}. \tag{1.17}$$

This equation states that the sum of the probability of the outcome of a measure must add up to 1. In fact,

$$1 = \sum_m p(m) = \sum_m \langle\psi| \Pi_m^\dagger \Pi_m |\psi\rangle = \langle\psi|\psi\rangle \tag{1.18}$$

In particular, this type of measurement is often called *projective measurement.* Π is determined by any set of ortogonal projector Π_m that satisfy the completeness relation. If $\Pi_m = |\phi_m\rangle\langle\phi_m|$ is a projector then

$$\begin{aligned} p(m) = \langle\psi| \Pi_m^\dagger \Pi_m |\psi\rangle = \delta_{mm'} \langle\psi| \underbrace{|\phi_m\rangle\langle\phi_m| \, |\phi_m\rangle\langle\phi_m|}_{\mathbf{\Pi_m = |\phi_m\rangle\langle\phi_m|}} |\psi\rangle = \\ = \langle\psi|\phi_m\rangle\langle\phi_m|\psi\rangle = |\langle\phi_m|\psi\rangle|^2. \end{aligned} \tag{1.19}$$

Expectation value

Generally, the outcome of a measure of an observable A in the state $|\psi\rangle$ is aleatory. The average value of the outcome of the measure of the operator A, namely its *expectation value*, is defined as

$$\begin{aligned}
\mathbf{E}_\psi(A) &= \sum_a \epsilon_a p(a) \\
&= \sum_a \epsilon_a \langle\psi| \underbrace{|\phi_a\rangle \langle\phi_a|}_{\Pi_a} |\psi\rangle \\
&= \langle\psi| (\sum_a \epsilon_a \Pi_a) |\psi\rangle \\
&= \langle\psi| A |\psi\rangle ,
\end{aligned} \tag{1.20}$$

where $A = \sum_a \Pi_a$. Usually the expectation value of an observable A is also denote with $\langle A\rangle$.

The measurement operation is *irreversible.* For any process to be reversible, it should be possible to obtain the *initial state* from the *final state.* During measurement, the information about the initial state is lost. Quantum mechanics provides some statistical information on the possible results of a measurement according to what is known as *Born's statistical interpretation* or known as *Born's rule.* By repeating measurements on copies of the system, it is possible to establish the probabilistic distribution of the results. The meaning of probability of an outcome must be understood according to the interpretation given by the theory of probabilities as relative frequency: the probability of a result is the ratio between the number of times the experiment is successful and the total number of experiments done (repeated for a sufficiently large number of times).

1.1.4 Entanglement

In this paragraph, we are going to formally introduce the concept of *entanglement*, properties that, by itself, differentiate classical and quantum computers. Entanglement is introduced when qubits are allowed to interact. *Interacting* systems cannot be written in the tensor product form as we seen for isolated system. If they are allowed to interact, the isolated system will contains both qubits together, and the states of the qubits may not be factorizable in the subsystems. When this happens, the qubits that are interacting each other are *entangled.* Entangled states own properties that we can not find in any other classical object. For the fact that they interact with other qubits, they do not own a individual state, and only the complete system has a well defined state. An example of state composed by entangled qubit is one of the Bell states [3]. If we take the the state \$ $|\beta_{00}\rangle$ defined as

$$|\beta_{00}\rangle = \frac{|00\rangle - |11\rangle}{\sqrt{2}}, \tag{1.21}$$

by following Equation 1.10, we can expand it as

$$|\beta_{00}\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix} . \tag{1.22}$$

It is clear that there is no decomposed composition of u_0, u_1, v_0, and v_1 such that the tensor product of $u = \begin{pmatrix} u_0 \\ u_1 \end{pmatrix}$ and $v = \begin{pmatrix} v_0 \\ v_1 \end{pmatrix}$ replicates $|\beta_{00}\rangle$. The state $|\beta_{00}\rangle$ is defined as *entangled state.*

Partial Trace

To identify whether a quantum state/system is made up by entangle qubits, we introduce the *Partial Trace* operatrion. The action of the partial trace is to reduce the whole system representation to one or more sub-systems. Given a quantum state represented by the density matrix ρ_{AB}, composed by the two subsystems A and B, the partial trace on sub-system A, is defined as

$$\rho_A = Tr_A(\rho_{AB}) = \sum_{i \in B} (I \otimes \langle i_B|)\rho_{AB}(I \otimes |i_B\rangle),$$

where i is a vector of the computational basis of the subsystem B. The obtained density matrix can also be defined as *reduced density matrix* for sub-system A. If we apply the partial trace on one qubit system, we obtain a simple *trace* operation i.e. the sum over the diagonal entries of the state density matrix. At this point, a quantum state is defined as *separable* if and only if it can written as mixture/linear combination of product states, i.e. product of the sub-systems reduced density matrices, as follows

$$\rho_{AB} = \sum_i p_i \rho_A^i \otimes \rho_B^i. \tag{1.23}$$

If a density matrix of a state cannot be written in the mixure form, then the state is *entangled.*

1.1.5 Entanglement measures

Along the whole thesis, multiple times the entanglement will be used as central metric in the algorithm developed. It is therefore necessary to introduce the associated measures that can be used to quantify entanglement or correlation between different component of the whole system.

Von Neumann Entropy

The *Von Neumann Entropy* can be used to measure the uncertainty of a quantum system. Given a quantum state described by the density matrix ρ_{AB}, the Von Neumann entropy $S(A)$ for the sub-system A, is defined as

$$S(A) = -Tr(\rho_A \log(\rho_A)), \tag{1.24}$$

where $\rho_A = Tr_A(\rho_{AB})$ and $S(A) \geq 0$. The Von Neumann entropy is the quantum equivalent of the classical *Shannon Entropy.*

Quantum Mutual Information (QMI)

Quantum counterpart of the classical *Quantum Mutual Information (QMI)*, it is used to measure the amount of correlation between two quantum states. It is defined starting as the sum of the single Von Neumann entropies of the sub-systems minus the entropy of the composed system, i.e.

$$I(A, B) = S(A) + S(B) - S(A, B), \tag{1.25}$$

where $I(A,B) \geq 0$ and $S(A), S(B), S(A,B) \geq 0$.

1.2 Quantum Circuit paradigm

After defining how a system of one single or multiple qubits is composed, naturally, we need to be able to manipulate them, evolve the state, perform the desired algorithm, and measure the outcome. The most common paradigm to deal with quantum algorithm is the Quantum Circuit paradigm. The *Quantum Gates* are the building blocks for quantum circuits and to each of them is associated an operation that evolves the state of the system. Quantum gates are unitary operators (i.e. $UU^\dagger = U^\dagger U = I$). In figure 1.1, a generic quantum circuit collects all the most common components used in standard execution: $\boxed{U}$ is a *single-qubit* gate, $\boxed{G}$ is a *multi-qubit* gate, x $\boxed{R(\theta)}$ is a *parametrized single-qubit* gate, $\boxed{CU}$ is a *controlled-operation*, each single horizontal — line represent a qubit, the double line ═ classical bit, while the 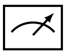 is the *measurement* operation. The gates are applied to the state from left to right, therefore, the temporal application of gates in the circuit model is reversed w.r.t. to the standard algebraic applications of operator. In particular, if we follow the same structure of Figure 1.1, the application of the state $|00\rangle$ will be $CU_{01} \cdot R(\theta)_0 \cdot G_{01} \cdot U_0 \, |00\rangle$.

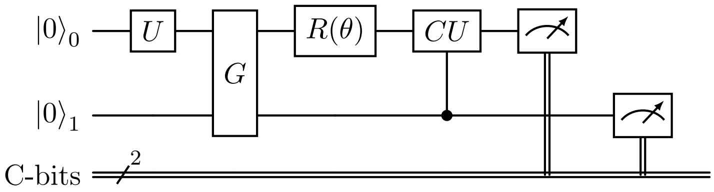

Figure 1.1: Example of a quantum circuit composed of the most common gates and operation that are present in quantum algorithms.

1.2.1 Single Qubit Gates

Here, we present some of the unitary operators that acts on a simple single-qubit system. Starting from the *Pauli's matrices*, namely X, Y, and Z, defined as

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \; Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \; Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \tag{1.26}$$

in which also the *Identity* matrix I can be added. which are known as *Pauli's matrices.* The set composed by the Pauli matrices $\{X, Y, Z, I\}$ is going to be of fundamental relevance in the following sections, in particular for the fact that, any operator that we want to measure and evaluate on a quantum computer, needs to be defined as a linear combination of *Pauli string*, i.e. tensor product of Pauli matrices. Pauli matrices in terms of manipulations, are associated with basic bit- and phase-flips.

We also introduce the *Hadamard gate*, which allows us to create a superposition of states:

$$H \equiv \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \equiv \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{1.27}$$

These type of transformations are *linear operators* i.e. function L: $\mathbb{C}^2 \mapsto \mathbb{C}^2$ is said to be linear if for each $\alpha_1, \alpha_2 \in \mathbb{C}$ and $|v_1\rangle, |v_2\rangle \in \mathbb{C}^2$ s.t.

$$L(\alpha_1 |v_1\rangle + \alpha_2 |v_2\rangle) = \alpha_1 L |v_1\rangle + \alpha_2 L |v_2\rangle$$

Other common operators are the generic rotations in the x, y, and z direction defined respectively by

$$R_X(\theta) = \begin{bmatrix} cos(\frac{\theta}{2}) & -isin(\frac{\theta}{2}) \\ -isin(\frac{\theta}{2}) & cos(\frac{\theta}{2}) \end{bmatrix}, \; R_Y(\theta) = \begin{bmatrix} cos(\frac{\theta}{2}) & -sin(\frac{\theta}{2}) \\ sin(\frac{\theta}{2}) & cos(\frac{\theta}{2}) \end{bmatrix}, \; R_Z(\theta) = \begin{bmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{bmatrix} \tag{1.28}$$

as well as *Phase* gate S and T-gate, respectively defined as

$$S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \; T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{bmatrix}. \tag{1.29}$$

The former three operations compose a continuous set of rotation that can be used to define any other single-qubit unitary. The listed above Pauli matrices are the generators of the Rotations gates.

1.2.2 Multi-qubit quantum gates

Operation on quantum register of multiple qubits are necessary to describe the transformation of composite system and they are essential for the introduction of *entanglement* in the system. In general, if our system is composed of n qubits, then applying a gate U to the j-th qubit corresponds to applying the following operation

$$\mathbf{I} \otimes \mathbf{I} \otimes \cdots \otimes U \otimes \cdots \otimes \mathbf{I} \otimes \mathbf{I}$$

to the entire composite system. In the circuit model, when we apply a gate to a single qubit of a series of qubits, the identity gate is considered to be applied to the remaining qubits even if it is not inserted in the circuits.

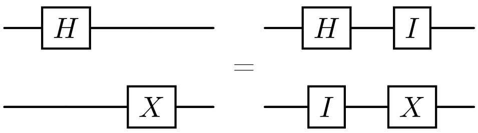

Figure 1.2: Each step is defined by an application of the gates to the system.

Controlled-NOT

One of the most important 2−qubit gate is the `CNOT`(controlled-`NOT`), which is the quantum counterpart of the classical `XOR`. The `CNOT`, acts on two qubits: the *target* and the *control.* If the control is 0 then the target is unchanged; if the control is 1, then the target is negated.

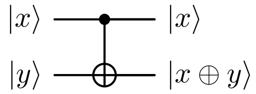

Figure 1.3: CNOT gate, $|x\rangle$ is the control qubit and $|y\rangle$ is the target qubit

Given two quantum states $|x\rangle$ and $|y\rangle$, we can see the `CNOT` as the transformation

$$|x, y\rangle \mapsto |x, x \oplus y\rangle\,,$$

where $|x\rangle$ is the control, $|y\rangle$ is the target and $\oplus$ is the classical `XOR` operation. This gate can also be represented in matrix[3] form

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix},$$

where we can that the `CNOT` maps the states $|00\rangle \mapsto |00\rangle\,, |01\rangle \mapsto |01\rangle\,, |10\rangle \mapsto |11\rangle$,and $|11\rangle \mapsto |10\rangle$.

Controlled-Unitaries

Given an arbitrary single qubit unitary operation U, a controlled-U gate is a two qubit operation, with the same structure seen for the `CNOT`. If the control qubit is in the state $|1\rangle$ then U is applied to the target qubit, otherwise the target qubit is unchanged. What we obtain by applying the $Controlled - U$ to the composite system is $|x, y\rangle \mapsto |x\rangle\, U^x\, |y\rangle$.

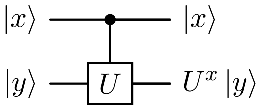

Figure 1.4: Controlled-U gate, $|x\rangle$ is the control qubit and $|y\rangle$ is the target qubit.

Toffoli Gate

A Toffoli gate is a 3-qubit quantum logic gate, that operates as a doubled controlled NOT acting on a single target qubit. It's a reversible gate, as all the one listed before, and is universal for classical computation, i.e. any Boolean function can be built using Toffoli gates. It enables implementing classical logic within quantum circuits and is a key component for error correction and reversible

---

[3] The reported CNOT matrix assume the ordering of the qubits in which the control qubit is the left-most, and the targer, the right-most. Software packages/ QSDK might assume different ordering (e.g. Qiskti assumes little-endian ordering).

computation. Generally denoted with CCNOT, with associated matrix and symbol,

$$CCNOT = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \tag{1.30}$$

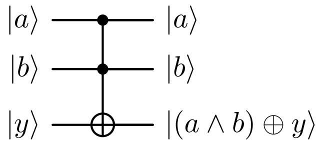

Figure 1.5: Toffoli (CCNOT) gate, $|a\rangle$ and $|b\rangle$ are the control qubits and $|y\rangle$ is the target qubit

1.2.3 Universal set

A quantum computer needs to be able to perform arbitrary single-qubit operations, e.g. H, S, T, R_x, R_y, and R_z, and introduce entanglement between qubits using at least one non-trivial multi-qubit gate, e.g. CNOT or CZ (Controlled-Z). If the set of gates that we are using contains gates allowing for both the manipulations, then , the set is a *universal gate set* for quantum computation, i.e. any n-qubit unitary transformation can be approximated to arbitrary accuracy. Amongst the previous presented gates, H, Pauli matrices, S, and the $CNOT$ gates are defined as *Clifford* gates. Clifford gates are such that they map Pauli matrices into Pauli matrices, and circuit composed by them, according to the Gottesman-Knill theorem [3], can be efficiently simulated with a classical computer. The T gate, rotations gate, and Toffoli gate, which are all[4] *non-cliffond* , are therefore fundamental to unlock the universality.

In this way, different combinations of them can lead to universality, in particular: Different gate combinations achieve universality [3]:

- **All single-qubit rotations** $+$ $CNOT$: If we can do arbitrary rotations on one qubit and entangle qubits with $CNOT$, we can implement any n-qubit unitary.
- H**,** S**,** T**,** $CNOT$: Widely used in theoretical developement of **fault-tolerant quantum computing** algorithm, as these gates are easier to protect with error correction.
- H**,** T**,** $CNOT$: $H + T$ cover arbitrary single-qubit unitaries and adding CNOT makes the set universal.
- **Pauli rotations** $+$ $CNOT$: Using rotations (R_x, R_y, and R_z), or equivalently (X,Z,H,T), any single-qubit unitary can be generated, and CNOT provides the entanglement needed.

---

[4] Rotations gate for $\theta = \frac{k\pi}{2}$ are Clifford.

- **Toffoli** + H: As demonstrated by Aharonov [4], Toffoli plus Hadamard is also universal. Toffoli handles all *classical logic*, and H injects the ability to form quantum superpositions.

- **Approximating arbitrary single-qubit rotations**: Using the **Solovay–Kitaev theorem**, discrete gate sets such as H, T can approximate any rotation (R_x and R_z) to arbitrary precision.

1.3 Quantum Devices

In this section, we briefly address the main technologies that, nowadays, represent the quantum computing ecosystem. Various physical implementations of qubits have been proposed, tested, developed and are currently under development. Each technology offers advantages, peculiarity, and limitations. It will be a short description for the main hardware technologies due the fact that hardware implementation is not the topic of this thesis.

- **Superconducting qubits (e.g., transmons)**: Superconducting circuits built from Josephson junctions dominate today's commercial devices (IBM, Google, Rigetti). They benefit from well-developed microfabrication processes, relatively fast gate times on the order of 10–100 ns, and straightforward on-chip scalability. However, they require operation at mK temperatures in dilution refrigerators, suffer from relatively short coherence times 50–500 μs), and are sensitive to fabrication variability and crosstalk as the number of qubits grows [5]

- **Trapped-ion qubits**: Ions such as $^{171}\mathrm{Yb}^+$ or $^{40}\mathrm{Ca}^+$ are held in electromagnetic traps and manipulated with laser beams. These systems exhibit some of the longest coherence times in any platform (seconds to minutes) and very high gate fidelities exceeding 99.9%, with all ions being nearly identical in frequency. The trade-offs include relatively slow gate operations (tens to hundreds of μ s), as well as engineering challenges in scaling to hundreds or thousands of ions due to the complexity of laser control and ion shuttling between trap zones. Top players in the sector are IonQ and Honeywell. A nice review can be found in the references [6]

- **Spin-qubits in semiconductors (e.g., quantum dots)**: Spin qubits leverage the spin of single electrons or holes confined in semiconductor quantum dots [7], often fabricated in silicon or III–V compounds materials. In this field, Intel and QuTech, are the leading vendor. They promise high-density integration compatible with established semiconductor manufacturing, and benefit from long coherence times in isotopically purified silicon. Nonetheless, their gate fidelities are still improving, and they remain sensitive to charge and nuclear-spin noise, while requiring advanced cryogenic control electronics for precise manipulation [8].

- **Neutral-atom / Rydberg atom qubits (e.g., optical tweezer arrays)**: Neutral atoms such as rubidium or cesium are trapped in arrays of optical tweezers or lattices and use Rydberg-state interactions for entangling gates. This approach offers natural scalability with flexible two-dimensional geometries and relatively long coherence times ranging from milliseconds to seconds. Present challenges include slower and less reliable two-qubit gates, which typically reach fidelities around 98–99%, and sensitivity to laser intensity noise and background gas collisions [9–12].

- **Photonic qubits (linear optics and integrated photonics)**: Photonic approaches encode qubits in single photons' polarization. They can operate at or near room temperature, experience negligible decoherence during transmission, and are well suited for quantum communication and networking. Yet, implementing deterministic two-qubit gates is difficult because photons do not naturally interact, leading to probabilistic gate schemes and heavy reliance on high-efficiency photon sources, detectors, and error-correcting schemes to achieve scalability [13–15]. Photonic qubits most famous players are Xanadu and PsiQuantum.

- **Topological qubits (e.g., Majorana-based qubits, still experimental)**: Topological qubits aim to encode information in non-local quasiparticle modes such as Majorana zero modes, theoretically offering intrinsic protection against many types of local noise. This inherent error resilience could dramatically reduce the overhead needed for quantum error correction [16, 17]. However, these systems remain in the early stages of experimental development, with no large-scale or fault-tolerant demonstration to date, and the reliable creation and braiding of topological excitations is still an open research challenge.

- **Quantum Annealers**: Even if is not properly a type of hardware, quantum annealing is a quantum computing paradigm applied to optimization problems, exploiting a similar resolution to Simulated annealing. On top of this, quantum annealer exploits quantum tunneling and superposition to find the global minimum of a given function, starting from the eigensolution of a know problem, and adiabaically evolving to the target problem [18, 19]. For quantum annealest, the Canadian company D-Wave Systems is the leading in the sector.

1.3.1 Noisy Intermediate-Scale Quantum Computers

Despite the presence of a lot players in the quantum hardware development field, fault tolerant digital quantum computer is still not available, and current intermediate scale devices are still effected by a considerable noise level. Although this limits the depth of the circuits that can be executed faithfully, these noisy intermediate-scale quantum (NISQ) devices are already been used in different field. NISQ devices are susceptible to errors, which can only be partially mitigated using error mitigation procedures. We report very briefly the main errors and techniques for dealing with them (for detailed and more in-depth explanation, check the associated references, and textbooks [2, 3, 20–22]).

Types of Errors

Physical qubits are prone to *bit-flip*, *phase-flip*, and combined *bit-phase* errors. A bit-flip error acts as a Pauli-X operator,

$$X|0\rangle = |1\rangle\,, \quad X|1\rangle = |0\rangle\,, \tag{1.31}$$

while a phase-flip corresponds to a Pauli-Z operator,

$$Z|0\rangle = |0\rangle\,, \quad Z|1\rangle = -|1\rangle\,. \tag{1.32}$$

A depolarizing channel [3, 20–22] acting on a single qubit can be modeled as

$$\mathcal{E}(\rho) = (1-p)\rho + \frac{p}{3}\big(X\rho X + Y\rho Y + Z\rho Z\big), \tag{1.33}$$

where p is the depolarization probability and ρ the qubit density matrix. Errors can occur during *storage* (decoherence), *gate application*, or *measurement*.

Noise

Noise describes the random, typically continuous interaction between a quantum system and its environment. The loss of phase coherence is characterized by the *dephasing time* T_2, while energy relaxation is described by T_1. For a pure dephasing channel the off-diagonal terms of $\rho(t)$ decay exponentially:

$$\rho_{01}(t) = \rho_{01}(0)e^{-t/T_2}. \tag{1.34}$$

Environmental fluctuations, control imperfections, and thermal excitations lead to *decoherence*, reducing the fidelity of quantum operations. Suppressing noise requires careful material engineering, shielding, and dynamical decoupling techniques.

Gate Errors

Gate errors occur during the application of quantum operations due to imperfect control pulses, cross-talk between qubits, or hardware limitations. If an ideal unitary operation is U, the implemented gate can be modeled as a noisy operation

$$\tilde{U}(\rho) = \mathcal{E}(U\rho U^\dagger), \tag{1.35}$$

where $\mathcal{E}$ is a quantum channel representing the error[5]. Common gate errors include *over-rotation*, *under-rotation*, and *phase errors*, which accumulate over successive operations and degrade the overall circuit fidelity:

$$\mathcal{F} = \mathrm{Tr}\left(\sqrt{\sqrt{\rho_{\text{ideal}}}\,\rho_{\text{actual}}\,\sqrt{\rho_{\text{ideal}}}}\right)^2. \tag{1.36}$$

Minimizing gate errors is critical for achieving high-fidelity quantum algorithms, and techniques *dynamical decoupling*, and calibration routines are often employed.

Measurement Errors

Measurement errors occur during the readout of qubit states, when the observed outcome differs from the actual qubit state due to detector noise, imperfect discrimination, or relaxation during measurement. If the true qubit state is ρ, the measured probability distribution p_{meas} can be modeled as

$$\vec{p}_{\text{meas}} = M\,\vec{p}_{\text{true}}, \tag{1.37}$$

---

[5] Mathematically, it is a completely positive, trace-preserving (CPTP) map acting on density matrices as $\mathcal{E} : \rho \to \mathcal{E}(\rho)$. If the system were perfect, this map would simply be the unitary evolution in Equation 1.35

where $\vec{p}_{\text{true}}$ is the ideal probability vector and M is a classical confusion matrix encoding the misassignment probabilities. For a single qubit, a typical readout error probability is

$$\epsilon = \Pr(0 \to 1) \approx \Pr(1 \to 0), \tag{1.38}$$

which reduces the fidelity of quantum algorithms, particularly when many qubits are measured simultaneously. Mitigation strategies include calibration of the confusion matrix and classical post-processing to invert the measurement bias.

Error Mitigation

Error mitigation does not change the underlying hardware but applies algorithmic or calibration techniques to reduce bias in measured observables. For example, *Zero-Noise Extrapolation* (ZNE) [23, 24] rescales the effective noise strength λ and extrapolates expectation values $\langle O \rangle(\lambda)$ to the zero-noise limit:

$$\langle O \rangle_0 \approx \lim_{\lambda \to 0} \langle O \rangle(\lambda). \tag{1.39}$$

Other methods include *Probabilistic Error Cancellation* [23], which reconstructs an ideal circuit as a quasi-probability mixture of noisy gates, and *Measurement-Error Mitigation*, which inverts the classical confusion matrix obtained from calibration.

Error Correction

Quantum error correction (QEC) [20, 21] encodes one logical qubit $|\psi_L\rangle$ into a subspace of n physical qubits, allowing detection and correction of errors without measuring the logical state directly. A simple three-qubit bit-flip code maps

$$|0\rangle_L = |000\rangle, \qquad |1\rangle_L = |111\rangle, \tag{1.40}$$

detecting any single-qubit bit-flip via parity checks. Fault-tolerant schemes such as the *surface code* [25] use topological redundancy to correct both bit- and phase-flip errors, enabling scalable quantum computation once gate and measurement error rates fall below a threshold (typically $\sim 10^{-3}$).

Chapter 2

Quantum Chemistry in a nutshell

As long as the quantum hardware is still sort of made up relative small number of qubits, noisy gates, and limited scale keep a modest problem still out of reach. For now, classical methods, powered by decades of refinement, remain the workhorse of computational chemistry. At its core, the problem is simple[1]: solve the Schrödinger equation for a collection of electrons and nuclei. In practice, however, the full solution is hopelessly out of reach, as the computational cost grows exponentially with system size. Faced with this challenge, generations of chemists and physicists have devised a hierarchy of approximations to make the impossible tractable. The Hartree–Fock method offered the first systematic approach, capturing the mean-field behavior of electrons but missing critical correlations. Post-Hartree–Fock methods such as configuration interaction and coupled-cluster theory added layers of precision, at the price of a steep increase in computational demands. In parallel, density functional theory revolutionized the field by trading exactness for efficiency, enabling simulations of thousands of atoms—though always with the nagging caveat of functional choices and limits to reliability.

These methods, refined over decades, now form a powerful toolbox. For the vast majority of chemical problems—from small organic molecules to complex reaction mechanisms—they deliver predictions accurate enough to guide experiments and inspire new discoveries. Computational chemistry has become not just a complement to the lab, but often a driver of innovation in materials science, drug discovery, and catalysis. And yet, the cracks remain. When electron correlation becomes strong, as in transition metal complexes, excited-state dynamics, or photosynthetic systems, classical approximations strain under their own weight. Problems that require both high accuracy and large system sizes remain stubbornly elusive, forcing researchers to choose between detail and scale.

The following chapter is a personal tentative to fit all the minimal theoretical requirement to deal with an electronic structure problem. It is written following two essential quantum chemistry textbooks, *Molecular Electonic–Structure Theory*[2] by Trygve Helgaker [26], and *Modern Quantum Chemistry* by A. Szabo and N. S. Ostlund [27]. The structure of the chapter is defined as follow: In Section 2.1, we introduce the fundamental problem of computational quantum chemistry, the *Molecular Electronic–Structure Problem*, passing from the first quantization definition, Section 2.1.1, through all the second quantized formalism, Section 2.2, ending by defining the Electronic

---

[1] Relatively simple once the thoery has been mastered.

[2] A.k.a. the *Purple bible* of computational quantum chemists.

Hamiltonian in second quantization, Section 2.2.5. In Section 2.3, we introduce the Hartree–Fock theory, giving an insight on how the self–consistent procedure works, in Section 2.3.3. After dealing with single determinant wavefunctions, we move to beyond Hartree-Fock methods, in Section 2.4, introducing the Configuration Interaction method, Section 2.4.1, Complete Active Space methods, Sections 2.4.2-2.4.3, passing by Density Matrix Renormalization Group theory, in Section 2.4.4, and concluding with other famous methods, but treated very quickly, such as Møller-Plesset perturbation theory, Coupled–Cluster, both in Section 2.4.5, and Density Functional Theory, in Section 2.4.6.

2.1 The Molecular Electronic Structure Problem

The fundamental goal of molecular quantum chemistry is to determine the electronic structure of atoms and molecules by solving the stationary Schrödinger equation

$$\hat{H}\Psi = E\Psi, \tag{2.1}$$

where $\hat{H}$ is the molecular Hamiltonian, Ψ the exact wave function and E the corresponding energy eigenvalue. For a system of N_e electrons and N_N nuclei, the first-quantized Hamiltonian in atomic units reads

$$\hat{H} = -\sum_{i=1}^{N_e} \frac{1}{2}\nabla_i^2 - \sum_{A=1}^{N_N} \frac{1}{2M_A}\nabla_A^2 - \sum_{iA} \frac{Z_A}{r_{iA}} + \sum_{i<j} \frac{1}{r_{ij}} + \sum_{A<B} \frac{Z_A Z_B}{R_{AB}}. \tag{2.2}$$

This Hamiltonian contains the kinetic energy of electrons and nuclei, the electron–nuclear attraction, the electron–electron repulsion, and the nuclear Coulomb repulsion. Even for small molecules, the exact solution of Eq. (2.1) is impossible analytically and grows exponentially in cost numerically [26, 28].

2.1.1 Born Oppenheimer Approximation (BOA)

Because nuclei are much heavier than electrons, they move on a slower timescale. The BOA exploits this separation by freezing nuclear positions while solving the electronic problem. Under BO, the Hamiltonian becomes

$$\hat{H}_e = -\sum_{i} \frac{1}{2}\nabla_i^2 - \sum_{iA} \frac{Z_A}{r_{iA}} + \sum_{i<j} \frac{1}{r_{ij}} + V_{NN}, \tag{2.3}$$

where V_{NN} is a constant for fixed nuclear coordinates. The electronic Schrödinger equation,

$$\hat{H}_e \Psi(\mathbf{r}_1, \ldots, \mathbf{r}_{N_e}) = E_e \Psi(\mathbf{r}_1, \ldots, \mathbf{r}_{N_e}), \tag{2.4}$$

defines the potential energy surface on which nuclei move.

2.2 Second Quantization Framework

While the first-quantized formulation expresses the many–electron wave function explicitly in terms of particle coordinates, the *second quantization* formalism provides a more compact and algebraically convenient representation of many-fermion systems. Instead of tracking the coordinates of individual electrons, it focuses on the occupation of one-particle states (spin–orbitals) and describes changes in occupancy through algebraic operators. This approach naturally incorporates the Pauli exclusion principle and simplifies the derivation of electronic structure methods [26, 28].

2.2.1 From First to Second Quantization

Let $\{\phi_p(\mathbf{x})\}$ be an orthonormal set of spin–orbitals, with $\mathbf{x} = (\mathbf{r}, \sigma)$ denoting both spatial and spin coordinates. A general antisymmetric N_e-electron wave function can be written as a linear combination of Slater determinants constructed from these spin–orbitals. Each determinant corresponds to a unique configuration of occupied orbitals, and the collection of all such configurations forms the *Fock space*. Second quantization provides a systematic way to operate on this space.

2.2.2 Fock Space, Antisymmetry and Slater Determinants

Given a finite set of M orthonormal spin-orbitals $\{\phi_p\}$, we define the fermionic Fock space as

$$\mathcal{F} = \bigoplus_{N=0}^{M} \mathcal{H}^{(N)}, \tag{2.5}$$

namely the direct sum of all N-electron Hilbert spaces. A many-electron basis state is written in *occupation number representation* (or *occupation number vector*) as

$$|n_1\, n_2\, \ldots\, n_M\rangle, \qquad n_p \in \{0, 1\}, \tag{2.6}$$

where n_p specifies whether spin-orbital ϕ_p is occupied.

Because electrons are indistinguishable fermions, the electronic wave function must be antisymmetric under particle exchange. If P_{ij} denotes the permutation operator that exchanges electrons i and j, any physical wave function must satisfy

$$P_{ij}\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_N) = -\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_N). \tag{2.7}$$

A compact way to enforce antisymmetry is the *Slater Determinant (SD)*. Given N spin-orbitals $\{\phi_1, \ldots, \phi_N\}$, the corresponding antisymmetric wave function is

$$\Phi(\mathbf{x}_1, \ldots, \mathbf{x}_N) = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(\mathbf{x}_1) & \cdots & \phi_N(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ \phi_1(\mathbf{x}_N) & \cdots & \phi_N(\mathbf{x}_N) \end{vmatrix}. \tag{2.8}$$

Exchanging two electron coordinates is equivalent to exchanging two rows of the determinant, which changes its sign, thus guaranteeing the correct fermionic transformation. Moreover, if two electrons attempt to occupy the same spin-orbital, two columns become identical and the determinant vanishes. Therefore, the Pauli exclusion principle is automatically enforced.

In second quantization, each SD corresponds uniquely to an occupation number vector in Fock space, as defined in Equation 2.5. For example,

$$|\phi_1\, \phi_2\, \phi_4\rangle \quad \longleftrightarrow \quad |110100\rangle.$$

2.2.3 Creation and Annihilation Operators

Transitions between occupation vectors are described by fermionic *creation* and *annihilation* operators, $\hat{a}_p^\dagger$ and $\hat{a}_p$, which respectively add or remove an electron in orbital p. Their action on a generic occupation vector is defined by

$$\hat{a}_p^\dagger|n_1 \dots n_p \dots n_M\rangle = (-1)^{\sum_{q<p} n_q} (1 - n_p) \, |n_1 \dots 1_p \dots n_M\rangle, \tag{2.9}$$

$$\hat{a}_p|n_1 \dots n_p \dots n_M\rangle = (-1)^{\sum_{q<p} n_q} \, n_p \, |n_1 \dots 0_p \dots n_M\rangle. \tag{2.10}$$

For instance, acting on $|1000\rangle$, the operator $\hat{a}_2^\dagger$ yields $-|1100\rangle$, where the negative sign arises from the antisymmetry of fermions. The operators obey the canonical anti–commutation relations

$$\{\hat{a}_p, \hat{a}_q^\dagger\} = \delta_{pq}, \qquad \{\hat{a}_p, \hat{a}_q\} = \{\hat{a}_p^\dagger, \hat{a}_q^\dagger\} = 0, \tag{2.11}$$

which ensure that no orbital can be doubly occupied and that the overall wave function remains antisymmetric under particle exchange, i.e. they enforce the Pauli exclusion principle at the operator level.

2.2.4 One- and Two-Electron Integrals

To express the electronic Hamiltonian in this operator form, the molecular spin–orbital basis is used to define the *one-electron* and *two-electron* integrals. The one-electron integrals,

$$h_{pq} = \int \phi_p^*(\mathbf{x}) \left(-\tfrac{1}{2}\nabla^2 - \sum_A \frac{Z_A}{r_A} \right) \phi_q(\mathbf{x}) \, d\mathbf{x}, \tag{2.12}$$

represent the combined kinetic energy and electron–nuclear attraction between orbitals p and q. The two-electron Coulomb integrals are defined as

$$(pq|rs) = \iint \frac{\phi_p^*(\mathbf{x}_1)\phi_q(\mathbf{x}_1)\phi_r^*(\mathbf{x}_2)\phi_s(\mathbf{x}_2)}{r_{12}} \, d\mathbf{x}_1 \, d\mathbf{x}_2, \tag{2.13}$$

where $r_{12} = |\mathbf{r}_1 - \mathbf{r}_2|$ is the electron–electron distance. These integrals depend only on the chosen one-electron basis and can be computed once at the beginning of a calculation.

2.2.5 Electronic Hamiltonian in Second Quantization

In terms of the creation and annihilation operators, the Born–Oppenheimer electronic Hamiltonian takes the compact form

$$\hat{H}_e = \sum_{pq} h_{pq} \, \hat{a}_p^\dagger \hat{a}_q + \tfrac{1}{2} \sum_{pqrs} (pq|rs) \, \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_s \hat{a}_r. \tag{2.14}$$

The first term describes all one-body processes (kinetic and electron–nuclear), while the second term accounts for pairwise Coulomb repulsion between electrons. Normal ordering of operators is usually implied, ensuring that creation operators precede annihilation operators.

This representation is completely general and forms the algebraic foundation for most modern electronic structure theories. The Hartree–Fock model corresponds to a single-determinant approximation to the exact wave function, while correlated post–Hartree–Fock methods (CI, coupled cluster, DMRG, etc.) emerge from systematic manipulations of Eq. (2.14).

2.3 Hartree–Fock Theory

The second–quantized Hamiltonian in Equation 2.14 provides the formal starting point for wave function methods. In practice, the simplest non–trivial approximation is obtained by restricting the N–electron wave function to a single Slater determinant. This leads to the *Hartree–Fock (HF)* method, which provides a mean–field description of electron correlation and constitutes the reference state for most post–HF approaches [26, 28].

2.3.1 Variational Principle and Hartree–Fock Ansatz

The electronic ground state energy satisfies the variational bound

$$E_0 \leq \frac{\langle \Phi | \hat{H} | \Phi \rangle}{\langle \Phi | \Phi \rangle}. \tag{2.15}$$

In the HF approximation, the trial wave function is constrained to be a single Slater determinant,

$$|\Phi\rangle = |\phi_1 \phi_2 \ldots \phi_{N_e}\rangle, \tag{2.16}$$

where the orbitals $\{\phi_p\}$ are orthonormal spin–orbitals to be optimized. Applying the variational principle under orbital orthonormality constraints leads to the Hartree–Fock equations [28].

2.3.2 Fock Operator and Roothaan–Hall Equations

Minimization of the energy expectation value yields the one–electron eigenvalue problem

$$\hat{F}\,\phi_p = \varepsilon_p\,\phi_p, \tag{2.17}$$

where $\hat{F}$ is the Fock operator. In a spin–orbital basis,

$$F_{pq} = h_{pq} + \sum_{rs} D_{rs} \Big[(pq|rs) - (pr|qs) \Big], \tag{2.18}$$

with D_{rs} the one–particle density matrix. The two–electron terms can be written as Coulomb and exchange contributions,

$$J_{pq} = \sum_{rs} D_{rs} (pq|rs), \qquad K_{pq} = \sum_{rs} D_{rs} (pr|qs), \tag{2.19}$$

so that $F = h + J - K$. The Coulomb term represents the average electron repulsion, while the exchange term arises purely from antisymmetry and has no classical analogue.

When the orbitals are expanded in a non–orthonormal atomic basis $\{\chi_\mu\}$,

$$\phi_p = \sum_{\mu} C_{\mu p}\,\chi_\mu, \tag{2.20}$$

Eq. (2.17) becomes the generalized eigenvalue problem

$$\mathbf{F}\mathbf{C} = \mathbf{S}\mathbf{C}\boldsymbol{\varepsilon}, \tag{2.21}$$

known as the Roothaan–Hall equation, where $\mathbf{S}$ is the overlap matrix and $\boldsymbol{\varepsilon}$ contains the orbital energies.

2.3.3 Self-Consistent Field Procedure

Since the Fock operator depends on the density matrix and the density depends on the orbitals, Eq. (2.21) must be solved self–consistently. The HF solution is thus obtained through a *Self–Consistent Field (SCF)* cycle:

1. guess the molecular orbitals (or density matrix),
2. build the Fock matrix $\mathbf{F}$,
3. solve Eq. (2.21) for new orbitals,
4. update the density matrix,
5. iterate until convergence.

Upon convergence, the HF determinant defines the best single–determinant approximation to the ground state. Its energy reads

$$E_{\mathrm{HF}} = \sum_{pq} D_{pq}\, h_{pq} + \frac{1}{2}\sum_{pq} D_{pq}\,(J_{pq} - K_{pq}) + V_{NN}. \tag{2.22}$$

HF includes *exchange* effects but neglects *dynamic electron correlation*, motivating post–HF methods.

2.3.4 Basis Sets and Basis–Set Limit

In practical calculations, the spin–orbitals are expanded in finite Gaussian or Slater basis sets $\{\chi_\mu\}$. Increasing the basis size systematically lowers the variational HF energy and approaches the *Complete Basis–Set (CBS)* limit. Correlated methods converge more slowly than HF with respect to the basis size, but HF still provides a controlled reference and a set of canonical orbitals widely used in post–HF theories.

2.4 Beyond Hartree-Fock Methods

Although Hartree–Fock provides the best single–determinant approximation to the ground state, it neglects electron correlation beyond exchange. Post–HF methods aim to recover this correlation by enlarging the space of accessible wave functions or by introducing more accurate electron–electron treatments.

2.4.1 Configuration Interaction

Configuration Interaction (CI) provides the most straightforward systematic improvement beyond the HF approximation. Instead of restricting the wave function to a single SD, CI expands it as a linear combination of many determinants constructed from a fixed orbital basis,

$$|\Psi_{\mathrm{CI}}\rangle = c_0|\Phi_0\rangle + \sum_{ia} c_i^a|\Phi_i^a\rangle + \sum_{ijab} c_{ij}^{ab}|\Phi_{ij}^{ab}\rangle + \cdots , \tag{2.23}$$

where $|\Phi_0\rangle$ is the HF reference, and $|\Phi_i^a\rangle$, $|\Phi_{ij}^{ab}\rangle$, etc. are singly and doubly excited determinants, and $\{c_I\}$ are linear coefficients determined variationally. This expansion allows the explicit mixing of electronic configurations and therefore accounts for electron correlation effects beyond the mean–field picture[28, 29].

The CI coefficients are obtained by solving the secular equation

$$\mathbf{Hc} = E\,\mathbf{c}, \tag{2.24}$$

where the matrix elements $H_{IJ} = \langle\Phi_I|\hat{H}|\Phi_J\rangle$ are computed using the Slater–Condon rules[28]. Because the CI wave function is linear and variational, its energy is an upper bound to the exact electronic energy in the chosen orbital basis.

The *Full Configuration Interaction (FCI)* limit includes all possible determinants generated by distributing N_e electrons among M spin–orbitals. It represents the exact solution of the non–relativistic Born–Oppenheimer electronic problem within that basis. However, the number of determinants grows combinatorially as $\binom{M}{N_e}$, leading to an exponential scaling in both memory and computational effort, which limits FCI to very small systems.

Truncated CI models, such as CIS (singles), CID(doubles), and CISD (singles and doubles), provide a practical compromise by retaining only a subset of excitations relative to the HF reference $|\Phi_0\rangle$. While they capture a significant portion of the correlation energy, they are not size–consistent, meaning that the energy of two non–interacting subsystems is not equal to the sum of their individual energies. This deficiency ultimately motivated the development of multireference and coupled–cluster methods. In particular, when several configurations contribute nearly equally to the ground state, as in bond–breaking or transition–metal complexes, a single HF determinant is an inadequate reference and a *multi–configurational* approach, such as *Complete Active Space Self–Consistent Field (CASSCF)* , becomes necessary.

2.4.2 Complete Active Space Configuration Interaction (CASCI)

A natural multi–reference extension of CI is the *Complete Active Space Configuration-Interaction (CASCI)* method. In CASCI, one selects an *active space* of m orbitals containing n electrons and performs a FCI expansion only within this space,

$$|\Psi_{\mathrm{CASCI}}\rangle = \sum_{I\in\mathrm{CAS}} C_I|\Phi_I\rangle, \tag{2.25}$$

where the determinants $\{|\Phi_I\rangle\}$ span all possible occupations of the chosen active orbitals. This construction retains full multi–reference flexibility while avoiding the exponential scaling of a global

FCI expansion. However, the orbitals are kept *fixed*, typically at the HF level. Consequently, although CASCI describes static correlation within the active space, its accuracy is limited by the choice of orbitals and it lacks full variational optimization. These limitations motivate the following method, the *Complete Active Space Self–Consistent Field (CASSCF)*, in which both the CI coefficients and the orbitals are simultaneously optimized.

2.4.3 Multiconfigurational Methods: CASSCF and MCSCF

The *Complete Active Space Self–Consistent Field (CASSCF)* procedure start from a CASCI wavefunction expanded on the subset of desired active orbital. CASSCF optimization proceeds through a two–step iterative process. For fixed orbitals, the CI step solves the eigenvalue problem within the active space, producing the optimal $\{C_I\}$. Next, for fixed CI coefficients, the orbital step optimizes the orbital rotation parameters to minimize the total energy. These two steps are alternated until self–consistency is reached, resulting in an energy that is stationary with respect to both orbital and configuration variations. The generalized version instead, *Multi-Configuration Self-Consistent Field (MCSCF)* approach expand the solution of the CI as a Truncated CI method, including a lower number of determinants compared to CASCI. A general CASSCF or MCSCF wavefunction Mathematically, CASSCF can be viewed as a generalization of HF in which the single determinant is replaced by a multi–configurational reference. It therefore captures static correlation exactly within the active space, while a portion of the dynamic correlation is recovered through the orbital relaxation from the outer space, between active and inactive orbitals. This recovery of dynamic correlation, however, is not fully complete because CASSCF treats electron correlation only within the active orbitals explicitly, and relies solely on orbital optimization to mimic the correlation with the external space. As a result, most of the dynamic correlation residing in the inactive and virtual orbitals remains unaccounted for. To address this shortcoming, post–CASSCF approaches, most notably multireference perturbation theories such as CASPT2 [30–33] and NEVPT2 [34–36], are employed. These methods provide a systematic, well–defined framework for capturing the missing dynamic correlation and thereby improving both the accuracy of energies and the reliability of potential energy surfaces. The accuracy of CASSCF depends critically on the choice of active orbitals, which should include all near–degenerate or chemically relevant orbitals.

The computational cost of CASSCF grows combinatorially with the size of the active space: a CAS(n, m), containing n active electrons in m active orbitals, scales roughly as $\mathcal{O}(m^n)$. This exponential growth limits conventional CASSCF to about 16–18 active orbitals even on modern hardware. To treat larger spaces, tensor–network approaches such as the *Density Matrix Renormalization Group (DMRG)* have been introduced, providing an efficient polynomial–scaling alternative.

2.4.4 Density Matrix Renormalization Group (DMRG)

The *Density Matrix Renormalization Group (DMRG)* was originally developed in condensed–matter physics [37, 38] and later adapted for quantum chemistry [39, 40]. Its power lies in efficiently representing large multi–configurational wave functions that would be intractable for traditional CI or CASSCF methods.

The FCI expansion scales exponentially with the number of active orbitals, since each orbital may be occupied or empty for each spin. DMRG replaces this explicit determinant expansion by a compact tensor network representation, the *Matrix Product Ansatz (MPS)*, whose number of parameters grows only polynomially with the number of orbitals. The MPS ansatz for a wave function in a basis of L spin–orbitals reads

$$|\Psi\rangle = \sum_{\{n\}} A_1^{n_1} A_2^{n_2} \cdots A_L^{n_L} |n_1 n_2 \dots n_L\rangle, \tag{2.26}$$

where $n_i \in \{0,1\}$ denote orbital occupations and the tensors $\{A_i^{n_i}\}$ carry internal "bond" indices connecting neighboring orbitals. The dimension of these internal indices, denoted D, controls the expressive power of the ansatz: $D = 1$ corresponds to a single determinant, HF determinant, while large D approaches the exact FCI limit.

DMRG exploits the physical observation that ground states of non–critical systems obey an *area law* of entanglement [41]: the entanglement entropy between subsystems grows only with the size of the boundary between them. Consequently, a modest bond dimension D is sufficient to accurately describe states with limited entanglement, as often encountered in molecules. The fact that it focuses computational effort on the most relevant parts of Hilbert space, truncating the rest based on density–matrix eigenvalues, becomes one of its key advantage.

DMRG determines the optimal MPS variationally by minimizing the expectation value $\langle\Psi|\hat{H}|\Psi\rangle$. This is accomplished through an iterative sweeping procedure: two neighboring tensors are optimized while keeping the rest fixed, the environment tensors are updated, and the sweep continues back and forth until convergence. The resulting energy is strictly variational and systematically improvable by increasing D.

For a given bond dimension D and number of orbitals L, the DMRG cost scales approximately as $\mathcal{O}(LD^3)$, much more favorable than the exponential scaling of FCI or CASSCF[3]. Typical chemical accuracy is achieved with D in the range 500–2000, depending on the degree of entanglement, and size of the molecular space.

In quantum chemistry, DMRG is commonly employed as an active–space solver within the SCF framework, leading to the *Density Matrix Renormalization Group Self–Consistent Field (DMRG-SCF)* method. Here, the DMRG algorithm replaces the FCI solver in the CASSCF procedure, allowing active spaces of 30–50 orbitals to be treated with high accuracy. The orbital optimization and DMRG sweeps are iterated to achieve mutual self–consistency, in analogy with the conventional CASSCF macro–cycle.

For a detailed explanation of the DMRG theory, please consult the reference [41].

2.4.5 Other Methods: MP2 and CC

Electron correlation can also be treated without explicit CI expansions. *Second—order Møller–Plesset Perturbation Theory (MP2)* provides a cheap correction to the HF energy and is based on Rayleigh–Schrödinger perturbation theory,

$$E_{\mathrm{MP2}} = \sum_{ijab} \frac{(ij|ab)\big[(ij|ab) - (ib|aj)\big]}{\varepsilon_i + \varepsilon_j - \varepsilon_a - \varepsilon_b}. \tag{2.27}$$

---

[3] For a more accurate scaling analysis check the Appendix 4.3.2

Coupled-Cluster (CC) methods use an exponential ansatz

$$|\Psi_{\text{CC}}\rangle = e^{\hat{T}}|\Phi_0\rangle, \tag{2.28}$$

with $\hat{T} = T_1 + T_2 + \cdots$ the cluster excitation operator [42, 43]. Truncations such as *Coupled–Cluster Singles–Doubles (CCSD)* or *Coupled–Cluster Singles–Doubles Pertubative Triples (CCSD(T))* [44] are size–consistent and the latter is widely considered the “gold standard” for weakly correlated systems.

2.4.6 Density Functional theory

Density Functional Theory (DFT) provides an alternative to wave function methods by expressing the ground–state energy as a functional of the electron density $\rho(\mathbf{r})$. The Hohenberg–Kohn theorems establish that the ground–state density uniquely determines all properties of a many–electron system and the exact ground–state energy is obtained by minimizing an energy functional of the density [45]. In practice, the *Kohn–Sham (KS)* formulation maps the interacting problem onto a fictitious system of non–interacting electrons that reproduces the same density, greatly simplifying the numerics [46].

The KS energy functional is written as

$$E[\rho] = T_s[\rho] + V_{ne}[\rho] + J[\rho] + E_{\text{XC}}[\rho], \tag{2.29}$$

where $T_s[\rho]$ is the kinetic energy of the KS reference determinant, V_{ne} is the nuclear attraction, J is the classical Coulomb term, and $E_{\text{XC}}[\rho]$ is the *Exchange–Correlation (XC)* functional, which must be approximated. Increasingly accurate approximations form the so–called Jacob’s ladder of DFT, ranging from the *Local Density Approximation (LDA)* and *Generalized Gradient Approximations (GGA)* to meta–GGA, hybrid, and double–hybrid functionals [47].

DFT offers a favorable accuracy–to–cost ratio, making it widely used for molecular structure, thermochemistry, and properties. For this reason, in Chapter 4 we will employ selected KS functionals as a reference for benchmarking correlated methods, including those based on CASSCF and DMRG.

Chapter 3

Tackling Quantum Chemistry problem with Quantum Computing

The two previous chapters have outlined the theoretical foundations of quantum computing and quantum chemistry, trying to briefly underlay principles and computational methods. First, quantum mechanics and the quantum circuit paradigm provide the mathematical and algorithmic framework for exploiting qubits as computational resources. Second, quantum chemistry describes the complexity of molecular systems, where diverse approximations—ranging from Hartree–Fock theory to advanced post-Hartree–Fock methods—have been devised to overcome the exponential scaling of the electronic structure problem.

Despite the maturity of classical quantum chemistry, its limitations become evident when dealing with strongly correlated systems or large active spaces. Quantum computing, though still in its early stages, offers a promising path to overcome some of these bottlenecks by embedding quantum subroutines within standard workflows. The challenge, therefore, lies not in replacing established methods altogether, but in designing hybrid pipelines where quantum processors complement classical algorithms in a targeted and efficient manner.

This chapter, introducing the topic of the thesis, focuses on the central idea of tackling quantum chemistry problems with quantum computing, serving as the bridge between theoretical background and the development of practical pipelines. First, the mapping of molecular problems to qubits is introduced, clarifying how the language of second quantization translates naturally into the quantum computing framework. Next, we put more attention to the main resolution algorithm used in our approaches, i.e *Variational Quantum Eigensolver (VQE)*, which has intermittently been considered a viable near–term approaches for estimating molecular ground states on Noisy Intermediate–Scale Quantum (NISQ) devices.

By establishing these connections, this chapter lays the groundwork for the second part of this thesis, which turns toward specific developments, such as active space selection, geometry optimization, single–point calculations, and multi–configurational self-consistent field methods, each tackled within a HQC paradigm.

The structure of the chapter is defined as follow: In Section 3.1, we introduce the tools to map the representation of electrons and molecular orbital form bit to the qubit space. We define what a*Fermion–to–Qubit (F2Q)* is in Section 3.1.1. Then, in Section 3.1.2, we report how, dealing with a restricted number of orbital, we can reduce the full molecular space representation into the active sites only. After this, we introduce out main solver tool, the Variational Quantum Eigensolver (VQE) algorithm in Section 3.2 and how we can define suitable circuit, i.e. the wavefunction of the problem, in the qubit space, with Section 3.3. Finally, in Section 3.4 we define the main steps an average quantum chemistry pipeline is composed of.

3.1 From Bits to Qubits

The previous chapter presented the basis and the main computational tools to tackle and face standard quantum chemistry calculations (classically) at different level of theory. Even the simplest procedure, as evaluating the energy of the $|HF\rangle$ state on a quantum computer, require different steps. Therefore, if the final objective, as we are going to show in the whole manuscript, is to perform a calculation on a quantum computer, so by exploiting a quantum algorithm, the first essential step is to performs a mapping between the Fermionic space to the qubit space by means of a Fermion-to-Qubit mapping method.

3.1.1 Fermion–to–Qubit mapping

As we seen in the previous section, a generic quantum chemistry wavefunction (or fermionic wavefunction) lives in the Fock space, $\mathcal{F}$, as defined in Equation 2.5. In this space, the anti-symmetrization principle, explained in Equation 2.7, and the commutation rules for Fermionic operators, Equations 2.11, strongly constrain the symmetries. Moving to a quantum computer, these symmetries need to be conserved and properly translated from the fermionic space to the qubit space. Therefore, if the final objective is to solve a task on a quantum computer, we need to be able to maintain the same symmetric structure we had in the fermionic space, in the qubit space. Contrary to Fermions, qubits are distinguishable bosons. In order to properly manipulate fermionic objects in the qubit space, *Fermion–to–Qubit (F2Q)* mapping is mandatory.

A *Fermion–to–Qubit (F2Q)* mapping is a transformation that translates the algebra of fermionic creation and annihilation operators into the algebra of qubit operators in such a way that the essential algebraic properties of the fermionic operators are preserved. Consider a system of N fermionic modes, described by the annihilation and creation operators

$$\{\hat{a}_p\}, \quad \{\hat{a}_p^\dagger\}, \quad p \in \{0, 1, \ldots, N-1\}, \tag{3.1}$$

satisfying the canonical anti-commutation relations

$$\{\hat{a}_p, \hat{a}_q\} = 0, \quad \{\hat{a}_p^\dagger, \hat{a}_q^\dagger\} = 0, \quad \{\hat{a}_p, \hat{a}_q^\dagger\} = \delta_{pq}. \tag{3.2}$$

These operators act on the fermionic Fock space $\mathcal{F}_N$, $\mathcal{F}(\mathbb{C}^N)$, and each basis state are defined as a occupation number vector

$$|n_0, n_1, \ldots, n_{N-1}\rangle, \quad n_p \in \{0, 1\}, \tag{3.3}$$

obtained by apply the creation operators on the fermionic vacuum state $|\mathbf{0}_f\rangle$.

Definition 1. ***Fermion–to–Qubit (F2Q) mapping:*** *A F2Q mapping is defined as a linear algebra homomorphism*

$$\mathcal{M} : \mathcal{A}_{fermion} \to \mathcal{A}_{qubit}, \tag{3.4}$$

where $\mathcal{A}_f$ *is the algebra generated by the fermionic operators acting on* $\mathcal{F}_N$*, and* $\mathcal{A}_q$ *is the algebra generated by tensor products of Pauli matrices acting on an* N*-qubit Hilbert space* $(\mathbb{C}^2)^{\otimes N}$*. Both spaces are* $\mathbb{C}^{2N}$*.*

The purpose of $\mathcal{M}$ is to provide explicit representations

$$\mathcal{M}(\hat{a}_p), \quad \mathcal{M}(\hat{a}_p^\dagger), \quad \mathcal{M}(\hat{n}_p = \hat{a}_p^\dagger \hat{a}_p) \tag{3.5}$$

as operators on qubits that reproduce the fermionic anti-commutation rules. By mean of a mapping, all physical observables and dynamics expressed in terms of fermionic operators are going to be conserved also in the qubit space. In general, $\mathcal{M}$ must satisfy the following key properties:

- **Faithfulness:** $\mathcal{M}$ is injective on the fermionic algebra so that no physical information is lost.
- **Preservation of algebraic relations:** $\mathcal{M}(\hat{a}_p), \mathcal{M}(\hat{a}_q^\dagger)$ satisfy the canonical anti-commutation rules.
- **Compatibility with the Fock space basis**: $\mathcal{M}(\hat{n}_p)$ acts diagonally in the computational basis $|0\rangle, |1\rangle$.

Thus, a F2Q mapping is a structure-preserving homomorphism that link the fermionic operator formalism of quantum many-body physics with the qubit-based operator formalism of quantum computing. Concretely, a mapping $\mathcal{M}$ assigns to each fermionic mode p a set of Pauli operators on a register of N qubits. Since qubits obey tensor-product commutation relations rather than fermionic anti-commutation, the map must encode the fermionc anti-simmetrization rule explicitly, often through nonlocal *parity strings*.

Different method exist to convert Fermionic operators to the bosonic space of the qubits, above all the most popular are Jordan-Wigner [48], Parity [49], Bravy-Kitaev [50], but also custom and ad-hoc mappings can be defined [51, 52].

Jordan-Wigner Mapping

We briefly explain the main mapping method used throughout the whole thesis, the Jordan–Wigner mapping [48]. The JW mapping allows for a more intuitive transformation, maintaining a almost 1–to–1 correspondence between Slater determinants, states in the computational basis in the qubit space, and creation/annihilation operators. Each active spin orbital is going to be mapped to one qubit and the encoding represent the occupation of that specific spin orbital, as for Slater Determinants. [1]The *vacuum* state is going to be defined in the same way, i.e.

$$|\mathbf{0}\rangle_f = |0\dots00\rangle_f \rightarrow |\mathbf{0}\rangle_q = |0\dots00\rangle_q \,. \tag{3.6}$$

The Hartree-Fock state of an H_2 two orbital system, $|HF\rangle = |0101\rangle_f$, is going to be mapped to the computational basis vector $|HF\rangle_q = |0101\rangle_q$. An excited determinant is going to follow the same

---

[1] The ordering of Fermionic modes and qubits is not fixed, changing from software to software and from publication to publication, but generally the two main orderings are *splitted*, $|0_\beta \dots 0_\beta 0_\alpha \dots 0_\alpha\rangle$, thus splitting alpha and beta spin orbitals, or *interleaved*, $|0_\beta 0_\alpha \dots 0_\beta 0_\alpha\rangle$. For this explanation, we stick to the first one.

structure, e.g. $|1001\rangle_f \to |1001\rangle_q$. The transformation applies to the fermionic operators. These operators, as already explained in the previous section, allows for creating (destroying) a particle associated with a specific spin orbital. In particular,

$$\mathcal{M}_{\mathrm{JW}}(\hat{a}_i^\dagger) = \hat{Q}_i^\dagger \prod_{j=0}^{i-1} Z_j \quad \text{and} \quad \mathcal{M}_{\mathrm{JW}}(\hat{a}_p) = \hat{Q}_i \prod_{j=0}^{i-1} Z_j, \tag{3.7}$$

where $\hat{Q}_i^\dagger = \frac{1}{2}(X_i - iY_i)$, $\hat{Q}_i = \frac{1}{2}(X_i + iY_i)$, and i iterates over the set of active spin orbitals. The $\hat{Q}$ operators can be associated to creation and annihilation operators in the qubit space, respectively. As well as the fermionic creation and annihilation operators act to modify the occupancy of the relative spin-orbital, the qubit operators act to change the state of the qubit related to the specific spin orbital. The full correspondence is finalized with a series of Pauli Z matrices that allow to compute the parity of the state and account for the fermionic anti-commutation between a and $a^\dagger$. Using Equations 3.7 in Equation 2.2, the effective Hamiltonian in the qubit space, $\hat{H}_q$, can be written as

$$\hat{H}_q = \sum_i^L c_i \hat{P}_i = \sum_i^L c_i \bigotimes_{j=0}^{N_{MO}} \hat{\sigma}_j^i, \tag{3.8}$$

where each $\hat{\sigma}_j \in \{X, Y, Z, I\}$ is a Pauli matrix acting on the j-th qubit, c_i is a coefficient relative to the i-th Pauli string $\hat{P}_i$, N_{MO} is the number of qubit or number of spin-orbitals, and L, proportional to N_{MO}^4, is the number of Pauli string composing the qubit Hamiltonian. Each Pauli string $\hat{P}_i$ is obtained by the multiplication of each term in the original Hamiltonian, $\hat{H}_{eff}$, once they get mapped in the qubit space.

We now know how to map a generic fermionic string (e.g. $\hat{a}_i^\dagger \hat{a}_j$ or $\hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l$) into the qubit-space equivalent. The missing link is how to relate a generic qubit Hamiltonian to a electronic structure second quantized Hamiltonian. This procedure is done through the coefficients c_i associated with each Pauli string $\hat{P}_i$. Each coefficient is proportional to the respective one- and two-body integrals defined from a classical quantum chemistry method, e.g.HF, CASSCF, DFT, and so on. Therefore, without entering in the merit of any particular method, suppose we now have a set of one- and two-body integrals (defined in second quantization) namely h_{pq} and $\Gamma_{pqrs} = (pq|rs)$, defined accordingly to Equation 2.12 and 2.13 in Section 2.2. We begin with the electronic Hamiltonian in second quantization, expressed over the full molecular orbital basis:

$$\hat{H} = \sum_{pq} h_{pq}\, a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} (pq|rs)\, a_p^\dagger a_q^\dagger a_s a_r \tag{3.9}$$

where h_{pq} are the one-electron integrals, $(pq|rs)$ the two-electron integrals (in chemist's notation), $a_p^\dagger, a_q$ fermionic creation and annihilation operators, as defined in Section 2.2.

3.1.2 Defining an Effective Hamiltonian

As we already introduced in the previous chapter, it is possible to focus on a subset of orbitals on which the FCI is solved, and perform what is called a *Complete Active Space (CAS)* method on this subset. Any CAS method, either CASCI, CASSCF, DMRG, or post-SCF method[2], require the exact

---

[2] For example CASPT2, or higher order CASPTn, as well as NEVPT

or trucated solution in the active space, and therefore, a scheme to reduce from the whole system to the region of interest is necessary. The active space procedure, which is going to be approached in the dedicated section, Chapter 4, allows to define a subset of molecular orbital, $\mathcal{A} \subset \mathcal{M}$, from the total molecular orbital space, $\mathcal{M}$. Thus, the MO basis can be partitioned into *core orbitals*, $\mathcal{C}$, orbitals that are going to be always doubly occupied, the *active orbitals*, $\mathcal{A}$, which are going to be treated more accurately either classically or by a quantum computing method, and the *virtual orbital*, $\mathcal{V}$, which are going to be always unoccupied. The core orbitals indexed by $u, v \in \mathcal{C}$, active orbitals indexed by $i, j, k, l \in \mathcal{A}$, and virtual orbitals indexed by $a, b, c, d \in \mathcal{V}$.

Following this subdivision, the effective one-body and two-body integrals can be defined, namely h_{pq}^{eff} and Γ_{pqrs}^{eff}, as well as the new core energy, E_{core}. Restricting to the active space, the *effective* Hamiltonian, H_{eff}, is defined as

$$\hat{H}_{eff} = E_{\text{core}} + \sum_{ij \in \mathcal{A}} h_{ij}^{\text{eff}} \, a_i^\dagger a_j + \frac{1}{2} \sum_{ijkl \in \mathcal{A}} \Gamma_{ijkl}^{eff} \, a_i^\dagger a_j^\dagger a_l a_k. \tag{3.10}$$

In order to obtain the new effective component, the following contraction are required. Assuming the core orbitals are always doubly occupied, their contribution to the energy becomes a constant scalar term

$$E_{\text{core}} = \sum_{u \in \mathcal{C}} h_{uu} + \frac{1}{2} \sum_{u,v \in \mathcal{C}} \left[2(uu|vv) - (uv|vu) \right]. \tag{3.11}$$

The one-electron integrals over the active space are modified due to the presence of the frozen core, leading to the effective one-body integrals as

$$h_{ij}^{\text{eff}} = h_{ij} + \sum_{u \in \mathcal{C}} \left[2(ij|uu) - (iu|uj) \right], \quad i, j \in \mathcal{A}, \tag{3.12}$$

so by reducing two-body interaction to one-body interaction restricted to the active space. The effective-two body tensors instead is just obtained as $\Gamma_{ijkl}^{eff} = (ij|kl)$, where the indices $i, j, k, l \in \mathcal{A}$, leaving outside all the indices non included in $\mathcal{A}$. Without building the effective Hamiltonian, the effective integrals can be now be used by a standard quantum chemistry pipeline in order to apply one of the methods listed in the main manuscript (e.g. CASCI, CASSCF, orDMRG).

3.2 Variational Quantum Eigensolver (VQE)

Variational Quantum Eigensolver (VQE) utilizes a HQC approach to estimate ground state energies of molecular systems by optimizing a parameterized quantum ansatz through iterative energy minimization. Since about a decade the method has been largely studied (and criticized) for its employment into near-term applications designed for quantum chemistry [53–59]. The VQE is generally used in place of a classical CI solver a-la CASCI for implementing Hybrid Quantum–Classical (HQC) pipelines for which, ideally, the resolution of the classical part is unfeasible. Its simplicity instead hides scalability challenges because of the growth of expressivity of the trial wavefunction with its depth, the flatness of the corresponding energy landscape, the intensive exchange of data between quantum ⟷ classical devices, and noisy quantum devices. The VQE went through ups and downs of popularity[3], nevertheless, the algorithm will be used as main resolution tool through-

---

[3] Recently, started to shine again thanks to Sample based algorithms.

out the whole thesis. In this section, we are going to present briefly the VQE algorithm in all its components. A exhaustive review on the VQE method can be found in the references [59].

3.2.1 Hybrid Variational Algorithm

The VQE algorithm bases is resolution on the *Rayleigh-Ritz Variational Principle* [60]. The principle states that given an Hamiltonian H and the associated groundstate energy E_0, for any trial wavefunction $|\psi\rangle$, the energy obtained by the trial state

$$E_0 \leq \langle\Psi| H |\Psi\rangle . \tag{3.13}$$

Therefore, the energy of any wavefunction is strictly lower bounded by the ground state energy E_0 of the system associated to H. The VQE method, exploiting the principle explained in Eq. 3.13, is meant to variationally optimize a parametrized wavefunction until it reaches, ideally, the optimal wavefunction i.e. the best approximation to the groundstate. The trial wavefunction is a *PQC*, denoted with $|\psi(\theta)\rangle$, where θ is the set of parameters that are optimized by the VQE procedure. If we now substitute the newly defined PQC in Eq. 3.13, we obtain

$$E_0 \leq E(\theta) = \frac{\langle\psi(\theta)| H |\psi(\theta)\rangle}{\langle\psi(\theta)|\psi(\theta)\rangle}. \tag{3.14}$$

Going into the detail of Eq.3.14, without loss of generality, we can define a PQC as

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta}) |0\rangle , \tag{3.15}$$

i.e. a parametrized unitary evolution applied to the qubit-vacum state, or in general, to a reference state (for quantum chemistry application, the starting guess is always the HF state). The parameters vector $\boldsymbol{\theta}$ is composed by a set of j real-valued parameter, where j is associated to the number of parametrized unitaries $U(\theta_j)$ that are used to composed the full $U(\boldsymbol{\theta})$. In particular, the PQC can be written as

$$U(\boldsymbol{\theta}) = \prod_{j=N-1}^{0} U(\theta_j). \tag{3.16}$$

After introducing also the ideal form of the trial wavefunction, the VQE algorithm can be defined as the task of finding the set of parameters $\boldsymbol{\theta^*}$ that minimizes the expectation value of a given operator, H, by iteratively improving the wavefunction, in order to find the best approximation of the ground state E_0, i.e.

$$E_0 \sim E(\boldsymbol{\theta}^*) = \min_{\boldsymbol{\theta}\in\mathbb{R}^j} \frac{\langle\psi(\boldsymbol{\theta})| H |\psi(\boldsymbol{\theta})\rangle}{\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle}. \tag{3.17}$$

If the operator we are trying to solve is non-degenerate, and the wavefunction is sufficiently expressive, then

$$|\psi_0\rangle = U(\boldsymbol{\theta}^*) |0\rangle \quad \text{with} \quad \boldsymbol{\theta}^* = \text{argmin}_{\boldsymbol{\theta}\in\mathbb{R}^j} \frac{\langle\psi(\boldsymbol{\theta})| H |\psi(\boldsymbol{\theta})\rangle}{\langle\psi(\boldsymbol{\theta})|\psi(\boldsymbol{\theta})\rangle}, \tag{3.18}$$

where $|\psi_0\rangle$ is the true groundstate of H.

3.2.2 Iterative procedure

As already mentioned, the VQE procedure is an iterative method that iteration by iteration improve the representation of the wavefunction to lower the energy, with the final objective of approximating as best as possible, the true groundstate of the Hamiltonian H. To be able to run correclty, the operator H need to be represented in the qubit space, and this procedure, as been already explained in the previous section. Therefore, we suppose to have already mapped our problem in the qubit space, and the Hamiltonian now is defined as The VQE algorithm is a Hybrid algorithm, composed by a quantum and a classical component:

$$\hat{H} = \sum_i c_i \hat{P}_i, \tag{3.19}$$

where $\hat{P}_i$ is a tensor product of $\{X, Y, Z, I\}$ Pauli matrices, and c_i is the associated coefficient in the linear combination. At this point, the VQE procedure proceed as follows:

1. **PQC preparation [Quantum Computer]** First, the shape and the type of the PQC need to be chosen. There are different type of wavefunctions (also called *Ansatz*) and the most common are going to be listed below. In general, on the quantum computer the state for a given set of parameters $\boldsymbol{\theta}$ is generated, i.e. $|\psi(\boldsymbol{\theta}\rangle$ is generated on the device.

2. **Operator measurement [Quantum Computer]** Now, for a fixed set of parameters $\boldsymbol{\theta}$, the operator is measured. As already introduced previously, the operator to be measure on a quantum device has been converted into a linear combination of Pauli strings. For each Pauli string, a measurement is performed

   $$\langle \hat{P}_i \rangle = \langle \psi(\boldsymbol{\theta})| \hat{P}_i |\psi(\boldsymbol{\theta})\rangle \tag{3.20}$$

   and exploiting the linearity of expectation values, the total energy is reconstructed by summing the partial measures into

   $$\begin{aligned} E(\boldsymbol{\theta}) &= \sum_i c_i \langle \hat{P}_i \rangle \\ &= \sum_i c_i \langle \psi(\boldsymbol{\theta})| \hat{P}_i |\psi(\boldsymbol{\theta})\rangle \\ &= \langle \psi(\boldsymbol{\theta})| \sum_i c_i \hat{P}_i |\psi(\boldsymbol{\theta})\rangle \\ &= \langle \psi(\boldsymbol{\theta})| H |\psi(\boldsymbol{\theta})\rangle . \end{aligned} \tag{3.21}$$

   Therefore, the process of measuring the energy of H is done by repeated measurement of each Pauli up to the desired precision, namely ϵ. The number of times, or *shots*, required to reach the precision ϵ[4] scales as $\mathcal{O}(\frac{1}{\epsilon^2})$, and this number is for each Pauli composing the operator. A generic quantum chemistry operator, contains a number of terms that scales as $\mathcal{O}(N^4)$. Measuring the expectation values of an operator within the desired accuracy is therefore one of the most prominent bottleneck in quantum computing application for quantum chemistry. In general, the task of improving the measurement outcomes and reduce its cost is an active

---

[4] For quantum chemistry, the desired accuracy ϵ is 1.6mHa, i.e. 1e-3 standard units.

research field. In particular, most used strategies are exploiting commutations rules for Pauli string and measuring commuting sets [61, 62], or using different measurement basis from the standard Pauli basis, as the Informational Complete Positive Operator Value Matrice (IC-POVMS) [63]

3. **Energy reconstruction and Optimization step [Classical computer]** After the evaluation of each Pauli string $\hat{P}_i$ is obtained, the whole energy, according to Eq. 3.21, is reconstructed on the classical side. Now, the energy is used as input as *cost function* evaluation for a classical optimizer and a single update of the parameter vector $\boldsymbol{\theta}$ is performed. The performance of the classical optimization algorithm are going to impact the overall VQE convergence. Different options exists such as Broyden–Fletcher–Goldfarb–Shanno (BFGS) [64], COBYLA [65], or Simultaneous Perturbation Stochastic Approximation (SPSA) [66].

4. At this point, the procedure is repeated by plugging the new parameters $\boldsymbol{\theta}'$, obtained from a single classical optimizer step, in step 1, until convergence criteria are met.

3.3 Ansatz design

An *Ansatz*, or *trial-wavefunction*, defines how a quantum state can be represented on a quantum computer. The shape and structure of the ansatz coincide with the representation of the wavefunction on the QPU.

Designing a good ansatz means representing an expressive wavefunction that can be able to fully describe the state that we want to encode on the quantum computer. In general, a well designed ansatz need to take into account that its expressivity is not free, and the more parameters the harder is to be optimized. Regardless of the approach used to construct the ansatz, increasing complexity in PQC results in longer and deeper circuits, which not only heightens the risk of error accumulation but also contributes to the emergence of Barren Plateaus [67], an issue characterized by an exponentially flat optimization landscape. Ansatz design is one on the most active field of research for VQE-like algorithms, and different constructions techniques have been defined, but we can categorize them in two major classes: *Hardware Efficient Ansätzes (HEA)* and *Physically Inspired*.

3.3.1 Hardware Efficient Ansätzes (HEA)

This first approach begins with wavefunctions directly constructed to leverage the characteristics of quantum hardware. This empirical approach, known as the Heuristic Ansatz [68–72], comprises repetitions of *Rotations blocks*, i.e. sequence of 1-qubit rotations gate, alternated with *Entangling block*, i.e. 2-qubits entangling gates. The most common type of HEA is the *ladder ansatz*. It is composed by a combination of R_x and R_z rotations, in its original definition [68], and CNOTs gate in a top-down fashion from the first qubit to the last, repreated for d times. An example of ladder-type ansatz is shown in Fig.3.1 It is designed without relying on information about the physical system, focusing solely on exploiting the quantum hardware's capabilities, While the Heuristic Ansatz better utilizes quantum hardware, it comes at the cost of losing the physical meaning associated with the variational ansatz. Generally, compared to the second family of circuits, the circuits appear to be more compact and shallower, providing a potential avenue to address scalability concerns. In this

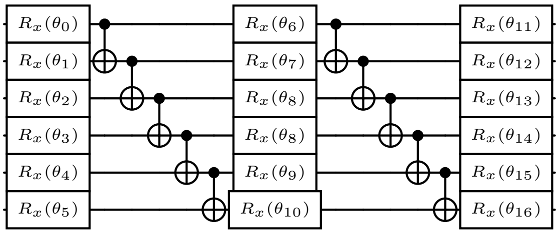

Figure 3.1: Generic depth 2 heuristic ansatz with entangling map in a ladder configuration, where $\{\theta_i\}$ are the variational parameters.

manuscript, in Chapter 5, a briding ansatz between the two types is defined, building an HEA based on information obtained from the underying physics of the system under investigation.

3.3.2 Physical or Chemically Inspired Ansätzes

The main task of this kind of ansatz is translating the same structure of the solution or Hamiltonian into the wavefunction. For quantum chemistry, the ansatz is generally composed by a sequence of fermionic excitations gates, carefully mapped into the qubit space, and applied according to some specific criteria to the reference HF state [73].

Unitary Coupled–Cluster ansatz

The most famous example is the *Unitary Coupled–Cluster (UCC)* method, which provides a unitary implementation of the classical Coupled-Cluster method [74–78]. The UCC ansatz is composed as

$$|\psi(\boldsymbol{\theta})\rangle = e^{T(\boldsymbol{\theta})-T^{\dagger}(\boldsymbol{\theta})}\,|HF\rangle \tag{3.22}$$

, where $T(\boldsymbol{\theta})$ is defined as a approximated expansion of the exact many-body function. In particualar, the expansion is generally stopped at the single and double excitation as follows:

$$T(\boldsymbol{\theta}) = \sum_{i\in\{1,2\}} \hat{T}_i \tag{3.23}$$

with

$$\hat{T}_1 = \sum_i \sum_a \theta_i^a \hat{X}_i^a \quad \text{and} \quad \hat{T}_2 = \sum_{i,j} \sum_{a,b} \theta_{ij}^{ab} X_{ij}^{ab} \tag{3.24}$$

where, i, j are occupied orbitals, a, b unoccupied orbitals, and X_i^a, X_{ij}^{ab} are excitation operators defined as product of fermionic operators, respectively, $a_a^\dagger a_i$ and $a_b^\dagger a_a^\dagger a_i a_j$. Each θ is a variational parameter. The operator $T - T^\dagger$ is anti-Hermitian, ensuring that the exponential remains unitary. This type of trucated UCC is named *Unitary Coupled Cluster Singles and Doubles (UCCSD)* [77, 79]. In general, the scaling of Unitary Coupled–Cluster (UCC)-like ansatz in terms of number of gate is not favorable to make them executable on real hardware, due to the amount of single and two-qubits gates required to encore each terms of the expansion above. For a system with N_{occ} occupied orbitals and N_{virt} virtual orbitals, the total number of variational parameters scales as

$$N_{\text{params}} = N_{\text{occ}}\, N_{\text{virt}} + \frac{1}{4}\, N_{\text{occ}}^2\, N_{\text{virt}}^2, \tag{3.25}$$

showing that the ansatz scales almost quartic in the number of spin orbitals in the system [80]. Each excitation term using a F2Q mapping becomes a unitary of the form

$$U_k(\theta_k) = e^{i\theta_k P_k}, \tag{3.26}$$

where P_k is a product of Pauli operators. The full ansatz is a sequence of such unitaries acting on the reference state,

$$U(\boldsymbol{\theta}) = \prod_k e^{i\theta_k P_k} |\Phi_0\rangle. \tag{3.27}$$

The UCCSD ansatz directly encodes physically meaningful single and double electron excitations, leading to a circuit depth that scales with the number of excitation operators rather than fixed hardware-efficient layers

Adaptive Ansätzes

Another physically inspired circuit that is widely used is the *Adaptive Derivative-Assembled Pseudo-Trotter ansatz Variational Quantum Eigensolver (ADAPT-VQE)* [81] is an iterative, adaptive extension of VQE designed to compute ground-state energies of many-body Hamiltonians with shallower, problem-tailored circuits. Like VQE, it minimizes the expectation value of a Hamiltonian but instead of choosing a fixed variational ansatz, ADAPT-VQE grows the ansatz operator-by-operator based, at least in its first implementation, on the gradient of the energy. The adaptive approach builds the parametrized circuit iteratively by selecting and appending unitary operators from a predefined operator pool. Starting from a reference state $|\psi_0\rangle$, the ansatz after m adaptive steps is defined as

$$|\psi^{(m)}(\boldsymbol{\theta})\rangle = \prod_{k=1}^{m} e^{\theta_k \hat{P}_k} |\psi_0\rangle, \tag{3.28}$$

where each $\hat{P}_k$ is an anti-Hermitian generator taken from the pool $\mathcal{P} = \{\hat{P}_j\}$. At a give iteration m, the commutator of the Hamiltonian, $\hat{H}$, and each the elements contained in the pool $\{\hat{P}_i\}$ are computed as

$$g_i = \left\langle \psi^{(m)} \middle| [\hat{H}, \hat{P}_i] \middle| \psi^{(m)} \right\rangle \tag{3.29}$$

on the reference wavefunction $\left|\psi^{(m)}\right\rangle$. This is equivalent to add the candidate operator $\hat{P}_i$ to the ansatz, set the variational parameter to zero and then compute the energy gradient w.r.t. to it. The operator that produces the largest gradient norm is the added to the circuit with the corresponding unitary, $e^{\theta_{m+1}\hat{P}_{j*}}$. This adaptive growth continues until all gradient magnitudes fall below a chosen threshold, indicating that no operator in the pool can further lower the energy. Multiple variations of the ADAPT-VQE algorithm have been implemented in recent years [72, 82–88]. This type of ansatz is going to be used in the extension of the WAHTOR algorithm to its adaptive and self-consistent version, explained in Chapter 6.

Physical inspired

For strongly correlated lattice models like the Heisenberg Hamiltonian, system treated lated in Chapter 5, specific ansätze have been developed to efficiently tackle the unique challenges presented by these systems. An example is the Hamiltonian Variational Ansatz (HVA) which leverages the compact structure of lattice model Hamiltonians, using fermionic operators to exploit low Pauli weight encodings [89–91]. This approach often results in shallower circuits compared to other ansatze like the *Unitary Coupled Cluster Singles and Doubles (UCCSD)*

In the thesis, the part of quantum computing resolution will be mainly tackled using VQE and HEA-like ansatz, limited to a number of qubit manageable with standard noiseless statevector simulators.

3.4 Hybrid Pipeline to Tackle Quantum chemistry Problem

We have now all the pieces to introduce the main topic of the whole thesis.

Electronic structure theory can be viewed as a workflow, in which each stage builds upon the quality of the previous one. The first task is to identify a chemically meaningful subset of orbitals: the active space. This step controls the balance between physical relevance and computational feasibility. Once the active space is defined, accurate single-point energies must be computed, ideally exploiting techniques that accelerate convergence and reduce the cost of state preparation. However, a description restricted to the active space captures only static correlation: dynamical correlation outside it must be recovered to achieve predictive power. Finally, the resulting electronic structure model acts as the engine for exploring the potential energy surface and optimizing molecular geometry[5]. Therefore the necessity of having a whole framework able to deal with all this type of calculations is high.

Thanks to the collaborations that took place during the PhD period, the whole standard pipeline has been faced to solve varying quantum chemistry application, ranging from standard benchmarking molecules, to industrial chemistry and biomedical chemistry. As already anticipated in the preface, for the whole length of the thesis, we are going to adopt a "workflow-oriented" perspective and investigate how quantum routines may be embedded into each of its stages. Therefore, rather than treating quantum algorithms in isolation, we follow the natural order of the electronic structure problem and integrate quantum components where they have the greatest impact, progressively building towards end-to-end hybrid pipelines. In particular, we resume the content of the thesis in

- **Active–Space Selection for Photodynamic–therapy**: Defining an active-space is generally a tedious task which can hardly be automatized due the different nature and mechanism to be described in quantum chemistry. The procedure becomes even more demanding in quantum computation when the sizes of the system or subsystem analyzed can becomes larger than the one tractable with standard CAS methods. Active–space design is valuable for tackling heavy–atom photosensitizers where spin–orbit coupling and near-degenerate states are key for Photodynamic processes. In these complexes, only a handful of valence orbitals (like metal d-orbitals and ligand π/π^* orbitals) dominate the photochemistry. Identifying these orbitals as the active space allows multireference calculations (like CASSCF) to describe excited–state charge transfer and intersystem crossing, crucial for light–driven generation of singlet oxygen in PDT [92]. In Chapter 4, we are going to introduce our active space selection method, *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS).*

---

[5] This step can also be seen as the first one and therefore its position is interchangeable.

- **Single–point resolution with shallow circuits for standard molecules**: Single-point calculations let you probe how different levels of theory (HF, MP2, CCSD, etc.) capture electron correlation for fixed nuclear geometries and fixed molecular orbital set. They are the basis of each algorithm, and every routine passes at least from a single calculations of this kind. Single–point calculation are therefore essential for testing new quantum or hybrid algorithms against "classical" quantum chemistry methods. From the other side, being also able to provide good starting point with a reduced number of quantum resource is essential, especially given the hardware limitation faced by NISQ devices. In Chapter 5, we are going to introduce and show the applications of our state–preparation/warm–start method, *Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA)*.
- **Multi–Configurational Self–Consistent Methods for Industrial chemistry**: Industrial processes mostly involve bond–breaking / forming, radical intermediates, and transition–metal catalysts, all cases where single-determinant methods fail. CASSCF is able to treat several configurations and optimizes orbitals simultaneously, providing reliable potential–energy profiles and electronic-structure insight along the reaction path. On top of that, CASSCF reveals the multi-configurational character of transition states and intermediates, enabling chemists to understand reaction mechanisms in catalysis and polymer growth. In fact, in Chapter 6, we are going to introduce our SCF procedure, *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)*, applied on generic molecules and for free-radical ethylene model polymerization reaction.
- **Geometry optimization for Drug Design**: *DompeKey (DK)*, originally designed for cheminformatics, can classify substructures and help map or prioritize candidate geometries (e.g., different tautomers, coordination modes in metal complexes) before running expensive quantum energy evaluations. Geometries are therefore essential for the classification, and providing more accurate and precise structure is crucial. Improving the strucures, requires a suitable geometry optimization procedure, which, essentially revolves around finding the most stable arrangement of nuclei (the equilibrium geometry) for a given electronic structure method. In Chapter 7, we introduce our structure optimization method, *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)*, employed to find the equilibrium structures for generic molecules and *DompeKey (DK)* [93] descriptors.

Once again, in the spirit of defining a HQC pipeline, all the methods, except the AEGISS, have been developed and introduced in the University of L'Aquila computational chemistry and quantum computing group code base, *Quantum @ L'Aquila (QuAQ)*. The whole code has been written from scratch and integrated with different Python libraries: QuAQ can manage different quantum chemistry backend, in particular, PySCF [94–96] and the recently introduced VeloxChem [97–99]. The chemistry backed is used for fetching one- and two-body integrals, compute molecular orbital guesses, and for standard quantum chemistry reference calculation. In term of QDK, the circuital model is based on Qiskit [100], and recently also integrated with CudaQ [101]. This distinction is also present in terms of simulation backend which is split between statevector simulators, both Qiskit and CudaQ, and also GPU simulators, thanks to CudaQ. In term of qubit operator and fermionic operator both have been redefined internally and *back–and–forth* conversion between Qiskit and QuAQ has been implemented. QuAQ is completely modular and can be either executed using a standard *input-file*, as most common quantum chemistry packages, or by manually creating the workflow. For the specific implementation details, AEGISS is not publicy available at the moment, as well as the QA–q–GeomOpt, both defined as standalone packages. In particular, QA–q–GeomOpt has QuAQ as main dependency and most of the functionalities are taken from it. Additionally, being WAHTOR–SCF and QA–q–GeomOpt considered *embarrassingly parallel tasks*[6], they have been developed in order to be *High–Performance Computing (HPC)* oriented, including both intra-node and multi-node parallelization. The last chapter of the thesis, Chapter 8, will be used to briefly present and recap the content of our code base, and how it is integrated with different chemistry and quantum computing backend.

With the development and the study of all this use cases, we have been able to build up an example of interacting pipeline between quantum computing and most common quantum chemistry problem, looking towards the possibility of defining new algorithms thanks to the modularity and the different building blocks. As a reminder, the whole pipeline, except the active space selection, has been meant to be applicable to NISQ devices as proof of concept, dealing with statevector simulations and manageable number of qubits.

---

[6] From the simulation point of view, their routines are fully and efficiently parallelizable.

Part II

Hybrid Quantum Classical Pipeline Development

Chapter 4

Active Space Selection for Photodynamic Therapy

This chapter addresses the first stage of the electronic structure workflow: the identification of a chemically meaningful active space. We introduce the *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS)*, which exploits orbital descriptors and quantum entanglement measures to automate this traditionally manual step, providing a reproducible and data-driven foundation for all subsequent electronic structure calculations. We explore the role of active space selection in modeling Photodynamic Therapy (PDT)-relevant systems, highlighting both the conceptual foundations and practical strategies that underpin accurate excited-state simulations. By bridging methodological advances with applications to photoactive complexes, we aim to demonstrate how systematic active space design can unlock new opportunities for understanding and improving PDT.

PDT relies on light-activated Photosensitizer (PS) to generate *Reactive Oxygen Species (ROS)* that selectively damage cancerous or diseased cells. The efficiency of this process depends critically on the accurate description of excited-state electronic structures, particularly in transition-metal complexes commonly employed as photosensitizers. These systems are characterized by strong static correlation and near-degenerate electronic states, making them challenging targets for conventional computational methods such as Density Functional Theory (DFT), which often fail to capture multi–configurational effects.

To overcome these limitations, multireference wavefunction-based approaches, most notably Complete Active Space Self–Consistent Field (CASSCF) and its extensions, provide a robust framework for treating electronic correlation in excited states. At the heart of these methods lies the definition of an active space—the subset of molecular orbitals explicitly correlated in the calculation. A carefully chosen active space is essential for balancing computational feasibility with the need to capture the key electronic transitions underlying photoactivation, intersystem crossing, and energy transfer processes central to PDT. However, defining such spaces is far from trivial: it often requires significant chemical intuition and is prone to inconsistency, particularly as molecular size and complexity increase.

In recent years, the importance of systematic and automated approaches to active space selection has become increasingly evident. Methods based on orbital energies, natural orbital occupation numbers, or entropy measures have been developed to reduce subjectivity and improve reproducibility. Yet, no universally accepted “black-box” solution exists, and many approaches remain limited when applied to large transition-metal complexes relevant to PDT. This challenge is further amplified by the growing interest in quantum computing for quantum chemistry, where the number of orbitals that can be treated is severely constrained by hardware limitations. Active space selection, therefore, is not just a technical step in quantum chemistry but more a decisive factor in predicting and optimizing the performance of candidate photosensitizers.

This chapter is composed as follows: In Section 4.1, we introduce some intuition why is necessary to defined proper active space for Complete Active Space (CAS)-based method analysis. After, in Section 4.2, we will recap the existing active space selection method present in literature, stopping on the most famous ones. The whole Section 4.3 will be used to introduce our Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS) method. Then, Section 4.4 will be dedicate to introduce how Photodynamic Therapy (PDT) works and the molecular system we are going to test. Then Sections 4.4.1, 4.4.2, and 4.4.3 present the results on the three molecular system treated, Trans-Cl, Ruby-3, and TLD-1433, respectively. Finally, Section 4.5, will be used to wrap up the conclusion and final considerations.

4.1 Necessity for active space methods

Conventional computational methods for modeling chemical mechanisms and properties are limited by both computational cost and the accuracy they can achieve, especially in systems with strong electron correlation or complex excited-state behavior. Single-reference approaches, most notably Density Functional Theory (DFT) [45, 46], are successfully employed for modeling a variety of molecular and extended systems due to their favorable cost-to-accuracy ratio [102, 103]. However, they can fail to capture essential static correlation effects and are inadequate for systems where the electronic structure is governed by multiple, often near-degenerate, electronic configurations [104]. These limitations are particularly pronounced in transition metal complexes, open-shell species, and systems with low-lying excited states, where single-reference methods exhibit poor predictive power for key phenomena such as photoactivation, bond dissociation, and charge-transfer processes [105].

To address these deficiencies, wavefunction-based (WF) approaches are the methods of choice since they can systematically include electron correlation beyond the mean-field level. Configuration Interaction (CI) techniques improve upon a mean-field reference by incorporating excited determinants in the wavefunction expansion using a linear operator ansatz; however, single-reference truncated CI methods (e.g., CISD) are not size-consistent and fail to capture static correlation [106]. By contrast, Full Configuration Interaction (FCI), while exact within a finite orbital basis, becomes computationally infeasible due to its factorial scaling with system size. In contrast, Coupled-Cluster (CC) methods [43, 107] offer a more balanced approach by employing an exponential cluster operator ansatz to account for electron correlation with controlled approximations, delivering high accuracy for ground and some excited states. Nonetheless, standard (truncated) CC methods are by construction single-reference in nature and thus turn out to be often inadequate for systems requiring a multi-configurational description [108]. In such cases, multi-reference methods, particularly the Complete Active Space Self–Consistent Field (CASSCF) [109, 110] approach and its post-CASSCF extensions (e.g., CASPT2 [33], NEVPT2 [34–36]), are indispensable tools in multi-configurational quantum chemistry [106]. These methods explicitly account for static correlation by constructing the wavefunction from a set of active orbitals spanning multiple configurations. This set, normally referred to as *active space*, defines the subset of electronic orbitals that are treated as correlated, and choosing it correctly is critical to both the accuracy and efficiency of the calculations. Despite recent advances in algorithms and high-performance computing, the exponential cost (as the number of active orbitals grows) of multi-reference methods severely limits their applicability to relatively small active spaces [111], underscoring the need for novel, scalable approaches to accurately treat strongly correlated quantum systems [112] such as, for example, Density Matrix Renormalization Group (DMRG) [113], Full-CI Quantum Monte-Carlo (FCIQMC) [114], (semi-stochastic) Heat-Bath CI [115], Adaptive Sampling CI (ASCI) [116] and Iterative Configuration Exansion (ICE) [117], to name just a few.

With the advent of quantum computing, multi-reference methods and quantum algorithms [118, 119], particularly those based on variational techniques [59], are poised to work synergistically, combining the best of both worlds. Quantum computers have the potential of, inherently, excel at handling the large, entangled wavefunctions required in multi-reference calculations, which are potentially computationally prohibitive for classical approaches. For this reason they hold significant promise for revolutionizing quantum chemistry owing to the ability to simulate molecular systems beyond the reach of traditional computational methods [120–124].

However, due to current hardware constrains (e.g. limited number of qubits, sensitivity to noise and decoherence) the size of accessible active spaces for molecular simulations on contemporary quantum hardware remains limited. In the same spirit as in CAS methods, one therefore has to make larger and more complex problems tractable on current quantum devices by choosing chemically relevant active spaces and map it onto qubits, e.g. Jordan–Wigner transformation [48]. While the mapping of fermionic operators to qubits is a crucial aspect of quantum simulations of electronic structure, a detailed discussion is beyond the scope of this paper. For recent developments, we refer the interested reader, for example, to the following works [51, 52]. A brief conceptual overview of the mapping procedure is provided in the Supplementary Information in Section A.

As alluded above, one of the most significant challenges in applying multi-reference methods to solve quantum chemistry problems, either classically or using quantum computing, is the selection of an appropriate active space. This selection process requires a careful balance between reaching chemical accuracy and computational feasibility. Determining the optimal active space remains a difficult task, often relying on heuristic or trial-and-error approaches that are either time-consuming or prone to error, or eventually both. To address this Achilles' heel, considerable research has been directed towards unbiased and least empirical active space selection leading to the development of a variety of methods. These range from simple energy-based selections, such as the orbital energy-based method (around the Fermi level) to more refined approaches using additional metrics. One such method involves Natural Orbital Occupation Numbers (NOONs) [125, 126], which leverage occupation numbers to identify and select frontier orbitals. Furthermore, various procedures have been employed to define active spaces in the most automated way possible. A refined and more advanced method, based on NOONs, namely ASS1ST, allows to build an active space with *quasi*-NOONs from a first-order perturbation theory calculation on a minimal space [127, 128]. Other relevant examples are the Atomic Valence Active Space (AVAS) [129] method, based on modeling the active space according to desired atomic orbital contributions, and Auto-CAS [130–132], which defined the active space exploiting to correlation metrics from approximated wavefunctions. While promising in specific scenarios, these approaches still face limitations, and no universally accepted "black box" solution currently exists.

A unified, efficient approach to active space selection is therefore of utmost importance, as it would make multi-reference methods more practical and scalable for relevant applications of technological and biological interest. Moreover, a robust method could greatly enhance the accessibility of multi-reference quantum algorithms, paving the way for more widespread use of quantum computing in areas such as drug discovery, material science, and environmental chemistry, where accurate simulations of excited states and complex molecular interactions are essential.

4.2 Background on Active Space selection

In the early 1990s, both CASSCF [109, 110] and Complete Active Space Second–Order Perturbation Theory (CASPT2) [33] methods became widely used, with CASPT2 among the first to provide quantitative results for molecular excited states of photochemical interest [133]. The CASSCF method optimizes a wavefunction over a chosen set of active orbitals, allowing accurate descriptions of strongly correlated systems. CASPT2 refines this wavefunction by adding dynamic correlation effects, leading to better energy predictions. As already mentioned in the Introduction, the main challenge of these methods is the need to carefully define an active space, in order to comprise the relevant excited determinants, which is directly related to the active orbitals that can take any occupation beyond their initial mean-field configuration. The effectiveness of CASSCF depends on selecting an appropriate active space, which requires deep chemical intuition and knowledge of the electronic structure. Automating this selection remains a challenge due to the need to balance essential electronic correlations with computational feasibility, avoid heuristic metrics, and address near-degenerate states or dynamic correlations.

To extend to active orbital spaces beyond limitations of the *exact* CAS-approach, ≈ CAS(24 electrons, 24 orbitals) [111, 134–136], DMRG approach [37, 38, 41, 137] constitutes an alternative and versatile option. It is a variational algorithm that optimizes a wavefunction expressed as a Matrix Product Ansatz (MPS).

Importantly, the bond dimension of the virtual bonds, that "connects" the tensors in the MPS and over which the summation runs to yield the scalar coefficient of the (approximated) FCI tensor, controls – along with other factors [138] – both the accuracy and computational cost. The integration of DMRG algorithms into quantum chemistry codes, such as, for example, PySCF [94–96] and OpenMOLCAS [139, 140], has significantly expanded the applicability of FCI and CAS methods. These advances enable accurate wavefunction calculations for systems above the standard FCI limit, counting approximated CAS solution of sizes that range between 20 to 100 orbitals [40, 141–147] and even extends to fully relativistic Hamiltonian simulations [148–151]. These advancements continue to push the boundaries of quantum chemistry, improving the accuracy and efficiency of electronic structure calculations for increasingly complex molecular systems across the periodic table of elements.

Yet, the selection of an appropriate active space remains, as mentioned above, perhaps *the* central and most critical step in multi-configurational calculations, as it directly affects the accuracy and feasibility of electronic structure methods. In the following section, we discuss the main methodologies developed for active space selection and their role in improving computational efficiency and accuracy. In the literature, a collection of different active-space selection methods can be found spanning the last few decades since the advent of multi-configurational approaches in the 1970's. Here, we report the most commonly used ones as of today, especially those from which the protocol presented in this work is inspired by:

Selection around Fermi-Level

The simplest and most straightforward is by using a specific number of orbitals above and below the Fermi level, including a predefined number of molecular orbitals closest in energy to it. For instance, the last N orbitals in the occupied space, and the first M virtual. This approach assumes the most relevant electronic excitations and correlations occur within this energy window, capturing orbitals with energies directly tied to the system's chemical or physical properties while avoiding unnecessary orbitals further away from the Fermi level. If instead the process requires specific orbitals that might be outside this range, a CASSCF calculation done on this space will surely miss the crucial components.

Natural Orbital Occupation Number

This method for selecting a candidate active space is based on constructing natural orbitals [152]. In particular, it relies on the *Natural Orbital Occupation Numbers (NOONs)* [125, 126]. As first proposed by Löwdin [152], the natural orbitals are obtained by computing the eigenvectors of the one-electron density matrix from an approximated electronic wavefunction (e.g., the Second—order Møller-–Plesset Perturbation Theory), and the relative eigenvalues. The eigenvalues correspond to NOONs. In particular, they are obtained by diagonalizing

$$\rho^{|\psi\rangle} = \sum_{pq} \rho_{pq}^{|\psi\rangle} = \sum_{pq} \langle\psi|a_p^\dagger a_q|\psi\rangle, \tag{4.1}$$

where $\rho^{|\psi\rangle}$ is the one-electron reduced density matrix for a given correlated wavefunction $|\psi\rangle$, $a_p^\dagger$ is a fermionic creation operator that creates one electron in orbital p, and a_q is the annihilation operator that destroys one electron in the respective orbital q, respectively. The active space is then composed by collecting all the natural orbitals which occupation number falls in the range between 0.02 and 1.98 (typical values for NOONs), generally closer to the former for virtual orbitals and closer to the latter for occupied ones. This approach has been proven to be effective, especially because partial orbital occupation is inherently related to the presence of correlation in the wavefunction $|\psi\rangle$. By contrast, main issues with this method are encountered in the case of molecular systems in which the number of frontier orbitals is large and they are close to each other in orbital occupation. For these systems, a general and fixed threshold does not guarantee the definition of balanced and consistent active spaces. Another critical issue is that the the NOONs are in most cases based on MP2 wavefunctions due to its favorable scaling with system size, which, however, might not be an adequate starting choice for the system under investigation.

ABC scheme

The ABC scheme [153] is an automatic procedure for selecting an active space in multi-reference calculations such as *Multi–State Complete Active Space Second–Order Perturbation Theory (MS-*

CASPT2) or *Multi–Configuration Paired Density Functional Theory (MC–PDFT)* [154, 155]. Its purpose is to choose the number and type of orbitals to include in the active space without requiring manual inspection or system-specific intuition. The scheme uses three parameters, denoted A, B, and C. The parameter A determines how many doubly occupied orbitals (holes) are included; the parameter B sets how many virtual orbitals (particles) are included; and the parameter C specifies the number of extra excited states to average over, to ensure balanced treatment of near-degenerate states. It requires three different steps, with three different calculations: one for identifying the key occupied orbitals, one for identifying the relevant virtual orbitals, and the last one, a state-average CASSCF calculation on the dominant hole/particle orbitals. Finally, the resulting multi-configurational wavefunction is used as a starting guess for an MS-CASPT2 or MC–PDFT to obtain excitation energies.

Atomic Valence Active Space (AVAS)

AVAS is an automated procedure that allows one to build active spaces starting from a list of atomic valence orbitals and a single determinant wavefunction [129]. Setting out from a (user) pre-defined set of atomic valence orbitals, it includes all molecular orbitals with a sufficiently large overlap of the predefined set. To this end, AVAS uses a simple linear algebra rotation between the set of occupied and virtual orbitals in order to maximize the desired atomic valence character. In particular, the aim is to isolate the part of the reference wavefunction which involves the set of preselected atomic orbitals from the part which is not related, and define the active space on the first one.

Single-orbital entropy

The single orbital entropy is a measure derived from quantum information theory, widely used in the context of quantum chemistry, to analyze and define active space. In particular, it is the fundamental metric used in the automated active space selection algorithm AutoCAS [130–132, 156]. Physically, the single orbital entropy, $S(1)_p$, quantifies the amount of correlation or entanglement of a given orbital p with respect to the rest of the system. The single orbital entropy $S(1)_p$ for a specific orbital p is defined as:

$$S(1)_p = -\sum_{\alpha=1}^{4} \omega_{\alpha,p} \ln \omega_{\alpha,p}, \tag{4.2}$$

where α iterates over the possible occupations of a given spatial orbital, and $\omega_{\alpha,p}$ are the eigenvalues of the 1-orbital reduced density matrix (1o-RDM) of orbital p. For example, in the context of DMRG, the 1o-RDM can be obtained from the density matrix of the MPS representing the approximated wavefunction of the system by tracing out all sites except the p-th one. Orbitals with high single orbital entropy are likely to play a key role in the description of the correlated electronic structure and should be included in the active space, whereas low-entropy orbitals are typically less relevant for static correlation and can be excluded from the CAS. The AutoCAS procedure uses $S(1)_p$, which has been obtained from a partially converged – that is, a low bond dimension – MPS wavefunction, to aid the active space selection. By constructing threshold plots, as defined in Ref. [131], active spaces can be easily identified.

An example of a threshold plot is reported in Fig. 4.1, where the number of active orbitals for a given single orbital entropy is on the y-axis, and on the x-axis, $S(1)_p$ relative to the maximum

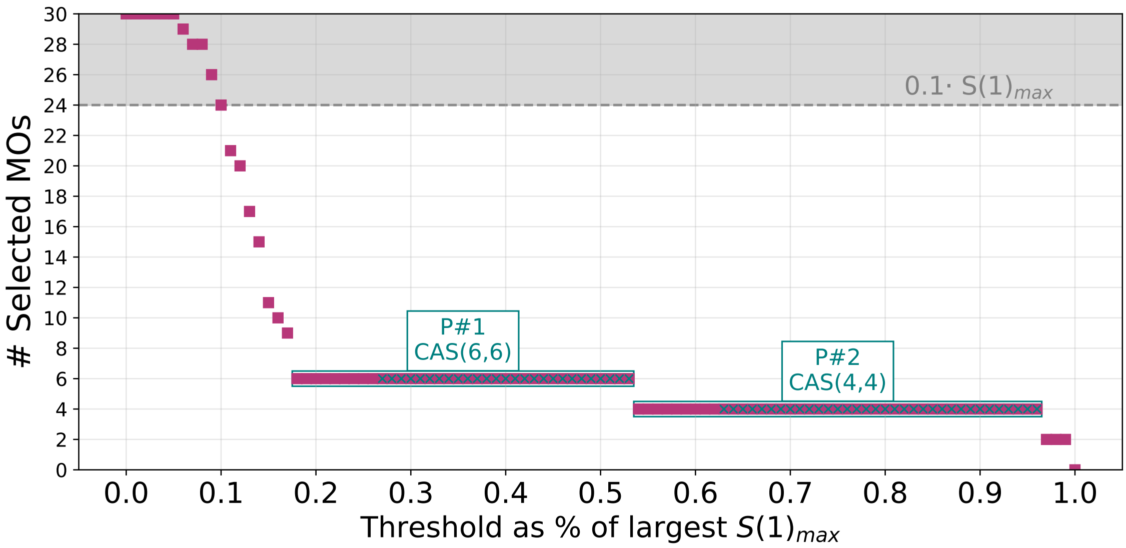

Figure 4.1: Examples of threshold plot for a C_6H_6 molecule in cc-pVTZ basis. The single orbital entropies have been extracted from a DMRG(30,30) calculation with a maximum bond dimension of 200. Each ■ corresponds to how many MOs have an $S(1)_p$ above the fixed fraction. The shaded area shows a 10% threshold. In this case, 24 orbitals have an entropy above this threshold. The "×"s displays when a plateau is identified. In the present case, two plateaus have been identified and they correspond to a CAS(6,6) and a CAS(4,4), respectively.

single orbital entropy, $S(1)_{max}$. In this way, the number of active orbitals that fall above a certain fraction of $S(1)_{max}$ can quickly be determined from a large set of candidate orbitals. Finally, an active space can be identified as a plateau in a threshold plot, corresponding to an invariance of active orbitals given a range of consecutive thresholds of the relative $S(1)_p$ value.

4.3 Atomic Orbital and Entropy Guided Inference for Space Selection

A novel active space selection method that integrates orbital entropy analysis with atomic orbital projections, drawing inspiration from key features of both AVAS and AutoCAS, and combining them within a unified workflow is defined here. The resulting framework can be directly interfaced with both quantum algorithms and classical multi-reference methods. A pictorial representation is given in Figure 4.2.

The proposed novel semi-automated framework combines the strengths of several active-space selection strategies to efficiently narrow the search for active orbitals and identify the most significant subset, optimizing a subsequent multi-configurational electronic structure modeling of the system of interest. It starts by selecting molecular orbitals (MOs) with the highest single-orbital entropies above a reasonable threshold. Then, atomic orbitals (AOs) linked to the fragments involved in the electronic state or transition of interest (e.g., $S_0 \rightarrow S_1$) are identified from the selected MOs. Hence, our approach combines two of the previous active space selection methods in order to define a *State-average/State-Specific* active space selection procedure that takes into account highly entropic orbitals (i.e., the most correlated ones) as well as accounts for the consideration of chemically and process-specific essential orbitals.

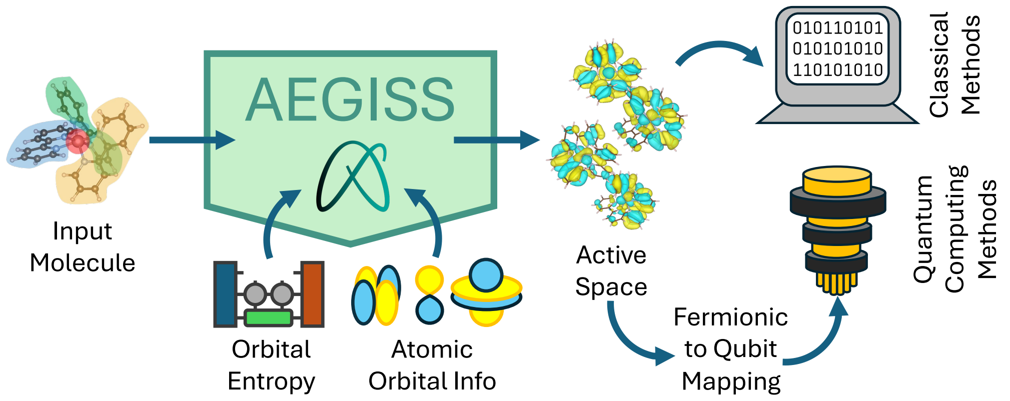

Figure 4.2: Simplified workflow of the AEGISS method showing the compatibility with both classical and quantum computing-based quantum chemistry methods.

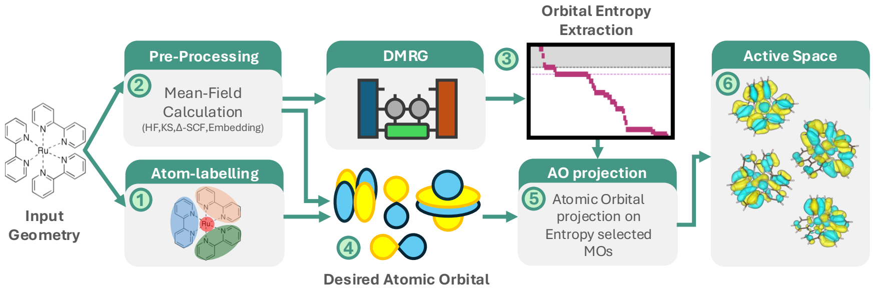

Figure 4.3: Pictorial representation of the AEGISS workflow. In ① , the molecule is possibly split into atom clusters, and the atomic orbitals of interest,④, are defined according to this subdivision. On the decorated geometry, ②, builds the initial guess molecular orbitals. An approximate DMRG calculation is then run to pre-select a subset of molecular orbitals, ③. On this set, the atomic orbital projection is performed, ⑤. Finally, ⑥, the active space is put together, refining the selection if required.

The active space selection procedure involves several key steps which are highlighted in Figure 4.3. First, the molecule is possibly divided into relevant clusters by labeling atoms, focusing on fragments where, for example, the excitation or chemical reaction of interest is assumed to occur. If nothing is known about the system, step 1 can be skipped. A starting guess for reference orbitals is obtained from a mean-field calculation, possibly within an embedding framework to account for solvation or any external effects that could have an impact on the mean-field solution. Next, single orbital entropies are computed by using DMRG to assess the relevance of a sufficiently large set of frontier orbitals, typically around 100 orbitals. This first selection helps to reduce the search space for the following step where an atomic orbital selection is performed next. Here, subsets of orbitals are chosen based on atomic-labeling and chemical relevance to the process or mechanism being studied, for example, near a reaction center or given their role in a bond formation/breaking process. The contributions of these selected atomic orbitals are then identified using a projection method on atomic orbitals built from a minimal basis-set. If multiple molecular orbitals have similar contributions and the obtained active space is not balanced, manual screening is required to refine the selection. Finally, additional post-processing steps may be applied to optimize or transform the chosen orbitals.

4.3.1 AEGISS workflow

We are now in a position to present in more detail the main steps of the proposed pipeline:

① **Atoms labeling**: The full molecule is going to be split in different clusters $\mathcal{C}_i$. An atom X belonging to one of the cluster i is going to be denoted with X_i . The combination of all the clusters defines the set $\{\mathcal{R}\}$, which represents the full molecule. A coarse labeling will allow us to describe the system from a more general point of view, while a finer one will allow us to select specific orbitals defined on a specific group of atoms or single atoms.

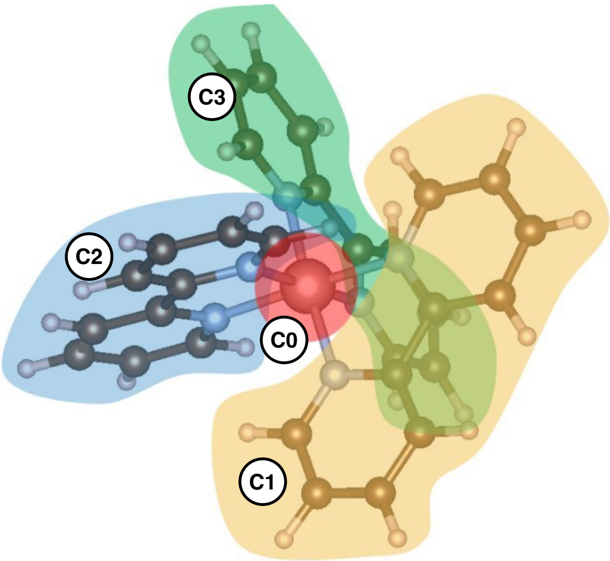

Figure 4.4: Visual representation of the labeling used for [Ru-*bpy*]$_3^{+2}$. Each of the color represent a different fragment, $\mathcal{C}_0$-$\mathcal{C}_3$, of the molecule. In this case, the complex has been split in the three ligands independently ($\mathcal{C}_1$ - $\mathcal{C}_3$) and the Ru(II) atom ($\mathcal{C}_0$).

② **Pre-processing/Starting guess**: In this step, we carry out the calculation of the guess MOs $\{\psi_p\}$, each of them associated with a ket $\{|p\rangle\}$, where p runs on both the occupied and virtual orbital. From the mean-field calculation, the MOs coefficient matrix $[\mathbf{C}]_{\mu p}$ is obtained. If the process under investigation is localized over a specific active region or fragment of the molecule $\mathbf{A}$, one could potentially apply fragment-based or embedding techniques to facilitate the active space selection, as well as dimensionality reduction like Effective Core Potential (ECP). Alternatively, transforming into Localized Orbitals or Natural Orbitals can help the selection of the window within which the subsequent single-orbital entropy calculation is carried out. This procedure will lead to a set of MOs of interest, $\{\psi_p^{\mathbf{A}}\}$, and the respective coefficient matrix, $[\mathbf{C}^A]_{\mu p}$. The size of the $[\mathbf{C}^{\mathbf{A}}]_{\mu p}$ is going to be $N_{AO} \times N_{\mathbf{A}}$, where N_{AO} is the number of AOs and $N_{\mathbf{A}}$ is the number of MOs in the active region $\mathbf{A}$. The environment part then is going to be left untouched by our proposed procedure. Hence, the remaining workflow will proceed and operate within only the active region $\mathbf{A}$. In this way, we can define the level of fragmentation that the Atomic Orbital projection is going to have (Step 5/6). If, for example, a Δ-SCF procedure is used in the starting guess mean-field optimization, the selection can be either defined using the same orbitals for each Δ-state MOs using a similarity measure (*State-Average active-space*) or by carrying out a specific selection for each Δ-state guess (*State-specific active-space*).

③ **Single Orbital Entropy extraction** At this step, we need to compute the $S(1)_p$ for each orbital $p \in \{\psi_p^{\mathbf{A}}\}$. The Single Orbital Entropy (SOE) can be extracted very easily by any wavefunction method as long as (elements of the) one- and two-*electron* RDMs are readily available [157]. Ideally, the electron correlation approach from which the single orbital entropies are to be extracted should be a computationally comparable cheap method, given the fact that the main correlations landscape emerges already with approximated methods.

(a) **Approximated wavefunction**: After the definition of the guess method and having obtained a reference molecular orbital set, an approximated wavefunction is computed. The electron correlation approximation should be chosen such that the main correlations and the single orbital entropies emerge even without running expensive calculations, that is, for example, within the context of DMRG without reaching full (energy) convergence wrt the corresponding required bond dimension of the MPS. In this step, the main pre-selection is carried out, by singling out only highly relevant MOs in terms of single orbital entropies. The set of MOs in the pre-selection step is going to be called *Large-CAS*. The *Large-CAS* can either be defined on the full MOs space, if the size of the system allows that, or on a meaningfully reduced portion of the same space. For the systems used as case studies in this work, given the size of the MO spaces, all the *Large-CAS* spaces have been selected by either including $\pm N$ orbitals above and below the Fermi-Level, or by a careful consideration of the orbital energies. Then, the wavefunction obtained is defined as

$$\begin{aligned} |\Psi^m_{\text{DMRG-CI}}\rangle &= |\Phi_{\text{occ}}\rangle \wedge |\psi^m_{\text{LCAS}}\rangle \wedge |\Phi_{\text{virt}}\rangle \\ &= |\Phi_{\text{occ}}\rangle \wedge \sum_{n=1}^{m} c_n \, |n\rangle \wedge |\Phi_{\text{virt}}\rangle \,, \end{aligned} \tag{4.3}$$

where $|\Phi_{occ}\rangle$ and $|\Phi_{virst}\rangle$ are the environment or in general the core part of the wavefunction which remains fixed (i.e., complementary to the large active space), $|\psi^m_{\text{LCAS}}\rangle$ is the approximated wavefunction obtained by the correlated method inside the *Large-CAS*, and ”$\wedge$" is an anti-symmetrized product operator. The notation follows the one used in Otten *et al.* [158]. The approximated wavefunction coefficients, $\{c_n\}$, and determinants/configuration state functions, $\{|n\rangle\}$, could, for example, be obtained by configuration sampling from an MPS. From the correlated subspace defined in Eq. 4.3, the relative density matrix

$$\rho^m = |\psi^m_{\text{LCAS}}\rangle\langle\psi^m_{\text{LCAS}}| \,, \tag{4.4}$$

can be defined, and it will be the description of the quantum state in the *Large-CAS* up to an approximation given by m. If the wavefunction has been obtained from a DMRG calculation, then both the creation of the density matrix of Eq.4.4 and the 1o-RDM can be easily and efficiently computed without the need for preceding configuration sampling (*vide supra*).

(b) **Entropy-based space-reduction**: As explained in Section 4.2, a first approximation on the orbitals that might define an active space can be defined by exploiting the SOE. In particular, for each orbital p, the 1o-RDM is computed as

$$\rho_p^m = \mathrm{Tr}(\rho^m)_{\overline{p}}, \tag{4.5}$$

where $\overline{p}$ is the set of all MO except th p-th. After computing the 1o-RDM for each orbital ψ_p, the respective single-orbital entropy $S(1)_p$ is computed according to Eq.4.2. The procedure ends with the selection of all the ψ_q such that $S(1)_q > \tau_E$, i.e. the fixed fraction of the found entropy. The threshold is defined as $\tau_E = \frac{S(1)_{max}}{10}$, where $S(1)_{max}$ is the highest single orbital entropy obtained. We are going to refer to the new set of MOs, $\{\psi_{E,p}^{\mathbf{A}}\}$, defined by the MOs coefficient matrix $[\mathbf{C}_E^{\mathbf{A}}]_{\mu p}$. The size of the obtained coefficient matrix $[\mathbf{C}_E^{\mathbf{A}}]_{\mu p}$ is $N_{\text{AO}} \times N_E$, where N_E is the number of MOs in the active region **A** with entropy value above the threshold.

④ **Atomic Orbital selection**: In this step, the AOs chemically relevant to the process/mechanism are identified. The information can be fetched from any prior-knowledge a.k.a. *chemical intuition* (reference work, reference calculations, by experience and/or theory). For each atomic cluster label, X_i, the information of the AO desired, "AO_s", is added, obtaining the combined labeling "$X_i\ AO_s$". Single or multiple sets of relevant AO-cluster labels, $\mathcal{D}$, are now built. Each $D \in \mathcal{D}$ is composed by identified groups of AOs and atoms that might contribute to the description of the mechanism. Same AO labels can be found in different groups.

⑤ **Atomic Orbital Projection**: For the AO projection, following the AVAS active-space selection procedure, an auxiliary minimal Atomic orbital basis (MinAO) is defined, with the coefficient matrix $[\mathbf{C}^{\mathbf{mAO}}]_{\nu p}$. The overlap $[\mathbf{S}]_{\nu\mu}$ matrix between the minimal basis set orbitals and the original active region MOs is defined as

$$[\mathbf{S}]_{\nu\mu} = \langle \nu | \mu \rangle = \int \phi_\nu^{\mathbf{mAO}}(\mathbf{r}) \phi_\mu^{\mathbf{A}}(\mathbf{r})\, d\mathbf{r}, \tag{4.6}$$

where $|\nu\rangle$ is one of the AOs in the minimal basis $[\mathbf{C}^{\mathbf{mAO}}]_{\nu p}$ and $|\mu\rangle$ is one of the AOs of the original active region $[\mathbf{C}^{\mathbf{A}}]_{\mu p}$. On the left side, the contraction on the AO basis is required. For each set D, we define the *Atomic-basis projector* $\mathbf{P}^D$, with shape $N_D \times N_{AO}$. The projector $\mathbf{P}^D$ does not have any particular structure; it simply trims the matrix to which it is applied to the rows/columns that satisfy some condition (similar to the one for selecting the highly entropic MOs). In this context, it selects only the rows associated with the AO basis function contained in the set D, by left-multiplication. The projector, once applied to the overlap matrix $[\mathbf{S}]_{\nu\mu}$, allows one to build the projected overlap matrix $[\mathbf{S}^D]_{\eta\mu}$, obtained as

$$[\mathbf{S}^D]_{\eta\mu} = \sum_{\nu \in D} [\mathbf{P}^D]_{\eta\nu} [\mathbf{S}]_{\nu\mu} \tag{4.7}$$

with shape $N_D \times N_{AO}$. Then, we determine $[\mathbf{O}_E^D]_{\eta p}$, which corresponds to the residual coefficient matrix of the projected overlap matrix $[\mathbf{S}^D]_{\eta\mu}$ and the entropy-selected active-region MOs, $[\mathbf{C}_E^{\mathbf{A}}]_{\mu p}$, that is,

$$[\mathbf{O}_E^D]_{\eta p} = \sum_{\mu} [\mathbf{S}^D]_{\eta n \mu} [\mathbf{C}_E^{\mathbf{A}}]_{\mu p}, \tag{4.8}$$

where η is one of the desired AO in the group D, μ is one of the AO in the extended original basis, and p is one of the MOs in the entropy-selected active-region. The contribution of each desired AOs in D computed for every MO, p, in the preselected set, is denoted as w_p^D and can be obtained as

$$w_p^D = \sum_{\eta \in D} [\mathbf{O}_E^D]_{\eta p}, \tag{4.9}$$

Figure 4.5 summarizes the whole procedure of this step.

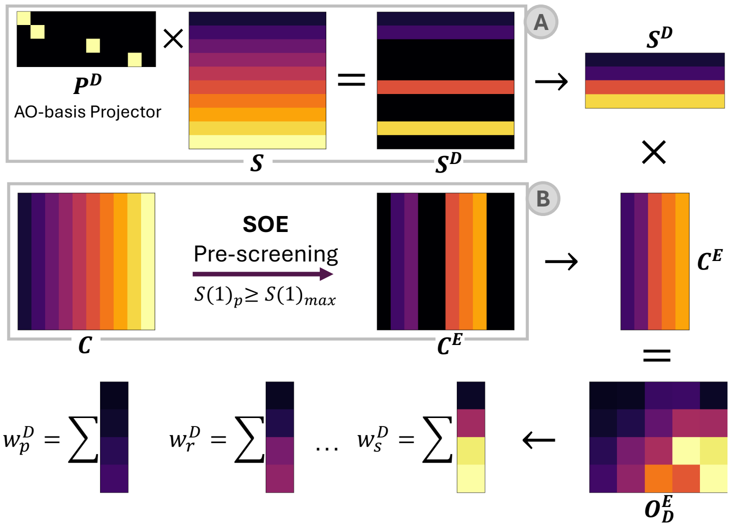

Figure 4.5: Schematic of the procedure used to obtain the weights w_x^D for each orbital $x \in \{\psi_E\}$ and the specific AO label group $D \in \mathcal{D}$. The section (**A**) is associated with the AO projection, while the section (**B**) to the SOE pre-screening.

Now, it is possible to either define a cutoff ϵ_D on the contributions values w_p^D, and select all the orbitals above, or to impose a threshold n_D, and select up to the n_D-th MO with the highest AOs overlap. The choice of the cutoff/threshold is done empirically; in particular, the former opens up for a more conservative selection and might require further analysis of the orbitals picked, whereas the latter uniquely selects only a number of MOs with the highest AO contributions.

⑥ **Active-Space construction**: At this point, all the reduction and selection techniques have been applied, and the remaining set of candidate MOs is identified. In particular, for each AO contribution group $D \in \mathcal{D}$, a set of orbitals $\{\psi_r^D\}$ is defined. Further refinement of the selected space can be carried out by applying the following strategies:

(a) **Post-Selection Screening**: Given the challenges in accurately selecting orbitals due to the complexity of the problem, the goal of our pipeline is to prioritize robustness over precision. To this end, we propose a more conservative strategy that ensures that no potentially important orbitals are excluded due to a narrow or overly precise criterion, which could overlook key contributions. This is particularly true when the system complexity increases. However, the set of orbitals selected $\{\psi_r^D\}$ for each group $D \in \mathcal{D}$, can be further analyzed and adapted in terms of size and balance. As a result, the final step of the workflow can include some level of visual inspection, which is still required to exclude Rydberg states or make decisions about whether to include specific orbitals, for example, tailored to a specific excitation of interest. Another useful tool, if the molecular system exhibits point-group symmetry, is the usage of irreps to discriminate which orbitals to include and which to exclude from the active-space, respectively.

(b) **Additional Post-Processing**: In addition to manual screening, further post-processing techniques such as Frozen Natural Orbitals (FNOs) [159–162] or localization of virtual orbitals can significantly enhance the efficiency and interpretability of the active space. FNOs are commonly used to reduce the size of the correlated space by identifying and excluding virtual orbitals with negligible occupation in a correlated wavefunction, typically based on an MP2 or CCSD calculation. This approach retains accuracy while lowering computational cost. Similarly, any localization of virtual orbitals can improve the chemical interpretability of the virtual space and reduce delocalization artifacts that complicate correlation treatment. Localized orbitals can also facilitate fragment-based embedding.

In terms of settings, the proposed workflow requires two types of inputs. The first is related to the simulation and computational setup. In particular, it requires a mean-field method, a basis-set (and ECP if needed, as well as solvation corrections), while the minimal basis-set for the projection is fixed to MinAO. The auxiliary basis set, following AVAS guidelines, can be chosen to be larger. For DMRG, the number of sweeps and number of states are required as input parameters. For the single orbital entropy pre-selection, the threshold is fixed to $\frac{S(1)_{max}}{10}$. Furthermore, the SOE extrapolation step is not tied to DMRG and alternative procedures such as SparQ [163] for other approximated wavefunctions are applicable as well. The second type of input is the one related to *chemical-intuition* which allows the workflow to be flexible according to user needs. This includes the definition of atomic-orbital labels and the AO projection cutoff (ϵ_D) or the threshold (n_D).

4.3.2 Detailed AEGISS scaling

Let N_{AO} be the number of atomic orbitals (basis functions) in the whole system, N_A the number of MOs kept for the active region A after fragmenting/transforming, N_E the number of entropy-preselected MOs (subset of A), k the number of orbitals in the *Large-CAS* treated by DMRG, m the DMRG bond dimension, s the number of DMRG sweeps, and $\sum_D N_D$ the total number of AOs included across all AO groups D:

- **Pre-processing (HF or KS-DFT)**: For naïve HF or KS, the dominant cost per SCF iteration comes from the Fock build, scaling as $\mathcal{O}(N_{AO}^4)$, for the former, and $\mathcal{O}(N_{elec}^3)$, for the latter. Any AO-to-MO transforms or orbital localization steps are typically $\mathcal{O}(N_{AO}^3)$.

- **Single-orbital entropies via DMRG**: Let k denote the *large-CAS* size in the approximate DMRG run. A two-site DMRG sweep has a cost of $\mathcal{O}(km^3 + k^2m^2)$, where $k\,m^3$ arises from local solutions and k^2m^2 from Hamiltonian contractions. The total cost is $s \times \mathcal{O}(km^3 + k^2m^2)$. The entropy extraction itself (from 1o-RDMs) is $\mathcal{O}(km^2)$–$\mathcal{O}(km^3)$, negligible compared to sweep costs. This stage often dominates the total runtime when k or m are large.

- **AO-projection**:The AO-projection involves building the cross-integral overlap matrix $\mathbf{S}$ scales as $\mathcal{O}(N_{mAO}N_{AO})$. For each AO group $D \in \{\mathcal{D}\}$, the projection $\mathbf{S}_D = \mathbf{P}_D\mathbf{S}$ is com-

> puted at $\mathcal{O}(N_D N_{AO})$ cost, followed by a matrix multiply $\mathbf{S}_D\mathbf{C}_E^A$ at $\mathcal{O}(N_D N_{AO} N_E)$. Weight computation and selection add $\mathcal{O}(N_D N_E)$ for summations. Summed over all groups D, the dominant projection cost is $\mathcal{O}((\sum_D N_D)N_{AO}N_E)$, much cheaper than DMRG.

With RI/DF, SCF pre-processing scales as $\mathcal{O}(N_{AO}^3)$. DMRG screening dominates for k and high m, scaling as $s \times \mathcal{O}(km^3 + k^2m^2)$. The AO-projection is relatively inexpensive, scaling as $\mathcal{O}(N_{AO}^2)$ for overlap storage and $\mathcal{O}((\sum_D N_D)N_{AO}N_E)$ for projection and weighting. For large systems with small *large-CAS*, SCF can dominate; for large *large-CAS*, DMRG dominates. To control scaling, it is advisable to use Resolution to Identity/Density-Fitting to keep SCF cubic, limit the Large-CAS to the size needed for reliable entropy estimates, and gradually increase m. Using symmetry-adapted or block-sparse DMRG can reduce effective prefactors.

To summarized, the cost is partly dominated by the initial mean-field step, either Hartree–Fock (HF) [27] or Kohn–Sham (KS) [164]. However, when the chosen *Large-CAS* contains many orbitals, the DMRG-based single-orbital entropy extraction can become the most expensive step. By contrast, the AO-projection step—dominated by the computation of the cross-basis overlap matrix and the subsequent MO projection remains negligible in comparison to the former steps.

4.4 Photodynamic Therapy

Photochemistry involves light-induced chemical processes that lead to changes in the (excited-state) molecular structures and drive chemical transformations. When molecules absorb light, they can transition to excited states, leading to ensuing processes like internal conversion, Inter–system Crossing (ISC), or radiative emission (fluorescence and phosphorescence) to dissipate the excess energy. These mechanisms are crucial in technologies such as solar energy conversion and biological processes like DNA photostability[165], photosynthesis, and Photodynamic Therapy[166]. In PDT, light activates a photosensitizer, initiating an ISC and energy transfer to oxygen, with the ultimate goal to generate Reactive Oxygen Species (ROS) that are capable of introducing apoptosis in cancer cells. Here, a small energy gap between singlet and triplet states (ΔS$_1$-T$_1$) enhances ISC and reverse ISC (RISC) rates, that can lead to improved ROS production and therefore an increase in therapy efficiency [92]. Different criteria must be met to define whether a candidate molecule is a good photosensitizer: Lack of dark toxicity, Solubility in aqueous vehicle, Chemically pure, easy to synthesize, High absorption in "PDT window" for deep light penetration, i.e. 700-900 nm, preferential accumulation in tumor; rapid elimination from the body, and also been able to generate sufficient singlet oxygen, even under hypoxic conditions.

Hence, understanding the molecular and dynamical nature of photochemistry related processes not only help to optimize them but also enable further research into non-invasive cancer treatments and, in addition, improve technologies like organic light-emitting diodes (OLEDs). Achieving chemical accuracy (∼1 kcal/mol) in excited-state modeling is a challenge in computational chemistry[167]. Calculating excitation energies, critical to photoactivated processes, requires consideration of geometric relaxation, vibrational energy, and solvent effects, complicating comparisons with experimental data. Excited states involve diverse electronic phenomena (e.g., π-π^* transition, charge transfer, generation of exciton pairs) that require accurate computational methods. Some systems, like transition metals and biradicals, demand multi-configurational or strongly correlated approaches. In such cases, active space selection is vital, as it determines which orbitals are included in the description of electronic states, ensuring accurate modeling of complex structures and interactions while also controlling the computational cost. In this study, we have selected three increasingly large molecules that could serve as models for PDT, including one currently in clinical trials. This progression from small to large molecules allows us to explore behaviors at different scales while demonstrating the effectiveness of our approach on increasingly complex problems. This helps us to better understand the underlying mechanisms and optimize our methodology for more challenging systems.

We selected three systems of increasing complexity and size: Trans(Cl)-Ru(bpy)$Cl_2(CO)_2$ (dubbed in this work as *Trans-Cl*, 27 atoms) [168], $[Ru(bpy)_3]^{2+}$ (dubbed in this work as *Ruby-3*, 61 atoms) [169], and rac-[Ru(dmb)$_2$(IP-3T)]Cl_2 (*TLD-1433*, 99 atoms) [92, 170]. The first two serve as simplified prototypes of photoactive drugs for PDT, with some literature available to guide our benchmarking of our semi-automated active-space selection approach. Trans-Cl [168], the smallest system, contains a single Ru(II)-coordinated ligand, making it useful for identifying key metal-to-ligand (M → L) transitions relevant to PDT. Ruby-3 [169] increases both size and complexity, acting as a precursor to many PDT candidates and offering insight into orbital contributions in larger systems. TLD-1433 presents the most challenging case due to its size and electronic structure complexity. With limited theoretical studies available, accurate wavefunction-based calculations remain unfeasible, leaving Time–Dependednt Density Functional Theory (TD-DFT) [171–173] as our primary reference [92, 170]. Although direct comparison with experimental data is beyond the scope of the present study, UV–visible spectra and theoretical results for Trans-Cl and Ruby-3 are available in the respective reference works and may be used for validation purposes. By contrast, for TLD-1433, both experimental and theoretical data remain very limited.

For all the system, the specific computational details are collected in the Appendix 1.

4.4.1 Trans-Cl

The ground state geometry of this molecule has been optimized with point group symmetry (C_{2v}) in gas-phase using B3LYP [174, 175] functional, a 6-31G*-type basis set on all atoms except Ru and a def2-TZVP [176] basis with quasi-relativistic Stuttgard-Dresden pseudo-potential [177] for the Ru atom. The Trans-Cl geometry optimization has been performed with ORCA [178]. The setup for the optimization has been chosen to be as close as possible to Ref. [168], which we consider as our reference data. Both in ORCA and on PySCF, the correct symmetry, C_{2v}, can be kept by increasing the symmetry threshold during the mean-field calculations.

While the basis set used in the reference work is ANO-rcc-VTZP [179, 180], in this work, for the Restricted-HF (RHF), CASSCF, and DMRG, we employed def2-TZVP basis-set on all atoms and def2-TZVP basis with quasi-relativistic Stuttgard-Dresden pseudo-potential for the Ru(II) atom, as done for the geometry optimization setup. In the reference article [168], the active space selected is composed of 14 electrons distributed over 13 molecular orbitals, CAS(14,13), including a balanced set of 8 bonding/antibonding π-orbitals placed on the ligands and the 5 *4d* orbitals of the Ru(II) atom. For this system, the main goal is to benchmark our approach against the reference paper. This is motivated by the fact that no procedure for the construction active-spaces is given by the authors.

Starting from the atomic-labeling step, the molecule can be split in three main clusters: $\mathcal{C}_0$

corresponding to the Ru(II) atom, $\mathcal{C}_1$ to the bipyridine ligand, $\mathcal{C}_2$ for two the CO groups, $\mathcal{C}_3$ for the two Chlorines. The clusters are highlighted in Figure 4.6.

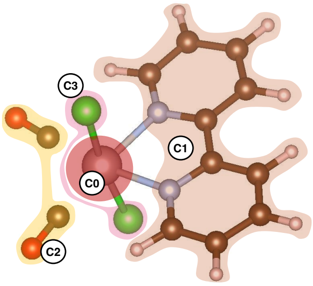

Figure 4.6: Trans-Cl highlighted clustering following the same labeling as explained in this section.

Thus, following the workflow as defined in Section 4.3.1, a mean-field calculation is performed and the initial-guess MOs are defined. For this system, initially, $N_{\mathbf{A}}$ is 658. In terms of occupation, the HOMO sits at the 80-th MO, counting a total of 160 electrons in the system. From the orbital occupations and the corresponding eigenvalues, we define a reduced search space cutting out core orbitals and high-energy virtual MOs. As a results, the size of the Large-CAS, N_{LCAS}, contains 90 MOs, leaving 32 orbitals as inactive core. The following setup leads to a DMRG with CAS(92,90), i.e. 92 electrons in 90 MOs.

Once the reference approximated MPS wavefunction is obtained, the $S(1)_i$ for each MO is calculated using Equation 4.2, and the entropy selection can be performed. For this system, the maximum single orbital entropy found is $S(1)_{max} = 0.226$, imposing an entropy threshold to $\tau_E = 0.00226$. In figure 4.7 we show the threshold plots defined for Trans-Cl. From this plot, it is possible to understand how many MOs are selected imposing different entropy thresholds. The one used for this molecule, and in the following work, is the standard 10% of the maximum entropy found. From the total pool of 90 orbitals, with the standard τ_E, the number of pre-selected MOs, N_E, is

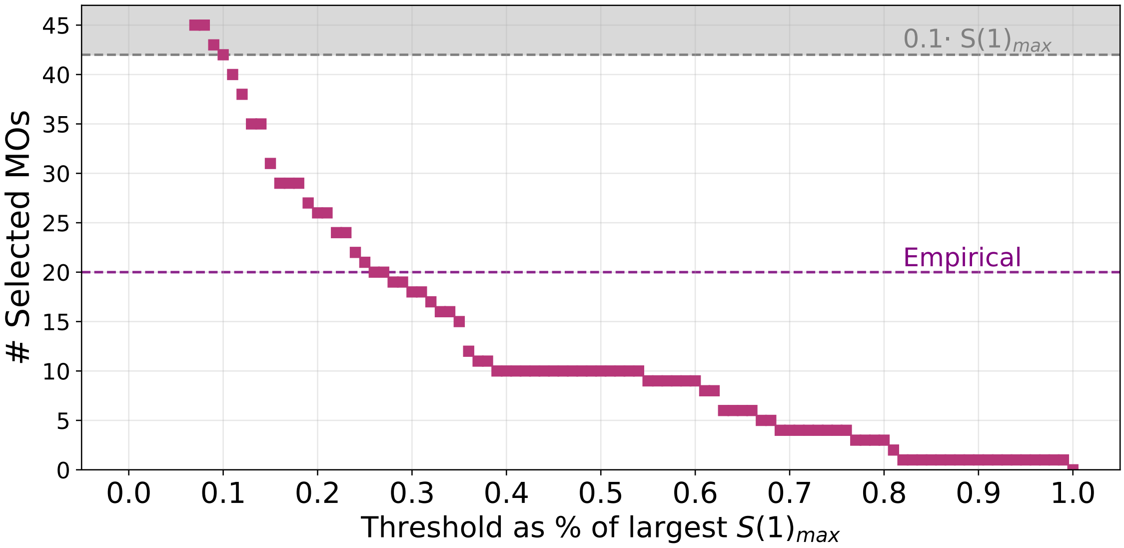

Figure 4.7: Threshold plot, defined according [130] , for Trans-Cl. In this case, the pre-selection step with DMRG single orbital entropy leads to the selection of 42 MOs. The approximated wavefunction is obtained on a CAS(92,90) with bond dimension 200.

reduced to 42. In Figure 4.8, it is possible to see the $S(1)_p$ associated to the 90 MOs considered in the DMRG calculation. With the pre-selection, the molecular orbital coefficient matrix of high entropy MOs is defined, $[\mathbf{C}^{\mathbf{A}}_{E}]$, and its shape is $N_{AO} \times N_E = 659 \times 42$.

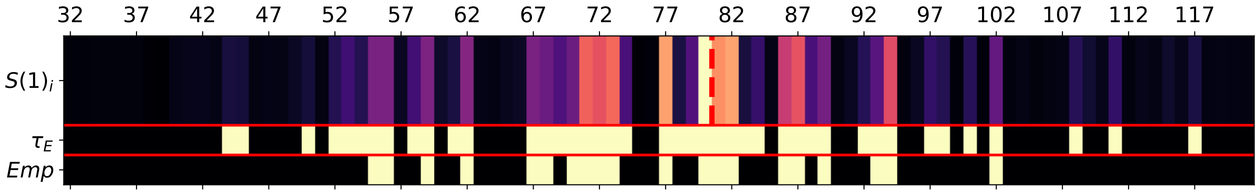

Figure 4.8: In this plot, the upper heatmap reports the value of the single orbital entropies, $S(1)_i$, for MOs. In the second frame, the pre-selected orbital with entropy above τ_E, and in the lower frame, the one with entropy above 0.06. The dashed line draws the HOMO-LUMO gap, set at MO number 80.

The AO labels of our interest for each Cluster $\mathcal{C}_i$ are the contributions of the five *Ru 4d* orbitals, as independent labels ("4dxz", "4dyz", "4dxy", "4x^2-y^2", and "4dz^2"). For recovering the π/π^*-orbitals we take the MOs having high *2p* orbitals contributions on the ligands. Specifically, the contributions of 2p$_z$ AOs from labeled Carbons and Nitrogen on the ligand, avoiding the contribution from the CO moieties.

ID	Cluster	AO Labels	O	V
L1	$\mathcal{C}_0$ (*Ru*)	4d$_{xz}$	1	
L2	$\mathcal{C}_0$ (*Ru*)	4d$_{yz}$	1	
L3	$\mathcal{C}_0$ (*Ru*)	4d$_{xy}$	1	
L4	$\mathcal{C}_0$ (*Ru*)	$4d_{x^2-y^2}$		1
L5	$\mathcal{C}_0$ (*Ru*)	4d$_{z^2}$		1
L6	$\mathcal{C}_1$ (*Bpy*)	[C$_1$ 2p$_z$, N$_1$ 2p$_z$]	4	4
L7	$\mathcal{C}_2$ (*CO*s)	None		
L8	$\mathcal{C}_3$ (*Cl*s)	None		

Table 4.1: Trans-Cl AO-labels of interest for each of the clusters. In this case, the atoms related to clusters $\mathcal{C}_2$ and $\mathcal{C}_3$ are not used.

After the identification of the AO-labels of interest, the AO-projection is applied on the pre-selected set of MOs, $[\mathbf{C}^{\mathbf{A}}_{E}]_{\mu p}$, and the relative contributions can be computed.

AS	Occupied							Virtual						
AS1	59	**67**	**68**	**70**	72	73	77	80	81	82	87	**93**	**100**
AS2	59	62	**70**	72	77	**78**	**79**	80	81	82	87	**93**	**100**
AS3	59	62	**67**	**68**	**70**	71	77	80	81	82	87	**93**	**100**
AS4	73	74	75	76	77	78	79	80	81	82	83	84	85

Table 4.2: Candidate active orbital spaces (AS) selected for the Trans-Cl complex. **Bold** MOs number are referred to Ru(II) 4d orbitals.

The minimal active space of interest is composed of only the five 4*d* Ru(II) orbitals. These have been chosen taking the first MOs for each Ru 4$_d$ contribution, independently of the occupation and is achieved by imposing a $n_{[\text{Ru } 4_d]}$ equal to 1, using labels group from L1 to L5.

The extended active spaces can be built including also π-orbitals centered on the bipyridine ligand. However, the inclusion of pi orbitals can be carried out in two different ways. In one case, the π/π^* orbitals are obtained by simply splitting the selection obtained from applying the workflow

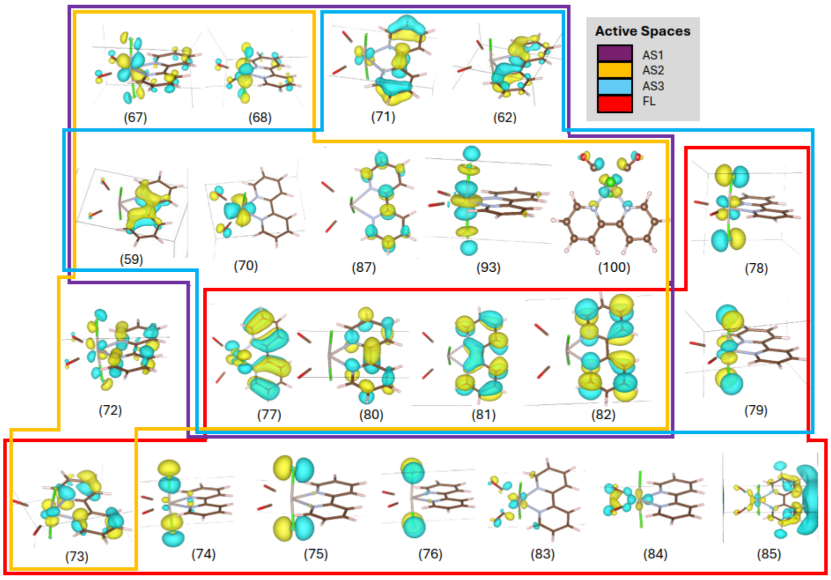

Figure 4.9: Pictorial representation of the Trans-Cl active spaces found and listed in Table 4.2

between occupied and virtual orbitals, imposing $n_{[\mathrm{C}_1\ 2\mathrm{p}_z]}$ equals to 4, for each subset respectively, according to group label L6. This space would correspond to AS3 in table 4.2.

The selection of the π/π^*-orbitals, can, however, be refined. In particular, this can been done by discarding all the MOs with $\epsilon_{[\mathrm{C}_1\ 2\mathrm{p}_z]} < 0.1$. In this way, the search pool of MOs have been reduced from 42 to 12. Finally, information about orbital symmetries, as provided in the paper has also been taken into consideration to further perfect the spaces found.

Following this procedure, the five *Ru 4d* orbitals are obtained quite easily with only an ambiguity regarding whether the HOMO and HOMO-1 are needed within the AS. We have generated two different active spaces with and without including the HOMO/HOMO-1 (AS1/AS2). This ambiguity comes from the fact that the HOMO and HOMO-1 are Ru $4d_{xz}$ and Ru $4d_{yz}$, respectively, and they present similar orbital character of HOMO-12 and HOMO-13 but lower entropy. A pictorial representation of the MOs involved in each of the AS can be found in Figure 4.9, this also includes the space one would obtain by taking orbitals around the Fermi level.

The active space used in the reference paper [168] can, hence, be reconstructed. In table 4.2 are collected three identified active spaces for Trans-Cl complex.

Spin-Averaged CASSCF benchmark

We performed some preliminary Spin-Averaged(SA)-CASSCF calculations to check the results consistency between our active spaces and the reference. We have not performed more accurate and complex calculations, such as CASPT2, because it is out of the scope of this work. We also obtained results with slightly different simulation conditions that are not completely comparable to the onse used in the reference article. By inspecting the excitations from CASSCF calculations of identified spaces and comparing them to the CASSCF results reported in the reference (using the same number of roots they used for singlets (12) and triplets (11)), we found that AS1 and AS2 give qualitatively similar results to the one reported in the work by Gonzalez et. al. In fact, in both these cases we find S_1 to be a predominantly metal-centered (MC) excitation involving d orbitals on the Ru(II) atom and S_9 to have a metal-to-ligand charge-transfer (MLCT) character. However, AS2 includes the HOMO and LUMO, which are typically found to be relevant orbitals in many applications and shows more sensible excitation energies compared to AS1. While energetically we cannot make direct comparisons with the paper, these findings suggest that AS2 is the best candidate active space for modeling Trans-Cl excitations. AS3 and AS4, on the other hand, give very different results in terms of excitation compositions. Additional TD-DFT benchmarking is reported in the Appendix 2. As a general trend, we see our calculations being in good agreement with the reference [168].

Singlet S_n	**AS1**	**AS2**	**AS3**	**AS4**
S_1	3.52084	3.4632	5.3783	3.5997
S_2	5.64656	3.5092	5.8276	3.649
S_3	5.75963	5.06	5.8517	4.3014
S_4	6.00366	5.2642	6.6143	4.5214
S_5	6.41763	5.3741	7.0279	5.42
S_6	6.71239	5.4494	7.8617	5.6879
S_7	7.22238	5.9283	7.8868	5.9645
S_8	7.37199	6.1987	7.9466	6.0366
S_9	7.51438	6.609	8.3037	6.2524
S_{10}	7.88445	6.6782	8.5376	6.8798
S_{11}	8.29203	7.1181	8.6899	7.0042

Table 4.3: Singlet excited state transition energies $\Delta(S_n - S_0)$, in eV, obtained from SA-CASSCF calculations on the Trans-Cl complex.

Triplet T_m	**AS1**	**AS2**	**AS3**	**AS4**
T_1	2.6606	2.7022	3.8196	2.7762
T_2	3.7442	2.8021	4.8487	2.9168
T_3	4.6312	3.8994	4.9330	4.0832
T_4	4.85	3.9288	5.2589	4.401
T_5	5.6106	4.1796	5.3183	4.5825
T_6	5.7031	4.4569	5.3635	5.2347
T_7	5.9157	4.6526	5.5326	5.2795
T_8	5.9889	4.9424	6.6266	5.7255
T_9	6.4157	5.6307	6.7555	5.7767
T_{10}	7.1017	5.9621	7.8793	6.1791
T_{11}	7.2738	5.9702	8.2575	6.2023

Table 4.4: Triplet excited state transition energies $\Delta(T_m - S_0)$, in eV, obtained from SA-CASSCF calculations on the Trans-Cl complex.

4.4.2 Ruby-3

The ground state geometry of this molecule has been optimized with symmetries in gas phase using RI-MP2 [181], def2-TZVP basis set on all atoms, and quasi-relativisitc Stuttgard-Dresden pseudopotential for the Ru(II) atom. Also in this case, we tried to reproduce the same simulation conditions as the reference paper[169], which we used as benchmark. In order to obtain the right symmetry, D_3, an additional step of geometry optimization has been further applied using AVOGADRO software [182]. For the RHF, CASSCF, and DMRG calculations, the basis-set used is slightly different from the one used in the reference work. In [169], ANO-RCC-VTZP on Ru(II) and all H atoms and ANO-RCC-VDZP without d functions on C and N atoms. As for the previous system, a def2-TZVP basis set on all atoms and def2-TZVP basis with quasi-relativistic Stuttgard-Dresden pseudopotential for the Ru(II) atom have been used.

Ruby-3 is a peculiar system, which it has been chosen for its highly symmetric structure and for being, ideally, the linking between Trans-Cl and TLD-1433. As for Trans-Cl, the objective for this system is to recreate the same active space as in the reference paper [169]. The reference active-space is composed of 16 electrons in 13 orbitals: five Ru 4d orbitals, a balanced set of three pairs of π/π^*-orbitals on the three bipyridines, and two σ-bonds between Ru and ligands. The cluster for Ruby-3 are defined quite easily: $\mathcal{C}_0$ is the central Ru(II) atom and the three bipyridines are associated to three independent clusters, $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$, respectively, as shown in Figure 4.10.

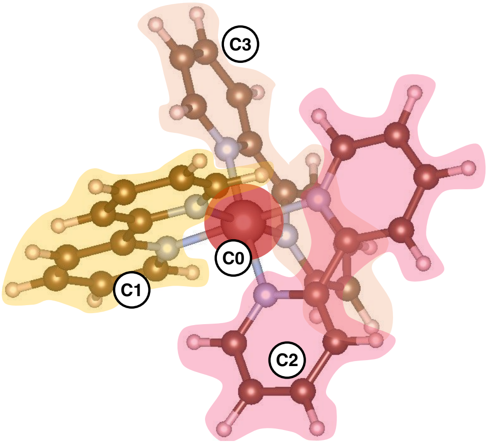

Figure 4.10: Ruby-3 highlighted clustering defined according to the atomic labeling explained in the current section. The Ru(II) atom and each of the bipyridine ligands are treated as independent clusters.

Unlike the previous system, there is no dominant or main plane along which we can identify the π/π^* MOs, and the high symmetry of the complex increases the difficulty related to the selection of a balanced set of them. Additionally, we are interested in determining excitation energies, thus, the active space should include MOs that are supposed to be involved in MLCT. The selection focus is also to identify σ-bonds between the central metal and the ligands, which are the most difficult to identify.

The initial mean-field calculation generates $N_{\mathbf{A}} = 1300$ MOs. The 130-th orbital is the HOMO with 260 electrons due to its positive charge, 2+. The DMRG has been carried out on a Large-CAS of 90 MOs, starting from the 100^{th} up to the 190^{th}, with a smaller bond dimension of 200. The following setup leads to a DMRG with a CAS(60,90), i.e. 60 electrons in 90 MOs.

The maximum single orbital entropy found for Ruby-3 is $S(1)_{max} = 0.177$, imposing a entropy threshold to $\tau_E = 0.00177$. Following the threshold plot, Figure 4.11, the total pool of 90 orbitals, with the standard τ_E, the number of pre-selected MOs, N_E, is reduced to 30. In Figure 4.12, it is possible to see how the relative $S(1)_p$ values are spread in the original 90 MOs.

In the absence of clear guidance from pre-existing literature, a reasonable initial choice for the active-space can be made based on chemical intuition and structural symmetry considerations. In the case of a Ru(II)-complex coordinated with bipyridine ligands, a minimal model could include: the valence orbitals of the transition metal, which for Ru(II) are the five $4d$ orbitals, the π bonding

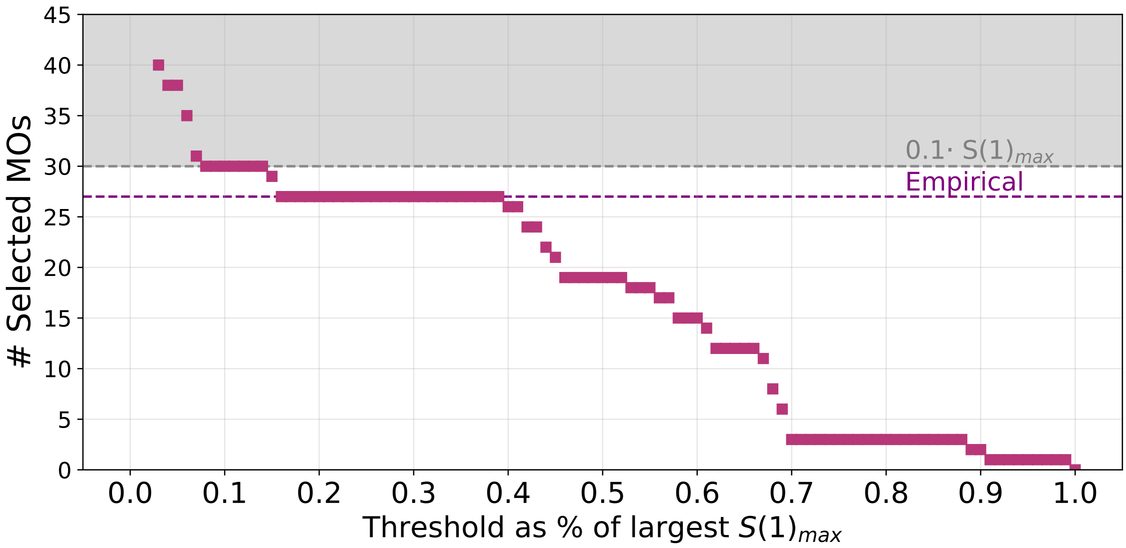

Figure 4.11: Threshold plot, defined according Stein et. al [130] , for Ruby-3. In this case, the pre-selection step with DMRG single orbital entropy leads to the selection of 30 MOs. The approximated wavefunction is obtained on a CAS(60,90) with bond dimension 200.

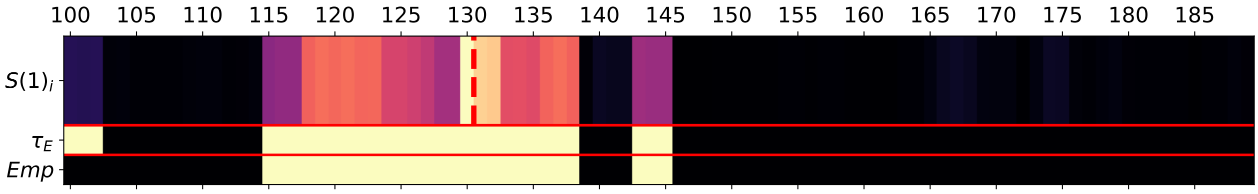

Figure 4.12: In this plot, the upper heatmap reports the value of the single orbital entropies, $S(1)_i$, for Ruby-3 MOs. In the second frame, the pre-selected orbitals with entropy above τ_E, and in the lower frame, the ones with entropy above 0.06. The dashed line draws the HOMO-LUMO gap, set at MO number 130.

orbitals from the bipyridine ligands and their corresponding π^* anti-bonding counterparts. Given the presence of three bipyridine ligands, three π bonding and three π^* anti-bonding orbital need to be included in the space. This leads to an initial active space of 11 orbitals, of which at least three are known to be virtual. However, in highly symmetric systems such as this one, the strong delocalization and mixing between metal d orbitals and ligand π orbitals complicates the a prior definition of the occupied or virtual space. For this reason, the active space selection is refined by separately inspecting the occupied and virtual manifolds, based on the dominant AO label desired of each orbital. According to literature source [169], a minimal yet satisfactory active space for this class of systems should contain at least 13 orbitals, exceeding the initially estimated 11. This extension is justified by the inclusion of two additional σ bonding orbitals.

Our analysis suggests that three of the five 4d orbitals of Ru(II) are occupied, while two are virtual. Notably, within the virtual space, as well as in the occupied space, a high degree of mixing between d- and π^*- type contributions is observed, which makes an unambiguous classification difficult. The refined target active space includes 13 orbitals, of which at least three are virtual and at least five are occupied. The remaining eight orbitals are selected based on their chemical relevance, and further analysis is required to determine their occupation status based on orbital character and wavefunction diagnostics. The desired AO labels are collected in Table 4.5.

ID	Cluster	AO Labels	O	V
L1	$\mathcal{C}_0(Ru)$	4d	3	2
L2	$\mathcal{C}_x\ (Bpy_x)$	[C_x $2p_\perp$, N_x $2p_\perp$, ...]	1	1
L3	$\mathcal{C}_1\ (Bpy_1)$	[C_1 $2p_y$, N_1 $2p_y$]	1	1
L4	$\mathcal{C}_2\ (Bpy_2) \cup \mathcal{C}_3\ (Bpy_3)$	[C_2 $2p_x$, N_2 $2p_x$, C_3 $2p_z$, N_3 $2p_z$]	1	1

Table 4.5: Ruby-3 AO-labels of interest for each of the cluster. The x in the labels is a placeholder for denoting all the cluster in $\{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3\}$. **O** and **V** are the number of orbitals selected for the specific AO labels.

The MOs appear to be very mixed and less localized (in term of AO contributions). This high mixing character is highlighted in Figure 4.13.

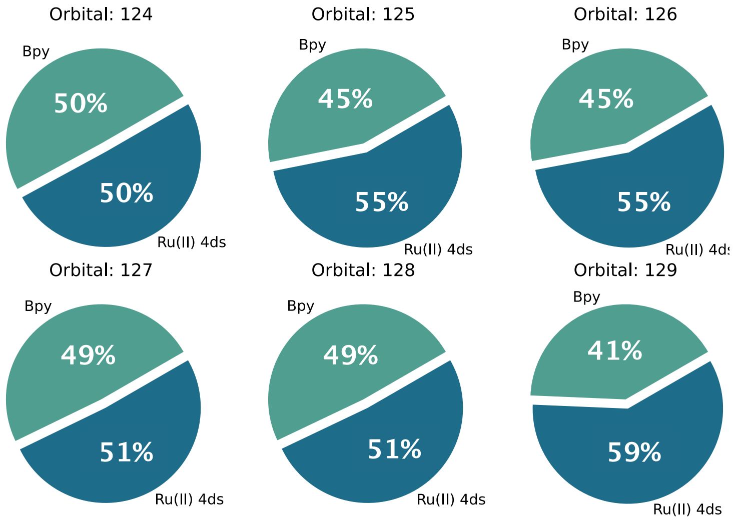

Figure 4.13: Molecular orbitals with highest Ru(II) 4d contribution. From here, it is possible to observe the composition of these orbitals, in particular, how the MOs are also highly composed of π-orbitals on the bipyridines.

In particular, for the occupied orbitals, there is no clear distinction between π and Ru(II) $4d$ orbitals, thus, the selection of the candidate orbitals in the occupied space has been defined selecting all the orbitals with $\epsilon_{[\mathcal{C}_0 \text{ Ru } 4d]} > 0.9$ (i.e. labels group L1). This selection leads to 6 *Ru 4d*s mixed with π-orbitals, defining a reasonable set of occupied orbitals. Then, for the three antibonding π^* orbitals, the required ones are composed respectively by a π-network on one ligand, L3, on two ligands (complementary to the first one), L4, and MOs spread on all three ligands, L2. In particular, L2 selects the MOs with carbon and nitrogen $2p_\perp$ contributions perpendicular to the plane of each respective ligand. From this, we select only the first MOs appearing. L3 and L4, instead, select two MOs whose complement each other, and even in this case, we select the first orbital for each label group.

Finally, to clearly describe MLCT, we needed to identify σ-orbitals. Unfortunately, from the explorative DMRG calculations and the AOs projections, no σ-orbital similar to the one in the reference paper was found. The σ-character is hard to identify using AOs projection, and from what we have observed, DMRG allows us to identify π-orbitals efficiently. Thus, by visual inspection of excluded molecular orbitals from the DMRG selection, we have identified two MOs that could potentially be the σ-orbitals used in the reference AS. In Table 4.6 we show the active space we have selected and investigated with CASSCF.

AS	Occupied								Virtual				
AS1	*110*	*111*	**124**	**125**	**126**	**127**	**128**	**129**	**130**	**131**	**132**	134	135
AS2 (FL)	**122**	**123**	**124**	**125**	**126**	**127**	**128**	**129**	**130**	**131**	**132**	133	134

Table 4.6: Candidate active orbital spaces (AS) selected for the Ruby-3 complex. **Bold** MOs number are referred to Ru(II) 4d orbitals, while *italics* for the σ-orbitals added in the post-processing screening.

The obtained active space is shown in Figure 4.14.

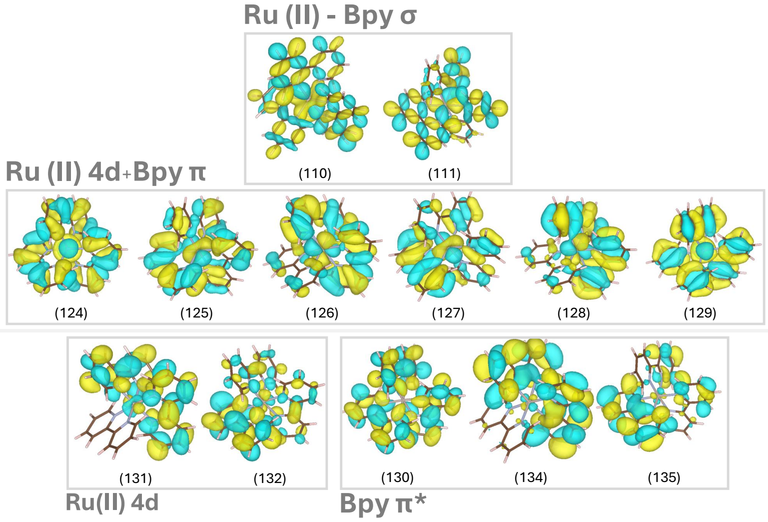

Figure 4.14: Candidate active space, AS1 in Table 4.6, identified with the presented procedures. The space is split between occupied mixed Ru(II) 4d and π orbitals, virtual Ru(II) 4d, virtual anti-bonding π^*, and the manually added σ-bonding orbitals.

Spin-Averaged CASSCF benchmark

As already pointed out for the previous complex, the scope of this preliminary investigation is, in fact, assessing our active space selection approach provides reliable results by comparing ourselves to reference active spaces. We performed preliminary SA-CASSCF calculations on 9 singlets and 9 triplets states to check the consistency between our active spaces and the reference. Because we compute results with slightly different simulation conditions, our energies are not exactly comparable to the ones in the reference article. Excitation energies with CASSCF and corresponding benchmarking DFT are reported in the Appendix in Section 3.

While both AS1 and AS2 have a balanced selection of orbitals and predict correctly the first singlet and triplet excitations to be predominantly composed by $dz^2 \rightarrow \pi^*$ transition, AS1 gives less mixed excitation state compositions. In particular, AS1 predicts the highest contribution to be $dz^2 \rightarrow \pi_1^*$, where π_1^* is the number 130 in Figure 4.14, in agreement with the reference paper [169]. All CASSCF, however, predict the correct energy order for T_1 being lower in energy than S1. AS1 results show very good agreement with B3LYP. For instance, for both S_1 and T_1 the energy differences with respect of B3LYP are of 0.058 eV and 0.004 eV, respectively. The AS2 (active orbitals around Fermi level) is considered not suitable for the type of physics that we are expected to simulate, in particular, it does not contain the σ-orbitals necessary to describe MLCT transitions. In comparison with TDDFT results, this space shows similar results as the ones computed with BP86 [183, 184] . This analysis led us to conclude AS1 to be the best candidate active space for the description of excitations in Ruby-3.

4.4.3 TLD-1433

For the TLD-1433 complex, an additional analysis of different possible conformations has been carried out, in order to correctly identify the ground state geometry. The conformations analyzed in the reference paper[170] are related to the orientation of the thiophene rings. This leads to two possible geometries with different energies, in particular, if the S atoms face the same direction or not. We found the structure with alternated rings is more stable. Also, the connection between the thiophene chain and the ligand coordinated to the Ru(II) atom can be modified. This describes how the H atom on the imidazole ring is oriented respectively to the S atom of the first thiophene ring. The H and S can point in the same direction or the opposite direction. In the reference paper[92, 170], only the first configuration has been analyzed, in the gas phase and water, for both the thiophene ring settings. As a preliminary study, in this work, a new local minimum has been found i.e. the structure with alternated ring and H facing in the opposite position w.r.t. the first sulfur ring. The ground-state geometries have been optimized with GAUSSIAN software[185] using PBE0 [186, 187] functional and def2-TZVP basis set using the Polarizable Continuum Model (PCM)[188, 189] solvent approach for water. All the ground-state geometry energies and analysis of the results can be found in the SI in Section E. For the RHF calculations, performed with PySCF, due to the increasing size of the system, the basis-set used is def2-SVDP [190] with cc-PVDZ-pp pseudo-potential for the Ru(II) atom. Additionally, the simulations for TLD-1433 have been carried out in water as solvent, using the PySCF implementation of the Conductor-PCM (C-PCM)[188, 189].

TLD-1433 contains, as substructure, the complete Ruby-3 complex, and it is asymmetrical as Trans-Cl. The atomic labeling chosen split the molecule in a similar way to the Trans-Cl Ru(II) complex, in particular, the interested clusters are the following: $\mathcal{C}_0$ is the Ru(II) atom, $\mathcal{C}_1$ is the bipyridine which extends into the thiophene chain, $\mathcal{C}_2$ is the imidazole ring connecting the bipyridine and the tail, $\mathcal{C}_3/\mathcal{C}_4/\mathcal{C}_5$ respectively each of the thiophenes, $\mathcal{C}_6$ / $\mathcal{C}_7$ the two remaining bipyridines. A coarser fragmentation can be defined by grouping all the components of the elongated ligand (bipyridine coordinated with the thiophene chain) in a unique cluster, $\mathcal{C}_1$, the two bipyridines as $\mathcal{C}_2$ and $\mathcal{C}_3$, and leaving the Ru(II) atom alone again, in $\mathcal{C}_0$. For this showcase, the second grouping has been chosen.

Thus, following the workflow as outlined in Section 4.3.1, after defining the desired atomic labeling, a mean-field calculation is performed and the initial-guess MOs are obtained. The initial search space, for the TLD-1433 molecule, $N_{\mathbf{A}}$ is 1559. The HOMO is the MO number 225, counting a total of 500 electrons. Working with the full orbital space to get the entropy-based preselection is impossible, thus, a first reduction of the search space is done. A reduced search space is defined by cutting out core orbitals and high-energy virtual MOs. The size of the Large-CAS, N_{LCAS}, is 100 MOs, leaving 185 orbitals as inactive core. The following setup leads to a DMRG with CAS(80,100), i.e. 80 electrons in 100 MOs.

Once the reference approximated DMRG wavefunction is computed, the entropy-based pre-

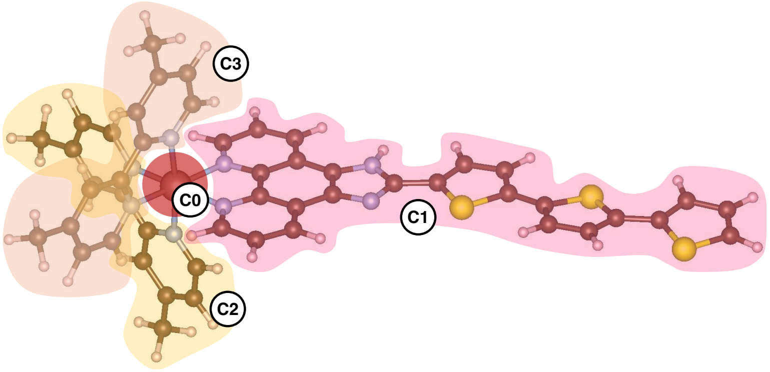

Figure 4.15: TLD-1433 molecule with highlighted cluster defined according to the coarse atom-labeling, as explained in this section. Each of the ligand (including the elongated one) is treated as an independent cluster, as the Ru(II) atom.

selection can be performed. The $S(1)_p$ for each MO in $\{\psi_p^{LCAS}\}$ is calculated using Equation 4.2. The maximum single orbital entropy found is $S(1)_{max}$=0.201, which corresponds to an entropy threshold of τ_E =0.0201 . The entropy-based selection allowed to cut down the initial reduced pool of 100 MOs to $N_E = 51$ pre-selected MOs. After the entropy selection, the coefficient matrix corresponding to high orbital entropy MOs is defined by $[\mathbf{C}_E^{\mathbf{A}}]$ with shape $N_{AO} \times N_E = 1559 \times 51$. In figures 4.16 and 4.17, it is possible to see how the number of MOs selected varies with different thresholds, in the former, and how the highly entropic MOs are spread in the molecular space, in the latter.

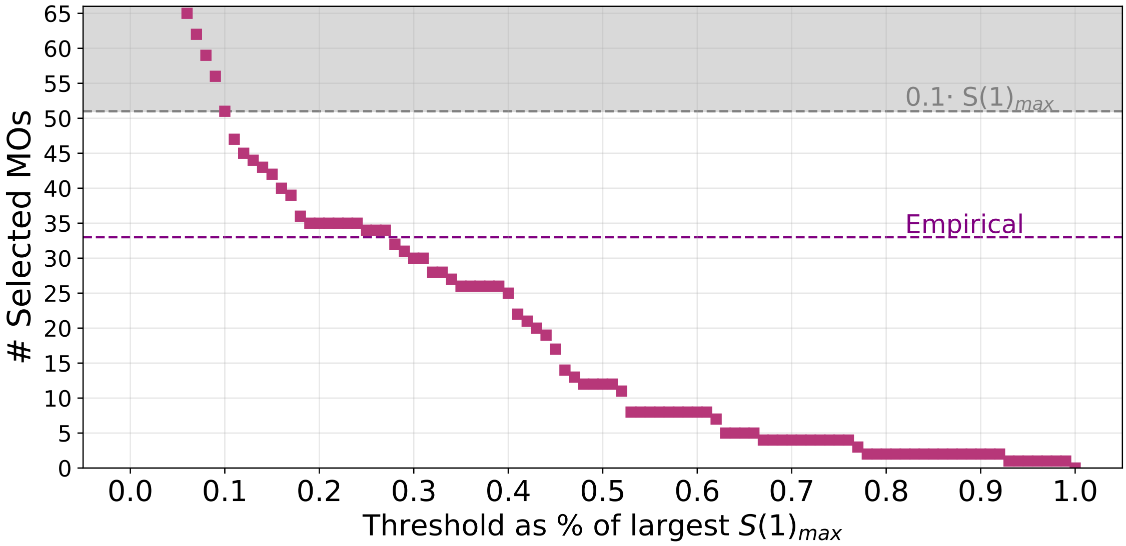

Figure 4.16: Threshold plot, defined according Stein et. al [130] , for TLD-1433. In this case, the pre-selection step with DMRG single orbital entropy leads to the selection of 51 MOs. The approximated wavefunction is obtained on a CAS(80,100) with bond dimension 500.

For this molecule, there is no reference in the literature about the active spaces that can be used to describe it using CASCI/CASSCF based method. Thus, a similar description of the active space defined for Trans-Cl is done, as first guess, and then enriching the representation scaling up to bigger active spaces, which may not be solvable using standard CASSCF/CASCI approaches. The objective here is to define a set of active spaces suitable to be used as effective Hamiltonian in

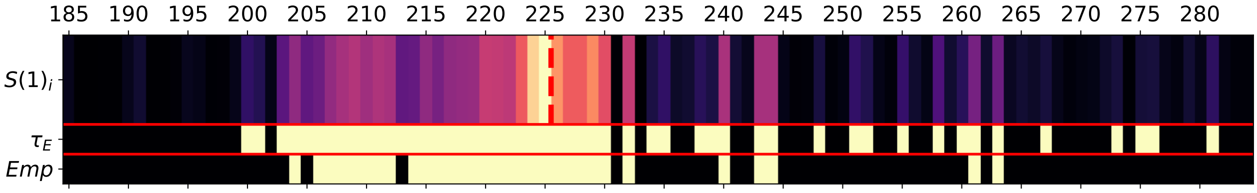

Figure 4.17: In this plot, the upper heatmap reports the value of the single orbital entropies, $S(1)_i$, for TLD-1433 MOs. In the second frame, the pre-selected orbital with entropy above τ_E, and in the lower frame, the one with entropy above 0.06. The dashed line draws the HOMO-LUMO gap, set at MO number 225.

a quantum computing application.

A balanced active space for describing Inter–Ligand (IL) and MLCT states should be composed of the 4d Ru(II) orbitals ($\mathcal{C}_0$), a balanced set of π/π^*-orbitals on the tail ($\mathcal{C}_1$), a set of π-orbitals coordinating the two non-substituted bipyradine ligands (referred to as $\mathcal{C}_2$ and $\mathcal{C}_3$) to the transition metal, and eventually a set of σ-orbitals for the description of the bonding between the thiophene tail and the coordinating bipyradine ligand.

ID	Cluster	AO Labels	O	V
L1	$\mathcal{C}_0$ (*Ru*)	4d	5	5
L2	$\mathcal{C}_1$ (*Tail*)	[C_1 2p_z, N_1 2p_z], [C_1 2p_z, S_1 3p_z]	1	1
L3	$\mathcal{C}_2$ (Bpy_1)	[C_2 2p, N_2 2p]	1	1
L4	$\mathcal{C}_3$ (Bpy_2)	[C_3 2p, N_3 2p]	1	1

Table 4.7: TLD-1433 AO-labels of interest for each of the cluster. **O** and **V** are the number of orbitals selected for the specific AO labels.

Following these guidelines on the type of MOs needed to be included in the active spaces, a set of size-scaling AO contributions can be defined. The resulting active spaces starts ideally from a **CAS(16,16)** composed by one π and one π^*-orbital from the $\mathcal{C}_2$ and $\mathcal{C}_3$ ligands, 5 occupied and 5 virtual Ru 4*d*s, and finally one π and one π^*-orbital from the tail. The desired active space follows a similar structure as the two previous systems, in which a balanced set of π/π^* dresses a good representative of Ru(II) 4*d*s orbitals. By further analyzing the molecular orbitals of the proposed active space above the composition of the final active spaces changed then slightly for the following reasons:

1. Due to high symmetry in the ligand orbitals, the same molecular orbital describing the π-bonding/antibonding orbital on $\mathcal{C}_2$ selected by AO projection also appears as the most contributing π-bonding/antibonding orbital for $\mathcal{C}_3$, and the same for the virtual space.

2. Due to the non-locality of RHF orbitals, the virtual orbitals appear to be defined as a combination of the AOs of our interest, even mixing $\mathcal{C}_2$ and $\mathcal{C}_3$ orbitals with Ru 4*ds*, thereby reducing the number of required orbitals.

Thus, the identified active space is reduce to be a CAS(14,13). The out-of-the-box active space contains one Rydberg state as virtual Ru(II) 4*d* contribution and the HOMO is not included. A refined active-space is instead obtained by increasing the threshold of selected orbitals to two for each set of π/π^* and to seven the number of Ru(II) 4*d*s. At this point, the refined space can be obtained by only considering the right amount of π/π^* orbitals including the HOMO in the occupied

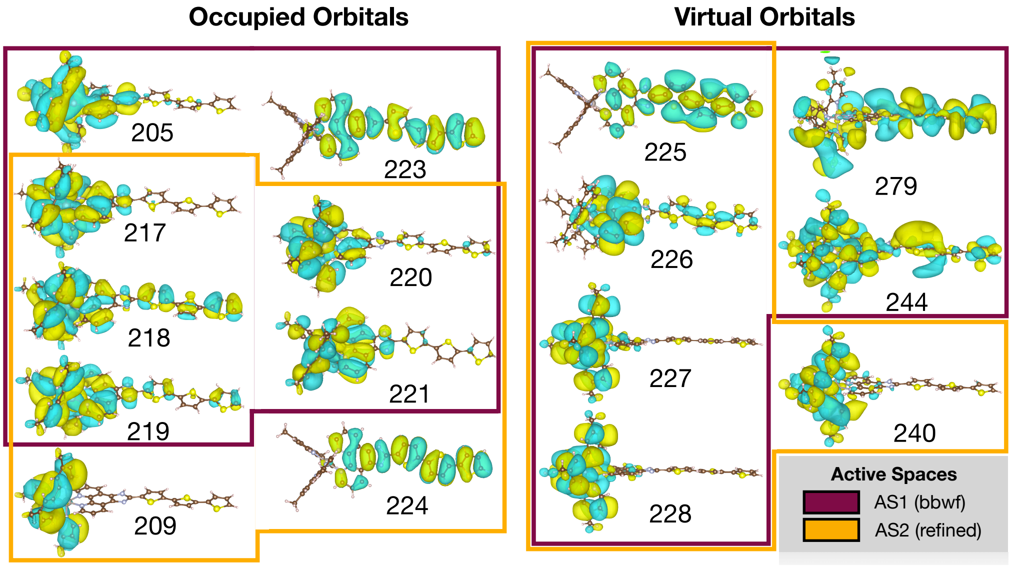

Figure 4.18: Pictorial representation of the TLD-1433 active spaces found and listed in Table 4.8

and removing the Rydberg from the virtuals. The defined active spaces defined for TLD-1433 are collected in Table 4.8. In Figure 4.18, we show the active orbitals selected in the active spaces identified.

AS	Occupied	Virtual
AS1 (bbwf) CAS(14,13)	205 217 218 219 220 221 223	225 226 227 228 244 **279**
AS2 (refined) CAS(14,13)	209 217 218 219 220 221 224	225 226 227 228 240 244

Table 4.8: Candidate active orbital spaces (AS) selected for the TLD-1433 complex. The obtained active-space corresponds to a CAS(14,13).

For TLD-1433, the calculations on the obtained active space are done with DMRG-CI due to highly computational cost required to run PCM-CASSCF calculation with PySCF. As explained in the computational setup in Appendix Section 1, the DMRG calculations have been carried out solving a state-average DMRG-CI on two singlet roots and then refining the result with two independent state-specific DMRG-CI. For the triplet state instead, a single DMRG-CI with target `spin=2` has been performed. After the converged MPS is found, the solvation correction is added by contracting the 1-RDM and the solvent potential.

AS	S_0	S_1	$\Delta(S_1 - S_0)$	T_1	$\Delta(T_1 - S_0)$	$\Delta(S_1 - T_1)$
AS1	-3597.605785	-3597.432341	4.719 (0.173444)	-3597.461174	3.935 (0.14461)	0.7846 (0.028833)
AS2	-3597.605786	-3597.432344	4.719 (0.173441)	-3597.461178	3.935 (0.14461)	0.785 (0.028834)

Table 4.9: S_0, S_1, and T_1 obtained from a PCM-DMRGCI calculation with CAS(14,13). Energies are given in Hartree. The gaps are given in eV and in brackets are reported the conversion in Hartree.

Given the absence of established computational benchmarks for this molecule based on wave function methods, we are also conducting a thorough DFT investigation, the results of which are provided in the Appedix Section 4. This detailed analysis offers a reliable starting point and a potential reference for future high-accuracy calculations.

4.5 Discussion and final remarks on active space selection

The AEGISS workflow proposed in this study offers a significant advancement in the semi-automated selection of active spaces for quantum chemistry, especially in the context of quantum computing. Its performance arises from the synergistic integration of two selection strategies: entropy-based metrics and atomic-orbital projections. By combining the physical rigor of orbital entanglement (via single-orbital entropy from DMRG wavefunctions) with chemically intuitive atomic-orbital contributions, the method effectively captures both static correlation and chemical specificity. This hybridization allows for the generation of active spaces that are both physically meaningful and tailored to the electronic transitions of interest, which is crucial for systems like transition-metal complexes where the interplay of metal-centered and ligand-based orbitals governs reactivity and excited-state dynamics.

Despite the workflow's attempt to reduce human intervention, certain stages, e.g. atomic orbital labeling and post-selection visual screening, still require chemical intuition. This human-in-the-loop aspect, while a strength in tailoring results, currently prevents full automation and could introduce subjectivity or variability. Future development should aim to mitigate these user-dependent bottlenecks. To reduce the complexity in the thresholding, a possible revision might be focused around the creation of a properly weighted measure that allows to have consistent values of AO projection with drastically different atoms in the clusters. Potential avenues for advancement include the integration of machine learning techniques to predict high-entropy orbitals from system topology or chemical descriptors, thereby reducing the dependence on full wavefunction sampling in cases where even approximate wavefunction approaches are computationally prohibitive. Prior studies addressing related challenges may provide a valuable foundation for guiding future developments in this direction [191]. Additionally, incorporating localization schemes or fragment-based embedding into the AO projection phase could further streamline the process and improve scalability. Another improvement might be related to the definition of the initial DMRG window, for which, relevant canonical HF orbitals can be too spread and fall outside the selected size. Correlated methods Natural Orbitals, e.g. MP2 or CCSD, can be used as starting guess from which the entropy preselection is performed. Additionally, the measure for the correlation-based pre-screening could also be defined according to other measures, analyzed by Ding et al., [192] and Evangelista [193]. On this line, during the preparation of this manuscript, another automated active-space procedure, namely Active-Space Finder (ASF) [194], has been presented and already employed. The procedure is similar to the one proposed in this manuscript, but differently from the present, it uses a clustering-like procedure on top of an approximated DMRG to defined the active spaces. It would be interesting to compare and, if possible, combine the two procedures [195].

Importantly, the AEGISS framework is well-positioned to interface with quantum computing workflows. Its capacity to generate state-specific and chemically adaptive active spaces aligns naturally with the needs of quantum algorithms like VQE [196] or Quantum Phase–Estimation (QPE) [3, 197], which require compact but expressive Hamiltonians. The method's output, a more chemically meaningful and entropy-aware active space, provides an ideal input for qubit-efficient encoding of molecular systems, as long as the effective Hamiltonian size is currently a bottleneck in quantum simulations. Moreover, the entropy-based pre-screening can be leveraged to prioritize orbital inclusion based on entanglement ranking, guiding qubit allocation and circuit depth considerations.

With the active space now defined in a principled way, we are ready to address the problem of preparing high-quality electronic states for hybrid quantum simulations. This task is going to be addressed in the next chapter.

Chapter 5

Quantum Information driven ansatz for State preparation

Despite its promise, VQE faces challenges in scaling to larger systems due to the expansion of the parameter space and the difficulties in implementing long circuits describing the variational wavefunctions and operators on noisy quantum devices. To overcome these hurdles, various improved algorithms have been proposed [79, 81, 85, 198]. In this respect, a fundamental role is led by the choice of the parametrized wavefunction that is going to be used as a trial wavefunction together with the VQE procedure. The shape and the structure of the ansatz are non-trivial and system-dependent, and in general, two types of ansätze can be defined. The first class exploits wavefunctions directly constructed to leverage the characteristics of quantum hardware. This empirical approach, known as the Heuristic Ansatz [55, 70–72, 199], comprises repetitions of blocks of parametrized rotations and entanglement gates. It is designed without relying on information about the physical system, focusing solely on exploiting the quantum hardware's capabilities. While the Heuristic Ansatz better utilizes the quantum hardware, it comes at the cost of losing the physical meaning associated with the wavefunction interpretation, with the advantage of considering significantly shallower circuits, providing a potential avenue to address scalability concerns. In contrast, the second family of ansätze involves translating classical Quantum Chemistry methods into the language of quantum computation. Regardless of the approach used to construct the ansätze, increasing complexity in PQC results in longer and deeper circuits, which not only heightens the risk of error accumulation but also contributes to the emergence of Barren Plateaus [67], an issue characterized by an exponentially flat optimization landscape. While adjustments in optimization techniques or error mitigation strategies do not directly resolve barren plateaus, employing shallow and adaptively structured ansätze has proven to be an effective countermeasure. Example of different ansätz have been presented in Section 3.3, and in this chapter, we will focus on the Hardware–Efficient Ansatz (HEA) family of circuits.

Shallow circuits that can also encode chemically relevant information and being build, for example, following the Mutual Information present in the system, as done with *Quantum Information Driven Ansatz (QIDA)* [200]. This approach allows to define quickly initial guess with a limited number of CNOTs. Obtaining energetically correct results with a good overlap with the ground state is not only useful as a starting guess for more complex ansatz construction protocol or other VQE-based procedures, but also as a trial wavefunction from which samples can be extracted. Recently, different works in which a non-VQE-based, mainly exploiting the quantum version of a Selected Configuration Interaction (SCI) approach, called Quantum SCI (QSCI) [201], Quantum Subspace Diagonalization (QSD) [202], or Sample-Based Quantum Diagonalization (SQD) [203], have been proving relevant industrial results and application of Quantum Algorithm to quantum utility-scale [203, 204]. In these works, the principal actor is the trial wavefunction from which the samples are drawn, which, for example, can be either a fixed Quantum Number Preserving (QNP) circuit [205] or a chemically inspired hardware-efficient Local Unitary Coupled Jastrow (LUCJ) [203]. Despite these advantages, being able to construct compact wavefunctions that can build effective trial wavefunctions that can be used to sample relevant determinants with the correct quantum symmetries is still a challenge.

Therefore, having established a well-defined active space, the next challenge lies in preparing accurate electronic states efficiently. This chapter presents the *Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA)*, a warm-start strategy that leverages information-theoretic filtering to construct compact variational forms and accelerate the convergence of hybrid VQE simulations. Starting from the idea behind the QIDA approach, MultiQIDA builds and optimizes ansätze layer by layer inspired by low-level preliminary classical calculations carried out on the system. Our state-preparation method aims to recover partially the missing correlation by exploiting a layered-structured ansätze, where each layer is selected based on the Quantum Mutual Information (QMI) values. Multi-QIDA extends the single threshold QIDA approach by including qubit-pairs that still present mid-high QMI values, that are excluded by the original method. Combined with a repeated iterative VQE routine, Multi-QIDA allows us to obtain effective shallow circuits that surpass the generic ladder-fashion ansatz in terms of performance and convergence properties. In addition, we explore the usage of a different correlator, namely an $SO(4)$ 2-qubit gate, which has several advantages over the traditional CNOT.

The chapter is therefore composed in the following way: In Section 5.1 we defined the original QIDA method and how to extrapolate Quantum Mutual Information (QMI) matrices from different sources, DMRG in Section 5.1.1 and SparQ [163] in Section 5.1.2. Then, in Section 5.2, we introduce our Multi-QIDA approach in each of its component: the extraction of single QIDA layers, in Section 5.2.1, how to reduce the number of correlators, in Section 5.2.2, the new *SO(4)* correlators, in Section 5.2.3, the costruction of each layer, in Section 5.2.4, and finally, the iterative-VQE procedure, in Section 5.2.5 In term of results, we have two sections split according to the type of system treated: Section 5.3 for Strongly–Correlated systems, and Section 5.4 for general molecules. A conclusive discussion is reporte in Section 5.5

5.1 Quantum Information Driven Ansatz (QIDA)

Quantum Information Driven Ansatz (QIDA) [200] is a heuristic ansatz approach used to design PQC with a topology that reflects the correlations that are contained in the quantum systems. After the calculation of the Quantum Mutual Information (QMI) matrix I, coupling pairs are identified by fixing a threshold $\mu \in [0, 1)$, splitting the qubits pairs into two separate sets:

$$\begin{cases} (i,j) \text{ is selected} & \text{if } I_{ij} \geq \mu, \\ (i,j) \text{ is discarded} & \text{if } I_{ij} < \mu. \end{cases} \tag{5.1}$$

The *selected* qubit pairs are used to design the empirical ansatz. In particular, the entangling block of the VQE is composed of two-qubit gates between each qubit couple $i, j \in$ *selected*, i.e. the coupling with QMI value greater or equal to the threshold chosen $I_{ij} > \mu$. In ref. [200], the entangling gate used, called *correlator*, is the CNOT gate. Any other two-qubit quantum gate can be used to correlate the two qubits.

QIDA follows the idea that the tunable threshold, μ, can provide the desired amount of correlation required to approximate the quantum state connecting only the spots that are more relevant for the approximate QMI reference. The algorithm does not necessarily connect all the qubits. Empirical ladder-fashion ansatz links all the qubits giving the possibility to span the whole Hilbert space at the price of having a larger entangling gate count, which will result in much more expensive quantum circuits. Also, ladder-fashion ansatz does not carry any information about the physical problem that is going to be studied. On the contrary, With QIDA, the topology of the correlation in the quantum systems is obtained by a classical ground-state pre-computation from which the QMI map is defined.

5.1.1 QMI from DMRG for spin lattices

For this application, the QMI matrices are computed exploiting MPS and the DMRG algorithm, which as already been explained in the introductory on quantum chemistry (Chapter 2). In this case, the Hamiltonian we are trying to solve is the one defined by Equation 5.10. After running the DMRG algorithm, the converged wavefunction $T_{\sigma_1,\ldots,\sigma_N}$, i.e. the best approximation of the ground-state at a fixed bond dimension, is obtained, where $\sigma_1 \ldots \sigma_N$ are the physical indices of the MPS. To compute the terms used in the QMI Equation 1.25, the RDMs for each couple of sites are required. Given the matrix $A_{i,j}^{i'j'}$, which can also represent a density operator, the partial trace is obtained by summing over the indices that we want to trace out.

$$\mathrm{Tr}_i(A_{ij}^{i'j'}) = \sum_i A_{ij}^{ij'} = A_j^{j'} \tag{5.2}$$

To compute ρ_{ij} which is define as a rank-4 tensor $\rho_{\sigma_i\sigma_j}^{\sigma_i'\sigma_j'}$, we need to trace out all the sites $\sigma_s \neq \{\sigma_i, \sigma_j\}$, which is equivalent to compute

$$\rho_{\sigma_i\sigma_j}^{\sigma_i'\sigma_j'} = \mathrm{Tr}_{\{\sigma_s \neq \sigma_i, \sigma_j\}} |\Psi\rangle\langle\Psi| \tag{5.3}$$

As shown in Equation 5.2, to trace out some sites, we need to sum over those indices, for the two-sites RDM, it is required to run over all the indices except i and j. If the MPS representing the quantum state is defined as $|\Psi\rangle = T_{\sigma_1\ldots\sigma_N}$, its density operator is

$$\rho = |\Psi\rangle\langle\Psi| = T_{\sigma_1\ldots\sigma_N} T^\dagger_{\sigma_1'\ldots\sigma_N'} = \rho_{\sigma_1\ldots\sigma_N}^{\sigma_1'\ldots\sigma_N'} \tag{5.4}$$

Now it is easy to see how to compute ρ_{ij} starting from ρ for each pair (i, j):

$$\begin{aligned} \rho_{ij} &= \sum_{\sigma_s, s\neq i,j} T_{\sigma_1\ldots\sigma_i\ldots\sigma_j\ldots\sigma_N} T^\dagger_{\sigma_1\ldots\sigma_i'\ldots\sigma_j'\ldots\sigma_N} \\ &= \rho_{\sigma_i\sigma_j}^{\sigma_i'\sigma_j'}. \end{aligned} \tag{5.5}$$

Finally, from each $\rho_{i,j}$ we can also compute the one-site RDMs ρ_i and ρ_j by tracing out the other index of two indices that resulted from the previous contraction (i.e. j and i respectively), completing all the necessary terms needed to compute the QMI value as defined in Equation 1.25. Exploting the properties of canonicalization of the MPS, the retieval of the one- and two-qubits QMIs is efficient due to the contraction to identities of the sites with indices $a \neq i$ for 1-qubit RDM, and with $a < i$ and $b > j$ for the 2-qubit RDM.

5.1.2 QMI with SparQ for molecular systems

Sparse Quantum–State Analysis (SparQ) is an innovative tool designed to efficiently compute key quantum information theory metrics for wavefunctions that are sparse in their definition space [163]. SparQ is particularly focused on wavefunctions derived from Post-Hartree-Fock methods and employs the Jordan-Wigner transformation to map fermionic wavefunctions into the qubit space. For excitation-based wavefunctions, this is done by applying excitations to the Hatree-Fock SD by the following equation:

$$\hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l |HF\rangle = \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l \prod_{s \in HF} \hat{a}_s^\dagger |\varnothing\rangle, \tag{5.6}$$

where the indexes $i, j \in Virtual$, $k, l \in Occupied$, and s is the set of occupied orbitals that defines the Hartree-Fock SD. The fermionic to qubit mapping is then directly applied to the excitation operators. This approach leverages the inherent sparsity of these wavefunctions to perform efficient quantum information analysis. This makes it possible to handle larger and more complex chemical systems than traditional methods such as the Density Matrix Renormalization Group (DMRG) even if sacrificing at times the quality of the wavefunction compared to the latter method. The sparsity of the wavefunction strictly depends on the Post-HF method in use, however, it is mostly related to these methods exploiting only a relatively low number of excitations in the overall Fock space. Given a computational cost linear in the number of states, SparQ is able to handle any wavefunction given by the current methods, proving an invaluable tool for the aim of the present work.

5.2 Multi–QIDA workflow

The main focus of the Multi-QIDA approach is the building of a systematic ansatz adding, step by step, a layer that will recover the missing correlation that is not caught using the first QIDA layer [200] as a standalone ansatz. The implementation of Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA) is a combination of three steps, described in the following sections. The procedure is composed of the following steps:

1. *QMI calculations*: The first step is to build the approximate QMI matrix exploiting the DMRG or SparQ method, depending on the nature of the system, as defined in [163]. (Section 5.1.1 and 5.1.2)

2. *QIDA-layers constructions*: Without discerning between classical and quantum correlation, we spilt the qubit-pairs based on a selected set of QMI values, namely *finesse-ratios*. For each range of QMI we obtain a QIDA-layer by performing a selection of only some relevant pairs. Therefore, it is a purely classical routine which, given an approximate QMI matrix, infers the adaptive layers of the VQE by sectioning the MI matrix by correlation levels.

3. *Circuit-composition*: the circuit made of the different entangler maps for each layer is constructed, with details varying based on the entangler block chosen, *CNOT* or *SO(4)*.

4. *Layer-wise incremental VQE*: Each QIDA-layer is independently optimized, and then the full circuit goes through a *relaxation* procedure. A varied version of the Layer-VQE [120].

Key aspects of Multi-QIDA's procedure are:

- **Layered Ansatz Construction**: Multi-QIDA constructs variational layers step-by-step, where each layer is informed by the QMI matrix, selecting qubit-pairs based on their QMI values. This incremental addition of layers allows for the capture of crucial correlations that single-threshold approaches might miss.

- **Efficient Resource Management**: The algorithm is designed to reduce the computational overhead typically associated with ladder-style heuristic ansatzes, especially standard Hardware Efficient Ansätze. By selectively entangling qubits with strong correlations, Multi-QIDA reduced the number of required entangling gates, effectively constructing shallower circuits without sacrificing accuracy. This approach could be particularly valuable for current quantum hardware, where circuit depth and gate count directly impact performance due to noise.

- **Improved Convergence and Accuracy**: The iterative approach embedded in Multi-QIDA allows for faster convergence to the ground-state energy with fewer optimization runs. It consistently outperforms traditional ladder ansatz methods by maintaining high precision with reduced mean energy deviation, demonstrating its effectiveness in calculating ground-state energies accurately. Benchmarks showed improvement in both energetic terms and in accuracy compared to other ladder ansatz.

- **Empirical Mitigation of Barren Plateaus compared**: Following the idea that a multi-layer construction of the ansatz may create a funnel in the parameter space that can guide the minimization process, we may argue that Multi-QIDA's iterative structure may also increase the probability of avoiding barren plateau. Barren plateaus are a common issue in variational quantum algorithms where the optimization landscape becomes flat, complicating the parameter optimization process. By breaking down the variational landscape into manageable steps and refining parameters in stages, Multi-QIDA achieves better results compared to Hardware–Efficient Ansatz (HEA) ladder ansatz in which the full parameters space is defined since the beginning.

In the next sections, we are going to guide you through all the component of the Multi-QIDA pipeline. The full Multi-QIDA version of the workflow can be found in Figure 5.1.

5.2.1 Chunking Procedure

The final goal of Multi-QIDA is to create a shallow-depth circuit for state preparation, improving the results obtained by the original QIDA method, in order to include a wider spectrum of different qubit-pair correlation. In the previous section, we have seen how QMI matrices can be obtained, either constructed from a DMRG MPS solution, or through SparQ, for spin-lattices and molecular system, respectively. Obviously, DMRG is absolutely suggested for obtaining QMI matrices

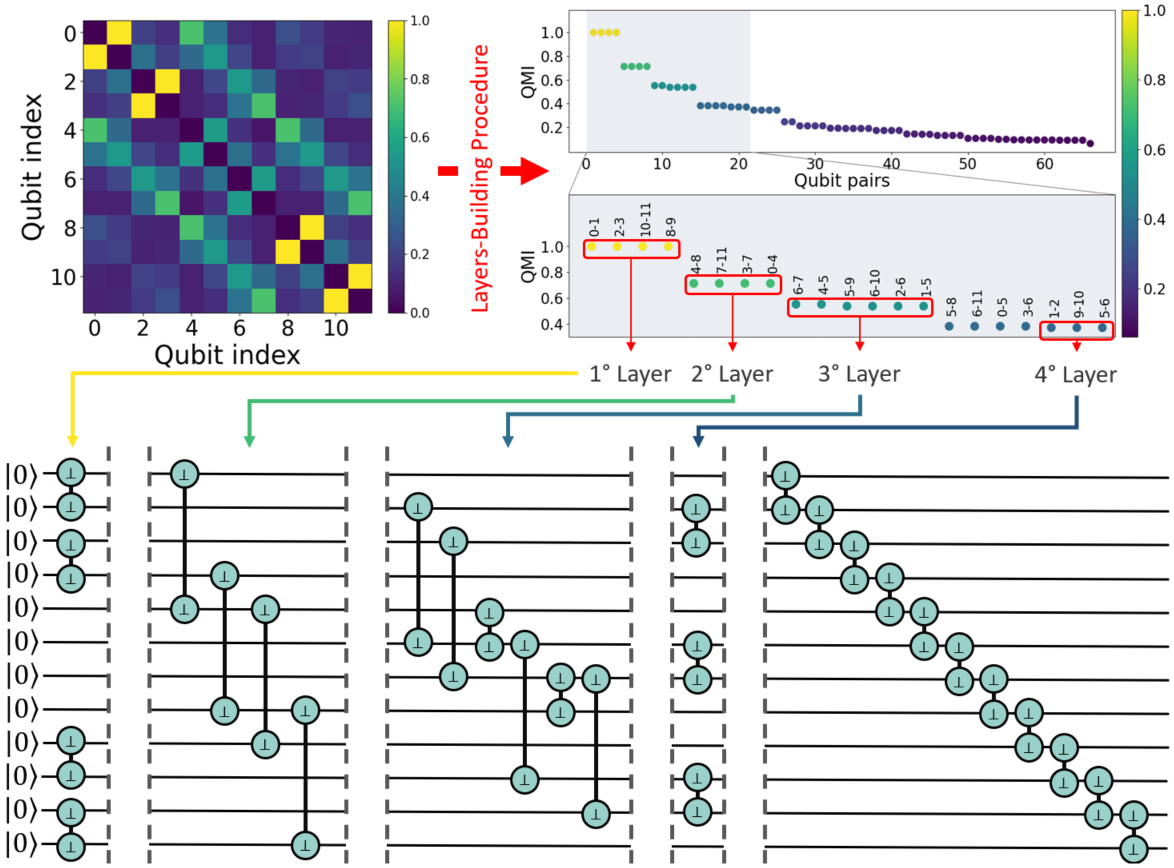

Figure 5.1: Pictorial representation of the Multi-QIDA method workflow. The reference wavefunction is used to build the QMI map on which is applied the Multi-QIDA Layers-Builder procedure as shown in Algorithms 1 and 2. The resulting collection of entangler maps is used to perform an Iterative-VQE, as defined in Algorithm 3.

in the case of strongly correlated system or unrestricted calculations, for which, RCISD or MP2 wavefunction may fail to represent the true ground state.

Once the QMI matrix I is generated, the first procedure that needs to be carried on is the definition of the different *QIDA-layers*. One single QIDA-layer is the application of the QIDA procedure i.e. selection of a mutual information cutoff μ and reduction of the number of correlators. With Multi-QIDA, instead, being a multi-threshold procedure, requires multiple applications of QIDA for different mutual information values. The selection is carried out on the list of descending QMI-value order qubit-pairs obtained from the QMI matrix $I_{u,v}$. This collection of thresholds $\bar{\mu}$, called *finesse-ratio*, represents how the full QMI matrix is split into different circuit layers. The selection of the finesse-ratio follows the structure of the distribution of qubit-pairs. The finesse-ratio is not computed automatically with a fixed step, it is instead empirically determined based on the distribution of the qubit-pairings. In general, the criterium is to select ranges of mutual information accordingly, for example, to the cluster that appears in the distribution. This clusters naturally appears in Spin lattices systems and the division has been chosen to follow these already present correlation structures. While for molecular system, by observing the distribution of the QMI spots, it can be noticed that it decays rapidly, concentrating on a minor group of highly/mid correlated spots (as Shown in Figure 5.2), and thus, as the lower the interval of correlation is chosen, the higher the number of pairings included in the chunks. Additional detail on the specific chunking and finesse-ratios selected for the system studied can be found in the result section.

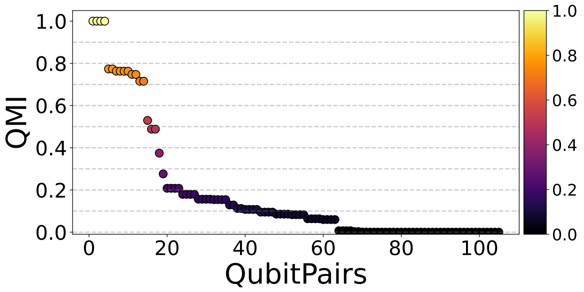

Figure 5.2: Example of the decreasing values of QMI for NH_3 molecular system. Details of the system are reported in the results section.

For each pair of consecutive finesse ratios, we define the range, l, in which the qubit-pairs are used to impose an edge on a graph G_l. This range corresponds to a Multi-QIDA layer. A QIDA entangling layer l-th is then composed by selecting the pairs q_i, q_j whose correlation I_{q_i,q_j} falls in the range $\bar{\mu}[l-1] > I_{q_i,q_j} \geq \bar{\mu}[l]$. Such pairs are then added to the candidate entangling list of the l-th layer, which represents the potential pairs to which then impose a correlator. At this point, a reduction of the number of correlators included in the mutual information chunk is going to be applied.

5.2.2 Selection criteria

As stated in [206, 207], one of the main components of the construction of PQC following the Multi-QIDA method, is the request for a selection criterion for the reduction of the number of entangling pairs that fall into each layer. The objective of this procedure is to exploit cross-entanglement built among the Multi-QIDA layer in order to reduce the number of correlators. In the two manuscripts[206, 207], different *selection criteria* have been defined, depending on the type of system to be solved. In particular, for Spin-lattices the criteria stops the production of layers once all the qubits have been touched at least one time by a correlator, and avoid the connection of pairs of qubits if already reachable by each other. For molecular system instead, the construction is more involved and exploits graph theory to build the reduced layers. Here, we report the specific procedure for each selection criteria.

Reduction criteria for Spin-lattices

For Spin-lattice the criteria according which a correlator, or set of correlators, is added to the circuit is by *coverage* i.e. if the correlators is trying to connect two qubits that are already connected, directly or exploiting indirect paths, the correlator is discarded. Initially, all the qubits are in the set of *free qubits*, meaning they have not been used in any entangling process and defining a disconnected graph using the qubits as nodes. A QIDA entangling layer l-th is then composed by selecting the pairs q_i, q_j whose correlation I_{q_i,q_j} falls in the range $\bar{\mu}[l-1] > I_{q_i,q_j} \geq \bar{\mu}[l]$. Such pairs are then added to the candidate entangling list of the l-th layer, which represents the potential pairs to which then impose a correlator. The criterion according to which the entangling gates are added to the current layer l is defined as follows: For each pair I_{q_i,q_j}, we check if at least on the two qubits q_i and q_j is contained in *free*, then the correlator associated to the pair is added to the

l-th layer, and the corresponding nodes are connected on the graph by an edge. Additionally, if the two qubits q_i and q_j are not in *free* and they are not connected in the graph, the correlator is added in the l-th layer. This means that the two qubits are contained in two different subset of qubits, and there is no path (or combination of previous layers correlators) that connects them. This procedure is iterated again until the set of free qubits is empty and all the qubits belong to the same unique subset, i.e. they are all connected. In other terms, the layers are added until all the qubits in the lattice have been covered by at least one entangling layer and they are reachable to each other by exploiting a combination of correlators chosen among all the layers. At the end of this building procedure, we define the entire collection of layers by L, which is then passed to the ansatz construction procedure. Algorithm 1 schematizes the procedure just explained.

Algorithm 1 Schematic outline of the Multi-QIDA Layers-builder

```
Require: QMI matrix $I_{ij}$, $\bar{\mu}$, $N_{qubits} > 0$
Ensure: List of entangling map $L$
$L \leftarrow$ empty list
$free \leftarrow \{n | \forall n \in N_{qubits}\}$
$G \leftarrow Disconnected\_Graph(free)$
$t \leftarrow 1$
$all\_connected = false$
while $(free \neq \varnothing) \wedge \neg(all\_connected)$ do
    $l \leftarrow$ empty list
    for $q_i, q_j \in \{\forall I_{i,j} : \bar{\mu}[m] > I_{i,j} \geq \bar{\mu}[m+1]\}$ do
        if $\neg(G.connected(q_i, q_j))$ then
            $l.add([q_i, q_j])$
            $G.add\_edge(q_i, q_j)$
            $free.remove(q_i, q_j)$
    $L.append(l)$
    $t \leftarrow t - \mu$
    $all\_connected \leftarrow G.is\_connected()$
$L.append($ladder_ansatz$(N_{qubit}))$
return $L$
```

Reduction criteria for molecular systems

For tackling the molecular system, we have exploited *Minimum Spanning Trees* (mST) to select which correlators are going to be used in each Multi-QIDA layer. We start by defining a *weighted graph*, $G = (V, E, w)$, where V is the set of the vertices, $E \in V \times V$ is the set of the edges, and $w : E \to \mathbb{R}$ is a function that maps each edge $e_i \in E$ to a real valued weight w_{e_i}. We can define a MST, $T \subseteq E$, which is a subset of the edges of the graph G that satisfies the following properties:

- *Vertices Span*: The subset of the edges T covers all the vertices V. Formally, for each pair of vertices $u, v \in V$, there exists a path in T that allows to reach v starting from u, and vice versa.
- *Tree structure*: T is an acyclic connected graph.
- *Total Weight Minimizing*: Given the weight function w, the sum of the weight in the subset T is minimized. Thus, if $T = \{e_1, e_2, \cdots, e_{|V|-1}\}$, then the total weight $w(T) = \sum_{e \in T} w(e)$

is the lowest possible value among all the available spanning trees that can be built from the graph G.

In the same way, a *Maximum Spanning Tree* (MST) can be defined following the previous properties by changing only the total weight values, which in this case, it has to be maximized. Thus, an MST is a subset $T \subseteq E$ that spans all the vertices V, it forms a tree, and the total weight $w(T)$ is maximized.

In this work, two different weight function has been used:

Maximum Correlation Spanning Tree (MCST): for each edge e the associated weight w_e is defined directly as the QMI value of the two vertices v, u that they are connected by the edge

$$w_e = w((u,v)) = I_{v,u}.$$

Thus, with this type of weight function, we are going to select the MST that collects the highest amount of correlation.

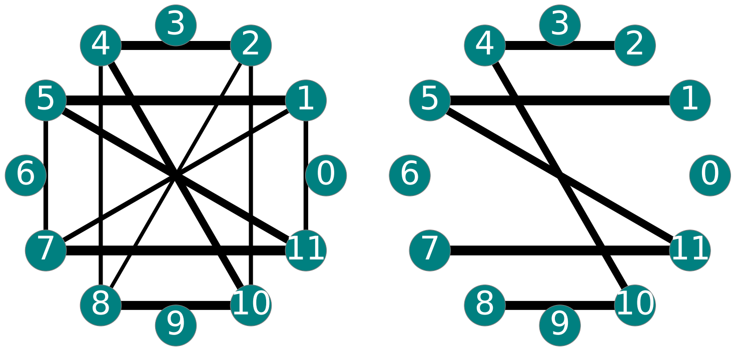

(a) **(b)**

Figure 5.3: Pictorial representation of the action of the correlators reduction. The width of each edge corresponds to the value of QMI between the two qubits, I_{uv}. In the two figures: (a) Original layer without any reduction applied. (b) Maximum correlation reduction applied leading to a QIDA-layer composed of the remaining edges.

Distance Reduction Spanning Tree (DRST): For each edge e, the relative weight is defined as

$$w_e = w((u,v)) = d(u,v),$$

where $d(u,v)$ is a *topology-based distance* function. The distance term is defined as the topological distance between qubits u and v, which can be seen as the number of edges that are included in the shortest path to connect these two qubits. From another point of view, this distance $d(u,v)$ can be defined as the minimum number of 2-qubits SWAP gates needed to make qubits u and j next-neighbour. For this specific case, the topology is linear, thus the distance function is defined as $|u-v|$. This weight function implemented on linear topology corresponds to the *empirical reduction* used in [200], i.e. the objective is to consider qubit-pairs as close as possible to the diagonal. Thus, it is required to build a Minimum Spanning Tree.

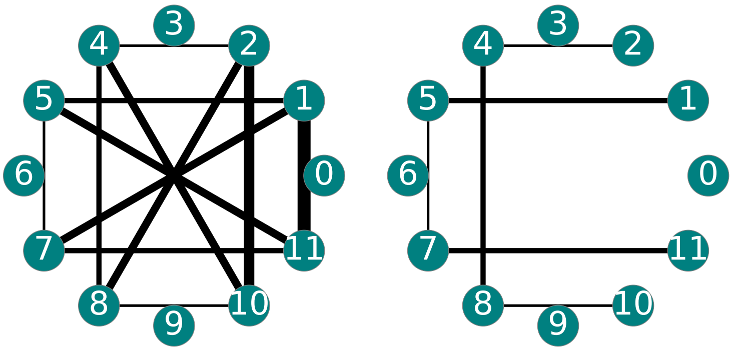

(a) **(b)**

Figure 5.4: Pictorial representation of the action of the correlators reduction. The width of each edge corresponds to the topological distance between the two qubits, $d(u, v)$. In the two figures: (a) Original layer without any reduction applied. (b) Empirical reduction leads to a QIDA-layer composed only of the identified edges.

In Figures 5.3 and 5.4, two examples of correlation reduction are shown. In particular, they represent how a candidate set of entangling gates, image (a) in both figures, gets reduced by the application of the selection criteria. The weights associated with each edge are defined randomly.

Algorithm 2 Schematic outline of the Multi-QIDA Layers-builder

```
Require: $I_{ij}$, $\bar{\mu}$, $N_{qubits}$, $w(\cdot,\cdot)$
Ensure: List of entangling map $L$
$L \leftarrow$ empty list
$m \leftarrow 0$
for $m \in [0, \cdots, len(\bar{\mu}) - 1]$ do
    $G \leftarrow (V = \{1, \cdots, N\}, E = \emptyset)$
    for $q_u, q_v \in \{\forall I_{u,v} : \bar{\mu}[m] > I_{u,v} \geq \bar{\mu}[m+1]\}$ do
        $G.add\_edge(q_u, q_v, w(u, v))$
    $T \leftarrow ComputeMST(G)$
    $L.append(T.get\_edges())$
return $L$
```

Selected all the pairings that fall in the selected QMI value range, for each u, v in the chunk, we define a weighted edge that connects vertex u to vertex v, with weight $w(u, v)$. The weight function $w(\cdot,\cdot)$ depends on the selection criteria preferred. Once all the edges are inserted in the graph, we create the MST (or mST), T_l, concerning the cost function used to reduce the number of correlators. From T_l, we can now retrieve the collection of the edges $E(T_l) \subseteq E(G_l)$ that will compose the entangler map for the l-th layer. At each step, the graph G_l is reset, ensuring that only the correlators contained within a given chunk are considered. This guarantees that the minimum spanning tree (MST) is constructed solely from the correlators relevant to that chunk, thereby preventing any edges from previous layers from being reused or exploited.For each tree T_l, the list of edges forming the MST (denoted as mST) is translated into a QIDA-layer. All QIDA-layers are then assembled into a complete set of layers, denoted as L . This collection L is subsequently passed to the layer-wise iterative VQE algorithm, which uses it as a blueprint to construct the adaptive quantum circuit. Algorithm 2 schematizes and generalized the procedure just explained, without assuming any selection criteria.

5.2.3 SO(4) Gates

Our approach has been also tested we considered a more general 2-qubit correlator with respect to the CNOT gate. This correlator can generate any unitary transformation belonging to the $\mathbf{SO}(4)$ group, and we call it *SO(4)* gate.

Differently from CNOTs-based ansätze, the parametrized *SO(4)* gates offer a tunable correlation, that can be chosen by the optimization.

On top of the variety of transformations that can be obtained by parametrized $\mathbf{SO}(4)$ matrices, the decision of this type of gates is also related to the nature of the system that we decided to treat. Given the time reversibility of the Heisenberg model Hamiltonians, we only need to include real transformations described by the group $\mathbf{O}(4)$ and, without lost of generality, we can restrict our attention to the group $\mathbf{SO}(4)$. To be explicit, the corresponding matrices consist of real numbers and these groups represent rotations in the four-dimensional space.

Following Theorem 3 in [208], every element of the group $\mathbf{SO}(4)$ can be realized by a circuit consisting of 12 elementary one-qubit gates and two CNOT gates.

A generic gate $U \in \mathbf{SO}(4)$ is composed by two generic one-qubit rotations $A, B \in \mathbf{SU}(2)$, four S gates and two R gates. It is also known that every matrix $M \in \mathbf{SU}(2)$ can be written as a composition of $R_z(\alpha)R_y(\theta)R_z(\beta)$ for some α, β and θ, while the R gate is defined as $R_y(\pi/2)$ and the S gate is obtained with $R_z(\pi/2)$. The U gate is then parametrized using the two sets of three parameters of gates A and B.

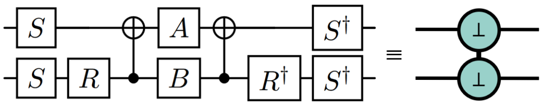

Figure 5.5: On the left, the circuit implementing a general $U \in \mathbf{SO}(4)$. The gates A and B are general $\mathbf{SU}(2)$ parametrized gates. On the right, the symbol we adopted in this work.

5.2.4 Circuit Building

Independently by how the previous procedure has been carried out, its final result is a list of QIDA-layer L. Before constructing the circuit, a final "ladder layer" is appended at the end of the L QIDA-layers. This layer arranges the remaining correlators in a top-down topology, allowing separate correlation groups to be interconnected. Functionally, the ladder layer acts as a final selection stage for qubit pairs that fall outside the specified finesse-ratio range. By applying MST-based selection to these leftover pairs, the resulting topology of this layer approximates that of a ladder both in term of structure and in CNOT cost.

As already introduced in the first section, the phenomenon on the barren plateau is related to the *circuit expressiveness*: the larger the variational space that the ansatz has to explore, the greater is the probability, when initialized according to a random distribution of the parameters, that the gradients of the cost function vanish. One of the possible solutions to avoid Barren Plateau is to employ short PQC [67]. By variationally optimizing only a portion of the search space with the restricted parameters, in some scenarios, shallow circuits present no flattened potential surfaces, with the drawback that the optimal solution could also not be obtained inside this constrained space. Following this concept, similarly to the Adapt-VQE [81], we decided to extend the shallow ansatz obtained by the QIDA method, adding layer-by-layer expressiveness to the PQC, without exploring the huge variational space at once. At each $l \neq 1$, we want the new layer L_l to be initialized to the identity to not disrupt all the previous $l-1$ rounds of optimizations once a new (yet) unoptimized layer is introduced.

To perform the proper construction, two different procedure are required based on the type of correlators we are using to build the whole circuit. The first implementation follows the standard combination of parametrized single-qubit rotation gates and $CNOT$ as entanglers [68], and it has been included in the analysis as a comparison to the next type of ansatz. In the following Multi-QIDA construction, this version of the Ansatz will be referred to as $QIDA^{CX}$. The second variant is obtained by entangling pairs of qubits using $\mathbf{SO}(4)$ gates, as in Fig. 5.5, which, as we have seen, guarantees the most general description of real two-qubit gates, such as Multi-QIDA implementation is referred in the following as $QIDA^{SO4}$.

CNOT V-shaped additional circuit

For the combination of R_y and CNOT, we can obtain the identity by imposing a V-shape structure to the entangling layer l, partly resembling the topology of a UCC excitation [74–76, 78]. The V-shape is obtained by a set of R_y rotations, the additional entangling layer, a central layer of R_y rotations, again the additional entangling layer but with reversed-order, completed with a last series of R_y rotations, as shown in Figure 5.6.

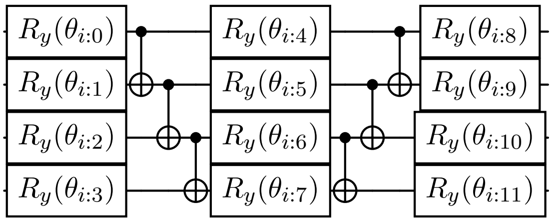

Figure 5.6: Circuit implementing a V-shape i-th additional layer with entangling map $[0,1]$,$[1,2]$,$[2,3]$. This circuit results into an identity, $\mathbb{1}$, when all the parameters $\theta_{i:j}$ are zero for all j.

SO(4) additional circuit

If the additional layer is built with $\mathbf{SO}(4)$, only $\mathbf{SO}(4)$ gates are used in the newly added circuit, and since the identity gate is already included in the parameterization of a general $\mathbf{SO}(4)$, shown in Figure 5.5, no further construction is needed to obtain it in this case. In summary, with both kinds of entangling maps (in Figure 5.7 and Figure 5.6), setting the parameters to zero guarantees an identity gate.

Parameter initialization to ϵ-Identity

In general, the initialization of the parameters of the additional layers with zeros, or in general with a circuit that results in an identity, ensures that the cost function being optimized remains unchanged after introducing this new layer. The parameter optimization can therefore proceed step by step, avoiding the risk of barrel plateau due to large expressivity. For spin-lattices, the initialization of the additional layer to identity, i.e. fully zero parameters,always leads the optimization to restart

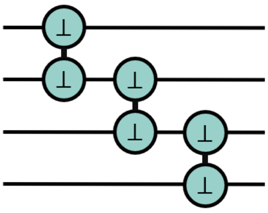

Figure 5.7: Additional layer with entangling map $[0,1]$,$[1,2]$,$[2,3]$ where each entangler is an $\mathbf{SO}(4)$ 2-qubit gate as describer in Figure5.5. This layer is initialized to identity, $\mathbb{1}$, at the beginning, with all the parameters of each entangler set to zero.

without any problem. In [207] istead, for dealing with molecular systems, the initialization to the identity of the additional layers turned out not being the right choice, due to the presence of local minima, preventing the optimization from proceeding. To overcome these limitations, we decided to initialize the new layer using an operator that is close to identity. This goal is achieved by adding additional layers which has a random offset from the identity. The set of initial parameters for the additional layer, $\bar{\theta}_l^i$, is created by randomly sampling values from a uniform distribution with mean 0 and standard deviation 0.1, i.e. $\bar{\theta}_l^i \in_R \mathcal{U}(0, 0.1)$. The value 0.1 has been empirically estimated by selecting the lowest value that allowed escaping the local minima of the previous QIDA-layer, while recovering in a few VQE iterations the previous energy after the addition of the successive layer.

5.2.5 Iterative–VQE

The selection of a suitable optimization routine is necessary when dealing with variational circuits in which the ansatz is built iteratively/adaptively and not directly optimizing the full wavefunction. As the name of this section suggests, we decided to perform the optimization of the full circuit iteratively along many steps, each of which will include two main phases: an optimization of the single layer L_l at the l-th step, and a relaxation procedure in which a global optimization of all the previous layers $\{0 \ldots, l-1\}$ is performed. To briefly resume a general procedure that is executed at each step l, exploiting non-fixed parameters quantum circuit, i.e. QC_{empty}:

1. Append the l-th layer to the QC_{empty};
2. Assign the optimal parameters up to the previous iteration $\bar{\theta}_{0,1,\ldots,l-1} = \bar{\theta}^*_{prev}$;
3. Initialize the parameters of the l-th layer to identity or an offset of the identity;
4. Find the optimal parameters of the l-th layer alone, $\bar{\theta}^*_l$;
5. Compose the total set of parameters up the the l-th layer as $\bar{\theta}_{tot} = \bar{\theta}^*_{prev} + \bar{\theta}^*_l$;
6. Use $\bar{\theta}_{tot}$ as starting parameters for a final VQE in which the variational wavefunction is defined by QC_{empty};
7. Once converged, the optimized set of the combined circuit, i.e. $\bar{\theta}^*_{tot}$ is obtained.

The starting parameters, denoted as $\bar{\theta}_l$, represent the unoptimized parameters associated with the l layer. With the vector of parameters $\bar{\theta}_l$, it is possible to define the unitary transformation to bring $|\Psi_l\rangle$ to the $l+1$-th state $|\Psi_{l+1}\rangle = U(\bar{\theta}_l)\,|\Psi_l\rangle$, where $|\Psi_l\rangle = U(\bar{\theta}^*_{prev})\,|\Psi_0\rangle$ and $\bar{\theta}^*_{prev}$ is the set of optimal parameters used up to the l-th layer. To optimize the $\bar{\theta}_l$ parameters we apply the VQE algorithm by minimizing

$$\langle \Psi_l | \, U^\dagger(\bar{\theta}_l) U^\dagger(\bar{\theta}^*_{prev}) H U(\bar{\theta}^*_{prev}) U(\bar{\theta}_l) \, | \Psi_l \rangle \, , \tag{5.7}$$

and we keep fixed $\bar{\theta}^*_{prev}$. After this first step VQE, we mark the optimal parameter vector by $\bar{\theta}^*_l$.

After the completion of optimization of the layer, L_l, the algorithm proceeds with a new VQE. In this second phase of the iteration l the algorithm encompasses all the layers $\{0 \ldots, l\}$ in a VQE that minimizes

$$\langle \Psi_0 | \, U^\dagger(\bar{\theta}_{tot}) H U(\bar{\theta}_{tot}) \, | \Psi_0 \rangle \, , \tag{5.8}$$

where $\bar{\theta}_{tot} = \bar{\theta}^*_{prev} + \bar{\theta}^*_l$. This optimization takes as a starting point the previous parameter vectors and gives as a result a new set of different parameter vectors, defined as $\bar{\theta}^*_{tot}$. This new set of parameters will be used as new fixed set of parameters for the next layer iteration. Algorithm 3 schematizes the Iterative-VQE procedures given an general entangler map,L, and the number of qubits in the circuit, N_{qubits}.

Algorithm 3 Iterative (Re)-Optimization routine

```
Require: $N_{qubits} > 0$, List of entangling map $L$
Ensure: Optimal parameters $\bar{\theta}^*$, Converged energy $E$
$QC_{empty}$ ←QuantumCircuit($N_{qubits}$)
for $l \in L$ do
    append($QC_{empty}, l$)
    $M$ ← n_params($QC_{empty}$)
    if $l$ is first layer then
        $\bar{\theta}_0 \leftarrow \theta_{0:0}, \ldots, \theta_{0:M} \in_R [0, 2\pi)$
        $\bar{\theta}^*_l$ ←VQE($QC_{empty}, \bar{\theta}_0$)
    else
        $\bar{\theta}_l \leftarrow \theta_{l:0}, \ldots, \theta_{l:M} = 0$ or $\bar{\theta}_l \in_R \mathcal{U}(0, 0.1)$
        $\bar{\theta}^*_l$ ←VQE($QC_{prev}, \bar{\theta}_l$)
        $\bar{\theta}_{tot} \leftarrow \bar{\theta}^*_{prev} + \bar{\theta}^*_l$
        $E_{tot}, \bar{\theta}^*_{tot}$ ←VQE($QC_{empty}, \bar{\theta}_l$)
    $\bar{\theta}^*_{prev} \leftarrow \bar{\theta}^*_{tot}$
    $QC_{prev}$ ←assign($QC_{empty}, \bar{\theta}^*_{prev}$)
return $E_{tot}, \bar{\theta}^*_{tot}$
```

5.3 Results for Strongly correlated systems

The first application in which Multi-QIDA has been tested is the strongly correlated systems. This section will focus first on the introduction of the physical system of Heisenberg-Spin lattices, then specific circuit construction, how the results between Multi-QIDA have been compared with standard HEA–ladder ansatz, and finally, the results.

5.3.1 Heisenberg Model

A general Heisenberg model Hamiltonian is defined as follows:

$$H = \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{J}_{i,j} \cdot \mathbf{S}_j + \sum_{i}^{N} \mathbf{h}_i \cdot \mathbf{S}_i, \tag{5.9}$$

where $\langle i, j \rangle$ indicates all nearest-neighbor spins pairs, $\mathbf{J}_{ij}$ is the interaction term between two adjacent spin sites, $\mathbf{S}_i$ and $\mathbf{S}_j$ are spins operators acting on sites i and j, respectively, and $\mathbf{h}_i$ is the local magnetic field acting on the i-th site, due to an external magnetic field or magnetic impurities. Expanding each spin operator $\mathbf{S}_i$ into the three components along x, y and z axis, the Hamiltonian H can be written as

$$\begin{aligned} H =& \sum_{\langle i,j \rangle} (J_{ij}^{xx} S_i^x S_j^x + J_{ij}^{yy} S_i^y S_j^y + J_{ij}^{zz} S_i^z S_j^z) \\ &+ h \sum_{i}^{N} S_i^z, \end{aligned} \tag{5.10}$$

where $\langle i, j \rangle$ indicates all nearest-neighbor spins in a 2D lattice with open boundary conditions, $J_{ij}^{\beta\beta}$ is the interaction term between two adjacent spin sites with $\beta \in \{x, y, z\}$, h is the external magnetic field and S_i^β is the β component of the $\frac{1}{2}$-Spin operator acting on the i-th spin site.

The Hamiltonian is written directly on the qubit space by converting each $\frac{1}{2}$-Spin operator into Pauli operator following

$$S_i^\beta = \frac{\hbar}{2} \sigma_i^\beta, \tag{5.11}$$

where $\sigma_i^\beta \in \{\sigma^x, \sigma^y, \sigma^z\}$ i.e one of the three Pauli matrices. The Hamiltonian to be measured on a quantum computer must be converted into a sum of Pauli strings. A Pauli string is defined as a tensor product of Pauli matrices. Each Pauli string corresponds to one of the coupling terms or the on-site term in Equation 5.10. For a general coupling term $J_{i,j}^{\beta\beta}$, the relative Pauli string is defined as

$$I_0 I_1 \ldots I_{i-1} \sigma_i^\beta I_{i+1} \ldots I_{j-1} \sigma_j^\beta I_{j+1} \ldots I_{N-1}, \tag{5.12}$$

where I_i are identity matrices.

To explore the validity of our Multi-QIDA approach we have considered different spin systems for the following Heisenberg Hamiltonian:

$$\begin{aligned} H =& \frac{J}{4} \sum_{\langle i,j \rangle} [\Delta(\sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y) + \sigma_i^z \sigma_j^z] - \\ &- \frac{h}{2} \sum_{i}^{N} \sigma_i^z. \end{aligned} \tag{5.13}$$

In Equation 5.13, the operator is defined in terms of three parameters: $J = J_{ij}^{xx} = J_{ij}^{yy} = J_{ij}^{zz} > 0, \forall \text{all } i, j \in \{1, \ldots, N\}$, whereas J is the spin coupling interaction, Δ as the Anisotropy term, and the external magnetic field h. In Table 5.1, we collect all the information about the parameters used for each system configuration with relative exact energies, E_{exact}, as well as the energy of the classical Néel state, $E_{\text{Néel}}$. E_{exact} obtained from exact diagonalization of the Hamiltonian, while $E_{\text{Néel}}$ is obtained according to Equation 5.16.

Size	*#Qubits*	J	h	Δ	E_{exact}	$E_{\text{Néel}}$
3×3	9	1.0	0.0	1.0	-4.749327	-3.00
2×6	12	1.0	0.0	1.0	-6.603472	-4.00
3×4	12	1.0	0.0	1.0	-6.691680	-4.25
3×4	12	1.0	2.0	1.0	-9.508473	-4.25
3×4	12	1.0	0.0	0.66	-5.338751	-4.25
3×4	12	1.0	0.0	0.1	-4.272670	-4.25

Table 5.1: Summary table containing all the configurations used in the simulations. E_{exact} is the energy obtained by exact diagonalization of the Hamiltonian of the system. The energy of the reference state, $E_{\text{Néel}}$, is obtained following Equation 5.16.

5.3.2 QMI manipulation and layer building

Here, we describe the full procedure for obtaining the complete set of entangling maps that are needed to compose a single Multi-QIDA ansatz. Starting from the DMRG reference MPS, we created the corresponding QMI map. We compute each term of Equation 1.25 by exploiting properties of the MPS wavefunction that allow us to compute RMD matrices, defined in Equation 5.5, efficiently. In Figure 5.8a and in the upper panel of Figure 5.1 (left), we show the QMI map obtained for the 3×4 isotropic Heisenberg Hamiltonian using DMRG.

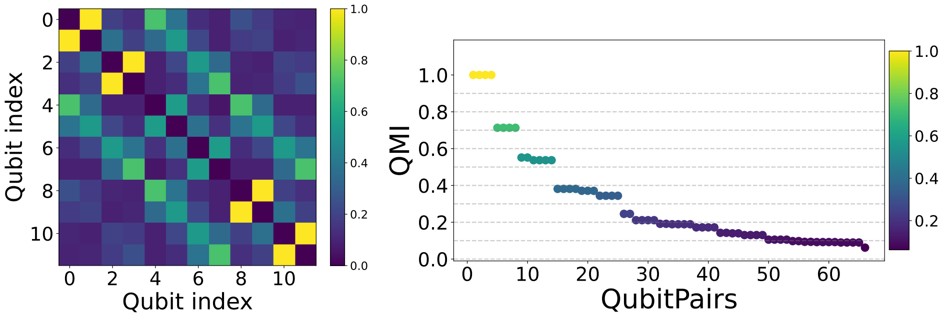

(a)

(b)

Figure 5.8: **(a)**: QMI map of 3×4 Isotropic Heisenberg model with $J=1.0$, $h=0$, and $\Delta=1.0$. Reference wavefunction obtained from converged DMRG calculation. **(b)**: Qubit-Pairing for the 3×4 Isotropic Heisenberg model with $J=1.0$, $h=0$, and $\Delta=1.0$. Dashed lines represent the values of the finesse-ratio $\bar{\mu}$ selected for this system.

From the QMI map, we have extracted the qubit-pairing and ordered them in descending mutual-information values. Then, we applied the layer-building procedure explained in Section 5.2.4, following Algorithm 1. In Figure 5.8b, we show a pictorial representation of the ordered set of coupling obtained by a QMI map. All the qubit-pairs are represented, and the color of each point is chosen accordingly with the value of QMI. The most correlated couplings are the ones with a QMI value near 1, light-tone color, followed by decreasing values of QMI corresponding to a color gradient toward a darker tone.

The first procedure of the Layers-builder procedure splits the correlators into QMI chunks, and this is done by imposing a finesse-ratio of $\bar{\mu}=[0.9,0.8,0.7,\ldots]$. With this finesse-ratio, it is clear that not all the chunks will contain qubit pairing e.g. in the ranges $[0.9, 0.8]$ or $[0.7, 0.6]$. In Figure 5.8b, it is possible to see how the QMI values are chunked according to the finesse-ratio selected. The QMI pairs that fall inside two dotted lines are the candidate coupling that may be selected to be inserted inside a QIDA-layer. Thus, it encloses also qubit-pairs that will be removed in the second phase of the Multi-QIDA layers-builder algorithm. These pairings are the ones that try to connect qubits that are already connected by the previous layer and thus, they are reachable to each other by exploiting the cross-correlation of the previous QIDA-layers.

From the QMI chunks obtained in the first procedure, we started building each QIDA-layer by applying the selection criteria explained in Section 5.2.4. Algorithm 1 stopped after the addition of the 7-th range of QMI values, i.e. 30% of correlation, which covered and connected all the 12 qubits.

The isotropic 3 × 4 model Hamiltonian required 17 of the 66 available qubit-pairs. The 17 correlators are split into the following layers :

- First Layer (4 correlators):[0, 1], [2, 3], [8, 9], and [10, 11].
- Second Layer (4 correlators): [0, 4], [3, 7], [4, 8], and [7,11].
- Third Layer (6 correlators): [1, 5], [2, 6], [4, 5], [5, 9], [6, 7], and [6, 10].
- Fourth Layer (3 correlators): [1, 2], [5, 6], and [9, 10],

and these pairings, i.e. the selected ones, are identified by the rounded box in the lower pane of Figure 5.1. Following the above composition, we obtain a PQC composed of 52 CNOTs for $QIDA^{CX}$ and 56 CNOTs for $QIDA^{SO4}$ having in total 5 QIDA layers. We notice that the $\#CNOT$ used in $QIDA^{CX}$ and $QIDA^{SO4}$ differs only for the number of correlators involved in the 1-st layer. This is due to the fact that for the $QIDA^{CX}$ ansatz, all the layers except the first are composed of the V-shape configurations of the layer, which doubles the count of CNOTs, while $QIDA^{SO4}$ counts two CNOTs for each gate used, including the first layer.

For the other system, the reference QMI maps are shown in Figure B.1, while the relative qubit-pairing is available in Appendix B.2. The resulting Multi-QIDA layers can be found in Appendix B.1.

5.3.3 Metrics and measures

We employed various metrics and measures to compare performances across different systems. In addition to the straightforward metric of the number of CNOTs, $\#CNOT$, other measures included are Absolute Quantum Energy, Relative Quantum Energy, Minimum Absolute Energy Deviation, and Minimum Relative Energy Deviation. While $\#CNOT$ is an unambiguous definition, other measures need a brief explanation:

Absolute Quantum Energy (AQE_i) for the i-th VQE run, i.e. the energy obtained from the i-th VQE run and the one obtained from Exact Diagonalization, E_{exact}.

$$AQE_i = \frac{E_i}{E_{exact}} * 100. \tag{5.14}$$

Relative Quantum Energy (RQE_i) i.e. the ratio between the difference of the i-th VQE run and the Néel energy and the difference between the ED energy and the Néel energy:

$$RQE_i = \frac{|E_i - E_{\text{Néel}}|}{|E_{exact} - E_{\text{Néel}}|} * 100. \tag{5.15}$$

The Néel state represents the ground state of the corresponding classical system, and it is obtained by setting spin in an alternated configuration between spin-up, ↑, and spin-down, ↓. The energy associated with the Néel state, $E_{\text{Néel}}$, is a constant that scales with the number of pairings and the value of the coupling term J_{zz}, i.e.

$$E_{\text{Néel}} = \sum_{\langle i,j \rangle} J^{zz}_{i,j}. \tag{5.16}$$

All the values relative to E_{exact} and $E_{\text{Néel}}$ are shown in Table 5.1.
Minumum Absolute Energy Deviation (MAED) represents the average deviation of the AQE_i from the best achieved AQE i.e. AQE_{best}, denoting the AQE value of the top-performing run. It is computed by summing the differences between the AQE of each VQE run and the AQE_{best}, then normalizing this sum by the total number of runs. It is expressed as:

$$MAED = \frac{\sum_{i=1}^{\#VQE} |AQE_i - AQE_{best}|}{\#VQE}. \tag{5.17}$$

Mininum Relative Energy Deviation (MRED) quantifies the average deviation of the relative quantum energy from the optimal relative energy E_{best}. Defined in the same way as the MAED but using RQE values:

$$MRED = \frac{\sum_{i=1}^{\#VQE} |RQE_i - RQE_{best}|}{\#VQE}. \tag{5.18}$$

Minimum Energy Deviation (MED) is defined as the two previous metrics, but using flat energies and not % values.

$$MED = \frac{\sum_{i=1}^{\#VQE} |E_i - E_{best}|}{\#VQE}. \tag{5.19}$$

5.3.4 Heuristic ansatz comparison

We carried out the comparison between Multi-QIDA configuration and ladder Heuristic ansatz defined by a layer of parametrized rotations $R_y(\theta)$, $N-1$ CNOTs in a top-down ordered configuration, followed by another set of parametrized rotations (see Figure 3.1). In a situation in which the number of CNOTs of a multi-QIDA configuration can not be compared directly with the exact depth of the ladder circuit, we decided to compare ladders composed of the previous and next depths. Thus, if the ladder contains fewer CNOTs than our approach, and the ansatz with depth increased by one, contains a higher number of CNOTs, then we take both ladders as a comparison. The number of CNOTs for ladder fashion circuits is defined as $(N-1) * d$, where N is the number of qubits, i.e. sites number, and d is the number of repetition of each layer with ladder displacement of the entangling gate i.e. the depth. We refer to the ladder ansätze as $(L)^{CX}_d$. For the 3×4 Isotropic Hamiltonian model, the number of sites is 12, which translates into 11 entangling gates per layer. The ladder ansätze we used as comparison contain 44 CNOTs for $(L)^{CX}_4$, 55 CNOTs for $(L)^{CX}_5$, and 66 CNOTs for $(L)_6^{CX}$. The cost in terms of CNOTs for all the configuration can be found in Appendix B.2.

5.3.5 Performance Analysis

Here, we present a comparison between our Multi-QIDA approach and the heuristic ladder ansätze for one specific system that we have studied, the 3 × 4 Isotropic Heisenberg Model Hamiltonian. In Table 5.2, we have collected all the results for this system configuration. As already introduced in the previous section, we carried out the comparison between ladder fashion ansätze and Multi-QIDA method circuits, in both CNOTs and *SO(4)* parametrized gates.

$(L)_d^{CX}$ - Heuristic ladder-fashion ansätze

For this type of ansätze, the results we obtained are very close between PQC at different depths. In particular, the two deeper circuits, with depths 5 and 6, present small variations in terms of mean result and best-performing run. The deeper circuit is the one that performs slightly better. For the absolute quantum error (AQE), the average is set to 84.37%, and the best-performing simulation reaches 90.36%, with a mean error from the best result of 5.99%. The relative quantum error (RQE) values are 57.17%, 73.58%, and 16.42%, corresponding to mean, best VQE, and deviation, in order. The ladder fashion circuits present on average a dispersion from the best-performing runs that is at least of 5.52% for the absolute quantum energy and at least 15.10% for the relative quantum energy. Looking at the mean energy deviation error (MEAD), we can see that the best-performing ladder incurs an error of 0.4.

$Lattice$	$Ansatz$	E_{avg}	E_{best}	AQE_{avg}	RQE_{avg}	AQE_{best}	RQE_{best}	MED	$MAED$	$MRED$
3×4	$(L)_4^{CX}$	$-5.56(\pm 24)$	-5.924784	83.03	53.49	88.54	68.59	0.368738	5.51	15.10
3×4	$(L)_5^{CX}$	$-5.65(\pm 14)$	-6.067645	84.36	57.14	90.67	74.44	0.422450	6.31	17.30
3×4	$(L)_6^{CX}$	$-5.65(\pm 25)$	-6.046642	84.37	57.17	90.36	73.58	0.401	5.99	16.42
3×4	$(QIDA)^{CX}$	$-6.200065(\pm 1)$	-6.200066	92.65	79.87	92.65	79.87	∼ 1e-6	∼ 2e-5	∼ 5e-5
3×4	$(QIDA)^{SO4}$	$-6.3612(\pm 42)$	-6.362179	95.061	86.46	95.08	86.51	0.000998	0.014918	0.040885

Table 5.2: Simulations results for the 3×4 Isotropic Heisenberg model Hamiltonian divided by ansatz layout. Results were obtained with $\#VQE = 50$ each, statevector simulation using Qiskit library, optimization algorithm BFGS with convergence tolerance 1e-6, and finesse ratio $\mu = 0.1$ for the layer selection. AQE_{avg} and RQE_{avg} are defined as the average performance metrics AQE_i and RQE_i over $\#VQE$ simulations. For the E_{avg}, the value in the brackets is the standard deviation.

$QIDA^{CX}$ - Multi-QIDA with CNOTS

We increased the mean AQE to 92.65% and the best-performing run too, with a value of 92.65. The mean absolute energy deviation that we commit using this ansatz is 2e-5%. Consequently, the RQE presents the same behavior, 79.87% for the average RQE and best VQE, and an error of 5e-5%. For the $QIDA^{CX}$ ansatz, the RQE_{avg} is increased by +22.70%, the RQE_{best} by +6.29%, with increased precision of seven orders of magnitude. In terms of MAED, we obtain an energy deviation of 1e-6.

$QIDA^{SO4}$ - Multi-QIDA with SO(4)

We obtained 95.061%, 95.08% and 0.149%, for average AQE, AQE_{best} and deviation, in order. The RQE results are instead 86.46%, 86.51, and 0.04% for mean RQE, best-simulation RQE and minimum relative deviation.In comparison to the ladders, for the $QIDA^{SO4}$, the enhancements that we obtained are +29.29% and +12.93%, for RQE_{avg} and RQE_{best}, respectively. The deviation from the minimum, $MRED$, is three orders of magnitude lower than the one obtained with the deeper heuristic ladder ansatz. We obtained a MEAD of 1e-3 for this ansätze configuration.

5.3.6 Convergence and precision

For a pictorial representation of the results shown in Table 5.2, a series of resuming violin plots can be found in Figure 5.9a.

Each violin represents the AQE and the RQE for a different ansätze configuration for the same system, tested in both ladder and QIDA ansatz settings. The lower tick represents the best VQE, and the higher tick represents the worst VQE. The black dots are associated with the average value of AQE and RQE. The width of the violins represents the frequency of the simulation's outcomes and they can be thought of as a rotated and smoothed histogram of several optimization procedures obtained with different initial parameters. From these plots, it is possible to see the difference in the performance between general heuristic ansätze and Multi-QIDA variational forms.

The ladder-fashion ansätze present average results that are generally distant from the minimum results, meaning that it is required to re-run the circuit different times before obtaining a satisfying result. Ladders can indeed encounter optimization paths that may converge to local minima that are far away from the best result that can be obtained. As you can see for the depth 4 ladder, only a few runs can reach similar results to the best-performing simulation for that configuration.

Multi-QIDA approaches instead guide the variational form in the right spot, providing good results with very high probability, thus, few independent optimizations are required to obtain these results. From Table 5.2, it is possible to see this behavior noting that the difference between the average case and best-performing simulation is almost zero. In terms of energy, $QIDA^{SO4}$ configuration provides the best choice. Looking at the errors committed by the two Multi-QIDA configurations, the $QIDA^{CX}$ ansätze maintain the optimal convergence and lower MRED. The $QIDA^{SO4}$ ansätze present a slightly larger inaccuracy with an error that is three orders of magnitude bigger than the CNOT implementation.

In Figure 5.9b, we show the convergence plot for the 3 × 4 isotropic Heisenberg system. Here, we have collected all the optimization trajectories for each simulation. Each lighter-colored line represents one VQE simulation, while we plotted in a darker color, the best-performing VQE and average VQE values. For these plots, it is possible to see how the VQEs simulations with Multi-QIDA ansätze tend to be more compact in terms of dispersion from the best results. Thus, we can say that with high precision Multi-QIDA-ansätzes drive almost all the VQE simulations to follow the same path as the best-performing VQE run.

In terms of the average number of iterations required for the different types of ansatz, it is clear that it is higher for the Multi-QIDA approach. This is due to the iterative optimization nature of the Multi-QIDA approach, against the standard VQE for the standard heuristic ladder ansatz which requires a single VQE routine. For the 3 × 4 isotropic Heisenberg Hamiltonian system, we

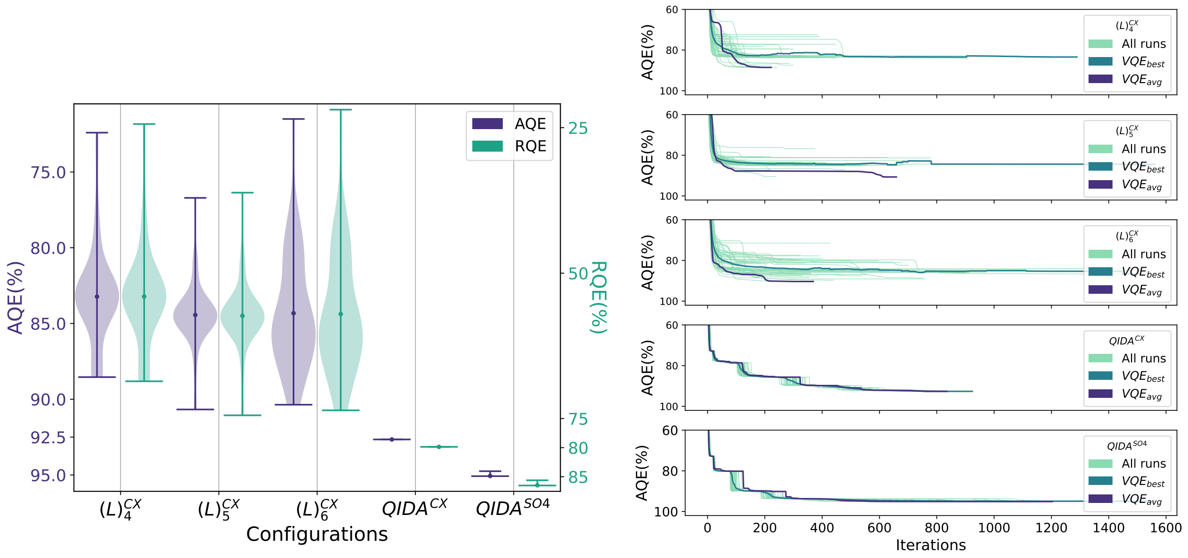

(a) **(b)**

Figure 5.9: **(a)**: 3×4 isotropic Heisenberg Hamiltonian with $J = 1.0$ VQE results. Comparison between AQE and RQE for the different ansätze settings. Starting from the left: HEA ladder depth 4 (44), depth 5 (55), and depth 6 (66), then, Multi-QIDA CNOT(52) and Multi-QIDA (56) *SO(4)* ansätze. In round brackets, the cost in terms of $\#CNOTS$. **(b)**: Representation of the optimization trajectories plotted by each VQE in terms of AQE for the 3×4 isotropic Heisenberg Hamiltonian with $J = 1.0$. From top to bottom: The first three plots represent ladder-fashion ansätze, then the $QIDA^{CX}$ circuit, and finally, the *SO(4)* implementation.

obtained that the number of optimization VQE steps for the three ladders ansatz is 338, 427, and 583, respectively, for depth 3, 4, and 5. For our Multi-QIDA ansatz, we reached 825 evaluations for the CNOT version and 1073 for the *SO(4)* implementation. The mean number of iterations required by the Multi-QIDA ansatz is at least double the evaluations required by the heuristic ladder ansätze. From Figure 5.9b, it is possible to note that, for Multi-QIDA configuration, a not negligible time is spent during the relaxation phase, in which the optimization is trying to re-optimize already converged circuits. In this situation, we could have implemented a different halting criterion for the relaxation phase, which could have helped reduce the time in which the optimization is almost converged. Using a fixed number of evaluations for the re-optimization routine or an additional threshold on the ΔE could help speed up the whole procedure.

The convergence and the trajectories plots for the other tested systems can be found in Figures B.3 and B.4, while the optimization iterations in Table B.3.

5.3.7 Complete results

In this section, we collect very briefly and concisely the results from all the systems configurations: 3x3, 2x6, and 3x4 isotropic Heisenberg model, 3x4 with $\Delta = 1/10$ and $2/3$ anisotropic term, and finally, the 3x4 isotropic with external magnetic field $h = 2$. Table 5.3 collects the results for the three isotropic configurations, while Table 5.4 shows the results for the 3×4 variations. For the Isotropic configurations remaining the 3×3 and the 2×6, our approach in the *SO(4)* settings performs better than the standard HEA in both the lattices. For the 2×6, both the Multi-QIDA configurations performed significantly better than the deeper ladder circuit $(L)_5^{CX}$ with an

Lattice	*Ansatz*	E_{avg}	E_{best}	AQE_{avg}	RQE_{avg}	AQE_{best}	RQE_{best}	*MED*	*MAED*	*MRED*
3×3	$(L)_4^{CX}$	$-4.22(\pm 14)$	-4.335326	88.81	69.61	91.28	76.33	0.117578	2.48	6.72
3×3	$(L)_5^{CX}$	$-4.30(\pm 13)$	-4.497591	90.56	74.37	94.69	85.61	0.196695	4.14	11.24
3×3	$(QIDA)^{CX}$	$-4.333(\pm 20)$	-4.336704	91.25	76.24	91.31	76.41	0.003003	0.063231	0.171668
3×3	$(QIDA)^{SO4}$	$-4.6029(\pm 11)$	-4.603961	96.92	91.63	96.94	91.69	0.001011	0.021294	0.057811
2×6	$(L)_4^{CX}$	$-5.42(\pm 15)$	-5.502422	82.07	54.52	83.33	57.71	0.083146	1.26	3.19
2×6	$(L)_5^{CX}$	$-5.57(\pm 11)$	-5.662730	84.31	60.21	85.75	63.87	0.095113	1.44	3.65
2×6	$(QIDA)^{CX}$	$-6.23776(\pm 31)$	-6.238618	94.46	85.95	94.48	85.99	0.000856	0.012955	0.032860
2×6	$(QIDA)^{SO4}$	$-6.265514(\pm 1)$	-6.265514	94.88	87.02	94.88	87.02	$\sim$ 5e-7	$\sim$ 8e-6	$\sim$ 2e-5
3×4	$(L)_4^{CX}$	$-5.56(\pm 24)$	-5.924784	83.03	53.49	88.54	68.59	0.368738	5.51	15.10
3×4	$(L)_5^{CX}$	$-5.65(\pm 14)$	-6.067645	84.36	57.14	90.67	74.44	0.422450	6.31	17.30
3×4	$(L)_6^{CX}$	$-5.65(\pm 25)$	-6.046642	84.37	57.17	90.36	73.58	0.401	5.99	16.42
3×4	$(QIDA)^{CX}$	$-6.200065(\pm 1)$	-6.200066	92.65	79.87	92.65	79.87	$\sim$ 1e-6	$\sim$ 2e-5	$\sim$ 5e-5
3×4	$(QIDA)^{SO4}$	$-6.3612(\pm 42)$	-6.362179	95.061	86.46	95.08	86.51	0.000998	0.014918	0.040885

Table 5.3: Simulation results divided by ansatz layout and system configurations for all the Isotropic Heisenberg Hamiltonian lattices. For the E_{avg}, the value in the brackets is the standard deviation.

Lattice	*Ansatz*	E_{avg}	E_{best}	AQE_{avg}	RQE_{avg}	AQE_{best}	RQE_{best}	*MED*	*MAED*	*MRED*
$h = 2$	$(L)_4^{CX}$	$-9.066(\pm 92)$	-9.184634	95.35	91.59	96.59	93.84	0.118396	1.25	2.25
$h = 2$	$(L)_5^{CX}$	$-9.068(\pm 87)$	-9.189169	95.37	91.62	96.64	93.93	0.121285	1.28	2.31
$h = 2$	$(QIDA)^{CX}$	$-9.286827(\pm 3\text{e-}4)$	-9.286827	97.67	95.79	97.67	95.76	$\sim$ 3e-10	$\sim$ 3e-9	$\sim$ 5e-9
$h = 2$	$(QIDA)^{SO4}$	$-9.4144(\pm 61)$	-9.426608	99.01	98.21	99.14	98.44	0.012195	0.128254	0.231911
$\Delta = \frac{2}{3}$	$(L)_4^{CX}$	$-4.80(\pm 16)$	-5.091959	89.93	50.61	95.38	77.33	0.290964	5.45	26.72
$\Delta = \frac{2}{3}$	$(L)_5^{CX}$	$-4.866(\pm 49)$	-4.934131	91.09	56.35	92.42	62.84	0.070614	1.32	6.49
$\Delta = \frac{2}{3}$	$(L)_6^{CX}$	$-4.89(\pm 17)$	-5.135422	91.52	58.42	96.19	81.33	0.249399	4.67	22.91
$\Delta = \frac{2}{3}$	$(QIDA)^{CX}$	$-5.1777(\pm 25)$	-5.181318	96.98	85.21	97.05	85.54	0.003624	0.067884	0.332874
$\Delta = \frac{2}{3}$	$(QIDA)^{SO4}$	$-5.232248(\pm 1\text{e-}2)$	-5.232248	98.01	90.22	98.01	90.22	$\sim$ 4e-9	$\sim$ 8e-8	4e-7
$\Delta = \frac{1}{10}$	$(L)_4^{CX}$	$-4.2599(\pm 40)$	-4.272528	99.70	43.64	99.99	99.37	0.013	0.295712	55.73
$\Delta = \frac{1}{10}$	$(L)_5^{CX}$	$-4.2629(\pm 37)$	-4.266284	99.77	56.71	99.85	71.83	0.003428	0.080234	15.12
$\Delta = \frac{1}{10}$	$(QIDA)^{CX}$	$-4.272574(\pm 3)$	-4.272579	99.99	99.57	99.99	99.60	$\sim$ 6e-6	0.000141	0.026570
$\Delta = \frac{1}{10}$	$(QIDA)^{SO4}$	$-4.272645(\pm 4)$	-4.272656	99.99	99.89	99.99	99.94	$\sim$ 1e-5	0.000249	0.046940

Table 5.4: Simulation results for the 3×4 Heisenberg model variations. Starting from the top: Isotropic with External magnetic field $h = 2$, Anisotropic with $\Delta = \frac{2}{3}$, and Anisotropic with $\Delta = \frac{1}{10}$. For the E_{avg}, the value in the brackets is the standard deviation.

increase in the average RQE of $+25.74\%$ and $+26.81\%$ for CNOT and *SO(4)* Multi-QIDA ansätze, respectively (shown in Figure B.3c). The RQE_{best} is increased by $+22.12\%$ by the $QIDA^{CX}$ ansatz and $+23.15\%$ for the $QIDA^{SO4}$. For the 3×3 lattice instead, the $QIDA^{CX}$ has the same performance of the $(L)_4^{CX}$ in terms of best-run VQE and worse performance for the $(L)_5^{CX}$ (Figure B.3a). On this lattice, the *SO(4)* implementation performs better than the deeper ladder ansatz, $(L)_5^{CX}$, obtaining an enhanced RQE_{avg} by $+17.26\%$ and a best RQE of $+6.08\%$. For the variants of the 3×4 lattice, the enhancement in terms of best-performing result is not very pronounced as the improvement in the average case. For the Anisotropic $\Delta = 2/3$ configuration, the RQE_{avg} is increased by $+26.79\%$ and the best runs by $+4.21\%$ for the CNOT ansatz. $QIDA^{SO4}$ instead lead to a $+31.80\%$ for the average case and $+8.89\%$ for the best-performing run (Figure B.1f). For the $\Delta = 1/10$ Anisotropic system, the CNOT Multi-QIDA ansatz increases the average case by $+42.86\%$ and the best one of $+27.77\%$. The *SO(4)* results are close to the previous one, changing only by a few decimals. Finally, for the Isotropic system with external magnetic field $h = 2$, $QIDA^{CX}$ obtains an increment in the average case of $+4.13\%$ and $+1.83\%$ for the best, and similar results of $+6.56\%$ and $+4.51\%$ for the *SO(4)* version.

The preprocessing information as QMI maps and qubit-pairings can be consulted in Figures B.1 and B.2, respectively. The composition and cost of each Multi-QIDA configuration are shown in Tables B.1 and B.2. Convergence violin plots are collected in Figure B.3, while the trajectories of the optimizations are shown in Figure B.4. The cost in terms of VQE evaluations are displayed in Table B.3.

5.4 Results for Molecular systems

To test the Multi-QIDA approach on molecular systems, we considered five different molecules that, once codified on the quantum computer, span the range between 8 to 14 qubits. The systems chosen are H_2O, NH_3, and BeH_2 in terms of full-size system, thus Full-CI level, while H_2O and N_2 with bigger basis have been studied at CASCI level. In Table B.4, we have summarized all the information related to the system under study.

5.4.1 Iterative Natural Orbitals molecules

One of the possible way in which the complexity of the circuit is reduced is exploiting Natural Orbitals (NOs). In particular, we use them as one-electron basis set functions for both VQE simulations and QMI calculations.

Given a wavefunction Ψ, the associated NOs are defined as the set of molecular orbitals (MOs) for which the one-body reduced-density matrix (RDM)

$$\rho_{u,v} = \langle\Psi| a_u^\dagger a_v |\Psi\rangle \tag{5.20}$$

is diagonal. From the diagonal terms of the RMD, suppose $\rho_{u,u}$, we can obtain information about the number of electrons in the u-th orbital and this value is denoted as Natural Orbital Occupation Number (NOON) of orbital u.

As claimed in [209], in the NOs basis, the resulting CI expansion of the state under study is composed by the minimal number of Slater Determinants. The reduction of the population of SDs of a reference state is reflected in an increased sparsity of the Quantum Mutual Information map. The simplification of the variational problem by means of Natural orbitals in quantum computing has been studied in different works [163, 200, 210]. For the *recursive* nature of the basis i.e. the NOs depend on the wavefunction, which is itself defined by the NOs, constructing a CI wavefunction in this basis is tricky. To address this, an iterative procedure known as Iterative Natural Orbitals (INO) [211] can be used, which aims to converge to MO where the wavefunction results in a diagonal one-body RDM. Achieving self-consistency with this iterative method is in principle costly, albeit the convergence rate is usually fast, so the process can be halted once a convergence criterion is achieved.

Therefore, for the first group of molecule, the iterative natural orbitals (INOs), have been employed as initial orbital guess. Only the frozen core approximation at the Hartree-Fock level has been used for the first group of three molecules, by freezing the first core orbital for each system.

All the systems in the first group have been analyzed with STO-3G basis set. One HF is computed, we have applied the procedure to obtain INOs for each system, starting from RCISD calculations. The INOs obtained are used as set of MOs that will define the true active space of the system, i.e. M_{act}.

Active space molecules

For the second group of molecules instead, a more fine selection is performed. For H_2O, the CAS(4,4) active space includes 4 electrons in 4 orbitals: two σ bonding (O-H) and two σ^* antibonding orbitals. Using the 6-31G basis set provides a moderate description of the molecular orbitals, accurately representing the bonding (HOMO) and antibonding (LUMO) levels. For this system, the selected orbitals will compose the set M_{act} with size 4. For N_2, the CAS(6,6) active space involves 6 electrons distributed across 6 orbitals: bonding $\sigma_g(2p_z)$, $\pi_u(2p_x)$, $\pi_u(2p_y)$, and antibonding $\sigma_u^*(2p_z)$, $\pi_g^*(2p_x)$, $\pi_g^*(2p_y)$. For the N_2, the set of M_{act} will be composed of 6 elements. Using the cc-pVTZ basis set, the orbital energies are more accurate due to better flexibility and polarization functions, resulting in a realistic HOMO-LUMO gap. In this case, the HOMO is degenerate, including both the $\pi_u(2p_x)$ and $\pi_u(2p_y)$ orbitals, as well as the LUMO, composed by $\pi_u^*(2p_x)$ and $\pi_u^*(2p_y)$. In contrast, the STO-3G basis set often predicts a different ordering, where the $\sigma_u^*(2p_z)$ orbital can incorrectly appear lower in energy than the π_g^*, leading to an underestimated HOMO-LUMO gap. With cc-pVTZ, the active space selection better captures the electron correlation essential for describing the triple bond in N_2. In minimal sets like STO-3G, improper orbital energy ordering may lead to inaccuracies in multiconfigurational calculations [109].

5.4.2 QMI recovery and layer building

For all the systems, an RCISD wavefunction has been used as a reference to build the QMI map. In particular, starting from the same set of orbitals used in the circuit, we used the PySCF RCISD solver in order to obtain the relevant information used by SparQ to build the QMI matrix. Then, for each system, by observing the distribution of the mutual-information pairs, the selection of the finesse-ratios is performed. Starting from the QMI matrices, shown in Figures B.5, the different finesse-ratios used are for H_2O INOs: [0.5, 0.3, 0.1], BeH_2 INOs: [0.7, 0.4, 0.35, 0.3, 0.2], NH_3 INOs: [0.75, 0.5, 0.25, 0.2], H_2O 6-31G CAS(4,4): [0.5, 0.20, 0.15], N_2 cc-pVTZ CAS(6,6): [0.80, 0.6, 0.4, 0.2].

The choice of the finesse-ratio is done accordingly to the criteria defined in Section 5.2.2, in particular, we avoid the creation of a highly populated candidate set, we cover all the qubits with at least one QIDA-layer, and we stop after reaching a layer below 0.2 of QMI value. The number of QIDA-layers is different for each system, but generally, we need to encode at least 3 layers in order to recover highly correlated pairs, mid-correlation, and low-lying correlations. As in the previous work, we completed the series of QIDA-layers with an additional ladder in order to join together disjointed groups of qubits.

5.4.3 Heuristic Ansätze Comparison

We decided to compare our approach with the most general variational wavefunction as the Hardware-Efficient Heuristic ansatz. The definition of the circuit is the same as previously done for Spin-

Lattices system. HEA are denoted with $(L)_5^{CX}$, while Multi-QIDA using the label $QIDA_{sel}$, where sel can be one of the two selection criteria defined in Section 5.2.2. In particular, the complete CNOTs count is shown in Table 5.5.

5.4.4 Metrics and Measures

#CNOTs	BeH_2	H_2O	NH_3	H_2O CAS(4,4)	N_2 CAS(6,6)
$(L)_d^{CX}$	66	55	65	35	66
$QIDA_{max}$	70	58	66	36	68
$QIDA_{emp}$	70	58	66	36	68

Table 5.5: Number of CNOTs used by each ansatz configuration. $(L)_d^{CX}$ denotes HEA, while $QIDA_{sel}$ to different Multi-QIDA ansätze. The number of CNOTs in Multi-QIDA configuration is equal due to the selection based on spanning trees.

The main metric used to compare different ansatz configurations is the number of CNOTS, #CNOT. We decided to not employ the measurement of the depth of the circuit, due the fact that for the Multi-QIDA method, the full circuit is composed by entangler maps that differs from layer to layer. Thus, a different number and disposition of the correlators, leads to an inhomogeneous metric. To measure the performance of the variational calculation, we used used *percentage correlation energy*, ϵ, defined as

$$\epsilon_i = 100 \cdot \frac{E_{VQE_i} - E_{HF}}{E_{FCI} - E_{HF}} \tag{5.21}$$

where E_{VQE_i} is the converged energy of the i-th simulation, E_{HF} is the Hartree-Fock SCF energy, while E_{FCI} is the exact solution, obtained by performing a diagonalization on the qubit Hamiltonian defined on the M_{act} orbitals, and selecting the lowest eigenvalue. The E_{HF} is instead directly obtained by the RHF solver of PySCF [94–96] package. For the system in which a specific active space is selected i.e. H_2O-CAS(4,4) and N_2-CAS(6,6), the reference exact energy correspond to the CASCI energy, $E_{CASCI} = E_{core} + E_{act}^{diag}$, where E_{core} is the energy contribution of inactive occupied orbitals, and E_{act}^{diag} is the lowest eigenvalue of the active Hamiltonian. Equation 5.21 can be redefined considering an active Hamiltonian as

$$\epsilon_{act} = 100 \cdot \frac{E_{VQE_i} - E_{HF}}{E_{CASCI} - E_{HF}}. \tag{5.22}$$

We have then computed the *Mean Correlation Energy Deviation* (MCED) which quantifies the average deviation of the correlation energy from the best performing simulation ϵ_{best}. It is obtained by summing up the difference between the correlation energies of each VQE simulation and the correlation energy of the best performing simulation, then, the sum is normalized by the total number of simulations. The MCED is formally defined as

$$MCED = \frac{\sum_i^{\#VQEs} |\epsilon_i - \epsilon_{best}|}{\#VQEs}, \tag{5.23}$$

where ϵ_i is the correlation energy for a specific VQE run, ϵ_{best} is the correlation energy of the best performing simulation, and $\#VQEs$ is the total number of simulations for a given ansatz configuration. All the results related to average and best-performing simulations for each ansatz configuration are collected in Table 5.6 for INOs systems, and in Table 5.7 for Active-Space systems.

We also computed four more quantities for each wavefunction computed by both Multi-QIDA and HEA circuits. We are interested in : **Fidelity with the true ground-state**, $\mathcal{F} = \langle \Psi_{GS} | \psi(\bar{\theta})_i \rangle$ where i is the i-th VQE simulation, **Projection along z-axis of the spin**,$\hat{S}_z$, **Spin squared** $\hat{S}^2$, and **Number of particles** $\hat{N}_e$. All the four additional properties measured, the values of best-performing results, are collected in Table B.5 for INOs systems, and in Table B.6 for Active-Space systems.

5.4.5 Performance analysis

In this section, the energetic comparison between HEA and Multi-QIDA, with also the deviation from the best performing VQE of both percentage correlation energy and absolute energy, are shown.

The results presented in Table 5.6 and Table 5.7 demonstrate that the proposed Multi-QIDA approaches, $QIDA_{\text{max}}$ (reduction of each QIDA-layer using a MST that maximizes the total QMI value) and $QIDA_{\text{emp}}$(reduction of each QIDA-layer using a mST that minimizes the distance between each qubit), consistently outperform the standard ladder ansatz $(L)^{d}_{\text{CX}}$ in terms of correlation energy and absolute energy across all tested molecular systems.

	BeH_2			H_2O			NH_3		
	$(L)_6^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$
ϵ_{avg}	21.25	**79.78**	**80.46**	-111.50	**89.52**	**89.71**	-71.13	**54.12**	**54.10**
	(14.55)	(10.22)	(9.54)	(284.70)	(2.31)	(2.31)	(248.89)	(0.36)	(0.36)
E_{avg}	-3.9146	-3.93490	-3.93510	-23.45440	-23.55380	-23.55390	-20.00260	-20.08510	-20.08510
ϵ_{best}	49.75	**90.9**	**84.03**	81.56	**91.51**	**91.55**	41.22	**54.95**	**54.95**
E_{best}	-3.92441	-3.93872	-3.93637	-23.54985	-23.55477	-23.55479	-20.07658	-20.08561	-20.08561
$M\epsilon D_{best}$	28.17	10.98	3.57	193.07	1.99	1.85	112.36	0.83	0.85
MED_{best}	0.00980	0.00380	0.00120	0.09550	0.00100	0.00090	0.07400	0.00050	0.00060

Table 5.6: BeH_2,H_2O, and NH_3 INOs system results. ϵ_{avg}, ϵ_{best}, and $M\epsilon D_{best}$ are given in percentage, while E_{avg}, E_{best}, and MED_{best} in Hartree.

INOs systems

In the simulations of BeH_2 , H_2O , and NH_3, Multi-QIDA circuits consistently achieve higher average percentage correlation energy (ϵ_{avg}%). In particular, for BeH_2, $QIDA_{\text{max}}$ achieves an ϵ_{avg} of 79.78%, a significant improvement over the 21.25% obtained using the ladder ansatz, with a lower standard deviation as well (10.22% versus 14.55%). Better results are obtained by $QIDA_{emp}$, which increases the ϵ_{avg} up to 80.46% Similarly, for H_2O, $QIDA_{\text{max}}$ obtains 89.52%, $QIDA_{emp}$ reaches a close 89.71%, whose compared to the negative correlation energy (-111.50%) from $(L)^{5}_{\text{CX}}$, indicating that the ladder ansatz optimization is failing to converge, leading to a state with energy higher than HF. In this case, the deviation is two orders of magnitude lower for the $QIDA$ circuit. For NH_3, while all methods show similar results, $QIDA_{\text{max}}$ and $QIDA_{\text{emp}}$ marginally outperform the ladder topology with ϵ_{avg} values of -54.12% and 54.10%, respectively, compared to the negative value of -71.13% for $(L)^{d}_{\text{CX}}$. Here, both $QIDA_{max}$ and $QIDA_{emp}$ obtain close values of percentage correlation energy deviation, close to zero. In terms of best VQE results, BeH_2, $QIDA$ obtains a clear increase of 40% over the ladders, whereas for the other systems, the increase is around 10/12%. In particular, for BeH_2, Multi-QIDA circuits do not fall far from the average case, as expected. Compared to the 49.75% obtained by the ladder, $QIDA_{max}$ reaches 90.90%, while

$QIDA_{emp}$ gets 84.03%. For the correlation energy of best performing VQE of H_2O, we can observe that the HEA reaches a good 81.56%, while $QIDA_{max}$ and $QIDA_{emp}$ get similar values, around 91%. Finally, for NH_3, the ladder reaches 41.22% correlation energy, while Multi-QIDA in both settings reaches 54.94%. In terms of percentage correlation energy deviation, $QIDA$ shows a lower dispersion w.r.t. the best-performing VQE, hitting a dispersion of two orders of magnitude lower than ladder ansatz for H_2O and NH_3. In general, we can find a consistent number of VQEs for the ladder topology, which due to the random initial parametrization, are guided in a completely wrong energetic solution, way lower than the HF energy.

	H_2O $CAS(4,4)$			N_2 $CAS(6,6)$		
	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_6^{cx}$	$QIDA_{max}$	$QIDA_{emp}$
ϵ_{avg}	55.42(27.46)	**82.36**(6.32)	**80.32**(9.18)	20.85(79.43)	**57.13**(6.53)	**58.48**(17.59)
E_{avg}	-6.62720	-6.62800	-6.62790	-11.44410	-11.46860	-11.46950
ϵ_{best}	92.17	**95.42**	**97.81**	79.99	**82.17**	**85.22**
E_{best}	-6.62828	-6.62837	-6.62844	-11.48403	-11.48551	-11.48757
$M\epsilon D_{best}$	36.75	13.06	17.49	59.13	25.04	26.75
MED_{best}	0.00110	0.00040	0.00050	0.04000	0.01690	0.01810

Table 5.7: H_2O 6-31G CAS(4,4) and $N2_2$ cc-PVTZ CAS(6,6) system results. ϵ_{avg}, ϵ_{best}, and $M\epsilon D_{best}$ are given in percentage, while E_{avg}, E_{best}, and MED_{best} in Hartree.

Active Region systems

The second group of simulations are related to the application of Multi-QIDA in more complex systems, which are considered by dividing orbitals into inactive and active space regions. For the H_2O CAS(4,4) system, $QIDA_{\text{max}}$ achieves an ϵ_{avg} of 82.36% while $QIDA_{emp}$ reaches a 80.32%, which are both enhanced performance with respect to the HEA ladder, which obtains on average 55.42% correlation energy. The same behavior can be found for the N_2 CAS(6,6), for which the standard ladder obtains on average 20.85% of correlation energy, while $QIDA_{max}$ is able to reach 57.13% and $QIDA_{emp}$, a slightly higher value of 58.48%. In terms of best-performing VQE results, we have that the results of HEA ladders and Multi-QIDA circuits are close, but in any case, the latter reaches slightly higher correlation energy.

Comparing instead the results obtained in terms of percentage correlation energy deviation w.r.t. to the best-performing VQE, we can notice that $Multi-QIDA$ behaves clearly better for complete systems, so in our case INOs systems, while it has a higher dispersion for Active space systems, which may be related to the fact that we are not including any kind of double excitations directly in the ansätze. Also, we can observe that there is no clear distinction between the two type of selection performed on the QIDA-layers, a difference that may be appreciated more if applied in the context of real hardware topology or real devices.

5.4.6 Convergence and Precision

Here, we briefly analyze the results presented in Figure 5.10a and in the Appendix Figures B.6a-B.7b, related to the precision of VQE runs, and Figure 5.10b and in the Appendix Figures B.8a-B.9b, related to the dispersion of the optimizations.

Each of the plots in the first group represents the energy, E, and the percentage correlation energy, E_{corr}, for every system and for all three ansatz configurations. In these plots, the results for each ansatz configuration are represented by a violin plot. The width of each violin is related to the frequency of the VQE outcomes. They are useful for assessing the consistency of the algorithm (how clustered or spread out the results are) and identifying trends, such as whether the algorithm reliably converges to a minimum energy or exhibits variability. At the two extremes of each violin are presented the worst and the best performing VQE results, while the central dot represents the average value.

We can notice that standard HEA with ladder-fashion connectivity presents results that are way lower than HF and thus the average value is strongly shifted from the best-performing VQE. In the systems in which this behaviour does not happen, the average result is anyway distant than the optimal one and this means that it is required to re-run the circuit even more time, compared to Heisenberg Model Hamiltonian, [206] before obtaining a satisfying result. Ladders quickly encounter and falls into local minima that are far away from the best result. An example is in Figure 5.10a, for which HEA can reach a good 81,54% of correlation energy but is heavily penalized on average because most of the runs falls below HF energy, while Multi-QIDA is able to maintain, as expected, a low dispersion around the best-performing VQE. We can further identify some difficulties of Multi-QIDA in correctly describing the NH_3 ground state and in particular, it get stuck around 50% of correlation energy, shown in Figure B.6b. But in general, we can see that even in the worst-performing VQE, for which the population is very small, the energy is still higher than the average case of HEA and in some cases higher that the best-performing ladder VQE. This last case is shown in Figure 5.10a. In Figure B.7b, we can see how, for the N_2 CAS(6,6) the performance of the two Multi-QIDA selection criteria is completely different, and in particular, the *maximum correlation* fails to compact the results towards the best-performing VQE but on the worst-performing, still reaching higher correlation energy than the average HEA circuits. Another example is for BeH_2, in which the best-performing for ladders reaches ∼50% correlation energy, while a very limited portion of VQEs for Multi-QIDA fails to follow the right variational path. Multi-QIDA approach, as already seen for spin systems, can guide the variational wavefunction in the right spot, in an iterative and adaptive way, without the requirement of building and optimizing the full variational space from the beginning. This behavior can also be noticed in the second group of plot, Figure 5.10b and in Appendix Figures B.8a-B.9b, in which we show the trajectories of the convergence of each VQE for all the systems. For HEA simulations, the upper subplot, it is clear that the population of VQE that are actually getting towards the right optimization path is very low compared to the one that diverges or gets stuck in local minima. For Multi-QIDA instead, it is possible to notice that the trajectories, even if they get perturbed at each additional layer, tend to be more compact and closer to the best-performing and average trajectory. From the best-performing VQE, which is the bolder trajectory, it is also possible to notice the quick recovery and restart of the optimization after being perturbed, allowing the escape from the previous local minima and without ending in a higher convergence point.

As in the previous work, we are aware of the higher computational cost required by Multi-QIDA to converge and to end the full optimization procedure. The average number of iteration is usually two/three times the number of iterations required by the corresponding HEA ansatz, and most of the optimization procedure is wasted in the relaxation procedure.

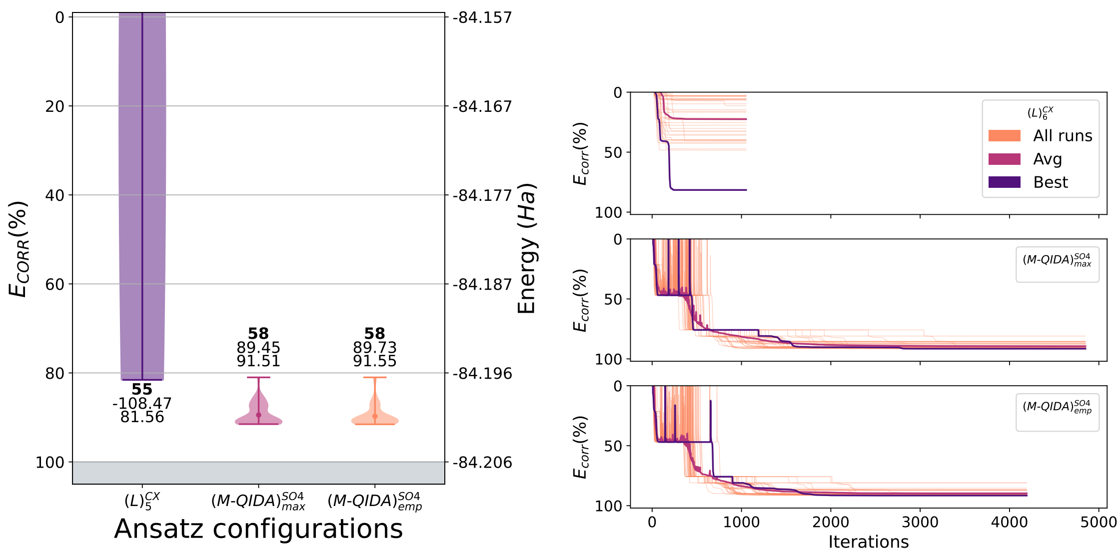

(a) **(b)**

Figure 5.10: (a):H_2O INOs system comparison between depth 5 ladder HEA and Multi-QIDA in both max and empirical configuration. The three numbers associated with each violin show the **number of CNOTs**, ϵ_{avg}, and ϵ_{best}, starting from the upper one, for each simulation setting.**(b)**:H_2O INOs system convergence trajectories for each of the 50 VQEs. In particular, starting from the upper plot $(L)^5$, Multi-QIDA with *max* selection criteria, and last Multi-QIDA with *emp* reduction.

5.4.7 Wavefunction Properties

Together with the measurement of the performance, we decided to define also metrics to evaluate the capability of the Multi-QIDA ansatz to satisfy symmetry constraints and fidelity w.r.t. the exact ground state. Given the fact that all the system studied are closed shell, the $\hat{S}_z$ and $\hat{S}^2$ are both zero, while the number of particles, $\hat{N}_e$, for each specific INOs system is: BeH_2=4 , H_2O=6, and NH_3=8 , for CAS system instead: H_2O $CAS(4,4)$=4 and N_2 $CAS(6,6)$=6. The properties analysis results are collected in Table B.5 for INOs systems, and in Table B.6 for Active-Space systems.

INOs Systems

On average, Multi-QIDA obtained an improvement on each of the property measured. For the Fidelity $\mathcal{F}$ with respect to the ground state, both the Multi-QIDA configuration recovered a slightly higher value, $\sim 1.30\%$,for BeH_2, a relevant increment of $\sim 10.8\%$ for H_2O, and a non-negligible $\sim 6.8\%$ for NH_3. For the spin symmetries, in both cases, Multi-QIDA has been able to improve the results or in general to not lower the quality of the result. In particular, for $\hat{S}_z$, Multi-QIDA obtain 0 in the two selection criteria, while the HEA gets 0.10343 for H_2O and 0.05541. In terms of $\hat{S}^2$ instead, the improvements have been obtained on all the three systems, in particular, for BeH_2, Multi-QIDA with *max* selection halved to 0.00111 the HEA value, 0.00261, while Multi-QIDA with *emp* reduce the value lower to 1e-3. For H_2O, the value for HEA is 0.22268 and Multi-QIDA configuration reduced it of two orders of magnitude. The best improvement have been obtained for NH_3 for which the value obtained by HEA, 0.13443, has been reduced by four orders of magnitude. For the number of particles, $\hat{N}$, the values have been refined to exact values only for NH_3. In terms of best-performing VQE, as we expected, also standard HEA is able to recover almost the same values of properties as Multi-QIDA.

Active Region systems

For this systems, the main improvement of Multi-QIDA with respect to HEA can be found mainly for the second system, N_2 CAS(6,6). On average, Multi-QIDA obtains an slightly higher fidelity compared to HEA and for N_2 closer $\hat{S}^2$ and $\hat{N}$ to the exact value. As before and as we expected, the best-performing VQE of HEA is able to obtain properties values closer to the one measure from a Multi-QIDA circuit.

5.5 Discussion and final remarks on state–preparation with shallow circuits

In this work, we introduced the Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA), a versatile framework designed to enhance quantum simulations of both lattice spin models and molecular systems. Multi-QIDA extends the single-threshold QIDA by leveraging multiple thresholds of Quantum Mutual Information (QMI) to identify and incorporate the most relevant correlations in a layered, compact quantum circuit. By systematically including mid- to high-QMI qubit pairs, Multi-QIDA captures richer correlations while maintaining shallow and resource-efficient circuits. A central innovation of Multi-QIDA is its layered ansatz construction. Each layer can act as an identity for specific parameter sets, enabling controlled circuit growth without unnecessary complexity. The inclusion of *SO(4)* correlators proved especially advantageous: these gates preserve essential wavefunction symmetries—such as total spin, spin projection, and particle number—while improving variational energy estimates and mitigating the barren plateau problem when parameters are properly managed. This design allows for both energy-efficient optimization and fidelity to the underlying physics of the system. For lattice spin systems (e.g., Heisenberg model Hamiltonians), Multi-QIDA consistently outperformed heuristic ladder ansätze, achieving higher precision in ground-state energy calculations, faster convergence, and significantly reduced average errors—often by two orders of magnitude—at only a moderate increase in cost-function evaluations. For molecular simulations, including H_2O, BeH_2, NH_3, and active-space models like N_2 CAS(6,6) and H_2O CAS(4,4), Multi-QIDA likewise surpassed hardware-efficient ansatz (HEA) approaches in accuracy, scalability, and symmetry preservation. These properties make it an excellent starting guess for adaptive algorithms such as ADAPT-VQE [81] or Quantum Selected CI (QSCI) based approaches [201]. Despite these achievements, several open questions remain. Future investigations will focus on: evaluating performance on larger, more strongly correlated systems, integration with adaptive or sampling-based approaches, such as ADAPT-VQE and QSCI, to further enhance accuracy without introducing optimization bottlenecks, adapting the method to noisy quantum devices with decoherence, and extension to alternative correlators (e.g., single/double-qubit excitations or Givens rotations).

Although Multi-QIDA enables a nice state preparation within the active space, or warm-start approach, most of the time, only the static-correlation inside the active space is not enough to describe properly a molecular system or a mechanism. In fact, also the dynamical correlation, coming from the complementary space, i.e. outside the active space, should be considered for an accurate description. The next chapter focuses on bridging this gap.

Chapter 6

Hybrid Quantum-Classical Multi-Configurational Self-Consistent Field for polymerization reaction

Many molecular systems of chemical, biological, and materials relevance cannot be described accurately by a single electronic configuration. In situations where several Slater determinants lie close in energy, their quantum mechanical mixing generates strong *static (or nondynamical) correlation*, which is fundamentally beyond the scope of single–reference methods such as Hartree–Fock, coupled cluster, or standard Kohn–Sham density functional theory. Typical examples include bond dissociation, transition metal complexes with partially filled d or f shells, polyradicals, spin-crossover compounds, magnetic materials, and electronically excited states involved in photochemistry. In such cases, multiple electronic configurations contribute with comparable weight to the true wavefunction, and the Born–Oppenheimer Potential Energy Surface (PES) may exhibit near-degeneracies, avoided crossings, or genuine conical intersections. Capturing these features requires *multi-configurational* wavefunction methods, where the electronic state is expanded as linear combination of Slater Determinant (SD) allowing the wavefunction to adapt flexibly to complex electronic arrangements.

Among such methods, *Complete Active Space Self–Consistent Field (CASSCF)* and its multi–configurational generalization, *Multi-Configuration Self-Consistent Field (MCSCF)* provide a balanced description of static correlation by treating a selected active space with full configuration interaction while simultaneously optimizing the molecular orbitals. This makes them the main tools for describing reliably excited states, bond rearrangements, and photochemical reaction pathways, and indispensable for applications ranging from vision and photosynthesis to catalysis. However, despite their conceptual elegance, the applicability of CASSCF and related approaches is fundamentally limited by the exponential growth of the active-space CI problem. In practice, classical algorithms rarely exceed active spaces of roughly (16–20) electrons in (16–20) orbitals, which is insufficient for many chemically realistic transition-metal or strongly correlated systems.

Quantum computing offers a promising path to overcome this bottleneck. By encoding the electronic wavefunction into qubits, a quantum computer can in principle represent and manipulate exponentially large CI spaces with only polynomial resources. Variational HQC algorithms, such VQE, enable the preparation of multi–configurational wavefunctions on quantum hardware, shifting the exponential cost of configuration interaction into a quantum state preparation that scales more favorably. When combined with classical orbital optimization in the spirit of CASSCF, this opens the door, at least theoretically, to *quantum multi–configurational methods* capable of treating larger and more realistic active spaces than what is classically tractable. Ultimately, this hybrid paradigm promises chemically accurate descriptions of strongly correlated electronic structure, improved prediction of excited-state reactivity, and faithful treatment of complex phenomena such as conical intersections, while retaining the interpretability and physical grounding of traditional multi-reference wavefunction theory.

In this chapter, we are going to present the path we followed to extend the *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR)* [210, 212, 213] algorithm to include a broader Self–Consistent Field (SCF) loop, in order to optimize not only the orbitals in the active space, but the full variational space, i.e. inactive occupied and inactive virtual orbitals. We are going to introduce to the *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)* algorithm, and how both from the theoretical and computational point of view, has been optimized to scale, also in High–Performance Computing (HPC) context. We tested our WAHTOR–SCF against standard molecular system, for which we described the full bond elongation up to dissociation, and also for industrially relevant reaction.

This chapter will be defined as follows: we will recap briefly the theory behind CASSCF/MCSCF approaches, in Section 6.1, passing by the explanation of already existing and notorious Hybrid Quantum–Classical algorithm. An brief introduction of the original WAHTOR procedure will be given in Section 6.2. We introduce the novel extension, i.e. additional self-consistency, in Section 6.3. In particular, we will dedicate its content to show the computational and optimization effort put to make the original procedure, and the extended one, scalable and HPC oriented. In Section 6.4, the benchmarks molecules and the simulation condition will be presented, while specific results are collected in Section 6.5. The industrially relevant molecules and associated results will be presented instead in Section 6.6. Section 6.7 will be dedicated to wrap up the results.

6.1 From CASCI to CASSCF: Classically and Quantum orbital–optimization

In Section 2.4 we already discussed very briefly about CASCI and CASSCF, but here, we will go a little bit more into the deatails of the two approaches.

Many molecular systems cannot be described adequately by a single Slater determinant, especially when strong static correlation is present. In such cases, the electronic structure is intrinsically multi–configurational, and the wavefunction must be expanded as

$$|\Psi\rangle = \sum_I C_I \, |\Phi_I\rangle, \tag{6.1}$$

where $\{|\Phi_I\rangle\}$ are Slater determinants built from a chosen molecular orbital (MO) basis. A common strategy to reduce the explosive cost of Full Configuration Interaction (FCI) is to restrict electron correlation to a selected subset of orbitals, forming a Complete Active Space (CAS). In a

CAS($N_{\text{act}}, M_{\text{act}}$), one performs FCI only among M_{act} active orbitals containing N_{act} active electrons:

$$|\Psi_{\text{CAS}}\rangle = \sum_{I \in \text{CAS}} C_I |\Phi_I\rangle. \tag{6.2}$$

Solving the projected Schrödinger equation

$$\hat{H}\,|\Psi_{\text{CAS}}\rangle = E\,|\Psi_{\text{CAS}}\rangle \tag{6.3}$$

defines the CASCI method. While CASCI captures static correlation within the active space, its energy and qualitative behavior depend sensitively on the underlying MO basis. Because CASCI optimizes only the CI coefficients and not the orbitals, it lacks invariance under orbital rotations. As a result, CASCI may produce distorted potential energy surfaces, incorrect excitation energies, or fail entirely near avoided crossings and conical intersections, where a balanced orbital description is essential. CASCI alone is not sufficient because it lacks orbital relaxation and can fail catastrophically, especially for excited states and bond rearrangements. CASSCF resolves this by a simultaneous optimization of both the orbitals and the CI expansion, yielding a balanced, state-consistent multiconfigurational reference for systems where single–reference methods break down.

The Multi-Configuration Self-Consistent Field (MCSCF) framework removes this deficiency by optimizing *both* the CI coefficients and the orbitals. The most widely used MCSCF family member is CASSCF, which minimizes the total energy

$$E = \langle \Psi(\mathbf{C}, \boldsymbol{\kappa})|\,\hat{H}\,|\Psi(\mathbf{C}, \boldsymbol{\kappa})\rangle, \tag{6.4}$$

where the orbitals are parameterized through a unitary rotation,

$$|\phi'_p\rangle = \sum_q U_{qp}|\phi_q\rangle, \qquad U = e^{-\hat{\kappa}}, \qquad \hat{\kappa} = \sum_{p>q} \kappa_{pq}(\hat{E}_{pq} - \hat{E}_{qp}), \tag{6.5}$$

where $\hat{E_{pq}} = \hat{a}^{\dagger}_{p\alpha}\hat{a_{q\alpha}} + \hat{a}^{\dagger}_{p\beta}\hat{a_{q\beta}}$.
Stationarity of the energy yields two coupled equations:

$$\hat{H}_{\text{CAS}}(\boldsymbol{\kappa})\,\mathbf{C} = E\,\mathbf{C} \qquad \text{(CASCI eigenvalue problem)}, \tag{6.6}$$

$$\frac{\partial E}{\partial \kappa_{pq}} = 0 \qquad \text{(orbital optimization)}. \tag{6.7}$$

Their self–consistent solution produces a multi–configurational reference that is smooth, orbital–invariant, and capable of describing degeneracies and non–adiabatic regions accurately.

6.1.1 Quantum counterpart of MCSCF methods

With current quantum hardware, diagonalization of large active spaces is out of reach, but VQE offers a scalable alternative to CASCI by preparing correlated wavefunctions on a quantum processor. However, if the orbitals are not optimized, a VQE calculation inherits the same limitations as CASCI. This motivates hybrid quantum–classical schemes that merge VQE with orbital optimization, in direct analogy to CASSCF.

The idea of a HQC orbital–optimization routine it has been already explored before, initially by Takeshita [214] and then, in a first form of completely Hybrid Quantum–Classical CASSCF [215]. In the second, approach, the quantum processor replaces the CASCI solver, while the orbital optimization loop remains purely classical, preserving the structure of the original algorithm. The energy depends on both the ansatz parameters $\boldsymbol{\theta}$ and on the orbital rotation parameters $\boldsymbol{R}$. The latter are optimized through a unitary orbital transformation, $\hat{U}_{\mathrm{OO}}$, defining the Orbital Optimization (OO),

$$\hat{U}_{\mathrm{OO}}(\boldsymbol{R}) = e^{-\hat{R}}, \qquad \hat{R} = \sum_{p>q} R_{pq}\left(\hat{E}_{pq} - \hat{E}_{qp}\right), \tag{6.8}$$

resulting in a transformed Hamiltonian

$$\hat{H}(\boldsymbol{R}) = \hat{U}^{\dagger}_{\mathrm{OO}}(\boldsymbol{R})\,\hat{H}\,\hat{U}_{\mathrm{OO}}(\boldsymbol{R}). \tag{6.9}$$

Given a converged wavefunction $|\psi(\boldsymbol{\theta})\rangle$, the quantum device provides the following one- and two-body RDMs as

$$\gamma_{pq} = \langle\psi(\boldsymbol{\theta})|\hat{E}_{pq}|\psi(\boldsymbol{\theta})\rangle, \qquad \Gamma_{pqrs} = \langle\psi(\boldsymbol{\theta})|\hat{E}_{pq}\hat{E}_{rs} - \delta_{qr}\hat{E}_{ps}|\psi(\boldsymbol{\theta})\rangle, \tag{6.10}$$

defining the electronic energy

$$\hat{E}(\boldsymbol{R}) = \sum_{pq}\gamma_{pq}h_{pq}(\boldsymbol{R}) + \frac{1}{2}\sum_{pqrs}\Gamma_{pqrs}g_{pqrs}(\boldsymbol{R}), \tag{6.11}$$

where $h_{pq}(\boldsymbol{R})$ and $g_{pqrs}(\boldsymbol{R})$ are the one- and two-body integrals respectively with a given set of orbital rotation parameters $\boldsymbol{R}$. Using the definition of electronic energy of Equation 6.11, we can then determine the orbital gradient w.r. the R_{pq} parameters

$$G_{pq}(\boldsymbol{R},\boldsymbol{\theta}) = \frac{\partial E(\boldsymbol{R},\boldsymbol{\theta})}{\partial R_{pq}} = \langle\psi(\boldsymbol{\theta})|\,[E^{-}_{pq},\hat{E}(\boldsymbol{R})]\,|\psi(\boldsymbol{\theta})\rangle \tag{6.12}$$

and Hessian

$$H_{pq,rs}(\boldsymbol{R},\boldsymbol{\theta}) = \frac{\partial^2 E(\boldsymbol{R},\boldsymbol{\theta})}{\partial R_{pq}\partial R_{rs}} = \langle\psi(\boldsymbol{\theta})|\,[E^{-}_{pq},[E^{-}_{rs},\hat{E}(\boldsymbol{R})]]\,|\psi(\boldsymbol{\theta})\rangle\,, \tag{6.13}$$

where $E^{-}_{pq} = \hat{E}_{pq} - \hat{E}_{qp}$. The commutator expressions in Equations 6.12 and 6.13 can be conveniently evaluated from the one- and two-body RDMs obtained on the quantum device, according to Equation 6.10. Substituting Equation 6.11 into the gradient definition, the orbital gradient becomes

$$G_{pq}(\boldsymbol{R},\boldsymbol{\theta}) = \sum_{rs}\gamma_{rs}\,[E^{-}_{pq},h_{rs}(\boldsymbol{R})] + \tfrac{1}{2}\sum_{rstu}\Gamma_{rstu}\,[E^{-}_{pq},g_{rstu}(\boldsymbol{R})], \tag{6.14}$$

showing that the gradient is entirely determined by contractions of the measured RDMs with the one- and two-electron integrals. Analogously, the orbital Hessian of Equation 6.13 involves second-order commutators that can be expressed through the same RDMs tensors, without requiring higher-order measurements from the quantum processor. After evaluating both the gradient and eventually the Hessian, the classical computer updates $\boldsymbol{\kappa}$ via a Newton–Raphson step,

$$\boldsymbol{\kappa}^{(n+1)} = \boldsymbol{\kappa}^{(n)} - H^{-1}G. \tag{6.15}$$

Following the presented procedure, multiple version of multi-configurational approaches appeared. In general, we can subdivide them in either using fixed ansatz [216–221], which are the majority, and adaptive approaches [222].

The first employs a *fixed ansatz* in a state-specific or state-averaged setting. An example, the State-Average Orbital-Optimized VQE (SA–OO–VQE) [223] targets multiple states simultaneously by minimizing a state-averaged functional, reproducing the role of SA-CASSCF and preventing root-flipping near degeneracies or conical intersections. In this scheme, the quantum computer provides reduced density matrices (RDMs), while the classical orbital step imposes the same self-consistency conditions as in SA-CASSCF. Additionally, the overall tendency is to used CC-like circuits like UCCSD or trucated versions. In some particular instances, also HEA circuits have been employed, like in *quantum CASSCF* [219] or in *Optimal Orbital VQE* [220].

The second family adopts *adaptive ansätze*. In ADAPT-VQE–SCF [222], the circuit is iteratively constructed while performing orbital optimization, yielding a state-specific analogue of CASSCF. At each macro-iteration, the quantum device prepares the correlated state and measures RDMs, and the classical computer updates the orbital rotations using the CASSCF Newton step.

In the following subsection, we introduce the WAHTOR–SCF algorithm, which combines qubit-space derivative measurements with different orbital optimizer algorithms, that can be paired with either Hardware–Efficient Ansatz (HEA) or adaptive ansätze to achieve hybrid CASSCF self-consistency.

6.2 Wavefunction–Adapted Hamiltonian Through Orbital–Rotation in a nutshell

Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR) method is a HQC algorithm that, unlike traditional variational approaches which adapt the wavefunction to a fixed Hamiltonian, adapts the Hamiltonian to better match a chosen variational ansatz. By exploiting the invariance of the electronic Hamiltonian under orbital rotations, WAHTOR shifts towards an orbital basis in which the target state is easier to represent with a shallow VQE circuit, thus improving convergence and reducing quantum resources.

The method applies a unitary orbital–rotation operator

$$U(\mathbf{R}) = e^{-i\mathbf{T}\cdot\mathbf{R}} = e^{-i\sum_{k<l} R_{kl}\hat{T}_{kl}}, \tag{6.16}$$

where $\mathbf{T} = \{T_{ij}\}$ are generators of orbital rotations, $i, j \in \boldsymbol{A}$ i.e. orbitals in the active-space, and $\mathbf{R}$ is the vector of rotation parameters. The molecular Hamiltonian is unitarily transformed as

$$H(\mathbf{R}) = U^{\dagger}(\mathbf{R})\, H\, U(\mathbf{R}). \tag{6.17}$$

For a given ansatz state $|\Psi(\boldsymbol{\theta})\rangle$, WAHTOR minimizes the energy

$$E(\mathbf{R}) = \langle\Psi(\boldsymbol{\theta})|\, H(\mathbf{R})\, |\Psi(\boldsymbol{\theta})\rangle \tag{6.18}$$

by computing the energy gradient

$$\frac{\partial E(\boldsymbol{R})}{\partial R_{ij}} = \sum_{ij} \frac{\partial h_{ij}(\mathbf{R})}{\partial R_{ij}} \langle a_i^\dagger a_j \rangle + \frac{1}{2} \sum_{ijkl} \frac{\partial g_{ijkl}(\mathbf{R})}{\partial R_{ij}} \langle a_i^\dagger a_j^\dagger a_k a_l \rangle, \tag{6.19}$$

where the expectation values are taken over $|\Psi(\boldsymbol{\theta})\rangle$. At $\mathbf{R} = 0$, the integral derivatives reduce to compact commutator forms, e.g.

$$\frac{\partial h}{\partial R_{ij}}\bigg|_{\mathbf{R}=0} = -i[T_{ij}, h_{HF}], \qquad \frac{\partial g}{\partial R_{ij}}\bigg|_{\mathbf{R}=0} = -i[T_{ij} \otimes I + I \otimes T_{ij},\, g_{HF}], \tag{6.20}$$

with analogous expressions for second derivatives.

The choice of classical optimizer is crucial, as WAHTOR does not seek the chemically optimal orbital basis (as in CASSCF), but instead a *VQE-compatible* one, similar to what happens in [220]. Different optimizers exhibit distinct behavior:

- **Steepest Descent (SD)**:

$$\mathbf{R}^{(k+1)} = \mathbf{R}^{(k)} - \alpha\, \nabla E(\mathbf{R}^{(k)}, \boldsymbol{\theta}), \tag{6.21}$$

  robust but slow and sensitive to step size α.

- **Newton–Raphson (NR)**:

$$\mathbf{R}^{(k+1)} = \mathbf{R}^{(k)} - H^{-1} \nabla E(\mathbf{R}^{(k)}, \boldsymbol{\theta}), \qquad H_{pq,rs} = \frac{\partial^2 E}{\partial R_{kl} \partial R_{rs}}, \tag{6.22}$$

  quadratic convergence but expensive due to Hessian evaluation.

- **BFGS (quasi-Newton)**:

$$H^{(k+1)} = H^{(k)} + \frac{y^{(k)} y^{(k)T}}{y^{(k)T} s^{(k)}} - \frac{H^{(k)} s^{(k)} s^{(k)T} H^{(k)}}{s^{(k)T} H^{(k)} s^{(k)}}, \tag{6.23}$$

  with

$$s^{(k)} = \mathbf{R}^{(k+1)} - \mathbf{R}^{(k)}, \qquad y^{(k)} = \nabla E^{(k+1)} - \nabla E^{(k)}, \tag{6.24}$$

  offering an optimal compromise for noisy hybrid environments.

- **Trust–Region (TR)**: In contrast to Newton or quasi–Newton methods, trust–region optimization restricts each orbital–rotation update to a neighborhood where the quadratic model of the energy is reliable. Around the current parameters $\mathbf{R}^{(k)}$, the energy is approximated by the second–order model

$$m_k(\Delta\mathbf{R}) = E(\mathbf{R}^{(k)}) + \nabla E^T \Delta\mathbf{R} + \tfrac{1}{2} \Delta\mathbf{R}^T H \Delta\mathbf{R}, \tag{6.25}$$

  and the step is obtained by solving the constrained problem

$$\min_{\Delta\mathbf{R}}\; m_k(\Delta\mathbf{R}) \qquad \text{s.t.} \qquad \|\Delta\mathbf{R}\| \le \delta_k, \tag{6.26}$$

  where δ_k is the trust–region radius. After each iteration, δ_k is enlarged or reduced depending on how well m_k predicts the true energy decrease, ensuring stable convergence even when gradients are noisy or Hessians are poorly conditioned. In the WAHTOR setting, this yields robust orbital updates and avoids unstable large rotations, making the method more reliable than pure Newton or steepest–descent schemes.

6.3 WATHOR-SCF: Adding Self-Consistency to WAHTOR

We introduce an extension of the WAHTOR framework, denoted *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)*, which removes the restriction that the orbital–rotation generators act only within the active space. Instead, the generator set $\mathbf{T} = \{T_{pq}\}$ is enlarged to span the full molecular orbital space, including inactive, active, and virtual orbitals. For giving a general definition, a generator $\hat{T}_{pq}$ will be defined as obtained from the cartesian product of the molecular orbital in the full space, i.e.

$$T_{pq} = (a_p^\dagger a_q - a_q^\dagger a_p), \tag{6.27}$$

where $p \in \boldsymbol{F}$ and $q \in \boldsymbol{F}$. This leads to a generalized orbital transformation

$$U_{\text{SCF}}(\mathbf{R}) = e^{-i \sum_{p<q} R_{pq} T_{pq}}, \tag{6.28}$$

with p, q defined as above, and a fully transformed Hamiltonian

$$H_{\text{SCF}}(\mathbf{R}) = U_{\text{SCF}}^\dagger(\mathbf{R})\, H\, U_{\text{SCF}}(\mathbf{R}), \tag{6.29}$$

analogous to a full MCSCF–style rotation instead of an active–space–restricted one. As in WAHTOR, the rotation parameters $\mathbf{R}$ are optimized by minimizing

$$E(\mathbf{R}, \boldsymbol{\theta}) = \langle \Psi(\boldsymbol{\theta}) |\, H_{\text{SCF}}(\mathbf{R})\, | \Psi(\boldsymbol{\theta}) \rangle, \tag{6.30}$$

using gradients and Hessians derived from commutator expressions of the form

$$\frac{\partial h}{\partial R_{pq}}\Big|_{\mathbf{R}=0} = -i[T_{pq}, h_{HF}], \qquad \frac{\partial g}{\partial R_{pq}}\Big|_{\mathbf{R}=0} = -i[T_{pq} \otimes I + I \otimes T_{pq},\, g_{HF}]. \tag{6.31}$$

By allowing rotations over the full orbital manifold, WAHTOR–SCF plays a role analogous to CASSCF rather than CASCI: the algorithm is no longer limited to adapting the Hamiltonian within the active space, but can reshape the entire molecular representation to optimally match the variational ansatz. There is still a caveat of the whole procedure. Even if WAHTOR–SCF employs generators acting on the full molecular orbital space, $\boldsymbol{F}$, the evaluation of energy derivatives still requires a projection onto the active space, $\boldsymbol{A}$. This is a direct consequence of the fact that the VQE wavefunction $|\Psi(\boldsymbol{\theta})\rangle$ and the corresponding reduced density matrices are defined only within the active space. Therefore, each derivative of the transformed Hamiltonian must be reduced to an *effective* operator acting in the active space before its expectation value can be computed. Formally, if $\mathcal{P}_{\mathbf{F}\to\mathbf{A}}$ denotes the operation to reduce from full orbital space, $\mathbf{F}$, to the active space, $\mathbf{A}$, as

$$\mathcal{P}_{\mathbf{F}\to\mathbf{A}}(H_{\text{SCF}}) \to H_{SCF}^{\text{eff}}, \tag{6.32}$$

i.e. applying the reduction procedure of Section 3.1.2, the relevant quantities entering the gradient

$$\frac{\partial E}{\partial R_{pq}} = \langle \Psi(\boldsymbol{\theta})| \frac{\partial H_{\text{SCF}}(\mathbf{R})}{\partial R_{pq}} |\Psi(\boldsymbol{\theta})\rangle \tag{6.33}$$

will be evaluated as

$$\frac{\partial H^{\text{eff}}_{\text{SCF}}}{\partial R_{pq}} = \mathcal{P}_{\mathbf{F}\to\mathbf{A}} \left(\frac{\partial H_{\text{SCF}}(\mathbf{R})}{\partial R_{pq}} \right). \tag{6.34}$$

Therefore, the effective derivative of the energy is given by

$$\frac{\partial E}{\partial R_{pq}} = \langle \Psi(\boldsymbol{\theta})| \frac{\partial H^{\text{eff}}_{\text{SCF}}}{\partial R_{pq}} |\Psi(\boldsymbol{\theta})\rangle, \tag{6.35}$$

computed for each orbital rotation generator T_{pq}. The associated $\partial E/\partial R_{pq}$, will be defined by the effective one- and two-body integrals, $\tilde{h}_{ij}$ and $\tilde{g}_{ijkl}$, where $i, j, k, l \in \boldsymbol{A}$. In practice, this means that the transformed one- and two-electron integrals are first computed in the full orbital basis, then contracted and projected down to an active-space Hamiltonian before their expectation values are evaluated using the active-space wavefunction, $|\Psi(\boldsymbol{\theta})\rangle$. The quantum resource requirements of WAHTOR–SCF procedure remain determined by the size of the active space, scaled by the number of the first- and second–order derivatives, while the classical orbital updates explore the full molecular orbital manifold.

In terms of overall procedure, the schematic workflow, as reported in Figure 6.1, is always the same: the single point evaluation, to obtain the reference wavefunction, $|\psi(\theta)\rangle$, is performed using any quantum algorithm available (in our case VQE and ADAPT-VQE). Once obtained the reference calculation, all the first- and second-order derivatives, according to the desired optimizer scheme, are defined on the full molecular space. The derivatives are then reduced to the active space, and evaluated on the reference wavefunction. From the first derivative, the gradient ∇H is composed, with length equal to the number of orbital rotation generators, N_T, and if needed, also the Hessian $\nabla^2 H$ is defined, in this case with shape $N_T \times N_T$. After the definition of the vector and matrices needed for the update step, the new set of parameters $\boldsymbol{R'}$ is obtained, and used to define the unitary orbital tranformation matrix. The unitary is then applied to the current step Hamiltonian, or specifically, to the current one- and two-body operators independently. Once the new set of integrals is defined, the procedure can start again from the beginning. The algorithm keeps looping between the two routines: CI solution, with VQE (or other algorithm), and the orbital optimization step, using different type of optimizer, until the convergence is met. In our case, the alting criteria is

$$E(\boldsymbol{R'}) - E(\boldsymbol{R}) = \Delta E < \delta, \tag{6.36}$$

where δ is a tunable energy difference threshold. In Figure 6.1, we present the essential difference between the two methods. It is possible to notice that the only component that differentiate between the two algorithm, WAHTOR and WAHTOR–SCF, is related to the orbitals on which the unitary tranformation, U, is defined. The upper panel of Figure 6.1, shows the standard WAHTOR approach, where the orbital on which the tranformation if applied are contained in the subset $\boldsymbol{A}$, i.e. the active-orbital, and also the effective Hamiltonian H_{eff} is defined on the same set. In WAHTOR–SCF instead, lower panel of Figure 6.1, the effective Hamiltonian is still defined on the subset $\boldsymbol{A}$, while the unitary, U, is defined on the orbital in the full molecular space, i.e. $\boldsymbol{F}$.

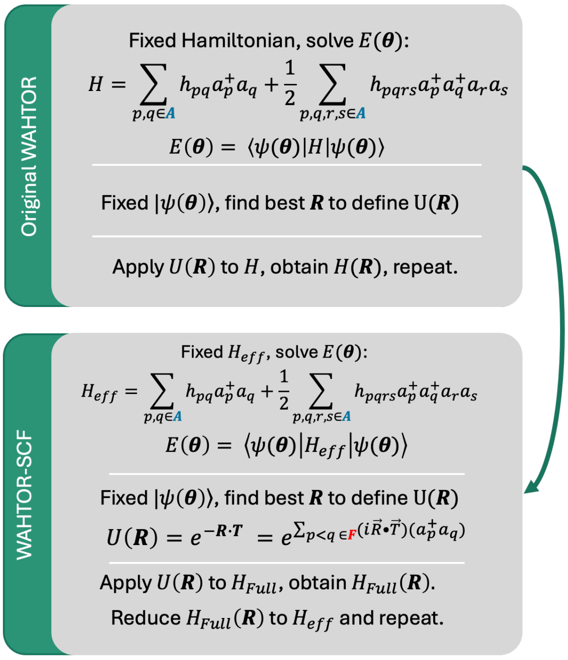

Figure 6.1: Schematization of both WAHTOR (upper panel) and WAHTOR-SCF (lower panel) algorithms.

From the theoretical point of view, WAHTOR–SCF approach is very close the the hybrid quantum algorithms listed above, like quantum CASSCF, Optimal Orbital VQE, ADAPT-VQE–SCF-SCF and OO–VQE. In all the cases, the CI solution is performed on the QPU, using either a fixed ansatz or an adaptive one, and then, once the single point is converged, it is used to evaluate the molecular orbital gradients w.r.t orbital rotation parameters. For ADAPT-VQE–SCF and OO–VQE, this step is obtained by evaluating 1- and 2-RDMs and then contract them to the analytical orbital rotation derivatives generators, completely from a classical point of view. In WAHTOR–SCF, instead, following the original idea of WAHTOR, the derivatives are evaluated as fermionic operators on the quantum computer. In terms of ansatz, we formalized two versions, with Hardware–Efficient Ansatz (HEA), namely HEA–WAHTOR–SCF, and with adaptive ansatz, namely ADAPT–WAHTOR–SCF. Additionally, we kept the same freedom of WAHTOR in the selection of the group of orbital between which the generators are constructed, implementing theoretical-based orbital generation reduction schemes, defined as state-of-the-art CASSCF or exploiting Point Group symmetries.

From the computational point of view instead, the implementation of such algorithm it everything but easy. Due to the necessity of measuring the derivatives as operators, defining a scalable and efficient algorithm is more that essential. Therefore, a huge effort has been spent into providing the best-performing implementation of the WAHTOR–SCF algorithm, starting from defining an efficient codebase, passing by the definition of memory efficient routines, and finally, building a fully parallelization infrastructure.

In the next section, we are going to present all the computational improvement that have been done on the QuAQ code base and to the original WAHTOR algorithm code, as well as the theoretical improvement and enhancement.

6.3.1 Terms, variables and associated cost

In order to fully explain the computational and memory requirements, a brief explanation of the nomenclature is needed.

Parameters:

- N_R: Number of MOs that are going to be considered in the SCF cycle. This is related to the size of the molecule, the atoms included, the basis-set, and whether FCA or FSA is used.
- N_A: Number of MOs included in the CAS. The selection of orbitals is done according to the methods explained in Section 4.2. In this situation, for solving a CASSCF procedure, generally the condition is $\{\psi_i\} \in \mathbf{A} \subseteq \mathbf{F}$, where i is a generic MO index, $\mathbf{A}$ is the subset of active MOs, and $\mathbf{F}$ is the total MOs space. Being an expensive procedure, and requiring ideally the resolution at ED level in the CAS, for big molecules or for extended basis-set, $N_A \ll N_R$.
- N_Q: Number of qubits used to define the circuit. The number of qubits is directly proportional to the number of orbitals in the CAS, in fact, $N_Q = 2 \times N_A$ (This is generally true, but it depends on the F2Q mapper used).
- *Double* entry for the integrals: each term in the integrals tensors ($\hat{h}$ for the 1-body and $\hat{\Gamma}$ for the 2-body) weights as a `np.float64` i.e. 64-bit precision floating-point (correspond to a standard C `double`). In the detail: 1 bit for the sign, 11 bits exponent, 52 bits mantissa.
- *Complex* entry for intermediate matrices: each term in a matrix representation weights as a `np.complex64` i.e. 32-bit for the real and 32-bit for the imaginary, both `np.float32`(each of them corresponds to a C `float`). In the detail: one bit for the sign, 8 bits exponent, 23 bits mantissa for imaginary and real part. This intermediate operators are related to the contraction of the integrals ($\hat{h}$, $\hat{\Gamma}$) with the orbital rotation generators (e.g. $\hat{T}_{ij}$ defined in the next section).

Main scaling: According to the original pipeline, the scalings that the procedure needs to take into account are:

- Orbital rotation generators $\boldsymbol{T}$: The number of generators, N_T, scales with the number of MOs in the total space considered for the SCF step. Without loss of generality, we consider it of the same size of the total MOs space $\mathbf{F}$ i.e. $dim(\mathbf{F}) = N_R$. Therefore, the number of generators (without considering any reduction, which are going to be introduce in the next sections) N_T is going to scale as $N_T \in \mathcal{O}(N_R^2)$. This scaling corresponds to consider all the combination $\boldsymbol{T} = \{\hat{T}_{ij} | i, j \in \mathbf{F}\}$.
- Parameters $\boldsymbol{\kappa}$ of the unitary $\hat{U}(\boldsymbol{\kappa})$: each parameter $\kappa_{ij} \in \boldsymbol{\kappa}$ is associated to the corresponding orbital rotation generator $\hat{T}_{ij} \in \boldsymbol{T}$. As consequence, $N_P = dim(\boldsymbol{\kappa}) = dim(\boldsymbol{T}) = N_T$.
- Number of first derivatives N_{dH} : The size of the gradient vector (or Jacobian ∇) scales 1-to-1 as the number of generators, i.e. $N_{dH} = N_T$.
- Number of second derivatives N_{ddH}: The size of the second derivative matrix (or Hessian ∇^2) scales, without taking into account any symmetry, with $N_{ddH} = N_{dH}^2 = N_T^2 \in \mathcal{O}(N_R^4)$.

- Number of operator kept in memory N_O: Without any optimization, the number of operator to be kept in memory scales as

$$N_O \in \mathcal{O}(1) + \mathcal{O}(N_R^2) + \mathcal{O}(N_R^4) = \mathcal{O}(N_R^4).$$

- Size of each operator $\hat{O}$: Each operator evaluated need first to be defined through the integrals and after as a qubit operator. As operator we include the original Hamiltonian $\hat{H}$, each of the first derivative operator $d\hat{H}$, and also the second derivative operators $dd\hat{H}$. All the operators are defined from the one- and two-body integrals, $\hat{h}$ and $\hat{\Gamma}$, respectively. The total memory footprint of a generic operator, originated from a chemistry calculation, is

$$M_O = dim(\hat{H}) = 2 \cdot \underbrace{(N_R \times N_R)}_{\hat{h}\ \text{size}} + 3 \cdot \underbrace{(N_R \times N_R \times N_R \times N_R)}_{\hat{\Gamma}\ \text{size}},$$

  where the factor 2 comes from the h_α and h_β, while the factor 3 comes from $\Gamma_{\alpha\alpha}$, $\Gamma_{\beta\beta}$, and $\Gamma_{\beta\alpha}$, which corresponds to the different one- and two-body integrals as defined in Section 2.2.5. Therefore, the size of an operator scales as $M_O \in \mathcal{O}(N_R^4) \cdot float$.

- Total memory usage M_T: Without any further reduction, the amount of memory required scales as

$$M_T \in M_O \cdot N_O = \mathcal{O}(N_R^4) \cdot \mathcal{O}(N_R^4) = \mathcal{O}(N_R^8).$$

- Number of Paulis L: Following the definition of Section 3.8, each operator, $\hat{O}$, once mapped from Fermionic space to the Qubit space by mean of a F2Q mapper, is going to be composed as a linear combination of Pauli strings

$$\hat{O} \xrightarrow{\text{F2Q}} \hat{O}_q = \sum_i^L c_i \hat{P}_i \quad \text{where} \quad \hat{P}_i = \bigotimes_{j=0}^{N_Q-1} \hat{\sigma}_i^j\,, \quad \sigma_i^j \in \{\hat{I}, \hat{X}, \hat{Y}, \hat{Z}\},$$

  where L is scales as $L \in \mathcal{O}(N_Q^4)$. In terms of composition, each operator is defined as a `qiskit.SparsePauli` object, in which each term $(c_i,\ \hat{P}_i)$ is composed by a coefficient, `np.complex128`, and `qiskit.PauliList` object.

6.3.2 Expectation values of operator and Operators generation

In order to make the procedure run smoothly, a series of optimization rounds have been performed, targeting different bottlenecks that were impacting the procedure. Firts of all, from the point of view of the backend, i.e. the QuAQ [224] codebase, was required to be cleaned and improved computationally. Related to how fast and efficient is the creation (transformation from integrals to `SparsePauliOp`) of each operator in the procedure. Given the fact that, for each iteration of WAHTOR and WAHTOR-SCF, the number of derivatives easily reach thousands or even more, a smooth qubit operator pipeline is essential. The critical point were mainly, the conversion between fermionic operators to qubit operators, and how the expectation values of an operator with a given trial state were done. An unique and simple solution has been found by substituting each qubit operator conversion from the simple dense representation, `to_matrix()`, to the sparse version, `to_matrix(sparse=True)`. In terms of memory, the storage is limited to only the entries different from zero, while in terms of construction efficiency, given the fact that each operator is a linear combination of tensor product of Pauli matrices, and each of them, if composed by p Pauli, can be represented by a sparse matrix with only $2p$ non-zero entries, the complete sparse construction becomes fairly quick. For smaller operator instead, less than 8 qubits, the sparse construction time is comparable with the dense one. The difference in terms construction times start to get noticeable when the number of qubits is greater than 14, as can be seen in Table 6.1 Additionally,

CAS(N_e,N_A)	Qubits	Dense	Sparse
CAS(4,4)	8	152 μs	208 μs
CAS(6,5)	10	3.06 ms	2.3 ms
CAS(6,6)	12	30.4 ms	7.19 ms
CAS(8,7)	14	897 ms	35.9 ms
CAS(8,8)	16	34.1 s	208 ms

Table 6.1: Time required to construct a dense and a spase operator from a `qiskit.SparsePauliOp`. Construction from a N_2 with cc-pVDZ basis with diffent CAS sizes. Each CAS(N_e,N_A) is defined by the number of electrons N_e and number of spatial orbitals N_A. The qubits are, as already mentioned, $2N_A$.

we also investigate the proper way to compute the expectation value, $\langle\psi| H |\psi\rangle$ efficiently in our range of application, i.e. a number of qubits between 4 and 16. In terms of evaluation methods we have analyzed: standard numpy matrix multiplication operator, `@`, numpy Einstein summation, `np.einsum`, and the combination of column vector product (left-side vector-matrix multiplication) with the concatenation operator, `np.vdot(` $|psi\rangle$ `,` H `@` $|psi\rangle$ `)`. Both `@` and combined `np.vdot` have been tested with dense and sparse, while `np.einsum`, only with dense operator. Already by using sparse matrices, in particular `csr_matrix` from `scipy.sparse`, the overall timing is about two order of magnitude faster, compared to methods using dense matrices.

Qubits	`@` Dense	`@` Sparse	`np.einsum` Dense	`np.vdot` + `@` Dense	`np.vdot` + `@` Sparse
4	0.000014	0.000161	0.000039	0.000007	0.000047
8	0.000245	0.000254	0.000551	0.000145	0.000110
10	0.001161	0.000491	0.004633	0.000975	0.000130
12	0.027064	0.001454	0.060489	0.013903	0.000507
14	0.499348	0.004050	0.729769	0.225560	0.004401
16	115.461459	0.119898	74.895082	45.548152	0.825535

Table 6.2: Time required to construct to evaluate the expectation value of a random vector, an operator in dense and sparse form with different contraction methods. Construction from a N_2 with cc-pVDZ basis with diffent CAS sizes, represented with the number of qubits.

It is clear by looking at Figure 6.2 and Table 6.2, that general sparse matrix based methods approaches results faster compared to dense ones, in the range of 8 to 16 qubits, with a peak improvement of almost 3 order of magnitude at 16 qubits. This is essential especially during the VQE optimization routine and for the evaluations of all the derivatives operator. Additionally, the whole expectation values operation has been further improved by exploiting properties of the derivative operators, discussed in the following section.

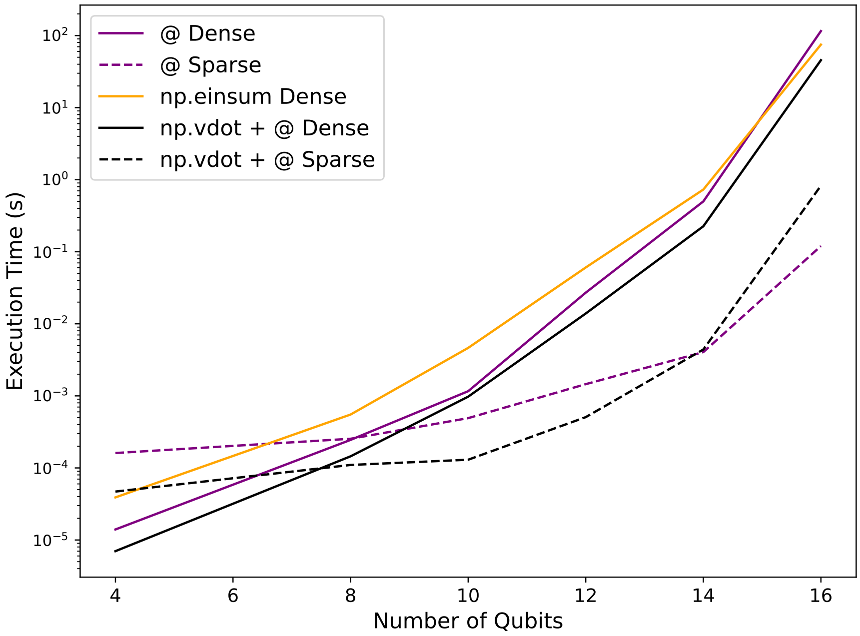

Figure 6.2: Expectation value comparison methods between sparse and dense operator contracted with a random $|\psi\rangle$ complex statevector.

6.3.3 Memory leakage

In multiple parts of the code, a non optimal management of the memory, especially due to massive usage of `copy.deepcopy()` and improper routines, was causing slowdowns, and, in the case of multi-processing and multi-node parallelization, also a memory blow-up, i.e. an infinite increase of memory usage. A correct management of the memory is essential when dealing with *embarrassingly* parallel tasks and the final objective is to have a code that keeps a linear usage of memory in the number of processes involved in the procedure. Python does not have a traditional, fully automatic garbage collector like Java. Instead, it relies mainly on reference counting to manage memory, each object keeps track of how many references point to it, and when that count drops to zero, the memory is immediately freed. However, this approach struggles with circular references, where two or more objects reference each other, preventing their reference counts from reaching zero. To handle such cases, Python includes a cyclic garbage collector, but it is limited and does not always efficiently clean up all cycles. Additionally, due to the Global Interpreter Lock (GIL) and Python's PyMalloc allocator, memory fragmentation can occur, and freed memory is not always returned to the operating system. These factors together cause Python to sometimes appear as if it has memory leaks or poor garbage collection performance.

`Deepcopy()` overuse

Each time a fermionic operator is defined starting from one- and two-body integrals, one of the essential step, as explained in Section 3.8, is the reduction from the full system to the active sub-system only, i.e. creating the effective operators. Given the fact that, in WAHTOR and consequently WAHTOR-SCF, each derivative (first and second order) are required to be evaluated on the wavefunction associated with the solution in the restricted space, a step of *integrals-cutting*[1]. In this step, before the code optimization, each of the tensor containing the one-body for each spin sector and the two-body tensors for all the three spin combinations ($\alpha\alpha$, $\beta\beta$, and $\beta\alpha$) were first copied, using `copy.deepcopy()` method, and after, the slices, i.e. the specific entries, associated with the active orbital space and the core orbital space, were extracted. The extrapolation was perfomed using the method `numpy.ix_()`, which create a mask that select only the entries desired. If defined as `numpy.ix_([a,b, ...], [c,d,...])` is going to create a mask for the entries obtained by the cartesian product of `[a,b, ...]` and `[c,d,...]`. When a tensor Γ is indexed using `np.ix_()`,

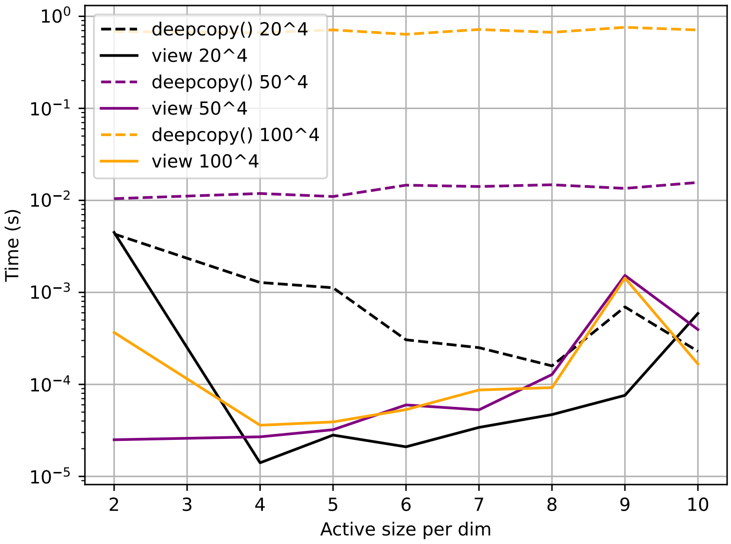

Figure 6.3: Comparison between extracting the two-body integrals entries associated with the active region. The size of the tensor is 20, 50, and 100, which are general size of the total molecular space we treat, and the active size is in the regime of orbitals we can deal with, i.e. $[2, 10]$. The extractions methods are done accordingly to the previous explanation: *deepcopy()* using the said method and *view* using the advanced-view.

NumPy constructs a broadcasted view or advanced-indexed copy depending on the indexing pattern. In the case of `np.ix_()`, the operation typically triggers advanced indexing, which results in the allocation of a new array rather than a view into the original memory buffer. Consequently, the resulting sub-tensor is already a distinct memory object, and subsequent contraction operations (e.g., via `np.einsum`, `np.tensordot`, or other tensor contractions) act on this new data block. Since tensor contractions inherently produce new output arrays, no aliasing occurs with the source tensor. Thus, performing a `deepcopy()` is redundant and inefficient. A shallow `copy()` would only be relevant if one intends to perform in-place modifications on the extracted sub-tensor prior to contraction. But this was not the case, therefore, a simple modification in the code, i.e. removing all the `deepcopy()`s, allowed for enhanced computational speed. From Figure 6.3 is it possible to notice the difference between the two extraction methods, which become more evident dealing with more than 20 molecular orbital in the full space. In the regime of 100 orbitals (and more), the reduction is almost four order of magnitude. This step, we repeat, is extremely important because is the core for reducing, not only the main molecular Hamiltonian $\hat{H}$, but all the first- and second-order derivatives, which can easily reach thousands of operators.

---

[1] Integrals cutting is referred to the reduction of the one- and two-body integrals from being spanned on the whole $\boldsymbol{F}$ molecular orbital space, to be reduce on the $\boldsymbol{A}$ active orbital space.

Derivatives calculation

The naive implementation of the derivative computation suffers from several inherent inefficiencies that limit its scalability for large Hamiltonian systems. The generator matrix T is highly sparse, containing only two nonzero elements per transformation block, corresponding to the antisymmetric pair $T_{ij} = i$ and $T_{ji} = -i$. Despite operating on such a simple two-element structure, the original method allocates full-dimensional temporary arrays for each nonzero entry:

```
temp = np.zeros((dim, dim), dtype=complex)
temp[gr, :] = T[r][c] * h[ham][gc, :]
temp[:, gc] = temp[:, gc] - T[r][c] * h[ham][:, gr]
```

Snippet 6.1: Integrals-Generator commutator before optimization.

The creation of `temp` incurs an unnecessary $\mathcal{O}(N^2)$ or $\mathcal{O}(N^4)$ memory cost, even though only a small subset of its entries (corresponding to the active indices i and j) are actually modified. Because NumPy's array creation semantics always initialize the full buffer in memory, each call to `np.zeros(...)` triggers a dense memory allocation and zero-fill operation regardless of sparsity. Consequently, the total runtime and memory footprint become dominated by allocation overheads rather than arithmetic operations. This inefficiency completely masks the potential computational savings that should arise from the generator's two-element sparsity.

Furthermore, each iteration reconstructs temporary arrays and performs duplicated slice updates, resulting in repeated cache invalidations and unnecessary data movement. These allocations create large contiguous memory buffers that are mostly unused, resulting in excessive memory consumption and heap fragmentation. Additionally, the repeated summation

```
comm_1 += temp
```

Snippet 6.2: Temporary commutator storing.

triggers internal temporary allocations within NumPy, amplifying memory pressure. Redundant computation of the antisymmetric pair (i, j) and (j, i), combined with Python-level loops over indices, further increases interpreter overhead and reduces cache locality. Because the arithmetic intensity is extremely low (only two active tensor contractions per commutator), these redundant memory operations dominate performance and make the algorithm *memory-bound* rather than compute-bound. As we can see from Figure 6.4, before the method optimization, the memory was keep increasing as the workers, or processes, were spawned and were finalizing their tasks. In the picture, the memory profiler shows that, after all the 96 workers (8 MPI ranks and 12 processes each), the memory usage is not cleaned, or at least reduced, but was keep increasing. The improved version has been defined to avoid all the issues above, removing all the Python-level loops and reducing the caching between operation, as well as making each evaluation constrained to the specific spin sector and type of integrals[2] Therefore, the refactoring restructures the computation of the single-restricted derivative to exploit the extreme sparsity of the generator matrix T and minimize memory allocations. The generator has exactly two non-zero elements per block, corresponding to the antisymmetric pair $T_{ij} = \frac{i}{\sqrt{2}}$ and $T_{ji} = -\frac{i}{\sqrt{2}}$. Instead of allocating a full temporary array for each nonzero entry, the algorithm pre-computes the nonzero indices once using

---

[2] i.e. one- or two-body integrals.

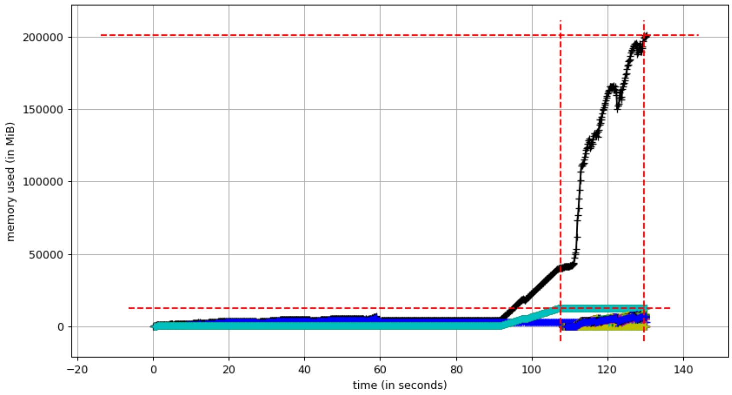

Figure 6.4: Memory blow-up observed as the step of derivative evaluation is started on multiple MPI-workers, each of which is spawning multiple processes.8 MPI workers, 12 processes for each worker.

```python
(r1, c1), (r2, c2) = zip(*np.nonzero(T))
```

Snippet 6.3: Optimized numpy in-place tensor operation

and then extracts the corresponding group indices and factors:

```python
gr1, gc1 = group_idx[r1], group_idx[c1]
gr2, gc2 = group_idx[r2], group_idx[c2]
Tij, Tji = T[r1, c1], T[r2, c2]
```

Snippet 6.4: Index matching.

For each Hamiltonian block $h[k]$, a pre-allocated commutator array `comm` of the same shape is created with a low-precision dtype (`np.complex64`) to reduce memory footprint. Updates to `comm` are performed entirely *in-place* using NumPy's broadcasting and `out=` argument:

```python
np.add(comm[gr1, :], Tij * h[k][gc1, :], out=comm[gr1, :])
np.subtract(comm[:, gc1], Tij * h[k][:, gr1], out=comm[:, gc1])
np.add(comm[gr2, :], Tji * h[k][gc2, :], out=comm[gr2, :])
np.subtract(comm[:, gc2], Tji * h[k][:, gr2], out=comm[:, gc2])
```

Snippet 6.5: Optimized numpy in-place tensor operation.

This approach ensures that only the active rows and columns corresponding to the nonzero generator entries are updated, avoiding any dense allocations for inactive portions of the tensor. Once the commutator is fully computed, it is multiplied by $-1i$ and yielded immediately:

```python
yield k, -1j * comm
del comm
gc.collect()
```

Snippet 6.6: Yielding and garbage collection.

On overall, the enhancement between old and the improved version of the derivative calculation method are

1. **Memory efficiency**: Peak memory usage scales with the size of a single Hamiltonian block ($\mathcal{O}(N^4)$ in the worst case) rather than the number of nonzero entries in T. No temporary arrays proportional to the full tensor are created.

2. **Streaming computation**: Using `yield`, blocks are returned as soon as they are computed, eliminating the need to store all results in memory simultaneously.

3. **In-place arithmetic**: Updates to `comm` avoid creating intermediate temporaries, improving cache locality and reducing pressure on the Python garbage collector.

4. **Exploitation of sparsity**: The algorithm touches only the rows and columns corresponding to the nonzero entries of T, scaling linearly with the actual sparsity rather than the full block size.

With the improved version of the derivative calculation, the memory blow-up has been contained, as it can be evinced from Figure 6.5, in which the same calculation as in Figure 6.4, has been performed. Here, the 96 workers, with the same setup as before, use a constant number of resources. The memory requirement remains constant for each worker, i.e. it scales lineraly with the number of processes spawned. In the particular case, the weight of the original Hamiltonian and each derivative is around 1.5Gb, for a total of an average memory occupation of 140Gb.

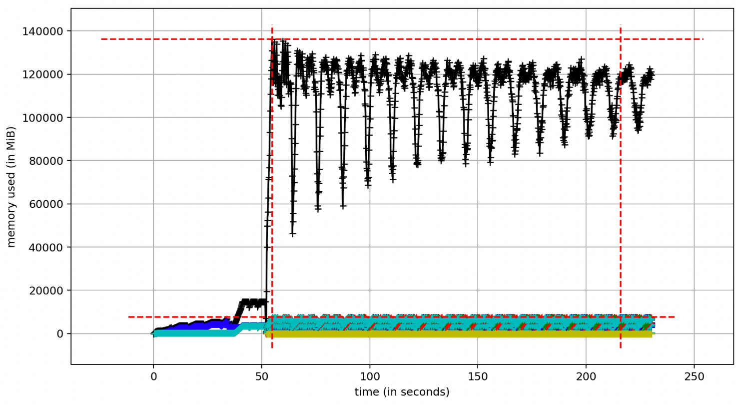

Figure 6.5: Memory blow-up contained after the optimization of the derivative evaluation routine. 8 MPI workers, 12 processes for each worker. The print is truncated after the first round of derivatives. Each drop in the memory usage correspond to the proper memory freed and garbage collection.

6.3.4 Generators pool

As theoretical-inspired optimization, we decided to follow the same grouping and space-reduction as normally performed in a CASSCF/MC-SCF procedure. This optimization is going to target the richness of the parameters space $\boldsymbol{\kappa}$ by removing the orbital rotation parameters that are not meant to rotate between each other, i.e. either forbbiden mixing, thanks to Poing group theory, or iso-energetic transformation, following electronic structure theory. The orbital rotation unitary matrix $\hat{U}(\boldsymbol{\kappa})$ is defined according to

$$\hat{U}(\boldsymbol{\kappa}) = e^{-i\boldsymbol{\kappa}\cdot\hat{\boldsymbol{T}}} = e^{-i\sum_{i,j\in\mathbf{F}}\kappa_{ij}\hat{T}_{ij}}, \tag{6.37}$$

Therefore, the reduction in the number of parameters is closely related to the number of generators in $\hat{\boldsymbol{T}}$. Here, we report the different policy that can be used and how they impact N_T.

- **ALL**: This is the default setup used also in the WAHTOR code. Given a specific group of orbitals, suppose $\mathbf{F}$, the set of generators, $\hat{\boldsymbol{T}}$, is defined by building

$$\hat{\boldsymbol{T}} = \{\hat{T}_{ij}|\ i \in \mathbf{F}, j \in \mathbf{F}, i > j\},$$

  leading to a number of generators N_T that corresponds to N_R^2. Intuitively, but mainly according to quantum chemistry knowledge, most of these rotations are going to be zeros due to the type of orbital they are trying to mix and the underlying symmetries.

- **CAS-LIKE**: Taking inspiration from standard MCSCF/CASSCF procedures, the complete orbital space $\mathbf{F}$ can be manipulated using the different components that build it i.e. $\mathbf{A}, \mathbf{O}$, and $\mathbf{V}$, respectively, active, occupied-inactive, and virtual-inactive orbitals. The generators are defined from the indices obtained by the combination of the three sets in pairs. In particular, the total set of generators is defined as

$$\begin{aligned} \hat{\boldsymbol{T}} &= \hat{\boldsymbol{T}}_O^A \cup \hat{\boldsymbol{T}}_O^V \cup \hat{\boldsymbol{T}}_A^A \cup \hat{\boldsymbol{T}}_A^V = \\ &= \{\hat{T}_{ij}|i \in \mathbf{O}, j \in \mathbf{A}\} \cup \{\hat{T}_{ij}|i \in \mathbf{O}, j \in \mathbf{V}\} \\ &\cup \{\hat{T}_{ij}|i \in \mathbf{A}, j \in \mathbf{A}, i > j\} \cup \{\hat{T}_{ij}|i \in \mathbf{A}, j \in \mathbf{V}\}, \end{aligned} \tag{6.38}$$

  where $\hat{\boldsymbol{T}}_O^A$ are the generators that rotate orbital between occupied and active, $\hat{\boldsymbol{T}}_O^V$ rotate between occupied and virtual, $\hat{\boldsymbol{T}}_A^A$ between active orbitals, and $\hat{\boldsymbol{T}}_A^V$ between active and virtual. This structure follows the same subdivision applied in a standard MCSCF/CASSCF calculation. It is possible to note how the rotations between occupied orbitals and between virtual orbitals are not considered due to the Brillouin's conditions. These conditions states that the occupied-occupied rotation does not change the determinant but only they redefine the occupied space, the virtual-virtual one do not contribute to the energy. Differently from CASSCF/MCSCF, the set of active-active rotations are required to improve the solution of the CAS space whenever the CI expansion is a truncation of the true ground-state. The total number of generators is defined as

$$\begin{aligned} N_T &= U_A^A \cup U_O^A \cup U_A^V \cup U_O^V \\ &= \underbrace{\frac{|\boldsymbol{A}| \cdot (|\boldsymbol{A}| - 1)}{2}}_{active-active} + \underbrace{|\boldsymbol{A}| \cdot (|\boldsymbol{O}| + |\boldsymbol{V}|)}_{\substack{occupied-active \\ active-virtual}} + \underbrace{|\boldsymbol{O}| \cdot |\boldsymbol{V}|}_{occupied-virtual} . \end{aligned}$$

  It is possible to notice how, the scaling N_T depends mainly from one term, the occupied-virtual rotations, which, especially for a in-active-space calculations, consist in a unbalanceness between the ratio $\mathbf{A} + \mathbf{O}$ and $\mathbf{V}$. For example, CAS(6,6) on C_6H_6 in cc-pVDZ leads to 18, 6, 90 orbitals, for $\mathbf{O}$, $\mathbf{A}$, and $\mathbf{V}$, respectively. Therefore, the $|\mathbf{A}| + |\mathbf{O}| \ll |\mathbf{V}|$, for most of the CAS calculations.

- **RED**: A minimal pool defined by removing one of the set of generators that contributes the most on the total number N_T. Starting from **CAS-LIKE**, the generators associated with the occupied-virtual rotations are removed. In particular, the total set of generators is going to be defined as

$$\begin{aligned}\hat{\boldsymbol{T}} =&\hat{\boldsymbol{T}}_O^A \cup \hat{\boldsymbol{T}}_A^A \cup \hat{\boldsymbol{T}}_A^V = \\ =&\{\hat{T}_{ij} | i \in \mathbf{O}, j \in \mathbf{A}\} \cup \{\hat{T}_{ij} | i \in \mathbf{A}, j \in \mathbf{A}, i > j\} \\ &\cup \{\hat{T}_{ij} | i \in \mathbf{A}, j \in \mathbf{V}\}.\end{aligned} \tag{6.39}$$

The total number of generators N_T is now composed only by three terms, the active-active, occupied-active, and active-virtual, as follows

$$\begin{aligned}N_T &= U_A^A \cup U_O^A \cup U_A^V \\ &= \frac{|\boldsymbol{A}| \cdot (|\boldsymbol{A}| - 1)}{2} + |\boldsymbol{A}| \cdot (|\boldsymbol{O}| + |\boldsymbol{V}|).\end{aligned}$$

This pool, especially for system in which the correlation in the virtual space is not dominant, leads to the same optimization results as the **CAS-LIKE** at a cheaper computational requirement.

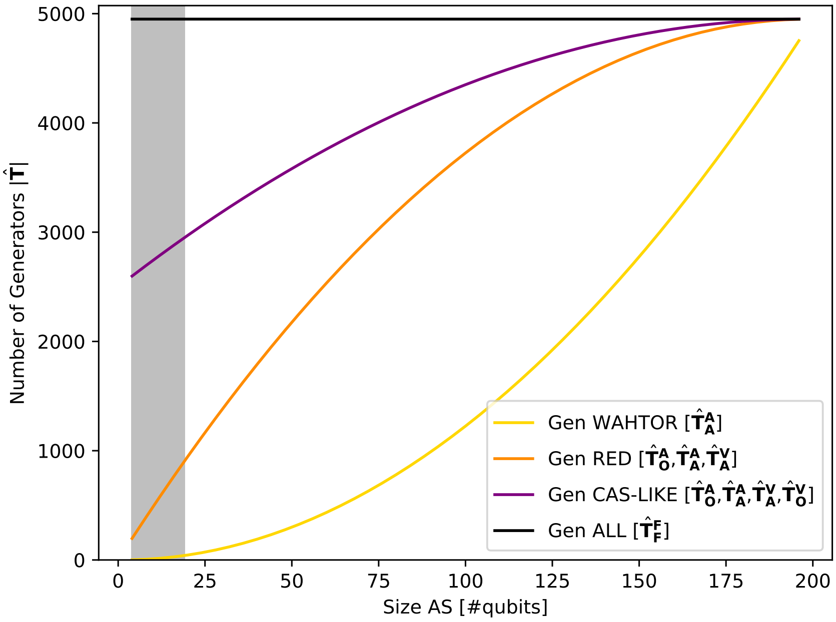

Figure 6.6: Ideal scaling of the number of generators (y-axis) as function of active orbitals, A, occupied orbitals, O, and virtual orbitals, V, with different pool reduction policies. The shaded area corresponds to the sizes of active space in the regime of use that can be reached by QuAQ. The values are obtained considering a balanced active space, equal number of occupied and virtual orbitals, starting from a CAS(2,2) with orbitals [49,50] as active space, in a system with 100 orbitals and 100 electrons.

In Figure 6.6, we present a graphical representation on how, the different type of pools, impacts the number of generators N_T. A dummy molecular space has been employed, specifically, a total of 100 MOs and 100 electrons for the full space. Then, for each value $i \in [1, 50)$, an active-space corresponding to a CAS($2i$, $2i$) is defined on the Fermi-level, i.e. picking i occupied orbitals and i virtual orbitals, for a total of $2i$ electrons. From the plot, it is possible to observe how the different pools impact directly the computational resources required to compute a single step of the orbital optimization procedure. The first line, *Gen WAHTOR*, corresponds to the scaling of the rotation in the active space, i.e. the old WAHTOR scheme without any in-active-space item clustering. The last one, *Gen ALL*, i.e **ALL** pool, and it corresponds to the direct extension of the WAHTOR scheme for the whole molecular space SCF procedure. Both of them, represent extreme setup of the procedure, an OO-VQE-like approach for the former, and a high-resource-hunger CASSCF for the latter. The remaining two lines, *Gen CAS-like* and *Gen Red*, correspond to **CAS-LIKE** and **RED** pool respectively, and they account only for chemically meaningful set of orbital rotation generators, leading to a reduction in the computational demand.

- **PG–RED**: In addition, as usually done by the standard quantum chemistry package, molecular information can be used to further reduce the number of generators involved in the orbital optimization procedure. A valid and robust selection criteria is to employ Point–Group (PG) symmetries (For a detailed and exhaustive discussion on Point-Group symmetries and related theory applied to molecular chemistry, see [225]). The **PG–RED** can be applied on top of each of the previous pool and it removes all the orbital–rotations generator that involves two orbitals with different irreducible representation. For example, couple with **ALL**, it will iterate over the pairs $i < j \in \boldsymbol{F}$ and keep all the generators $\hat{T}_{i,j}$ such that $\chi(\phi_i) = \chi(\phi_j)$, where ϕ_i and ϕ_j are the associate molecular orbitals, and $\chi(\phi_i)$ is the irreducible representation associated with orbital i.

6.3.5 Pauli Recycling

Returning back to the original WAHTOR paper, the analysis of the individual Pauli strings was required. We wanted to keep the same structure and improve also this functionality. In the original WAHTOR code, the procedure for evaluating the operators was, from a pure simulation point of view, completely inefficient. Due to the assumption that the simulation is only carried on with a noiseless statevector simulator, measuring for each VQE micro-step all the Pauli strings composing the Hamiltonian becomes prohibitive, especially for a number of qubits greater than 10. Therefore, as first changed in the procedure, only for the VQE optimization step, the Hamiltonian $\hat{H}$ is saved and used as a unique `csr_matrix` sparse matrix. After the optimized wavefunction is found $|\psi\rangle$, the Hamiltonian $\hat{H}$ is reconstructed as single Pauli strings, as a standard qubit operator would be. Starting form this expectation values, we improved the whole derivative evaluations. Therefore, to efficiently evaluate derivative Hamiltonians without constructing full matrix representations, each derivative is represented as a Pauli string list and wavefunction evaluations are recycled. This approach avoids the exponential memory scaling associated with explicit Hamiltonian matrices and minimizes redundant computations across multiple derivatives. The whole procedure is based on the assumption that each derivative of $\hat{H}$ is composed by a subset of the Pauli string of the original Hamiltonian. If the Pauli strings are more that the one in the Hamiltonian, the whole procedure might not be as efficient as expected.

As already explained in the previous sections, each derivative is computed from the original Hamiltonian contracted with the orbital rotations commutator, and then truncated/cutted to fit the active orbital space, i.e. the effective Hamiltonian derivatives. The truncated Hamiltonians are then converted to a list of Fermionic operators through `quick_FermionicOp`, which build the Fermionic operator corresponding to the derivative Hamiltonian starting from the integrals. If the the derivative is already composed by only zero elements, it gets discarded. If they contain non-zero elements, these Fermionic operators are subsequently mapped to qubit Pauli operators via the

Jordan-Wigner transformation using `QubitOperator`, producing a simplified list of nonzero Pauli strings without constructing the full qubit Hamiltonian.

To avoid repeated evaluation of the same Pauli strings on the reference state $|\psi\rangle$, expectation values are cached. In the main procedure, when Pauli recycling is enabled, the Pauli strings and their coefficients are extracted, and a Numba-accelerated cache is built as:

```python
CODE = np.zeros(128, dtype=np.int64)
CODE[ord(’I’)] = 0b00
CODE[ord(’X’)] = 0b01
CODE[ord(’Y’)] = 0b10
CODE[ord(’Z’)] = 0b11

def encode_pauli_string(pauli_str):
    result = 0
    for i in range(len(pauli_str)):
        c = pauli_str[i]
        result |= (CODE[ord(c)] << (2 * i))
    return result

def build_numba_pauli_cache(pauli_strs: list[Pauli],
                            psi: np.ndarray) -> dict[int, complex]:
    cache =  dict()
    for p in pauli_strs:
        key = encode_pauli_string(str(p))
        cache[key] = np.vdot(psi, p.to_matrix(sparse = True) @ psi)

    return cache
```

Snippet 6.7: Pauli recycling utils methods.

Here, `build_numba_pauli_cache` encodes each Pauli string into an integer key,representing the hashing of the Pauli string, using `encode_pauli_string` and stores the expectation value

$$\mathrm{cache}[P] = \langle\psi|P|\psi\rangle \tag{6.40}$$

in a dictionary. The integer encoding is achieved via a simple bitwise mapping of ASCII characters (`I=00`, `X=01`, `Y=10`, `Z=11`) into integer keys, which enables rapid dictionary lookup and Numba acceleration. Once the cache is built, the expectation vector for each derivative Hamiltonian is obtained by contracting the cached values (plain expectation values) with the corresponding Pauli coefficients using `get_expectation_vector_numba` :

```python
    def get_expectation_vector_numba(cache: dict[int, complex],
                                     query_paulis: list[Pauli],
                                     query_coeffs: list[complex]) -> np.ndarray:

    vec = 0.0 #np.zeros(len(query_paulis), dtype=np.complex128)
    for i in range(len(query_paulis)):
        key = encode_pauli_string(str(query_paulis[i]))
        vec += cache.get(key, 0.0) * query_coeffs[i]
    del query_paulis, query_coeffs
    return vec
```

Snippet 6.8: Expectation value reconstruction with Pauli recycling.

This corresponds to effectively computing

$$\text{vec} = \sum_i c_i \langle\psi|P_i|\psi\rangle \equiv \sum_i c_i \cdot \texttt{cache.get(Pi)} \tag{6.41}$$

without creating full Hamiltonian matrices or intermediate dense arrays. The resulting vector is directly used in the evaluation of derivative Hamiltonians, eliminating the need to track individual coefficients during contraction.

This Pauli-string approach provides significant memory and computational advantages. Memory usage scales linearly with the number of nonzero Pauli strings rather than the full Hilbert space dimension, allowing simulations of large qubit systems. Redundant computation is minimized, as each Pauli expectation value is computed once and reused across all derivative evaluations. The contraction step touches only the active Pauli strings, avoiding operations on the full Hilbert space and reducing intermediate memory allocation. Explicit deletion of intermediate objects, such as derivative fermionic operators (`fermop`), derivated integrals (`xh`), and on-fly qubit operators (`q_temp`), combined with calls to `gc.collect()`, prevents memory accumulation over iterative derivative computations. Additionally, the use of integer encoding and Numba acceleration reduces Python interpreter overhead, enabling high-throughput evaluation. The general derivative calculation method with and Pauli recycling is reported in Appendix 1.2.

In term of best performance, the methods shines when dealing with Hamiltonians and derivatives which active space is bigger than 8 qubits. In Table 6.3, we report the computational time to perform a single WAHTOR-SCF iteration between two regimes: small derivatives (6 qubit) and mid sized ones (12 qubits), with and without Pauli recycling. It is clear that, with smaller effective operators, the diffence in the performance is almost negligible, due to the fact that, as reported in Table 6.1, for small spaces the time to construct and evaluate the operators is very low.

Molecule	**Qubits**	Total Derivatives	**No Recycle**	**Recycle**
N_2	12	230	618	6.78
C_2H_4+H	6	11934	1696.4	1406.04

Table 6.3: Different timings, in seconds, obtained by using OpenMP parallelization and recycling of Pauli strings. 24 cores used for each computation. The two systems are N_2 cc-pVDZ basis (28 total orbitals) and the transition state of model free radical polymerization of ethylene (110 total orbitals, explained in the results Section 6.6.1).

Additionally, the Pauli recycling framework fits perfectly the regime in which, once the VQE is completed, the Pauli strings composing the Hamiltonian have already been evaluated on the final wavefunction, and therefore, the statistics on the average expectation values, shot based measurement, or noisy aware experimental data, can be used in place of the `cache_pauli`, and evaluate all the derivatives using the same error as on the Hamiltonian.

6.3.6 Workers definition for parallelization

The distributed evaluation framework employs two distinct worker types, `fermionic_worker` and `pauli_worker`, each responsible for computing first- and second-order derivatives of the Hamiltonian with respect to the generator matrices T within their respective algebraic representations.

Both worker routines share a unified computational architecture designed for modularity and controlled memory utilization during derivative evaluation. Each worker receives a tuple of generator descriptors and indices

$$(g_i, T_i, i, g_j, T_j, j),$$

which define the operator pair for which first-order and, if requested, second-order derivatives are computed. The execution begins with the construction of the first derivative Hamiltonian dH by calling the differential wrapper, as defined in the previous section, which applies the generator T_i to the reference Hamiltonian H. The resulting operator blocks are collected into a local dictionary `h_first`, ensuring structured access to the Hamiltonian partitions throughout the evaluation.

Once dH has been assembled, it is immediately contracted with either the cached Hamiltonian, with Pauli recycling, or the system wavefunction, depending on the computational backend, producing the first-order energy derivative. If the `second_order` flag is enabled, the same sequence is repeated using T_j as the secondary generator, acting this time on the intermediate dH to yield the second derivative Hamiltonian ddH. The evaluation of dH follows the same path, generating the mixed or diagonal second-order derivative contribution. Both derivative levels are returned as tuples of generator indices and scalar derivative values, formatted for collective reduction by the higher-level parallel routine.

```python
def fermionic_worker(args, psi: np.ndarray, num_qubits: int,
                     masks: dict,
                     core: list | np.ndarray,
                     virtual: list | np.ndarray,
                     h: dict[tuple[str, str], np.ndarray] | None = None,
                     second_order: bool | None = False) -> tuple[int, float, int| None, float | None]:

    jac_i, hess_ij = None, None
    orbs_g_i, t_i, i_idx, orbs_g_j, t_j, j_idx= args
    e = single_dh(T=t_i, g=orbs_g_i, h=h)
    h_first = {}
    for k, block in e:
        h_first[k] = block
    jac_i = fermionic_dh_evaluation(h_first, psi, masks, num_qubits, core, virtual)
    if second_order:
        e = single_dh(T=t_j, g=orbs_g_j, h=h_first)
        h_second = {}
        for k, block in e:
            h_second[k] = block
        hess_ij = fermionic_dh_evaluation(h_second, psi, masks, num_qubits, core, virtual)
        return (i_idx, jac_i, j_idx, hess_ij)
    else:
        return (i_idx, jac_i, None, None)
```

Snippet 6.9: Simplified `fermionic_worker` for evaluating the first- and second-order derivatives.

Both the implementation, `fermionic_worker` and `pauli_worker` are fully reported in Appendix 1.1 and Appendix 1.2, respectively.

A central aspect of both workers is the explicit management of temporary data and memory persistence. After each derivative stage, all intermediate Hamiltonian dictionaries (`h`, `h_first`, `h_second`) are cleared explicitly using element-wise deletion followed by total object deallocation through `del`. When large operators are stored as memory-mapped arrays, these mappings are closed and dereferenced immediately after use to prevent file locking and memory growth across processes. In addition, both workers invoke Python's garbage collector via `gc.collect()` at key synchronization points to enforce deterministic cleanup and avoid residual allocations when thousands of derivative tasks are processed in parallel. This explicit memory control, coupled with the use of temporary dictionaries and stateless evaluation calls, ensures reproducibility and scalability of the derivative computation even in high-dimensional operator spaces or under constrained shared-memory environments.

It is noticeable that, given the fact that computing the second-order derivative is not possible without keeping track of the first derivative Hamiltonian, `h_first`, then number of *wasted* single derivative evaluation are equal to the number of second derivative. Therefore, each second-order evaluation for generators T_i and T_j produces the value for the first-derivative w.r.t. T_i and the second-derivative, meaning that there is going to be an additional of each first-derivative for every second-order one. Given N_T the total number of generators, the second-derivative evaluations are $\frac{N_T \times (N_T+1)}{2}$, i.e. the upper-triangular Hessian and the diagonal entries, and the same number of first-derivative evaluation, due to the caveat explained above. So far, the total number of evaluation has been reduced by reducing the number of generators, N_T, exploiting iso-energetic orbital rotation removal or Point-Group symmetry, but additional screening might be performed with further research.

6.3.7 Intra-node and Multi-Node parallelization

Once all the previous optimization have been implemented, a proper parallelization scheme can be put in place. WAHTOR-SCF, as well as the previous version, is a procedure perfectly suitable to be fully and efficiently parallelized. Each derivative is an independent calculation performed on the original Hamiltonian, and once the expectation value is obtained, any intermediate object can be deleted. Therefore, there is not actual connection between the derivatives, if properly implemented, and each of them can be treated in parallel.

The evaluation of the Jacobian vector and Hessian matrix is distributed through a hybrid parallelization scheme combining MPI-based inter-process communication with intra-node parallel execution via `joblib`. This design enables scaling across both compute nodes and cores, optimizing memory locality and minimizing communication overhead during the evaluation of large numbers of Hamiltonian or derivative terms.

In terms of rank specific objects, the `rank==0` store all the information about:

- **System object**: original set of integrals[3],`starting_h`, the original set of molecular orbital coefficients,`initial_mo_coeff`, the updated integrals, `h`, the total unitary tranformation, `U_tot`, the set of active orbitals, `active_space`, the core orbitals, `core`, and the virtual orbital, `virtual`.
- **Circuit object**: the empty circuit, `ansatz`, associated optimal parameters, `params`, and the associated statevector, `psi`.

---

[3] For original we mean Hartree-Fock integrals

- **SCF info**: as well as all the information required to compute the derivatives like the mask associated to each group of orbitals, `mask_h`, the type of derivative evaluation method, `pauli_worker`[4] or `fermionic_worker`[5].

The other ranks, only when needed, receive from the broadcast the objects relative to the evaluation method, `eval`, the sub-set of derivative that are required to be computed, `batch`, and the statevector, `psi` or the cached Pauli evaluations, `pauli_cache`, on which the derivative are evaluated.

At initialization, the root process (`rank == 0`) constructs the full task list containing all single- and double-generator combinations. Each task encodes the pair of generator indices, their corresponding labels, and a set of orbital group identifiers. The complete list of tasks is then broadcast to all ranks using `comm.bcast`, ensuring consistent task distribution across processes. The total workload is partitioned into chunks of approximately equal size, computed as `len(tasks) // size`, with a small correction for remainder tasks to maintain load balance across ranks. Each MPI process thus receives a unique subset of tasks and proceeds independently on its assigned portion. A `comm.Barrier()` synchronization call ensures that all processes begin computation simultaneously.

Within each MPI rank, a second level of shared-memory parallelization is achieved via `joblib`, using a multiprocessing backend. Each rank executes a local batch of `n_jobs = wahtor_workers`, corresponding to the number of CPU cores allocated to that rank. The `Parallel` call wraps the `partial_eval` function—either `pauli_worker` or `fermionic_worker`—which handles the physical evaluation of the derivative contributions. The two-level parallel structure (MPI for distributed nodes and `joblib` for local cores) enables efficient scaling up to large cluster environments, while still maintaining efficient CPU utilization on shared-memory nodes. In Figure 6.7, the double parallelized structure is shown. The `partial_eval` function is built dynamically depending on the simulation mode. When Pauli recycling is enabled, the evaluator uses precomputed cached expectation values (`cached_h`) constructed earlier via the Pauli cache. In this case, only the Pauli coefficients and masks (`core`, `virtual`, and `h_mask`) are broadcast to the workers, avoiding redundant transfer of large Hamiltonian matrices. When Pauli recycling is disabled, a standard fermionic evaluation

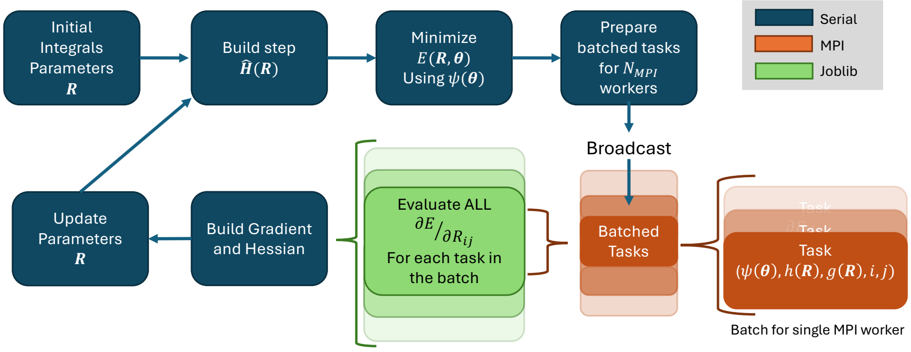

Figure 6.7: Resuming scheme of the double parallelization structure for WAHTOR-SCF. All the control tasks and updating checking are performed sequentially, while the construction and evaluation procedure of the gradient and Hessian entries are performed in the double parallelized framework. First all the tasks are generated, then batched for each MPI rank, and internally, the intra-node (or worker specific parallelization) is performed to obtain the evaluations. After that, all the results are collected by the master, or `rank==0`.

---

[4] Using Pauli recycling procedure

[5] Using standard construction of fermionic derivative operator, mapping into qubit space, and evaluation as a whole matrix.

path is used instead. Each evaluator is broadcast to all MPI ranks via `comm.bcast`, ensuring every process holds an identical function definition and context.

For systems with large numbers of fermionic modes (`fermionic_modes > MAX_FERM_MODES`), the Hamiltonian data are handled via NumPy memory-mapped files to prevent exceeding the available system memory and avoid the 2Gb MPI broadcast problem[6]. In this case, the root process creates temporary on-disk storage for each Hamiltonian block ($h_1^{\alpha}, h_1^{\beta}, h_2^{\alpha\alpha}, h_2^{\beta\alpha}, h_2^{\beta\beta}$) using `np.memmap`, and the file metadata (filenames, shapes, and dtypes) are stored in a distributed dictionary object (`self.info`). The metadata are then broadcasted to all ranks, and the `MemmapManager` utility ensures that each process accesses the same data files locally in read-only mode. This strategy prevents memory duplication across MPI ranks while maintaining concurrent read access to the same dataset, enabling very large Hamiltonian tensors to be processed even under tight memory constraints, and also reduce the memory footprint of each task.

Each rank independently launches its `joblib` workers through the context manager, ensuring isolation between processes and threads. Within each rank, all local results are aggregated into a list of partial Jacobian and Hessian contributions. Once computation completes, the rank-local results are gathered on the root process using `comm.gather`. The root process then reconstructs the global Jacobian and Hessian arrays, filling the symmetric entries where necessary when second-order derivatives are computed. Finally, the resulting arrays are broadcast back to all ranks to synchronize state and ensure consistent access to the final data structures across all processes.

This hybrid MPI–`joblib` approach achieves near-linear scaling across both nodes and cores by distributing task generation and derivative evaluation while keeping memory usage and inter-process communication minimal. MPI handles the coarse-grained distribution of independent Hamiltonian derivatives, while `joblib` manages fine-grained local parallelism for each batch of tasks. The use of memory mapping and broadcasted evaluator functions minimizes data replication, achieving efficient parallel resource utilization even for systems with extremely large orbital spaces.

Performance Scaling and Computational Efficiency

The hybrid parallelization scheme exhibits near-linear scaling for moderately large systems, where the total number of generator pairs and Hamiltonian blocks substantially exceeds the number of available computational processes. In the strong scaling regime, where the problem size is fixed and the number of MPI ranks increases, performance is primarily limited by inter-process communication and I/O latency arising from memory-mapped Hamiltonian blocks. Since each MPI rank operates on an independent subset of generator combinations, communication overhead remains minimal during computation, with synchronization required only at task broadcasting, memory map initialization, and result gathering phases. Consequently, the computational throughput is largely determined by the local evaluation efficiency of each rank and the intra-node load balancing handled by `joblib`.

In the weak scaling limit, where the problem size grows proportionally with the number of ranks, the approach maintains a high parallel efficiency provided that the ratio between the number of generators and ranks remains sufficiently large to saturate all available cores. The `joblib` multiprocessing backend effectively overlaps CPU-bound operations with minimal scheduling overhead, as each derivative evaluation is an embarrassingly parallel task with negligible interdependence. This

---

[6] For MPI version < 4, the maximum allowed broadcasting object size is 2Gb. This can be solved either wrapping the `MPI.comm` with `mpi4py.util.pkl5`

ensures that, within each rank, all cores are fully utilized throughout the evaluation of local task chunks. The hybrid model thus achieves a good compromise between distributed memory scalability (through MPI) and shared memory concurrency (through `joblib`), while avoiding excessive inter-rank communication or data replication.

6.3.8 Adaptive and Self–Consistent WAHTOR (ADAPT–WAHTOR–SCF)

Finally, after the assessment of the efficiency of the whole infrastructure, integrating an adaptive CI solver inside the WAHTOR–SCF routine is straight-forward. We define the *Adaptive Hardware–Efficient Ansatz Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (ADAPT–WAHTOR–SCF)* algorithm in the same way as ADAPT-VQE–SCF, with the only difference that, once again, the derivative operators are measured at each step as fermionic operators. The Adaptive Hardware–Efficient Ansatz Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (ADAPT–WAHTOR–SCF) algorithm combines an adaptive ansatz construction with Hamiltonian shaping through orbital rotations. At each macro–iteration, an ADAPT criterion (e.g gradient norm or energy drop) is evaluated on the operator pool; if the threshold is satisfied, a single generator is selected and appended to the circuit. The VQE parameters are then optimized to convergence, yielding a fixed wavefunction $|\Psi(\boldsymbol{\theta})\rangle$. With the circuit frozen, the orbital–rotation derivatives operators are evaluated and used to update the Hamiltonian via the WAHTOR–SCF transformation, After the Hamiltonian update, the adaptive circuit is relaxed again so as to best represent the new Hamiltonian. In ADAPT–WAHTOR–SCF, we have therefore two pools, the operator pool $\{\hat{\mathcal{P}}_i\}$, which contains all the excitation-gate eligible to be added to the circuit, and the orbital–rotation generators pool, $\{\hat{T}_{pq}\}$. The two pools are independent but both evaluated at fixed reference wavefunction $|\Psi(\boldsymbol{\theta})\rangle$. The ADAPT–WAHTOR–SCF can be therefore summarized as:

1. Start from the reference wavefunction $|\Psi(\boldsymbol{\theta})\rangle$

2. Select the operator based on the desired metric. For example, selecting the operator $\hat{\mathcal{P}}_i$ such that it maximizes the gradient w.r.t. the energy as

$$\left.\frac{\partial E(\boldsymbol{R})}{\partial \theta_i}\right|_{\theta_i=0} = \langle\Psi(\boldsymbol{\theta})|\,[\hat{\mathcal{P}}_i, \hat{H}(\boldsymbol{R})]\,|\Psi(\boldsymbol{\theta})\rangle\,, \tag{6.42}$$

   where $\hat{\mathcal{P}_i}^\dagger = -\hat{\mathcal{P}}_i$. If all the gates are above the metric threshold, γ_{exp}, the circuit is not expanded and it keeps the previous wavefunction.

3. Optimize the circuit in order to minimize $E(\boldsymbol{R})$ and obtain Minimize the energy using the new circuit $|\Psi(\boldsymbol{\theta})\rangle \to |\Psi(\boldsymbol{\theta'})\rangle$.

4. Use the newly obtained wavefunction $|\Psi(\boldsymbol{\theta'})\rangle$ to evaluate the effective energy derivative w.r.t. to orbital–rotation parameters $\boldsymbol{R}$, defined as fermionic operators, and obtain the new parameters $\boldsymbol{R'}$.

5. Define the unitary $U(\boldsymbol{R'})$ and update the Hamiltonian $\hat{H}(\boldsymbol{R}) \to \hat{H}(\boldsymbol{R'})$.

6. Relax the wavefunction for the new Hamiltonian $\hat{H}(\boldsymbol{R'})$ before restarting again.

This procedure is repeated until both the adaptive metric and the SCF orbital conditions are simultaneously converged. In this way, the ansatz grows only when necessary while the Hamiltonian is continuously rotated toward a representation that is maximally compatible with the evolving ADAPT wavefunction.

6.4 WAHTOR-SCF applied to general molecules

In this section, we collect the simulation results obtained by HEA–WAHTOR–SCF and ADAPT-WATHOR-SCF agains model molecular systems, for different setups. In particular, we have targeted a subset of small molecules, LiH, H_2O, N_2, and C_2H_4, along their potential energy curve under geometrical tranformation, i.e. bond dissociation / atomization. On top of this, we also evaluated the performance of WAHTOR-SCF with different circuit configurations, as well as different optimizer results, and different orbital–rotation generators pool.

6.4.1 Molecular system treated

Here, we collect all the molecular system treated. Given the fact that all the molecules have been analized along the PEC of single–bond, double–bond, or symmetric–bond elongation, the specific geometry will be reported during the discussion of the results. For all the system, the unit used are Å, except for H_2O, for which *bohr*s are used.

LiH: Li–H Dissociation

The dissociation of LiH serves as an ideal initial test case for a multi–reference calculation due to the small active space required and the suitability for comparison with Full Configuration Interaction (FCI) using reasonably large basis sets. To accurately describe its dissociation into the ^{2}Li and ^{2}H fragments, a minimal CAS(2,2) active–space (i.e. 4 qubits) is sufficient. The active space consists of three determinants among which the covalent one, where both active orbitals are singly occupied, increasingly dominates the Multi–Configurational Self–Consistent Field (MCSCF) wavefunction as one moves beyond equilibrium geometry toward the dissociation limit. We employ Pople's 6-31G [226] basis for both atoms and we use CISD natural orbitals as a starting point. CISD natural orbital have been obtained starting from PySCF RHF orbitals. The naturalization is necessary for the VQE calculation in order to get the desired active orbitals around Fermi level.

H_2O: Symmetric O–H bond elongation

The symmetric elongation of O–H bonds in H_2O is one of the most frequently used test cases for investigating the interplay of different levels of electron correlation effects [227–233]. The minimal active space required to properly describe the full PEC is a CAS(4,4) (i.e. 8 qubits), which includes two pairs of bonding and antibonding orbitals belonging to the following IRREPs of the C_{2v} point group: $3A_1$, $4A_1$, $1B_2$, and $2B_2$ [229]. We use the 6–31G basis for both atoms and the O-H bonds are stretched simultaneously keeping the H–O–H angle fixed at the equilibrium value of 110.6° [229]. When using PySCF, the C_2 subgroup of C_{2v} was used to define the initial RHF orbitals so as to maintain the active orbitals at the Fermi level for the subsequent VQE computation. We identify five O–H bond distances spread across the CASSCF PEC and compared the result of our WAHTOR-SCF apporach compared to the classical benchmark.

C_2H_4: C=C bond dissociation

The double bond dissociation in C_2H_4 requires a minimal active space of CAS(4,4) (8 qubits) to properly describe the full PEC. The active space includes two pairs of bonding and antibonding orbitals around the Fermi level, one of them being the C=C π-bond and the other is the C=C σ-bond. We fix the H–C–C bond angles at 120°, C–H bond at 1.089Å and the H–C–C–H dihedral angles at 0° or 180° depending on the cis or trans position of hydrogens respectively [234]. The cc–pVDZ basis was employed for our computations. D_{2h} point group symmetry was used to generate the initial RHF orbitals and we identify five C=C bond distances distributed across the CASSCF PEC to compute our quantum CASSCF results. The point group symmetry was also used to select the proper unitary generators which rotate the orbitals, using **PG-RED**.

N_2: Triple N=N bond elongation

The triple bond dissociation in N_2 is one of the more challenging applications for any multireference theory [235]. The minimal active space required to properly describe the full PEC is a CAS(6,6) (12 qubits), which includes three pairs of bonding and antibonding orbitals around the Fermi level, two of them being π bonds and the other is a σ bond. The cc–pVDZ basis was employed for our computations. As before, we employed the D_2H point group symmetry to generate the initial RHF orbitals and identify five N≡N bond distances spread across the CASSCF PEC to compute our current WAHTOR–SCF results. The point group symmetry was also used to select the proper unitary generators which rotate the orbitals, using **PG-RED**.

6.4.2 General simulation setup

All the molecules tested have been solved, from the classical point of view, using the quantum chemisty python package PySCF[94–96]. The reference CASSCF calculations have been obtained using `mcscf.CASSCF` solver from PySCF. For LiH, the CISD natural orbitals have been obtained from a PySCF CISD calculation (`ci.CISD`), diagonalizing the 1-RDM obtained from the correlated wavefunction, and applying the resulting unitary tranformation to the RHF initial orbitals. For H_2O, the active-space, as explained in the previous section, has been defined according to the core and active orbital irridicible representation. In particular, the selection of the active space is done using the method `mcscf.sort_by_irreps`. All the integrals for generating the fermionic operators are extracted from PySCF `scf.RHF` mean-field calculation, except LiH, which as already been commented. The associated fermionic operators are defined in a QuAQ custom format, while the corresponding qubit operator, for compatibility with the quantum circuits (`qiskit.QuantumCircuit`), are defined as `qiskit.SparsePauliOp`. For the quantum simulation, the *QuAQ* package, built as a Qiskit [236] wrapper, is used to design the quantum circuits, both for hardware efficient ansatz (HEA) and ADAPT–VQE. Both the VQE and ADAPT–VQE routines, as well as the (ADAPT–)WAHTOR-SCF implementation are again implemented in *QuAQ* package. For the VQE and ADAPT-VQE routine, the optimizer chosen is L–BFGS–B with a gradient tolerance set to 1e-6 and it has been carried out using `Scipy.minimize` method and its optimizer implementation.

For the PEC analysis, shown in Section 6.5, all the single-points have been obtained using Trust-Region as optimizer for the WAHTOR–SCF orbital optimization update step. For the generators pool chosen, if available, the **PG–red**[7] has been always preferred, and applied on top of a **CAS-like** pool. The chemical accuracy value in the results is chosen to be 2mHa.

For the generators pool analysis, in which we compared the outcomes obtained from HEA-WAHTOR-SCF under different level of reduced orbital–rotations pool, we focused on the equilibrium geometry of H_2O in 6-31G basis. On overall, a reference number of generators used for each of the system is reported in Table 6.4. In particular, for all the system the combination used is **CAS-LIKE** with **PG-RED** applied on top.

System	$\mathbf{N_{act}}$ **(qubits)**	$\mathbf{N_{occ}}$	$\mathbf{N_{virt}}$	$\mathbf{N_T^{ALL}}$	$\mathbf{N_T^{CAS-LIKE}}$	$\mathbf{N_T^{RED}}$
LiH	2 (4)	1	8	55	27	19
H_2O	4 (8)	3	6	78	60	42
C_2H_4	4 (8)	6	38	1128	410	182
N_2	6 (12)	4	18	378	219	147

Table 6.4: Number of orbitals included in each set of active (**act**), occupied (**occ**), and virtual (**virt**). The number of generators $\mathbf{N_T^x}$ according to the method $\mathbf{x} \in \{$**ALL**, **CAS-LIKE**, **RED**$\}$.

6.4.3 HEA–WAHTOR–SCF setup

For the HEA-WAHTOR-SCF simulations we use a hardware efficient ansatze with two kinds of entangling gates, CNOTs and *SO(4)* gates (which representation can be found in Section 5.2.3). Given the idea behind WAHTOR and WAHTOR–SCF to adapt the Hamiltonian to the wavefunction, the simulations have been performed using the simplest form of HEA, i.e. using ladder-fashion ansatz. As already explained multiple time through the whole thesis, the ladder ansatz is composed by two blocks: the *entangling layer* in a top-down fashion, starting from the first alpha orbital, crossing all the alpha orbitals, all the beta orbitals, and connecting the last beta orbital[8]. The entanglement block is than interleaved with a layer of *single–qubit rotation* gates, spanning all the qubits. The ladder circuit is then repeated d times to get the final VQE ansatz, where d depends on the active space required to study the potential energy curve (PEC) of a particular molecular system. In terms of correlators/entangling gate, as already mentioned, we used both CNOTs and *SO4* gate, for the rotation gates instead, we employed single–qubits R_y gates. All the circuits are initialized to the Hartree-Fock determinant using X gates in correspondence with the occupied orbitals.

In terms of circuit cost, each *SO4* gates cost 6 parameters and 2 CNOTs. A depth d CNOTs ladder circuit will have a cost of $N_q*(d+1)$ parameters and $d*(N_q-1)$ CNOTs, while a *SO4* ladder circuit $6*(N_q-1)*d+N_q*(d+1)$ parameters and $2*d*(N_q-1)$ CNOTs (double the amount in a CNOT-based ladder). A recap of the circuital resources used for each system and configuration is shown in Tables 6.5 and 6.6.

In terms of depth, we imposed that the depth of the circuit d, between CNOTs–based and *SO(4)*–based, is the same in terms of number of CNOTs, therefore, the former have double the number of layers compare to the latter.

---

[7] Point–Group reduciton

[8] Ladder-fashion circuit representation can be found in Section 3.1

System	**qubits**	**Layers** (d)	**Number CNOTs**	**Number Params**
LiH	4	4	12	20
H_2O	8	8	56	72
C_2H_4	8	8	56	72
N_2	12	12	132	156

Table 6.5: CNOTs-based circuit specifications.

System	**qubits**	**Layers** (d)	**Number CNOTs**	**Number Params**
LiH	4	2	12	48
H_2O	8	4	56	208
C_2H_4	8	4	56	208
N_2	12	6	132	480

Table 6.6: *SO(4)*-based circuit specifications.

Given the heuristic nature of the circuit, a single run is not meaningful, each point of the PEC has been solved by 50 indipendent VQE runs. Each VQE has been started with a random initialized circuit which parameters are picked in the range $(-\pi, \pi)$. This initialization procedure is also shared with the *SO4*-based ladder circuit.

6.4.4 ADAPT–WAHTOR–SCF setup

For the ADAPT–WAHTOR–SCF simulations, Qubit–ADAPT [72] version of the ADAPT–VQE algorithm have been used as ADAPT–CI solver. The Qubit–ADAPT pool, $\mathcal{P} = \{\hat{P}_i\}$, can be defined starting from all the allowed fermionic excitations

$$\begin{aligned} \hat{A}_{ij} &= \frac{1}{2}(X_iY_i - Y_iX_j)\prod_{p=i+1}^{j} Z_p \\ \hat{A}_{ijkl} &= \frac{1}{8}(X_iY_jX_kX_l + Y_iX_jX_kX_l \\ &+ Y_iY_jY_kX_l + Y_iY_jX_kY_l - X_iX_jX_kY_l \\ &- Y_iX_jY_kY_l - X_iY_jY_kY_l)\prod_{p=i+1}^{j} Z_p \prod_{p=k+1}^{l} Z_p \end{aligned} , \tag{6.43}$$

assuming $i < j < k < l$ spin–orbital (or qubit) indices. After creating the fermionic pool, the qubit–pool is obtained by splitting as independent operators all the Pauli strings composing each excitation operator $\hat{A}$. Therefore, the qubit–ADAPT pool will be composed by

$$X_iY_j \quad X_iX_jX_kY_l \quad Y_iY_jY_kX_l \tag{6.44}$$

including all the possible indices combinations. The metric according which the excitations contained in $\mathcal{P}$ are picked is the energy gradient w.r.t. the excitation parameter, i.e. the gate with highest $\frac{\partial E}{\partial \theta_i}|_{\theta_i=0}$ in absolute value. Both double and single excitations are inserted in the pool, and the *expansion threshold*, γ_{exp}, to decide whether to perform an expansion of the circuit or using the current one, is set to 1e-5. If a gate is added, the previous parameters are kept unchanged and the new parameter is set to 0. If no gate is added, the wavefunction is kept unchanged and only refined on the rotated Hamiltonian $\hat{H}(\boldsymbol{R'})$. The whole ADAPT–VQE and ADAPT–WAHTOR–SCF is `joblib` parallelized (for the moment only intra-node parallelization), in particular, the same number of workers, i.e. joblib `n_jobs`, is used for both the WAHTOR–SCF orbital–rotation gradients evaluation and ADAPT–VQE excitation gradients. Given the small sizes of the systems and active spaces, the number of jobs used 12.

6.5 WATHOR–SCF results on general molecules

In order to validate both our approached, HEA–WAHTOR–SCF and ADAPT–WAHTOR–SCF, in terms of performance, efficiency, applicability, and reliability for NISQ computers by taking into account both the circuit complexity, in terms of parameters and number of gates, and the accuracy compared to the classical counterpart, i.e. CASSCF. Once again, in order to ease the analysis and feasibility of the method, we consider exclusively statevector noiseless simulations, without including any shot-based statistic or hardware noise models. In the next sections, we report in order the resutls of HEA–WAHTOR–SCF, the adaptive version, ADAPT–WAHTOR–SCF, and the specific analysis results on different generators pool and optimizers.

6.5.1 Results for HEA–WAHTOR–SCF

In this section, the performance of the WAHTOR-SCF algorithm is analysed on four molecular systems of increasing electronic complexity: LiH, H_2O, C_2H_4 and N_2. For each case, the potential energy surface (PES) obtained with two different entangling circuits — a CNOT-based ladder ansatz and an *SO(4)*-based ladder ansatz — is compared with the CASSCF reference. A chemical-accuracy threshold of 2 mHa is adopted throughout. Distances are reported in Å, except for LiH where the original profile is plotted in Bohr. In addition to the energy profiles, a statistical analysis based on 50 independent VQE optimisations is used to quantify the stability of each ansatz in terms of correlation recovery, circuit expressibility, and standard deviation with respect to initial conditions. In the sections, we will report an analysis of the results obtained, which a selection of figure. For the complete results, please check the Appendix C which contains all the information about all the simulations, and all the plots.

LiH

The LiH dissociation, described in a CAS(2,2) active space, represents the simplest benchmark in this study and is dominated by a single-reference character. Both the CNOT and *SO(4)* ansätze yield PES curves that lie essentially on top of the CASSCF reference, remaining well within the 2 mHa window over the entire stretching coordinate (see Figure 6.8 and for both in Figure C.1). As reported in Table C.1, both circuits recover nearly 100% of the correlation energy, although the CNOT ansatz displays extremely small standard deviations (on the order of 10^{-7}), while *SO(4)* exhibits slightly larger fluctuations. This indicates that, in a weakly correlated regime, the additional expressibility of *SO(4)* provides no real benefit over a shallower CNOT structure. The LiH test therefore establishes a baseline: when correlation remains essentially dynamic and the electronic manifold is compact, both ansätze are more than adequate, and the more compact CNOT circuit can even be regarded as the more stable option due to its reduced parameter space.

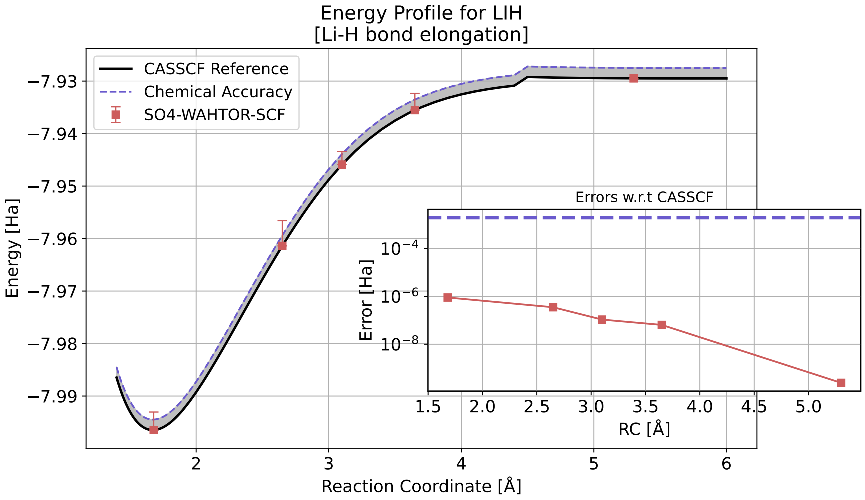

Figure 6.8: Comparison between our *SO(4)*–WAHTOR–SCF approach and CASSCF for the LiH Li-H bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the *SO(4)*–WAHTOR–SCF solution with associate error bars, from the 50 VQEs. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our *SO(4)*–WAHTOR–SCF approach.

H_2O

The symmetric O–H bond elongation in water introduces a moderate multireference character, providing a clearer separation between the capabilities of the two ansätze. As shown in Figures 6.9 and both in Figure C.2, both circuits match the CASSCF surface at equilibrium, but differences emerge as the bond is stretched. The CNOT curve begins to deviate from the reference, exceeding the chemical-accuracy threshold in the dissociation region, whereas the *SO(4)* ansatz maintains a smoother profile with reduced oscillations and a significantly smaller error. Statistical results from Table C.2 confirm this behaviour: while the CNOT ladder remains stable, it consistently recovers a smaller percentage of the correlation compared to *SO(4)*, which reaches values above 99.99% with moderate variance. This system highlights a first key observation: once static correlation becomes relevant, the limited expressibility of CNOT is no longer sufficient to guarantee chemical accuracy across the entire PES, whereas the *SO(4)* ladder provides a more robust representation of the correlated wavefunction.

C_2H_4

The C=C stretching in ethylene is characterised by π-electron delocalisation rather than strong bond-breaking correlation. In this scenario, both ansätze are capable of reproducing the CASSCF reference along the full coordinate, as illustrated in Figure 6.10 and both in Figure C.3. However, the *SO(4)* curve exhibits a slightly smoother deviation profile and a smaller maximum error with respect to the reference, remaining comfortably within chemical accuracy. The data in Table C.3 show that both circuits achieve high correlation-recovery percentages, although *SO(4)* retains a small advantage while maintaining reasonable statistical stability. This suggests that, in systems

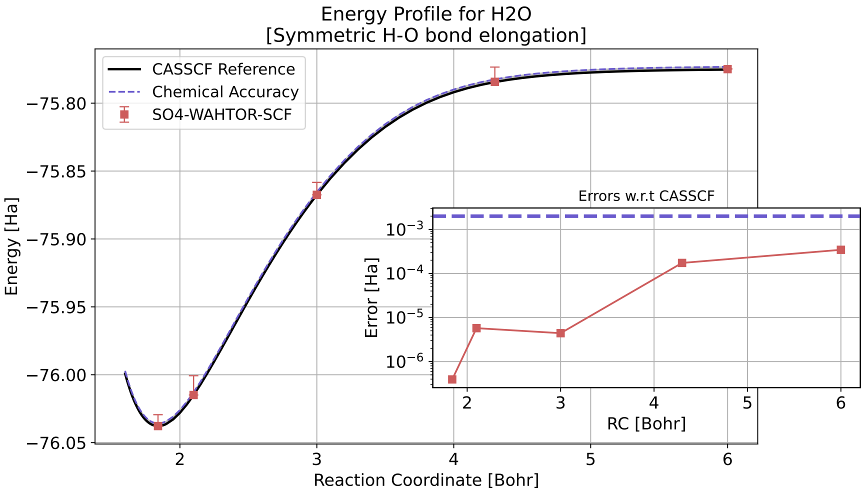

Figure 6.9: Comparison between our *SO(4)*–WAHTOR–SCF approach and CASSCF for the H_2O Symmetric H–O bond elongation.The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the *SO(4)*–WAHTOR–SCF solution with associate error bars, from the 50 VQEs. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our *SO(4)*–WAHTOR–SCF approach.

where electron delocalisation dominates, both ladder structures remain effective, but the larger two-qubit algebra of *SO(4)* continues to offer a more flexible optimisation landscape and a more uniform error distribution.

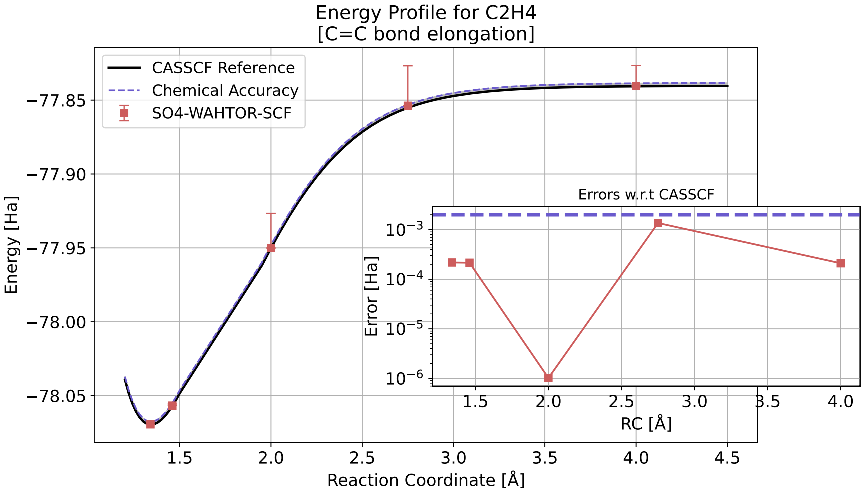

Figure 6.10: Comparison between our *SO(4)*–WAHTOR–SCF approach and CASSCF for the C_2H_4 C=C bond elongation.The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the *SO(4)*–WAHTOR–SCF solution with associate error bars, from the 50 VQEs. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our *SO(4)*–WAHTOR–SCF approach.

N_2

The dissociation of molecular nitrogen represents the most demanding and challenging benchmark due to its well-known strong multireference character. The CNOT ansatz fails to maintain chemical accuracy away from equilibrium and exhibits significant deviations from the CASSCF curve, as seen in 6.11 and both in Figure C.4. In contrast, the *SO(4)* ansatz is able to preserve the qualitative and quantitative features of the dissociation curve, significantly reducing the error throughout the stretching region. The statistical results in Table C.4 reveal large standard deviations and insufficient correlation recovery for the CNOT circuit, whereas *SO(4)* attains up to 96% of the reference correlation with far more controlled fluctuations. This case provides the clearest evidence that shallow CNOT-based parametrization are inadequate in the presence of strong static correlation, while a richer entangling structure is essential to avoid variational collapse and maintain a physically meaningful PES.

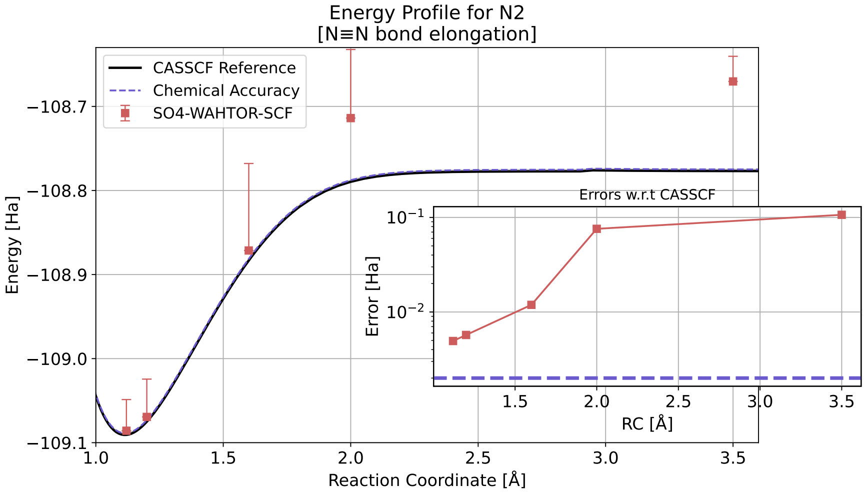

Figure 6.11: Comparison between our *SO(4)*–WAHTOR–SCF approach and CASSCF for the N_2 triple N≡N bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the *SO(4)*–WAHTOR–SCF solution with associate error bars, from the 50 VQEs. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our *SO(4)*–WAHTOR–SCF approach.

Once again, we report all the results of the VQE for each molecule in Appendix C, specifically, Tables C.1 - C.4.

6.5.2 Results for ADAPT-WAHTOR-SCF

In this section, we analyse the performance of the ADAPT-WAHTOR-SCF algorithm on the same molecular set considered in the previous chapter: LiH, H_2O, C_2H_4 and N_2. For each system, the potential energy surface obtained with the adaptive ansatz is compared to the CASSCF reference and to the heuristic ladder circuits previously employed. A chemical-accuracy threshold of 2 mHa is retained throughout. Distances are reported in Å, except for LiH where the original profile is plotted in Bohr. Notably, only the region prior to dissociation is reported for some molecules, as the ADAPT-WAHTOR-SCF procedure fails to converge in the asymptotic limit.

LiH

The LiH profile (Figure 6.12) demonstrates the ideal scenario for the adaptive strategy. ADAPT-WAHTOR-SCF converges rapidly to the CASSCF reference and reconstructs the potential energy curve with negligible error, remaining well within chemical accuracy across the sampled coordinate.

Owing to the predominantly single-reference nature of LiH, only a few operators are selected, resulting in a compact ansatz. This confirms the conceptual strength of the method: when correlation is weak and chemically local, the adaptive mechanism can efficiently identify the minimal operator set required to represent the state.

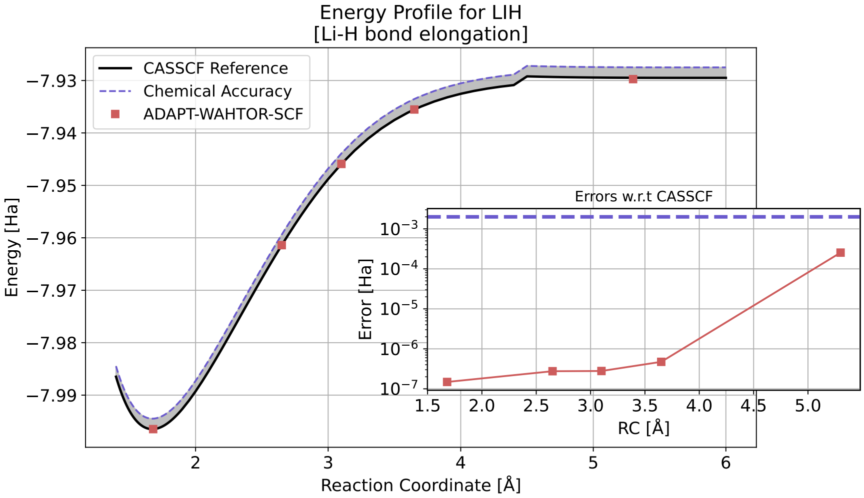

Figure 6.12: Comparison between our ADAPT–WAHTOR–SCF approach and CASSCF for the LiH Li-H bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the ADAPT–WAHTOR–SCF solution. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our ADAPT–WAHTOR–SCF approach.

H_2O

For the symmetric O–H bond elongation (Figure 6.13), ADAPT-WAHTOR-SCF again delivers accurate, smooth energy curves up to the onset of static correlation. Compared to heuristic circuits, the adaptive ansatz achieves similar accuracy with a significantly reduced number of parameters. This indicates that, in moderately correlated regimes, the adaptive selection is advantageous. In fact, it avoids unnecessary rotations, shortens the variational search, and preserves a more favourable optimization landscape. Up to this region, the ansatz growth remains controlled, and convergence is stable.

C_2H_4

A qualitatively similar behaviour is observed for the C=C stretching in ethylene (Figure 6.14). ADAPT-WAHTOR-SCF successfully recovers the underlying CASSCF profile in the non-dissociative region and maintains chemical accuracy. As in H_2O, delocalized correlation is well captured without excessive circuit growth. This confirms that the adaptive scheme remains efficient as long as the electronic structure is not strongly multi–reference and the orbital optimization steps remain well-conditioned.

N_2

While ADAPT-WAHTOR-SCF converges close to equilibrium, the algorithm fails to reach convergence in the dissociative region. This result contrasts with the corresponding single-point ADAPT-

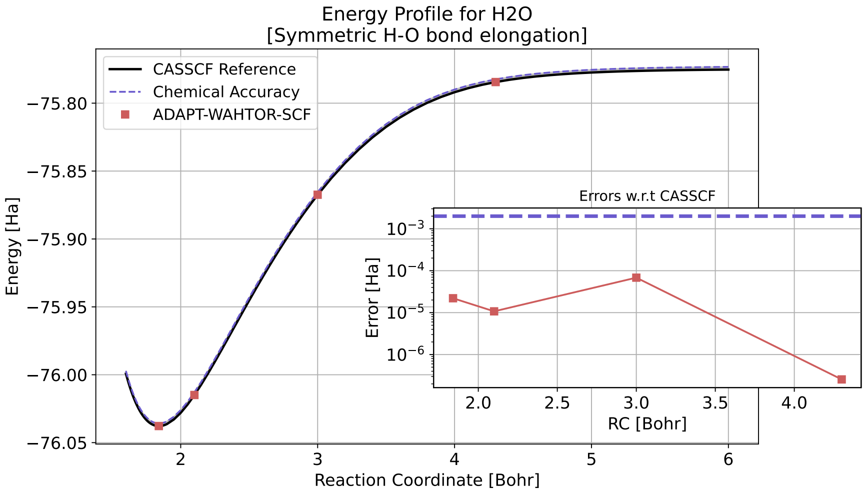

Figure 6.13: Comparison between our ADAPT–WAHTOR–SCF approach and CASSCF for the H_2O symmetric H–O bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the ADAPT–WAHTOR–SCF solution. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our ADAPT–WAHTOR–SCF approach.

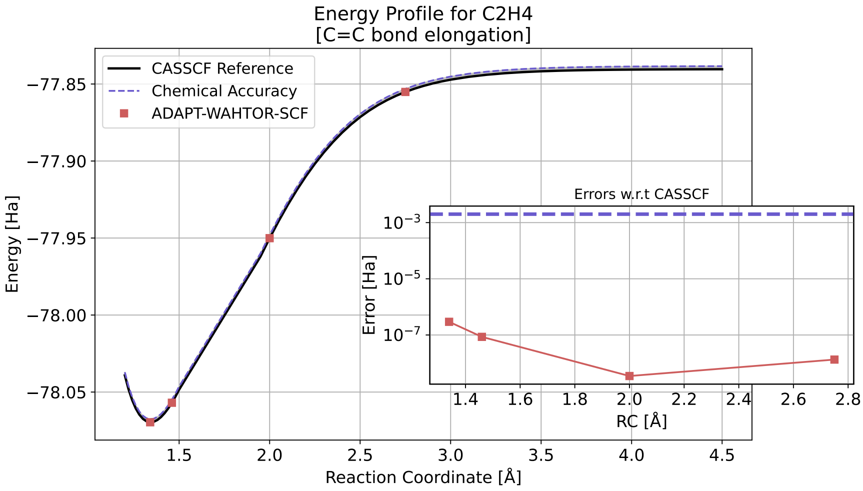

Figure 6.14: Comparison between our ADAPT–WAHTOR–SCF approach and CASSCF for the C_2H_4 C=C bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the ADAPT–WAHTOR–SCF solution. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our ADAPT–WAHTOR–SCF approach.

VQE on a CASCI reference, which is in fact able to converge even at large internuclear distances. Therefore, the failure is not intrinsic to the adaptive ansatz itself, but rather originates from the combined instability of orbital optimization and ansatz growth in strongly multi–reference regimes. In this region, the SCF equations become ill-conditioned, the orbital response amplifies small variational fluctuations, and the optimization landscape deteriorates, preventing convergence of the full self-consistent cycle.

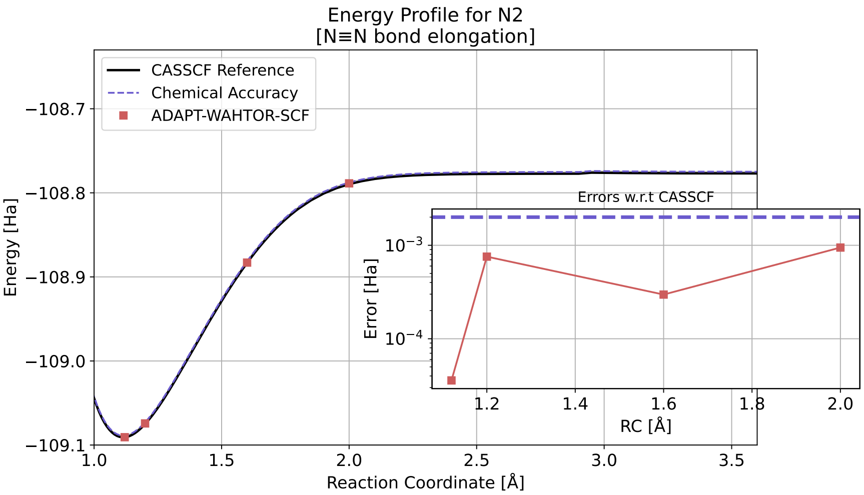

Figure 6.15: Comparison between our ADAPT–WAHTOR–SCF approach and CASSCF for the N_2 triple N≡N bond elongation. The black solid line corresponds to the CASSCF reference energy, the chemical accuracy (2mHa) is denoted with a dashed line, as well as a shaded area between the two values. The squares correspond to the ADAPT–WAHTOR–SCF solution. In the smaller panel instead, we report only the error w.r.t. to CASSCF of our ADAPT–WAHTOR–SCF approach.

6.5.3 Comparison between heuristic and ADAPT–WAHTOR–SCF

Table 6.7 summarizes the number of CNOTs and total parameters required to achieve chemical accuracy for representative geometries. ADAPT produces dramatically more compact ansätze than the heuristic circuits, particularly when compared to the SO(4) ladder. For LiH, the adaptive method requires only 3 parameters, whereas the heuristic circuits employ between 20 and 48. For H_2O and C_2H_4, the adaptive ansatz remains significantly smaller, confirming the efficiency of the operator-selection mechanism. In N_2, ADAPT still produces a more compact representation than SO(4), and it converges, whereas the heuristic ansatz does not. This comparison highlights a central conclusion: adaptive ansẗze excel in efficiency and compactness, but when coupled with orbital optimization, their robustness breaks down in the dissociation limit, where static correlation becomes dominant. In the Appendix C, the exact resource counts, in terms of CNOTs and parameters is collected in an unique Table C.5.

		CNOT Ladder		**SO4 Ladder**		**qubit-ADAPT**	
Molecule	**Geometry**	**CNOTS**	**Params**	**CNOTS**	**Params**	**CNOTs**	**Params**
LiH	3.10 Å	12	20	12	48	10	3
H_2O	3.00 a.u.	56	72	56	208	96	16
C_2H_4	2.00 Å	56	72	56	208	108	18
N_2	1.20 Å	132	156	132	480	258	43

Table 6.7: Comparison between heuristic ladder and ADAPT–VQE ansatzes for the molecular systems that we have studied. The number of decomposed CNOTs in each circuit along with total number of parameters needed to reach chemical accuracy is reported. For each molecule, we choose a representative geometry where all ansatzes could reach chemical accuracy (except the heuristic ansatz in N_2). The trend is similar for rest of the geometries.

Across all four systems, a clear trend emerges. For weakly correlated molecules such as LiH, both ansätze perform equally well, with the CNOT circuit even providing superior numerical stability due to its smaller parameter space. As correlation increases, however, the gap widens in favour of *SO(4)*: in H_2O and C_2H_4 the improvement is moderate but systematic, whereas in N_2 the difference becomes critical, with *SO(4)* being the only viable option for maintaining chemical accuracy along the dissociation coordinate. These results illustrate a fundamental trade-off: CNOT offers robustness and compactness, but limited expressibility; *SO(4)* offers superior accuracy and transferability at the cost of a larger parameter manifold. This trade-off provides a natural motivation for the next step, namely the development of an adaptive ansatz construction. ADAPT-WAHTOR-SCF outperforms heuristic variational circuits in terms of compactness and near-equilibrium accuracy, but remains less robust in the presence of strong multi-reference character. The dissociation failure in N_2 suggests that, although adaptive ansätze are fundamentally powerful, their combination with orbital optimization requires further stabilization strategies. This motivates the developments discussed in the next chapter, where we explore adaptive approaches capable of retaining compactness while restoring robustness in strongly correlated regimes.

6.5.4 Different optimizer results and generators pool

In order to assess the efficiency and the improvement of the method proposed in this work, and a completely variational approach. For the former results, two second-order optimization strategies are used, Newton–Rapshon and Trust-Region, while for the latter, the optimal orbital rotation parameters, $\mathbf{R}'$, are obtained variationally using BFGS. The comparison has been carried out on the second system studied, the H_2O with symmetric dissociation of O-H with $\angle \mathrm{HOH} = 110.6°$ using 6-31G basis at equilibrium (1.84 Bohrs). We compare each of the optimizer listed above **BFGS**, Newton-Raphson (**NR**), and Trust-Region (**TR**), with all the three possible generators pool **ALL**, **CAS-LIKE**, and **RED**. Point–group symmetry screening, **PG-RED** was used for each pool of generators. In table 6.9, the average number of iterations (**iter**) is shown, averaged w.r.t the number of chemically accurate runs (**G**) among the 50 executed. For the parameterized wavefunction, the same settings of the previous simulation (Section 6.4.2) has been used, in particular selecting a *SO(4)* ladder ansatz with depth $d = 4$.

	BFGS		**Newton-Raphson**		**Trust Region**	
Gen type	**G**	**iter**	**G**	**iter**	**G**	**iter**
ALL	17	25.35	0	0	25	37.6
CAS	38	13.71	5	35.2	37	31.22
RED	17	20.59	3	24	22	32.64

Table 6.8: Comparison between different optimizers on the same H_2O instance. The unitary generators for orbital rotations has been screened using point group symmetry. We show the average number of iterations required by OS-WAHTOR among simulations that converge within chemical accuracy w.r.t. CASSCF. **G** is the number of successful runs, **iter** is the average number of orbital optimizer cycles for each successful run.

The comparison between the three optimization strategies (BFGS, Newton–Raphson and Trust-Region) shows clear differences in both robustness and efficiency. Overall, BFGS proves to be the most reliable optimizer. It achieves the highest number of chemically accurate convergences for all generator pools (**ALL**, **CAS-LIKE** and **RED**), with and without **PG-RED** on top, while also requiring a relatively low and stable number of orbital–optimization iterations. The Trust–Region method also exhibits robust convergence, with a success rate close to that of BFGS, even if it generally requires more iterations on average, indicating slower progress toward the energy minimum.

	BFGS		**Newton-Raphson**		**Trust Region**	
Gen type	**G**	**iter**	**G**	**iter**	**G**	**iter**
ALL	15	19.93	0	0	33	32.8
CAS	20	26.5	1	167	39	32
RED	14	21.86	0	0	41	30.2

Table 6.9: Comparison between different optimizers on the same H_2O instance. All unitary generators for orbital rotations has been used, without considering symmetry. We show the average number of iterations required by OS-WAHTOR among simulations that converge within chemical accuracy w.r.t. CASSCF. **G** is the number of successful runs, **iter** is the average number of orbital optimizer cycles for each successful run.

This is mainly due to the fact that, BFGS does not require to build any derivative operators, but simply search for the optimal set of parameters associated to each orbital–rotation generator, and directly minimizes the unitary transformation iteratively. In contrast, the Newton–Raphson optimizer is the least stable and efficient among the three. Trust–Region frequently fails to converge, and even in its successful runs the number of iterations can become excessively large, reflecting difficulties in handling the underlying optimization landscape. This poor performance can be attributed to the use of an *approximated CI*[9] solution at each macro-iteration, which introduces noise in the energy and Hessian evaluations. Since Newton–Raphson is highly sensitive to inaccuracies in the curvature information, its quadratic convergence is lost, leading to unstable behavior. Quasi-Newton methods such as BFGS are instead more resilient to these inaccuracies, explaining their superior performance in this setting. Therefore, BFGS offers the best balance between convergence stability and computational efficiency, Trust-Region provides a slower but still dependable alternative, and Newton–Raphson is too unstable to be practical, despite its theoretical appeal as a second-order method. This behavior is consistent across both symmetry-screened and unscreened generator pools, confirming the superior robustness of quasi-Newton schemes in variational orbital optimization.

6.6 WAHTOR–SCF applied to Industrial chemistry

Polyethylene is among the most widely produced polymers worldwide, with an annual output exceeding 100 million tons [237, 238]. Its versatility, durability, and low manufacturing cost have made it indispensable in sectors such as packaging, automotive components, coatings and advanced materials. The industrial relevance of polyolefins has driven continuous innovation in polymerization methods since the mid-20th century, spanning radical high-pressure technologies, Ziegler–Natta catalysis, Phillips catalysts, and, more recently, metallocene and post-metallocene systems [237, 239–243].

Modern industry requires not only efficient polymerization protocols, but also fine control over microstructure, structurization and molecular-weight distribution [238, 239]. These structural variables govern the macroscopic processability and mechanical response of polyethylene materials. In this context, a profound understanding of the elementary steps of polymer growth is essential, especially at the level of the electronic structure of reactive intermediates and transition states.

Historically, high-pressure radical polymerization has played a key industrial role in the pro-

---

[9] Incorrectly defined as *approximated CI*. The VQE expansion might also contains determinants that a CI expansion would not.

duction of low-density polyethylene (LDPE), and its fundamental chemistry remains relevant for chain-growth processes [244, 245]. However, a major breakthrough in polyolefin science came with the advent of coordination catalysts developed by Ziegler and Natta in the 1950s, which opened the way to stereoregular polymers with unprecedented control over chain architecture [239]. The elementary step governing Ziegler–Natta catalysis is the insertion of an olefin into a metal–carbon bond, typically described by the Cossee–Arlman mechanism [243]. Despite decades of study, the electronic structure of this step remains nontrivial, as it often involves transient multi-reference character.

The original objective of this project was to address such electronically complex processes in Ziegler–Natta chemistry using hybrid quantum–classical approaches. However, the full catalytic system was deemed too electronically and structurally demanding for a first quantum implementation. For this reason, we selected a reduced prototype, the free-radical polymerization of ethylene, preserving the central chemical pattern of *olefin activation and C–C bond formation* while drastically reducing the computational burden [246]. This simplification enables a systematic validation of the hybrid multi-configurational framework developed in the previous chapters, serving as a stepping stone toward future simulations of realistic Ziegler–Natta catalysts.[10]

6.6.1 Free–Radical Model Polymerization of Ethylene

Free-radical polymerization of ethylene proceeds through initiation, propagation, and termination steps [238, 247, 248]. In contrast to coordination catalysis, the reactive center is here a carbon-centered radical, and chain growth occurs through successive $\pi \to \sigma$ bond conversions. This mechanism is emblematic of high-pressure LDPE production [244, 245], and provides a clean minimal platform for studying the electronic structure of olefin activation.

In the model investigated, a hydrogen radical triggers chain initiation, forming a carbon-centered radical that adds subsequent ethylene monomers. The reaction coordinate features two key transition states,one for initiation and one for propagation, each exhibiting partial π-bond cleavage and formation of a new σ bond. This chain reaction continues as radicals repeatedly add to new monomers, growing a long polymer chain. These configurations display diradical-like character that cannot be reliably captured by single-reference methods, thereby motivating a multi-configurational treatment [33, 109, 110, 133]. Regarding the model reaction investigated, to the best of our knowledge, it has been studied, with ab-initio method DFT and MP2, only for the activation step in Nassabeh *et al.* [247].

The elementary steps of the model reaction are reported in Figure 6.16.

We will use a minimal active space containing only those orbitals which are directly involved in the bond breaking/making process for each structure in this reaction. Since in our termination step this minimal active space has just two orbitals for the triplet reactant and transition state, which is a single ROHF determinant, there is no need for a CASSCF computation for this step. We focus on the initiation and propagation steps, where we will use three active orbitals for the doublet reactants (R1 and R2) and doublet transition states (TS1 and TS2). These three active orbitals are the π, π^* along C-C double bond in R1, R2 and the non-bonding orbital containing the single electron on H or C_2H_5 radical. For TS1 and TS2, the three orbitals correspond to the σ, σ^* orbitals

---

[10] This research activity was supported by ENI (Ente Nazionale Idrocarburi).

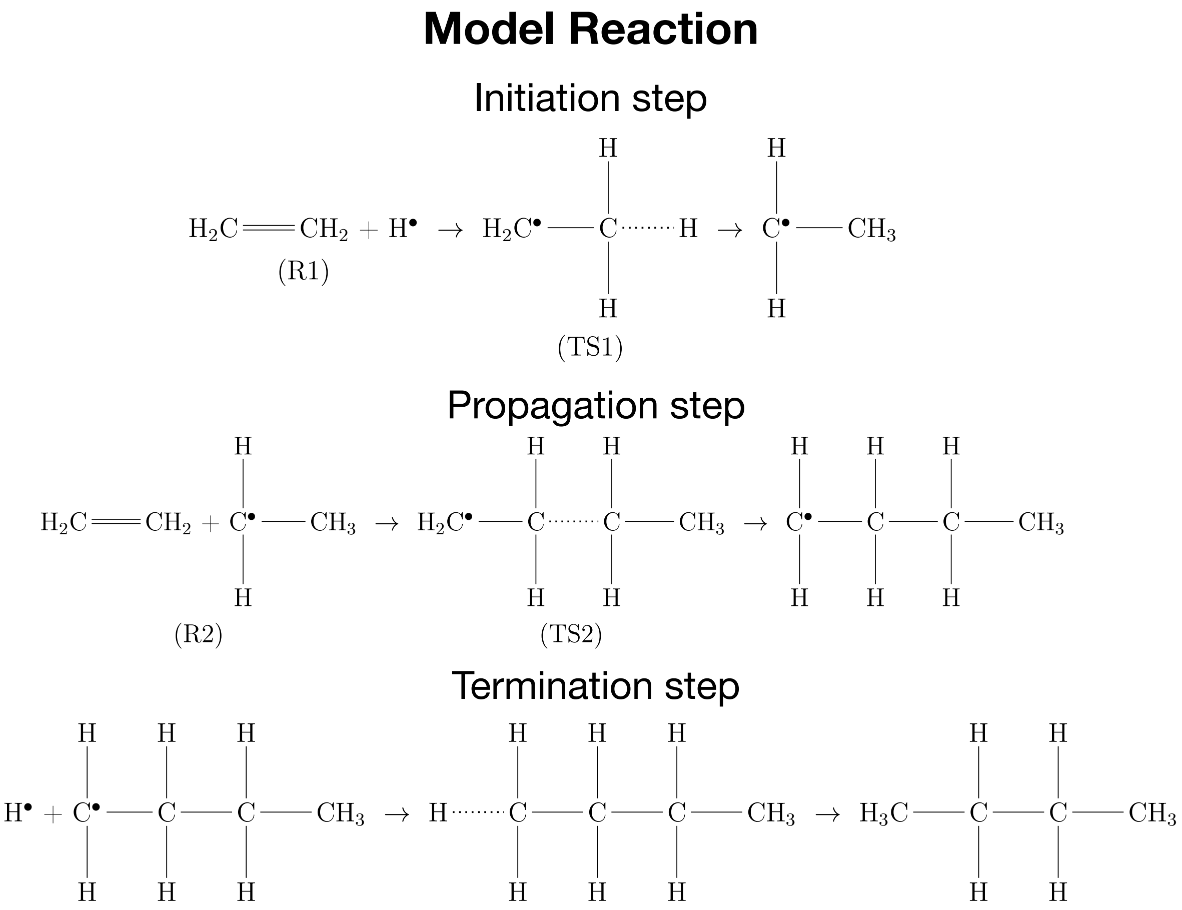

Figure 6.16: Three steps of the model reaction of free-radical polymerization of ethylene.

along the newly formed C-H/C-C bond respectively along with the non-bonding orbital having the lone electron.

Since this model reaction has no literature benchmarks, we find out optimized geometries of the reactants and transition states by performing a NEB-TS computation in ORCA 6.0 [178] at the CCSD(T) level of theory using the cc-pVDZ basis. These geometries are shown in Appendix Figure C.5 and Figure C.6. We computed the activation energies required to reach the transition state from the reactants in the initiation and propagation steps of this model reaction, i.e. R1 to TS1 and R2 to TS2. The reaction profiles for the first two steps at CASCI and CASSCF level using (3,3) CAS with the CCSD(T) optimized geometries are shown in Figure 6.17 and Figure 6.18.

6.6.2 WAHTOR–SCF on Model reaction

In this section we present the results of our WAHTOR–SCF procedure on the model reaction. In terms of WAHTOR–SCF variant, given the reduced size of the active spaces for initial, propagation, and termination step, we opted for HEA–WAHTOR–SCF. In particular, we employed the *SO(4)* variant of WAHTOR–SCF, with depth $d = 3$, total number of parameters 114, and total number of CNOTs 30. Every reaction step has been solved by running a single VQE run given the small size of the active space and the accuracy of the single point VQE as reported in Appendix Figure C.7. The VQEs are optimized using L-BFGS-B `scipy` optimizer. For WAHTOR–SCF specific settings, we

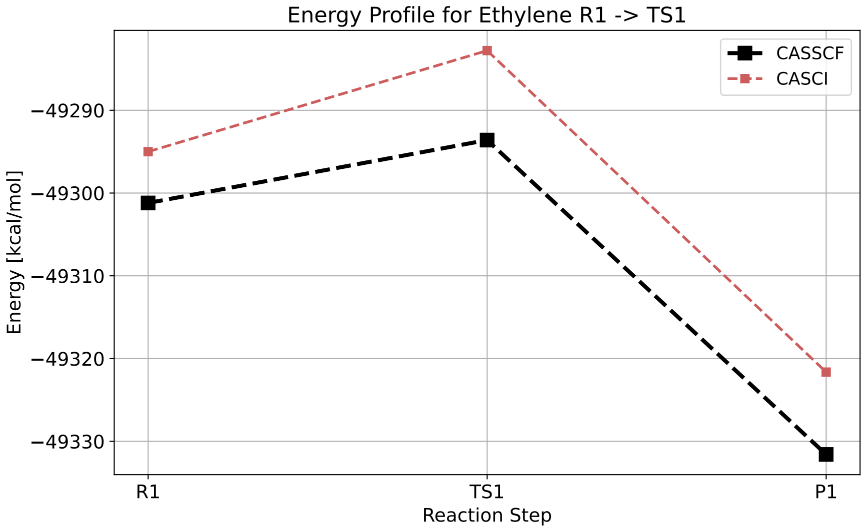

Figure 6.17: Reaction profile in the initiation step. Minimal CAS(3,3) used.

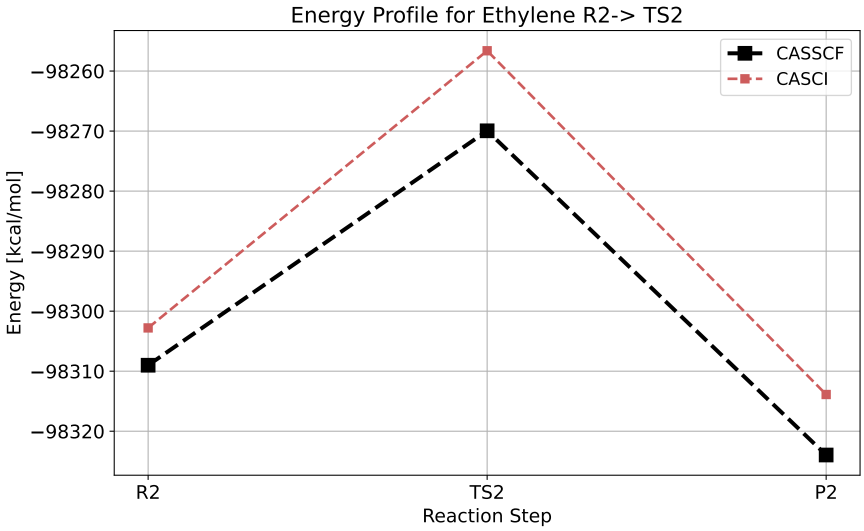

Figure 6.18: Reaction profile in the propagation step. Minimal CAS(3,3) used.

used Trust-Region second-order optimizer for the orbital–rotation optimization update and **RED** (minimal) generators pool, i.e. occupied-active, active-active, and active-virtual.

From the simulation, according to Table 6.10, the best VQE performace is in good agreement with the CASSCF benchmark calculation. In fact, for each of the structure the value recovered by WAHTOR–SCF is within chemical accuracy w.r.t. the corresponding CASSCF reference calculation.

For each of the structure, in Table 6.11, we report all the energies evaluated by out hybrid methdo, WAHTOR–SCF, and the reference CASSCF calculations. From the reported data, it is possible to find good agreement between our procedure and the classical benchamarks. Additional, as we introduce in the previous section, we are interested in the activation energies. Activation energies are obtained by computing the energy difference between a transition state, TS$_i$, and the energy of the respective reactant, R$_i$. We compared with the reference calculations reported in Figure 6.17 and Figure 6.18, specifically the CASSCF calculation. As reported in Table 6.12, our

Structure	N_R	N_T	**CASSCF**	**WAHTOR–SCF**	%E$_{corr}$
R1	53	153	-78.566463	-78.566365	99.66
TS1	53	153	-78.554354	-78.554253	99.69
P1	53	153	-78.614882	-78.614721	∼
R2	101	297	-156.665376	-156.665279	99.66
TS2	101	297	-156.603140	-156.602794	99.14
P2	101	297	-156.689209	-156.689209	∼

Table 6.10: Correlation energy recovered by QASSCF w.r.t. "classical" CASSCF for the initiation and propagation steps in the model polymerization reaction. N_R is the number of orbital to be rotated while N_T is the number of orbital–rotation generators. **CASSCF** and **WAHTOR–SCF** are the energy recovered by the respective method in Hartree, while %E$_{corr}$ is the percentage correlation energy recovered by WAHTOR–SCF w.r.t. CASSCF. P1 and P2, the product states, are still obtained with WAHTOR–SCF but the procedure took only one iteration.

Structure	**CASSCF**	**WAHTOR–SCF**
R1	-49301.200	-49301.138
TS1	-49293.601	-49293.538
P1	-49331.583	-49331.482
R2	-98309.008	-98308.947
TS2	-98269.954	-98269.737
P2	-98323.963	-98323.881

Table 6.11: Comparison of the energies obtained with our hybrid approach, WAHTOR–SCF, against the benchmark calculation with standard CASSCF in kcal/mol for each of the structure and reaction step.

HEA–WAHTOR–SCF method is in good agreement with the reference calculation, matching exactly the initiation activation energy (TS1 - R1), and producing a propagation step (TS2 -R2) activation energy slightly above the benchmark one, specifically 0.16 kcal/mol above. From Table 6.12 we see that our quantum version to compute CASSCF energies is well within the limit of chemical accuracy (∼1 kcal/mol) required in quantum chemical calculations. In Figure 6.19 and Figure 6.20, a complete overview on the benchmark datapoint, our quantum approach simulation results, and the relative error, can be found.

Reaction	**ΔE**	**ΔE**	**TS**		**R**	
Step	**CASSCF**	**WAHTOR**	**CASSCF**	**WAHTOR**	**CASSCF**	**WAHTOR**
Initiation	7.60	7.60	-49293.601	-49293.538	-49301.200	-49301.138
Propagation	39.05	39.21	-98269.954	-98269.737	-98309.008	-98308.947

Table 6.12: Activation energy required in the initiation and propagation steps of the model polymerization reaction. All units are reported in kcal/mol.

6.7 Discussion and final remarks on hybrid self-consistent procedure

The results presented in this chapter, two hybrid self-consistent-field schemes were implemented and tested: the *HEA-WAHTOR–SCF* and the *ADAPT-WAHTOR–SCF* . Both approaches integrate a classical orbital optimization cycle with a quantum configuration interaction solver, i.e.

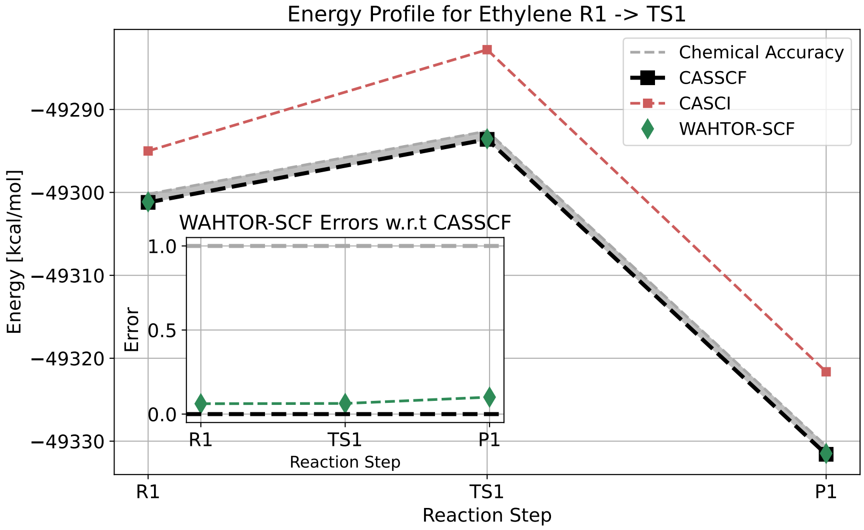

Figure 6.19: Complete comparison between our WAHTOR–SCF method and the quantum chemistry benchmarks with CASCI and CASSCF for the initiation step. Values in the main panel are reported in kcal/mol as well as the errors in the in-box plot. Only R1 and TS1 are reported due the fact that the P1 is composed by a single determinant, i.e. the energy corresponds to ROHF mean field on that structure.

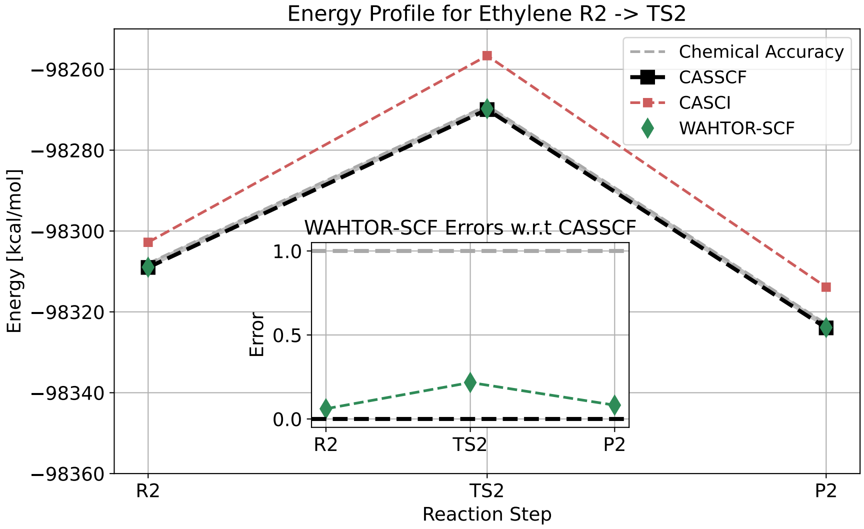

Figure 6.20: Complete comparison between our WAHTOR–SCF method and the quantum chemistry benchmarks with CASCI and CASSCF for the propagation step. Values in the main panel are reported in kcal/mol as well as the errors in the in-box plot. Only R2 and TS2 are reported due the fact that the P2 is composed by a single determinant, i.e. the energy corresponds to ROHF mean field on that structure.

VQE, extending the original WAHTOR procedure. The HEA-WAHTOR–SCF employs a fixed hardware-efficient ansatz optimized variationally across iterations, while the ADAPT-WAHTOR–SCF introduces adaptiveness. Together, these two variants provide complementary routes to balance expressibility, convergence rate, and efficiency.

WAHTOR–SCF has been benchmarked across a diverse set of molecular systems, and in all cases, our approach reproduced conventional CASSCF energies within the chemical accuracy, validating both the parameter optimization strategy and the HQC Self–Consistent Field loop. The HEA-WAHTOR–SCF method demonstrated smooth and predictable convergence for small and moderately correlated systems, while the ADAPT-WAHTOR–SCF variant exhibited reasonably superior accuracy and faster convergence for strongly correlated cases.

Our method then was applied to a representative industrially motivated process: the free-radical polymerization of ethylene. This model system captures the fundamental bond-breaking and bond-forming events underlying olefin polymerization and provides a rigorous test of the hybrid algorithms on a chemically meaningful potential energy surface. Using a compact CAS(3,3) active space including the π/π^* (or σ/σ^*) orbitals and the singly occupied non–bonding orbital, both the initiation and propagation steps were successfully modeled. The computed potential energy profiles obtained with WAHTOR–SCF algorithm reproduced classical CASSCF results with excellent agreement, differing by less than 0.2 kcal/mol in activation energies. Quantitatively, the activation barriers were found to be 7.6 kcal/mol for initiation and 39.0 kcal/mol for propagation, with the latter representing the rate-determining step in the radical chain mechanism. The nearly identical results between HEA-WAHTOR–SCF and classical CASSCF confirm that the hybrid methods not only preserve the underlying electronic structure but also deliver quantitatively accurate potential energy surfaces. These results demonstrate, despite considering a small sized case, that the quantum variational solvers can capture the essential multi-reference character associated with $\pi \to \sigma$ conversion and open-shell spin polarization without introducing numerical artifacts.

In conclusion, the combination of algorithmic development and chemical validation presented in this chapter highlights the versatility of the hybrid framework. The HEA-WAHTOR–SCF and ADAPT-WAHTOR–SCF methods form two robust strategies that can be tailored to the complexity of the target system: the former optimized for computational efficiency, the latter for systematic accuracy and ansatz compactness. Obviously, our procedure can be refined and make it more suitable for large scale applications, especially from the point of view of single point solutions, i.e. finding more expressive HEA, and including additional reduction techniques, such as *Commuting–Pauli groups*, for reducing even more the computational effort required by both the classical and quantum side of the procedure. Additionally, for the adaptive version of WAHTOR–SCF, more work can be done in order to analyze which operator pool fits better in our pipeline and if the evaluation of the gradients in the ADAPT-VQE step can be recycled also for the derivative evaluation, especially when scaling to bigger active space. Finally, it would be interesting to see how our WAHTOR–SCF approach performs against longer reactions and bigger system, hopefully being able to target the original Ziegler-Natta reaction. To conclude, as said in Section 2.4.3, a MCSCF-like approach is able to fully recover, if expressive enough, the static correlation in the active space, but only a portion of the dynamical correlation of the whole system. We should therefore focus also on the introduction of perturbative corrections, in order to perform CASPT2 [34–36] or NEVPT2 [30–33]

Now, having defined our hybrid procedure for dealing with both static and dynamical correlation, having incorporated both static and dynamical correlation, we can now move to the last piece of the pipeline. The structure optimization, which generally is the first step for starting the whole calculation, here is treated in the last chapter of the manuscript, due to the time frame in which the project has been started. The geometry optimization task is finally explained in the next chapter.

Chapter 7

Quantum Geometry optimization for Drug Discovery

Geometry optimization is a essential step of molecular modeling, as it determines the equilibrium structure from which physical / chemical properties and molecular descriptors can be derived. Traditionally, this process has been carried out entirely on classical computers, using analytical or numerical derivatives of the energy to relax nuclear positions toward the minimum. However, the exponential complexity of accurate electronic structure calculations makes these optimizations challenging for strongly correlated systems or large molecular assemblies. Hybrid classical–quantum algorithms have emerged as a promising solution. Ideally, quantum computers can be used to evaluate molecular energies, and then, relying on classical routines to guide the optimization. In particular, the central difference approach provides a practical way to approximate nuclear gradients directly from quantum-evaluated energies, thus making geometry optimization feasible even in the Noisy Intermediate–Scale Quantum (NISQ) era.

Within this framework, optimized quantum geometries offer a unique opportunity to enrich molecular representation. While *DompeKey (DK)* provide a powerful set of substructure-based descriptors for chemical space mapping, machine learning, and virtual screening, their current formulation is largely independent of geometry. By integrating equilibrium structures obtained from Hybrid Quantum–Classical optimization, DKs can be augmented with structural fidelity that captures subtle conformational and electronic effects often missed by purely topological descriptors. This extension introduces a new descriptor variant in which quantum-refined geometries serve as an input layer, allowing DompeKeys to encode not only connectivity and functional groups but also the precise arrangement of atoms in three-dimensional space.

In this chapter, we introduce our proof–of–concept integration of Hybrid Quantum–Classical structure optimization algorithm, namely *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)*. QA-q-GeomOpt exploits HEA circuit, finite difference method, and a classically–inspired *potentially quasi-adiabatic* routine, to search the equilibrium geometry. We demonstrate both the applicability of the method on on small DKs, i.e. Amine, Imine, and Aldehyde. Additionally, we show the implementational side, in terms of quantum computing simulation backend, classical`Qiskit` [236] and `CUDA-quantum` [101] CPU statevector simulators, as well as `CUDA-quantum` GPU backend. Procedure which adaptability to different quantum chemistry and simulation backend can be interchanged, in a fully parallelized environment. The implementation demonstrates the feasibility of incorporating HQC geometry optimizations into cheminformatics workflows, but much remains to be refined.

The chapter will be defined as follows: In Section 7.1 we will introduce they theory behind geometry optimization, i.e. how to extrapolate energy gradients w.r.t. atomic coordinates and forces. After that, in Section 7.2, we will briefly recap which HQC techniques exist in literature, splitting between how to evaluate forces, navigate the potential energy curve, and additional improvements. We will finally introduce our *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)* procedure, in Section 7.3, dedicating completely Section 7.4 for the implementational side. After defining the framework, we will present the application on DompeKeys, in Section 7.5. Finally, Section 7.2 will be used for discussion and consideration.

7.1 Molecular gradients for variational and fully variational wavefunctions

Within the Born–Oppenheimer Approximation, the molecular forces are defined as the total derivative of the electronic energy with respect to the nuclear coordinates,

$$\mathbf{F} = -\frac{dE}{d\mathbf{x}}. \tag{7.1}$$

Let $\boldsymbol{\theta}$ denote the variational parameters of the wavefunction (CI coefficients or, in a hybrid quantum algorithm, the circuit parameters), and let the orbitals be parametrized by a separate set of rotations $\boldsymbol{\kappa}$. By applying the chain rule one obtains

$$\frac{dE}{d\mathbf{x}} = \left.\frac{\partial E}{\partial \mathbf{x}}\right|_{\boldsymbol{\theta},\boldsymbol{\kappa}} + \frac{\partial E}{\partial \boldsymbol{\theta}}\frac{d\boldsymbol{\theta}}{d\mathbf{x}} + \frac{\partial E}{\partial \boldsymbol{\kappa}}\frac{d\boldsymbol{\kappa}}{d\mathbf{x}}. \tag{7.2}$$

For a variational wavefunction, the stationarity condition

$$\frac{\partial E}{\partial \boldsymbol{\theta}} = 0 \tag{7.3}$$

eliminates the CI (or circuit) response. This is true for both *variational* (CASCI, truncated CI, VQE with fixed ansatz) and *fully variational* (HF, CASSCF) electronic structures. The key distinction lies in the second response term. Fully–variational methods also satisfy

$$\frac{\partial E}{\partial \boldsymbol{\kappa}} = 0 \tag{7.4}$$

because orbital rotations are optimized at every geometry. In this ideal case, Equation 7.2 reduces to the Hellmann–Feynman [249, 250] form (with basis-set Pulay contributions [251, 252] only if the atomic orbitals themselves depend on $\mathbf{x}$),

$$\frac{dE}{d\mathbf{x}} = \underbrace{\left\langle \Psi \left| \frac{\partial \hat{H}}{\partial \mathbf{x}} \right| \Psi \right\rangle}_{\textit{Hellmann-Feynman}} + \underbrace{\left\langle \frac{\Psi}{\partial \mathbf{x}} \left| \hat{H} \right| \Psi \right\rangle + \left\langle \Psi \left| \hat{H} \right| \frac{\Psi}{\partial \mathbf{x}} \right\rangle}_{\text{Pulay Forces}}. \tag{7.5}$$

They needs to be accounted every time the basis set depends explicitly on the atom coordinates.

For variational but not fully–variational methods, such as CASCI or VQE with non-relaxed orbitals, the energy is stationary only with respect to $\boldsymbol{\theta}$. The orbital response term survives and gives rise to *Pulay forces*,

$$\mathbf{F}_{\text{Pulay}} \propto \frac{\partial E}{\partial \boldsymbol{\kappa}} \frac{d\boldsymbol{\kappa}}{d\mathbf{x}}, \tag{7.6}$$

which account for the fact that the underlying one–electron basis (and therefore the Fock and Hamiltonian matrices) changes when the nuclei move. Physically, Pulay terms correct the Hellmann–Feynman force to ensure that the electronic state remains variationally optimal with respect to the orbital manifold. This is why CASCI, truncated CI and fixed–ansatz VQE do *not* admit a pure Hellmann–Feynman gradient: even though the wavefunction is variational in the CI (or circuit parameters) space, it is not variational in the orbital rotation space [26].

7.2 Hybrid Quantum–Classical Implementation in literature

Hybrid Quantum–Classical (HQC) approaches to molecular geometry optimization replace the electronic–structure step of a classical optimizer with a variational quantum solver, while keeping the nuclear update on the classical side. Early demonstrations adopted the simplest workable loop by obtaining a variational ground state at a trial geometry using VQE, approximate nuclear forces via finite differences of nearby single points, and update coordinates with a standard optimizer. Although straightforward, this architecture immediately raised two methodological issues central to robust geometry optimization on near-term hardware: *force quality*, i.e., whether gradients are formally consistent with the Hellmann-–Feynman/Pulay [249–252] framework in a variational setting, and *landscape navigation*, i.e., how to maintain stable convergence on a nonconvex variational energy landscape across a sequence of geometries.

7.2.1 Force–Evaluations

A first family of works addressed *force evaluation*. On the theoretical side, hybrid derivative formulations clarified how Hellmann–Feynman contributions and wavefunction–response (Pulay–like) terms should be handled within variational circuits and multi-state settings, providing a route to analytic forces beyond naive finite differences [253, 254]. While several studies still rely on finite differences for practicality on NISQ devices, these analises established the conditions under which gradients are formally correct (state consistency at displaced geometries) and how response terms can be incorporated in principle, thereby informing the design of more robust geometry update loops.

7.2.2 PES exploration

A second strand focused on *closing the geometry loop* with quantum resources. Delgado *et al.* formulated a variational algorithm in which circuit and Hamiltonian parameters are co–optimized to locate equilibrium structures directly on the Born–Oppenheimer surface, demonstrating end–to–end geometry optimization beyond isolated energy evaluations [255]. Complementary efforts explored quantum–driven molecular dynamics and PES traversal using variational states with explicit gradient handling, concluding that consistent forces are essential for stable structural updates [253, 254]. For electronically challenging regions (near–degeneracies, conical intersections), HQC algorithms have been proposed to navigate multi–surface features and track minima in the presence of strong state interactions, illustrating how hybrid strategies can be extended to more intricate PES topologies relevant to relaxation paths [256].

7.2.3 Beyond VQE–based

Beyond fixed–ansatz VQE, *ansatz–adaptive and active–space* strategies are increasingly used to stabilize PES navigation. ADAPT–style procedures dynamically grow the ansatz along a scan, mitigating over–parameterization and helping maintain accuracy across geometries; while many ADAPT studies emphasize energies, their operator–selection logic is expressly motivated by smooth coverage of the PES and has been evaluated in that context [257, 258]. Also, as we already seen in the previous chapter, quantum–CASSCF/active–space workflows and quantum embedding have been pushed to make geometry optimization viable for larger molecules by employing quantum resources to treat accurately only chemically relevant subspaces, which is particularly attractive when structural relaxation requires multi–reference character [219, 259]. In fact, as suggested by Mizukami *et al* [217], employing fully variational algorithms HQC algorithm will favor the resolution of the structure optimization task on quantum devices, for example, using the Orbital–Optimized UCC procedure [216, 217]. Recent embedding formulations explicitly target geometry optimization/co–optimization by coupling fragmentation (DMET[1]–like) with variational solvers to reduce qubit counts and classical–quantum nesting costs, making routine for structural refinement closer to practical scales [260].

From these developments, a general tendency appears. For near–term devices, *finite–difference forces* remain the most common pragmatic choice in geometry loops, but must be deployed with care (ensuring state consistency at displaced geometries and controlling stochastic/shot noise) to avoid bias in updates. *Analytic–gradient frameworks* are maturing and clarify how to include response terms, positioning HQC geometry optimization on firmer theoretical ground as hardware and measurement strategies improve. On the *navigation* side, the nonconvex nature of variational landscapes necessitates warm–starting and, when needed, multi–start or branching to escape local minima, while *ansatz adaptivity* and *embedding* provide scalable routes to maintain accuracy across a sequence of geometries.

7.3 Quasi-Adiabatic Hybrid Classical-Quantum Geometry Optimization

As briefly introduced in the previous sections, the geometry optimization procedure has the objective to identify the equilibrium nuclear configurations by iteratively descending the Potential Energy Surface (PES) toward stationary points of the electronic ground–state energy. Once the electronic energy is obtained through a variational quantum algorithm at a given molecular geometry, nuclear coordinates are updated based on estimates of the energy gradient until convergence criteria are satisfied, e.g. molecular gradients or bond/angle not refined anymore. In this framework, forces

---

[1] Density Matrix Embedding Theory

are evaluated from single–point electronic structure calculations at displaced geometries, and an effective combination of energy evaluation, gradient estimation, and wavefunction preparation is required to achieve both accuracy and stability along the optimization trajectory.

In this section, we present our proof–of–concept Hybrid Quantum–Classical method for structure optimization, namely *geomopt*. The method presented here integrates three key components. First, nuclear forces are computed through finite–difference approximations, leveraging variationally optimized electronic energies at slightly perturbed geometries. Secod, an *adiabatic* wavefunction–tracking strategy is employed to accelerate convergence by reusing the optimized state from the previous step whenever the PES evolves smoothly. Third, a complementary *non–adiabatic* exploration strategy is introduced to prevent variational optimization from becoming trapped in local minima. The combination of these elements provides a proof–of–concept algorithm that is flexible, reliable, scalable and capable of achieving reasonably accurate equilibrium geometry for different molecular systems. The overall procedure is shown in a schematic and simplified diagram in Figure 7.1. The following subsections detail these three components of the geometry optimization protocol.

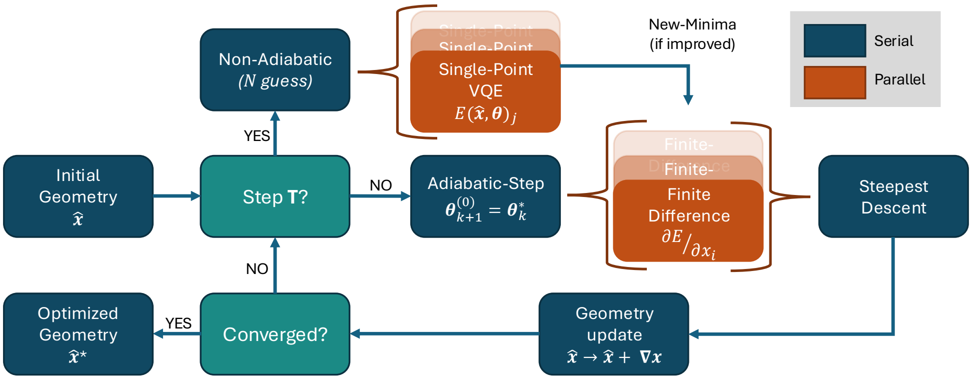

Figure 7.1: Quasi-Adiabatic Quantum Geometry–Optimization pipeline scheme.

7.3.1 Finite–Differences approach

In the context of molecular geometry optimization, the evaluation of nuclear forces traditionally relies on the computation of energy gradients with respect to nuclear coordinates, $\partial E/\partial x_i$, which can be estimated using finite-difference formulas. In particular, given an initial geometry $\mathbf{x}$, for a given coordinates x_i, the corresponding derivative in the finite difference framework is defined as

$$\frac{dE}{d\mathrm{x}_i} = \frac{E(\mathrm{x}_i + \delta) - E(\mathrm{x}_i - \delta)}{|2\delta|}, \tag{7.7}$$

where δ is the nuclear displacement applied on the i-th nuclear coordinate, x_i. This strategy implicitly incorporates both Hellmann–Feynman [249, 250] and Pulay-like [251, 252] effects, because the total energy difference reflects all contributions coming from the change in the electronic problem, regardless of their origin. The finite–difference strategy provides a simple and implementation–friendly estimate of the forces, at the price of requiring multiple energy evaluations and ensuring that each displaced geometry is sufficiently reoptimized.

Within a VQE framework, the molecular Hamiltonian $\hat{H}(\mathbf{x})$ is constructed from a given nuclear configuration $\mathbf{x}$, and the electronic ground state is approximated by a parametrized wavefunction $|\psi(\boldsymbol{\theta};\mathbf{x})\rangle$, whose parameters $\boldsymbol{\theta}$ are optimized according to the variational principle (introduce in Section 3.13). Once the optimal parameters $\boldsymbol{\theta}^*$ are obtained, the ground-state energy is given by

$$E_0(\mathbf{x}) = \frac{\langle\psi(\theta;\mathbf{x})|\,\hat{H}(\mathbf{x})\,|\psi(\theta;\mathbf{x})\rangle}{\langle\psi(\theta;\mathbf{x})|\psi(\theta;\mathbf{x})\rangle}. \tag{7.8}$$

A naive finite–difference implementation of the nuclear gradient proceeds by introducing a displacement $\pm d_i$ along the i-th coordinate and constructing the perturbed Hamiltonians $\hat{H}_\pm = \hat{H}(\mathbf{x} \pm d_i)$. The derivative is then approximated as

$$\frac{d\tilde{E}(\mathbf{x})}{dx_i} = \frac{\tilde{E}_+ - \tilde{E}_-}{2d_i} = \frac{\langle\psi(\boldsymbol{\theta}^*;\mathbf{x})|\,\hat{H}_+\,|\psi(\boldsymbol{\theta}^*;\mathbf{x})\rangle - \langle\psi(\boldsymbol{\theta}^*;\mathbf{x})|\,\hat{H}_-\,|\psi(\boldsymbol{\theta}^*;\mathbf{x})\rangle}{2d_i}, \tag{7.9}$$

where the same optimized wavefunction $|\psi(\boldsymbol{\theta}^*;\mathbf{x})\rangle$ is used for all three geometries (central, + and −). This is a reasonable approximation under the assumption that the wavefunction $|\psi(\boldsymbol{\theta}^*;\mathbf{x})\rangle$ is close enough to the true groundstate.

However, this approach is theoretically inconsistent, in fact, by construction, $|\psi(\boldsymbol{\theta}^*;\mathbf{x})\rangle$ is the ground state of $\hat{H}(\mathbf{x})$, but it is generally *not* the ground state of $\hat{H}_\pm$. As a consequence, the resulting forces might neglect wavefunction–response contributions and are not consistent with the Hellmann–Feynman framework [249, 250], nor with the Pulay corrections [251, 252] that arise when the basis (or ansatz) depends on nuclear coordinates. The theoretically correct finite–difference expression requires two independent ground–state optimizations,

$$|\psi_+\rangle = \mathrm{GS}[\hat{H}_+], \qquad |\psi_-\rangle = \mathrm{GS}[\hat{H}_-], \tag{7.10}$$

which yield

$$\frac{dE(\mathbf{x})}{dx_i} = \frac{E_+ - E_-}{2d_i} = \frac{\langle\psi_+|\,\hat{H}_+\,|\psi_+\rangle - \langle\psi_-|\,\hat{H}_-\,|\psi_-\rangle}{2d_i}, \tag{7.11}$$

representing the true finite–difference gradient. This formulation is fully consistent with the variational principle and incorporates the correct electronic response to nuclear displacements. In practice, however, performing two fully independent VQE optimizations for each displaced geometry may be computationally demanding and potentially unstable, as the two optimizations can follow different variational trajectories and converge to states of unequal quality [261].

A practical strategy is therefore to initialize both optimizations from the central–geometry wavefunction parameters, i.e., setting $\boldsymbol{\theta} = \boldsymbol{\theta}^*$ for the + and − calculations. This *warm–start* approach significantly reduces computational cost while ensuring that the perturbed Hamiltonians are minimized with respect to their own electronic wavefunctions, thus recovering a consistent definition of the finite–difference gradient and avoiding the theoretical inconsistency of the naive scheme. The warm start strategy provides a physically sound balance between accuracy and efficiency, and it represents the appropriate implementation for finite-difference force evaluation within variational quantum simulations of molecular systems.

7.3.2 Adiabatic Search for Accelerated Geometry Optimization

In molecular geometry optimization with VQE, when structural changes between successive optimization steps are sufficiently small, it is advantageous to adopt an *adiabatic* search strategy, in which the converged wavefunction from the previous geometry is used as the initial guess for the next single-–point calculation. This idea is loosely inspired by the concept of adiabatic continuity on a PES which intuitively states that *if the molecular geometry is perturbed only slightly, the ground state at the new geometry is expected to remain close, in parameter space, to the ground state at the previous geometry.* Given the optimal parameters $\boldsymbol{\theta}_k^\star$ at geometry $\mathbf{x}^{(k)}$, the next VQE iteration is initialized as

$$\boldsymbol{\theta}_{k+1}^{(0)} = \boldsymbol{\theta}_k^\star. \tag{7.12}$$

For sufficiently small nuclear displacements $\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\| \ll 1$, the energy can be expanded as

$$E(\boldsymbol{\theta}_{k+1}^\star; \mathbf{x}^{(k+1)}) \approx E(\boldsymbol{\theta}_k^\star; \mathbf{x}^{(k)}) + \mathcal{O}(\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|), \tag{7.13}$$

which justifies a warm–start initialization, since the optimal solution at step $k+1$ lies near $\boldsymbol{\theta}_k^\star$ in parameter space, enabling faster convergence. Operationally, the adiabatic strategy reuses $\boldsymbol{\theta}_k^\star$ as the initial guess for the VQE optimization at geometry $\mathbf{x}^{(k+1)}$, thereby reducing the number of classical optimization iterations.

By avoiding repeated random or Hartree–Fock initializations, the adiabatic strategy keeps the classical optimizer in the correct basin of attraction on the variational landscape, typically reaching convergence in a small number of iterations and reducing the total computational cost by up to an order of magnitude. This behavior reflects the empirical smoothness of the PES and the expected continuity of the electronic ground state for small structural deformations. However, the adiabatic strategy remains intrinsically local i.e. it efficiently follows a single branch of the PES, but cannot prevent convergence to sub–optimal local minima in the presence of strong landscape non-convexity or qualitative changes in the electronic structure. For this reason, the search remains adiabatic only when the PES evolves smoothly, and we introduce an additional non–adiabatic step, to check whether we are in this condition, trying to ensure robustness beyond the local PES topology. A pictorial representation of the adiabatic procedure is depicted in panel Figure 7.2(a).

7.3.3 Non–Adiabatic Search for New Local Minimum Identification

To overcome the limitations imposed by the non-convex nature of the variational optimization landscape, in which VQE procedures can become trapped in suboptimal local minima [261], we incorporate a *potentially non–adiabatic* search strategy designed to improve minima exploration. At predefined intervals T along the geometry optimization, instead of warm–starting from the previously converged wavefunction, the algorithm generates an ensemble of initial parameter configurations

$$\{\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(N)}\}, \tag{7.14}$$

obtained either by perturbing the previous solution or by fully randomized initialization in the range $(-\pi, \pi)$. Each candidate is variationally optimized, and the retained state is chosen according to

$$\boldsymbol{\theta}_k^{\star} = \underset{j\in\{1,\dots,N\}}{\arg\min}\; E(\boldsymbol{\theta}^{(j)}; \mathbf{x}^{(k)}), \tag{7.15}$$

corresponding to the branching energy rule

$$E_{\text{branch}}(\mathbf{x}^{(k)}) = \min_{j} \left\{ E(\boldsymbol{\theta}^{(j)}; \mathbf{x}^{(k)}) \right\}. \tag{7.16}$$

This branching mechanism enables, if exploiting enough expressive and accurate circuit, and ideal global exploration of the PES and mitigates the risk of converging to suboptimal local minima, particularly in the presence of complex landscapes or quasi-degenerate states.

Operationally, the non–adiabatic step samples distinct regions of the parameter landscape by launching N independent VQE optimizations,

$$\boldsymbol{\theta}^{*(j)} = \arg\min_{\boldsymbol{\theta}} E(\boldsymbol{\theta}; \mathbf{x}^{(k)}), \qquad j = 1, \dots, N, \tag{7.17}$$

and selects the lowest–energy solution as the new reference state for continuing the geometry optimization. After this branching phase, the algorithm reverts to the adiabatic mode to efficiently refine the geometry, reactivating non–adiabatic exploration every T steps. The term *potentially non–adiabatic* reflects the fact that the evolution remains predominantly adiabatic, but temporarily deviates from the adiabatic path whenever a more favorable minimum is discovered. This hybrid adiabatic/non–adiabatic strategy balances global robustness with local efficiency, making it well–suited for challenging molecular systems where following a single adiabatic trajectory is insufficient to guarantee reliable ground–state tracking across the VQE. A pictorial representation of the non–adiabatic procedure is depicted in panel Figure 7.2(b).

It is worth emphasizing that the geometry, $\mathbf{x}^{(k)}$, remains fixed during the non–adiabatic branching step. While the adiabatic update explores the PES along nuclear coordinates, the non–adiabatic procedure explores instead the variational landscape $E(\boldsymbol{\theta}; \mathbf{x}^{(k)})$ at fixed geometry, in oder to escape local minima of the VQE energy surface.

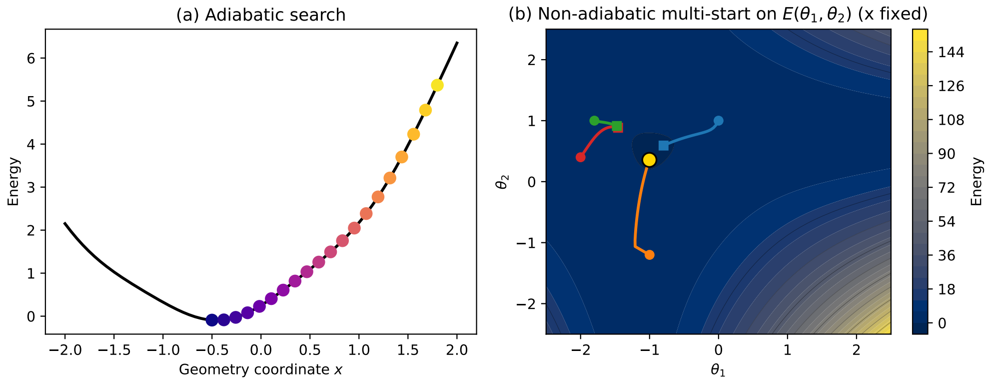

Figure 7.2: Adiabatic versus non–adiabatic search during VQE–driven geometry optimization. In panel **(a)**, a warm–start strategy follows a single valley of the PES, yielding fast local convergence for small structural changes. In panel **(b)**, when the adiabatic strategy risks becoming trapped in a local minimum, a non–adiabatic multi-start exploration is performed at fixed nuclear geometry: several candidate trajectories are launched on the variational energy landscape $E(\theta_1, \theta_2)$ and the best solution is selected.

7.4 Framework structure

In this section, we report the implementation and coding part of the algorithm, diving them between how we deal with the chemistry side, i.e. the generation of the molecules and associated operators, and the computational side, i.e. how we simulate the circuit, how the evaluation is computed, and also, the technical side in terms of parallelize routines.

7.4.1 Quantum Chemistry Backend

Aside form the computational setting, in terms of convergence thresholds, number of macro– and micro–iteration, frequency of *potentially-non-adiabatic* searches or update step, for citing some of them, the only inputs that the *QA–q–GeomOpt* algorithm should (at least) ideally receive are the initial coordinates $\mathbf{x}$, the basis–set, and the type of circuit used. From the point of view of the circuit, the next section will address it. In terms of chemistry, we wanted a code, and therefore an algorithm, that was agnostic from the quantum chemistry backend used.

For the proof–of–concept, we defined the whole pipeline using two different chemistry backend: *PySCF* [94–96], which was both constrained due to the *QuAQ* code base dependency on it but also for its accessibility, and the recently introduced *VeloxChem* [97–99], which is started to be considered on of the main competitor of the former, chosen for its HPC and GPU oriented nature. In terms of implementation, both PySCF and VeloxChem are Python packages, what changes is the underlying library of routines. The former is written almost entirely in Python (87%), while the remaining code is C++ (12%), and an smaller Fortran (0.2%) component. Contrary, the latter is entrirely written in C++ (97.8%) which as been wrapped into a Python exposed layer (2.2%), making it way more faster and scalable compared to the former. Additionally, VeloxChem can be seen as an efficient SCF and DFT provider which is then integrated with other HPC oriented library, again Python wrapped C++ packages, specialized in Multi-Configurational and CI methods, *Multi–Psi* [262], and Møller-Plesset (MPn) and propagator based method[2], *Gator* [264].

In term of implementation point of view, the main difference between the two chemistry backend is in the extraction of the one– and two–body integrals. In particular, for PySCF, using the CASCI constructor `mcscf.CASCI` and the two methods `get_h1eff` and `get_h2eff`, we can extract the effective one- and two-integrals, as well as the core energy. Then, passing by Qiskit-Nature `ElectronicEnergyProblem`, we use the method `from_raw_integrals` to obtain the `FermionicOperator` associated to the electronic Hamiltonian to be solved. For VeloxChem instead, we can only extrapolate the one– and two–body integrals of the whole molecular space, therfore, we employ QuAQ effective operator builder to obtain the effective Hamiltonian. Additionally, PySCF allows to build molecules both with *Cartesian coordinates*, defined by a list of tuple `(atom, x,y,z`, and *Internal coordinates* or `z-matrix`, defined in our case by a `template string` in which the coordinates are replaced by placeholder and a dictionary containing the value associated with each placeholder coordinate. VeloxChem instead only works with cartesian coordinates.

---

[2] Such as Algebraic Diagrammatic Costruction (ADC) [263] scheme and more involved variant.

7.4.2 Quantum Computing Simulation backend

For the circuit simulation, we implemented two different quantum computing backed: Qiskit statevector [236], and CudaQ [101], statevector and GPU simulator. The whole backbone is defined and integrated using some revised QuAQ functionalities, adapted for the geometry optimization task, like the VQE algorithm, the symmetry penalty operators ($\hat{N}_e$, $\hat{S}_z$, and $\hat{S}^2$), and the effective operator construction. The underlying management of the fermionic operators and qubit operators is defined using Qiskit-Nature [100] `FermionicOperator` and Qiskit `SparsePauliOp`, respectively. If the backend chosen for the simulation is Qiskit, then, no additional transformation is required. If the backed is CudaQ, then, all the qubit operators are converted, using `QuAQ.tools.converters`, from `SparsePauliOp` to CudaQ `SpinOperators`. Main difference between the two implementation is the circut object. In fact, Qiskit use a `QuantumCircuit` dataclass, which, once instantiated allows for further ciruit manipulation. CudaQ instead, uses a `cudaq.kernel` decorator applied to a python method that, give a set of parameters and the specific of the circuit, build it every an expectation value is called, from scratch. The second structure is more constrained, requiring also more carefulness in the setup of method, especially for parallelization. Additionally, in order to perform an expectation value, contrary to Qiskit in which a statevector can be converted into a complex array and contracted to the matrix representation of the qubit operator, CudaQ has a specif method, namely `cudaq.observe`, which requires an operator defined as `SpinOperator`, a circuit kernel, and the set of parameters and arguments required to build the associated circuit.

7.4.3 Intra-node parallelization

Force calculations via finite differences require energy evaluations at displaced geometries. Given a geometry with $3N_{\text{atoms}}$ cartesian degrees of freedom, the requirement of single evaluations is $6N_{\text{atoms}}$ additional VQE calculations when using finite differences. Given a geometry in internal coordinates define on $3N-6$, generally, degrees of freedom, the number of VQE drops to $6N-12$. We implement parallel evaluation strategies to distribute these calculations across available computational resources. The parallel framework divides force component calculations among independent workers, each performing complete VQE optimizations for their assigned displaced geometries. In particular, a single task is defined by: initial geometry, `cords`, coordinate on which apply the displacement, `i`, value and sign of the dispacement, `d` and `sign`, respectively, the method for building the displaced geometry, `builder`, the evaluation method, `evaluation`, the initial variational parameters, `psi_0`, and the optimization strategy, `optimization_procedure`. We report here the structure of a gradient worker

```python
def single_diff_x_i_reoptimized(
args: tuple[int, float],
builder: callable,
evaluation: callable,
optimization_procedure: callable,
x: list[float] | np.ndarray,
psi_0: np.ndarray,
*,
is_sparse: bool | None = False):

i, sign, d = args
```

```
x_d = x.copy()
x_d[i] += sign*d
H_d = builder(x_d, partial=True)
m = H_d.simplify()
if is_sparse:
    m = m.to_matrix(sparse = True)
_, e_d, _ = optimization_procedure(eval_method=evaluation,
                                   op = m,
                                   initial_parameters = psi_0)
return ((i, sign), e_d / (2*d))
```

Snippet 7.1: Single perturbed geometry energy evaluator

Here, every argument except the task specific info, position and displacement, are fixed using Python `partial` wrapping method.

At this point, for all the coordinates, both cartesian and internal, and for all the displaced direction, positive and negative, the associated task is defined as a tuple `(i, sign, d)` and evaluated in parallel using `joblib` as multi-processing manager. Then, all the evaluation are collected, and the gradient vector is defined.

Additional parallelization is implemented during the `non-adiabatic` new minima search, which is applied at the beginning and then after evert T step. In this case, the worker has a more compact and simpler structure compared to the one for the single perturbed energy evaluation. In particular, each worker receives a fresh new set of random parameters, `initial_parameters`, the operator to minimize, `op`, and the evaluation method, `eval_method`. Internally then, it calls a `scipy.optimize.minimize` instance as follows

```
def simplified_VQE( eval_method: callable,
                op: np.ndarray | csr_matrix,
                initial_parameters: list[float] | np.ndarray = None) -> tuple[
np.ndarray, float]:
    exp_method = partial(eval_method,op=op)
    if len(initial_parameters) == 0:
        return [], exp_method([]), 0
    res = minimize(fun = exp_method,
             x0=initial_parameters,
             method='L-BFGS-B',
             bounds=[(-np.pi, np.pi)] * len(initial_parameters),
             options={'disp': True, 'maxiter': 1000, 'maxcor':50})
    params = res.x

    return params, res.fun, res.nit
```

Snippet 7.2: Parallel non-adiabatic VQE worker

Once again, given the number of workers, each of them receive a task that is only composed by the set of initial parameters, while the other arguments are fixed using Python `partial`.

7.5 Structure Optimization for DompeKey Molecular Descriptors

The ability to translate molecular structures into machine–readable descriptors is the one of the main structural pillar of cheminformatics and computer–aided drug design. A descriptor can be seen as a bridge between chemical space and computational analysis, with the objective to enable tasks such as similarity searching, database indexing, machine learning model training, and the exploration of structure—activity relationships. Over the years, a wide variety of molecular representations have been developed, ranging from simple one-dimensional strings (e.g., SMILEs [265], InChI [266]) to graph–based encodings of connectivity and complex three–dimensional coordinate sets. Among them, the most used are molecular fingerprints [267], which encode the presence or absence of chemical substructures into bit vectors. While powerful and computationally efficient, classical fingerprints suffer from certain drawbacks such as the possibility to be time–consuming to generate for very large databases, that often rely on hashing procedures that may lead to collisions, and might not be always transparent or easily interpretable.

Dompé researchers, in order to address the previously listed limitation, have introduced *DompeKey (DK)* [93], a novel family of substructure–based descriptors designed to combine efficiency, interpretability, and broad applicability. DKs are part of EXSCALATE, Dompé's end–to–end drug discovery platform, and were designed to provide a flexible, scalable and chemically meaningful representation of molecules. DKs are built from a set of 1064 SMARTS [268] strings, each encoding a chemical feature or fragment of pharmaceutical relevance. These range from simple functional groups and pharmacophoric points to complex scaffolds such as amino acids, heterocycles, or toxicophores.

A distinctive feature of DKs is their hierarchical architecture. Substructures are organized across five levels of complexity:

- **Level 0** includes well–defined, highly complex structure such as amino acids, natural product fragments, or pharmacologically relevant scaffolds.
- **Level 1** captures ring systems, such as pyridines or imidazoles.
- **Levels 2 / 3** encode functional groups, with Level 2 offering more specific descriptions that account for substituents and chemical environments (e.g., distinguishing between aromatic vs. aliphatic amines).
- **Level 4**: represents the simplest units, such as individual atoms with defined roles (sp^2 carbons, hydrogen bond donors or acceptors, etc.), which can be interpreted as pharmacophoric points.

This layered design ensures that DKs can describe molecules at multiple levels of resolution, from general structures to precise functional details. Importantly, it also means that chemists and modelers can select subsets of descriptors depending on the task, balancing granularity with computational cost.

The applications of DKs are diverse and demonstrate their versatility. In chemical space mapping, DK descriptors have shown strong discriminative power, clustering molecules by structural classes more effectively than many traditional fingerprints. This makes them particularly useful for organizing large compound collections and visualizing diversity. In machine learning, DK have been successfully employed as inputs to predictive models of biological activity and metabolism, yielding performance comparable to or better than standard descriptors. Again, thanks to their hierarchical structure, they allows to trace back predictions to meaningful chemical features.

DKs have proven useful for virtual screening and substructure searches. Because they are based on explicit SMARTS patterns rather than hashed identifiers, queries can be designed to directly capture the presence of functional groups relevant to a target class or toxicity alert. Also applied to the prediction of drug metabolism and toxicity, demonstrating their ability to capture reactivity–related features more effectively than classical atom–typing schemes. Additionally, DKs are designed for efficient precomputation and storage, i.e. they are defined such that they can index databases containing tens of trillions of compounds without loss of resolution.

7.5.1 DompeKeys selected

Functional groups such as imines, amides, and aldehydes play a central role in both structural chemistry and medicinal design. Their presence strongly influences molecular conformation, hydrogen-bonding capacity, and electronic distribution, which in turn modulate pharmacophoric behavior and chemical reactivity. Within the DK fragment dataset, these structures occur frequently and represent chemically diverse yet computationally accessible systems.

Imines

Imines are characterized by a C=N bond, formally derived from the condensation of carbonyl compounds with primary amines. They serve as key intermediates in many biochemical transformations, including enzymatic transamination[3]and Schiff-base formation[4], and appear in drug scaffolds where reversible C=N bonding can modulate conformational flexibility or metal-binding properties [269–271]. Their moderate size make them ideal prototypes for evaluating how quantum methods handle delocalized π-systems during geometry optimization. As first test molecule, we selected the *Methylene Imine* [272] defined as $H_2C{=}NH$ [273] with initial geometry optimized at mean–field level with Restricted-HF and STO-3g basis set, fetched from the cccbdb database. The initial geometry is shown in Figure 7.3, the cooridinates in Appendix D, and the bonds/angle in Table 7.2.

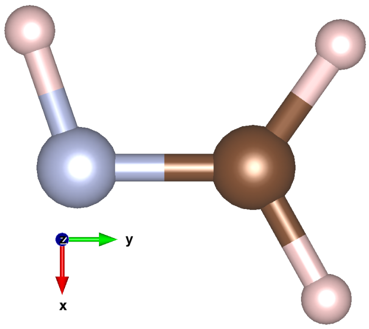

Figure 7.3: Methylene Imine initial geometry optimized with Restricted-HF and STO-3g basis set.

Aldehydes

Aldehydes feature a polar carbonyl group with an electrophilic carbon center and a strongly polarized C-H bond, which makes them central building blocks in oxidation, condensation, and polymerization reactions, relevant also in the context of drug metabolism and covalent inhibitor design [274].

---

[3] Transamination is a chemical reaction that transfers an amino group.

[4] Organic compound with structure $R^1R^2C{=}NR^3$, where $R^{1/2/3}$ are organic groups.

In the DK set, aldehydes provide a baseline to assess whether our HQC geometry optimization procedure can reproduce consistent trends for polar molecules. As test molecules, for the aldehydes, we selected *Acetaldehyde*, defined as $CH_3CH{=}O$ with initial geometry optimized at mean-field level with Restricted-HF and STO-3g basis set, fetched from the cccbdb database. The initial geometry of acetaldehyde is shown in Figure 7.4, the cooridinates in Appendix D, and the bonds/angle in Table 7.3.

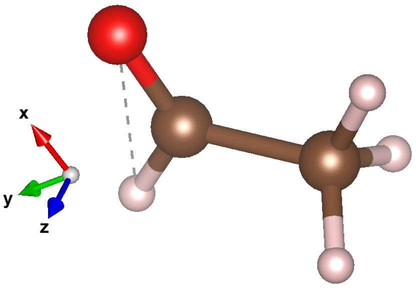

Figure 7.4: Acetaldehyde initial geometry optimized with Restricted-HF and STO-3g basis set.

Amides

Amides are among the most ubiquitous functional groups in organic and biological chemistry, essential in the construction of peptides and proteins. Their stability arises from strong resonance between the carbonyl and the adjacent nitrogen lone pair, leading to partial double-bond character and restricted rotation about the C–N bond [275–278]. These features test quantum algorithms against shallow potential energy surfaces and small energy gradients, typical in biological systems. For this functional group, we have selected the *Acetamide* which is an organic compound with the formula CH_3CONH_2, which is derived from ammonia and acetic acid, commonly known for its usage as industrial solvent [279]. The initial geometry is optimized at mean-field level with Restricted-HF and STO-3g basis set, fetched from the cccbdb database. The initial geometry of acetamide is shown in Figure 7.5, the cooridinates in Appendix D, and the bonds/angles in Table 7.4.

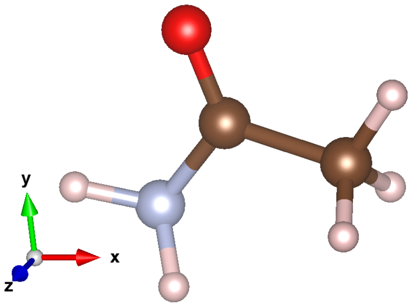

Figure 7.5: Acetamide initial geometry optimized with Restricted-HF and STO-3g basis set.

Together, these three families of compounds span a representative range of bonding situations, from delocalized π-conjugation in imines and amides to polar σ–π separation in aldehydes, providing an ideal chemical landscape for benchmarking the behavior of HQC structure optimization algorithms across distinct electronic regimes.

7.5.2 Results on small DK cases

In this section, we present the preliminary results obtained with our QA–q–GeomOpt procedure against some of the smallest DKs of interest, as we already introduced in the previou section, i.e. methylene imine, acetaldehyde, and acetamide. Given the fact that the structure refinement depends strongly on the possibility to describe as precise as possible the largest portion of the molecular system, we define for all the DKs a moderate sized active space consisting of 6 electrons in 6 orbitals, i.e. CAS(6,6), and the active orbitals have been selected around the Fermi level. For all the DKs the basis set chosen is cc-pVDZ and the integrals on which the electronic Hamiltonian are built, according to Equation 3.12, are defined from Restricted-HF orbitals.

In terms of circuital model selected, in spirit of the whole thesis, we selected HEA for all the molecules. In particular, CNOT-based ladder-fashion ansatz with different depth according to the DK studied. As defined in Section 3.3.1, we used a depth d ladder ansatz as reported in Table 7.1. The initial set of parameters has been selected starting from the usual randomized initialization in the range $(-\pi, \pi)$ and HF initialization with X gates in correspondence of occupied orbitals in the HF determinant.

DK	**Depth**	*CNOTs*	**Params**
Methylene Imine	4	44	60
Acetaldehyde	4	44	60
Acetamide	4	44	60

Table 7.1: Circuital resources required by our structure optimization procedure for the resolution of the three DompeKeys. As reported in the paragraph, the active space is a CAS(6,6) selected around Fermi-level corresponing to 12 qubits.

All the simulations have been run using 50 VQE as initial non-adiabatic search, 50 VQE trials, with a frequency step of 3 macro-iterations. As optimizer, we empliyed L-BFGS-B with `maxcor=50` and bounded parametrization in the range define above, i.e. $(-\pi, \pi)$.

For all the DKs, the initial geometry has been fetched from cccbdb database, and in particular, we started from a geometry optimized in minimal basis set, STO-3g, at mean-field level with Restricted-HF. In order to assess the accuracy of our method, we compared the bond–length and angles between the initial geometry, with the one obtained with our QA–q–GeomOpt procedure, and the exact solution obtained by optimizing the structure using CASCI with the same setup. Additionally, to validate the final structure, we also compared to the optimal distances and angles obtained by CASSCF, CCSD, and DFT with B3LYP Exchange–Correlation functional. These additional calculations provide a consistent comparison across different levels of electronic-structure theory: CASSCF representing a fully variational multi–configurational method with orbital relaxation, CCSD a non-variational correlated wavefunction approach that closely approximates the CCSD(T) gold standard, and B3LYP a widely adopted hybrid density-functional method whose accuracy for molecular geometries is well established, particularly when used with correlation-consistent basis sets such as cc-pVDZ.

In order to assess the agreement between the geometries obtained from the VQE optimization and the CASCI reference, the numerical precision of bond lenghts and bond angles is evaluated within commonly accepted chemical accuracy limits. For reliable electronic-structure methods with medium to large basis sets, deviations smaller than approximately 0.02 Å in bond lengths and 1°

in bond angles are considered chemically insignificant, while differences below 0.05 Å and 3° still indicate very good structural agreement. Therefore, all the comparisons discussed in this work are interpreted within these established precision criteria.

Methylene Imine

Table 7.2 reports the optimized structural parameters of methylene imine as obtained from different electronic-structure methods. The initial geometry was generated at the restricted–HF level with the minimal STO-3G basis, while subsequent optimizations were carried out using the cc-pVDZ basis set to ensure a consistent level of correlation treatment. In particular, the CASCI, QA–q–GeomOpt (QGO), and CASSCF calculations employed a CAS(6,6) active space constructed from orbitals selected around the Fermi–level, as reported in the previous paragraph. The QA–q–GeomOpt procedure reproduces the CASCI equilibrium structure with excellent accuracy, bond–length deviations remain below 0.02 Å, and bond-angle differences are within 1°, which are both well inside the thresholds commonly regarded as chemically indistinguishable. The deviations between the QGO and CASCI geometries are minimal, with bond-length differences ranging from 0.003 to 0.017 Å and bond-angle variations below 1°, confirming the agreement between the two metods. The small residual variations observed between QGO and CASCI geometries can be attributed to finite-difference numerical noise and the limited stochastic precision intrinsic to variational quantum algorithms. CASSCF results, which include additional orbital relaxation, show only slightly larger discrepancies, while the CCSD and B3LYP optimizations yield marginally elongated C–H and N–H bonds, a well-known tendency of these correlated and DFT approaches, respectively.

Measure	HF	CASCI	QGO		CASSCF	CCSD	B3LYP
C1-H1	1.09100	1.09099	1.08998		1.08997	1.10584	1.10687
C1-H2	1.08902	1.08895	1.08747		1.08688	1.10154	1.10232
H3-N1	1.04818	1.01299	1.01307		1.03730	1.03159	1.03143
N1-C1	1.27270	1.27362	1.27411		1.29291	1.28181	1.27103
H1-C1-H2	115.483	116.8113	116.1245		116.9035	116.1042	115.6657
H1-C1-N1	125.4003	124.2247	124.3875		124.6718	125.2326	125.7248
H2-C1-N1	119.1164	118.9640	119.4880		118.4248	118.6631	118.6094
H3-N1-C1	109.1299	110.0790	110.2613		108.2197	108.7566	109.9548

Table 7.2: Bond lenght and distances between methylene imine starting geometry obtained with **HF**, our QA–q–GeomOpt procedure (**QGO**), the reference **CASCI** results, with also **CASSCF**, **CCSD**, and DFT with **B3LYP** XC functional. Bond lenght are given in Ångstrom while angles in degress.

Relative to the HF geometry obtained with the minimal STO-3g basis, the correlated QA–q–GeomOpt and CASCI structures exhibit slightly longer bond lengths (by up to 0.03 Å) and small angular adjustments within 2–3°, reflecting the inclusion of electron-correlation and basis-set effects and although these differences are modest in magnitude, they represent a quantitative improvement in the structura description.

Acetaldehyde

Table 7.3 summarizes the optimized bond lengths and bond angles for acetaldehyde obtained from different electronic-structure methods. As explained, in the previous paragraph, the reference structure was computed at the CASCI level using a CAS(6,6) active space constructed from orbitals around the Fermi-level and the cc-pVDZ basis set. The QA–q–GeomOpt optimization reproduces the CASCI equilibrium geometry with high quantitative accuracy, i.e. bond–length differences lie within 0.002–0.009 Å, and angular deviations remain below 1°, indicating an almost perfect correspondence between the two approaches. Compared to the initial restricted-HF/STO-3G geometry, both CASCI and QA–q–GeomOpt yield slightly shorter C–O and longer C–C and C–H bonds, consistent with the inclusion of electron-correlation effects and the improved flexibility of the cc-pVDZ basis. CASSCF geometries, which include additional orbital relaxation, remain close to CASCI but display marginally larger bond distances, wheres CCSD and B3LYP produce the expected systematic bond elongation typical of correlated and density-functional treatments, respectively. Overall, the QA–q–GeomOpt geometry of acetaldehyde closely mirrors the multi–configurational CASCI reference, demonstrating the reliability of the variational quantum optimization in capturing both local and global structural features of a correlated molecular systems.

Measure	**HF**	**CASCI**	**QGO**		**CASSCF**	**CCSD**	**B3LYP**
H1-C1	1.10421	1.10371	1.10146		1.10720	1.12094	1.12394
C1-O1	1.21714	1.20786	1.20691		1.21705	1.21182	1.20836
C2-C1	1.53649	1.50463	1.50988		1.50600	1.51321	1.50672
C2-H2	1.08499	1.09042	1.08952		1.08893	1.10056	1.09889
C2-H3	1.08692	1.08068	1.08090		1.09321	1.10467	1.10426
C2-H4	1.08692	1.08068	1.08091		1.09321	1.10467	1.10426
H1-C1-O1	121.4454	119.3901	119.1317		120.5802	120.6787	120.6682
O1-C1-C2	124.2802	124.6258	124.8443		124.6598	124.5321	124.8185
H1-C1-C2	114.2744	115.9841	116.0240		114.7600	114.7892	114.5132
C1-C2-H2	110.4813	109.9179	110.2597		110.4073	110.1830	110.6801
C1-C2-H3	109.9005	108.4688	109.0364		109.6603	109.6107	109.6199
C1-C2-H4	109.9005	108.4688	109.0367		109.6603	109.6107	109.6199
H2-C2-H4	109.2066	109.6758	109.3818		109.8665	110.0639	110.1517
H2-C2-H3	109.2066	109.6758	109.3691		109.8665	110.0639	110.1517
H3-C2-H4	108.1017	110.6100	109.7439		107.3255	107.2590	106.5287

Table 7.3: Bond lenght and distances between acetaldehyde starting geometry obtained with **HF**, our QA–q–GeomOpt procedure (**QGO**), the reference **CASCI** results, with also **CASSCF**, **CCSD**, and DFT with **B3LYP** XC functional. Bond lenght are given in Ångstrom while angles in degress.

Acetamide

Table 7.4 presents the optimized structural parameters of acetamide obtained with the QA–q–GeomOpt procedure in comparison with reference CASCI results and several conventional electronic-structure methods. All correlated calculations were performed with the cc-pVDZ basis set, and the multi–configurational methods employed a CAS(6,6) active space constructed from orbitals near the Fermi level. The QA–q–GeomOpt optimized geometry shows very close agreement with the CASCI reference, with bond–length deviations ranging between 0.002 and 0.015 Å and angular differences within approximately 1–2°, indicating that the variational quantum optimization is able to reproduce the correlated equilibrium structure with quantitative fidelity. Relative to the Restricted–HF/STO-3g starting geometry, both CASCI and QA–q–GeomOpt predict a slight shortening of the C=O and N−H bonds and a moderate elongation of the C−C linkage, reflecting the inclusion of electron–correlation and polarization effects. CASSCF results, which include full orbital relaxation, remain consistent with CASCI while showing small systematic shifts in bond distances. The CCSD and DFT/B3LYP optimizations exhibit the expected trend toward marginally longer covalent bonds, a characteristic behavior of these correlated and density-functional approaches.

Measure	**HF**	**CASCI**	**QGO**		**CASSCF**	**CCSD**	**B3LYP**
H1-C1	1.08474	1.08591	1.08902		1.08666	1.09852	1.09676
H2-C1	1.08663	1.08487	1.08707		1.09290	1.10448	1.10329
H3-C1	1.08614	1.08420	1.08739		1.09158	1.10318	1.10182
C1-C2	1.54143	1.51552	1.53160		1.51423	1.52285	1.52218
C2-O1	1.21765	1.20648	1.19802		1.22014	1.21688	1.21936
C2-N1	1.44476	1.37806	1.39166		1.37771	1.38502	1.37398
N1-H4	1.02750	0.99997	1.00357		1.00072	1.01559	1.01416
H5-N1	1.02686	0.99918	0.99937		0.99902	1.01365	1.01200
H1-C1-H2	109.0395	109.3409	108.6537		109.2067	109.2583	108.7627
H1-C1-H3	109.3694	109.4019	109.1225		109.7395	109.8135	109.7416
H2-C1-H3	108.4740	108.9098	108.1855		107.9441	107.9688	107.8558
H1-C1-C2	109.7055	108.8634	110.7309		109.2747	109.0451	108.7953
H2-C1-C2	110.3794	109.9810	110.2019		110.3046	110.3490	110.1023
H3-C1-C2	109.8469	110.3274	109.8916		110.3493	110.3890	111.5415
C1-C2-O1	124.4914	122.4708	122.7274		123.2819	123.4083	123.1885
C1-C2-N1	113.5655	115.6444	114.7677		114.5210	114.1977	114.6845
O1-C2-N1	121.8485	121.8406	122.4252		122.1628	122.3386	122.1204
C2-N1-H4	110.8919	114.3867	113.0765		115.1177	114.4135	116.9742
C2-N1-H5	112.0851	118.2653	114.3868		118.4189	117.9016	121.2092
H4-N1-H5	109.2865	114.8427	110.8682		115.1578	115.0918	117.6910

Table 7.4: Bond lenght and distances between acetamide starting geometry obtained with **HF**, our QA–q–GeomOpt procedure (**QGO**), the reference **CASCI** results, with also **CASSCF**, **CCSD**, and DFT with **B3LYP** XC functional. Bond lenght are given in Ångstrom while angles in degress.

7.6 Discussion and final remarks on hybrid structure–optimization

In this chapter, we introduced how geometry optimization is performed following from the classical point of view by giving an insight on how molecular gradient are evaluated for both variational and fully–variational approached. After that, we gave an brief overview on what has been done in quantum ecosystem directly, and also indirectly, for the structure optimization procedure. In particular, we dedicated one section on the main step of the task i.e. the evaluation of the energy derivative w.r.t. nuclear displacement on the quantum devices, then, we described the state–of–the–art procedure for both navigating the PES and how to evaluate accurate single point calculations, including corrections beyond simple *a-la* CASCI solutions. We then started to introduce our proof–of–concept *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)* procedure, which mixes HEA with adiabatic and non–adiabatic search strategies to optimize structure towards chemically accurate ones.

The benchmarking results demonstrate that the proposed QA–q–GeomOpt approach reproduces reference CASCI geometries with remarkable precision across different molecular systems. Bond–length deviations consistently fall within 0.002–0.02 Å, and bond–angle differences are below 1–2°, which are well within the limits of chemical accuracy and essentially indistinguishable from the CASCI reference. When compared to the initial Restricted–HF/STO-3g geometries, our approach and CASCI both capture the expected correlation–driven effects, slightly longer covalent bonds and subtle angular relaxations, while maintaining full internal consistency with higher–level multi–configurational and coupled–cluster results. CASSCF optimizations, which include orbital relaxation, yield nearly identical equilibrium structures, whereas CCSD and DFT/B3LYP predict the typical marginal elongation of bond distances characteristic of these correlated and density-functional approaches. Overall, our QA–q–GeomOpt method achieves quantitative agreement with CASCI and performs similarly to post-Hartree–Fock techniques, confirming its capability to deliver chemically accurate molecular geometries within a Hybrid Quantum–Classical framework.

Obviously, a lot of work can be done in three main part of our procedure: first, improving the single point evaluation, shifting from simple HEA to more robust circuits, such as UCCSD ansatz or ADAPT-like [81] ansätze, as well as introducing post-processing procedures such as sub-space diagonalization-based method [201, 280]. The second improvement can be done on the theoretical side, meaning the replacement of finite-difference approach to analytical molecular gradients. Following this path, we can either decide to defined all the derivatives as fermionic operator, such as in the WAHTOR procedure, or, as done in ADAPT-SCF-VQE / OO–VQE, building the one- and two-body RDMs and using it to evaluate the gradients/Hessian entries, as illustrated by T. Helgaker [26]. Finally, from the computational point of view, additional integration and testing on actual quantum computing backend, as well as , from the simulation side, to improve the parallelization, moving from only intra-node parallelization, to multi-node and multi-GPU integration.

The following chapter concludes the thesis by reflecting on the hybrid workflow developed so far and outlining directions for future quantum-assisted molecular modelling.

Part III

Quantum @ L'Aquila

Chapter 8

QuAQ: Quantum @ L'Aquila

All the algorithms presented in this thesis have been implemented withing a dedicated quantum computing package, namely *Quantum @ L'Aquila (QuAQ)* [224], developed with in mind the objective of integrating a full quantum chemistry pipeline directly interfaced with the quantum computing environment. With the exception of Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS) [281], the active space selection algorithm, also designed for "*classical*" quantum chemistry, and in a private package, all procedure are natively implemented in our codebase.

QuAQ has been developed from scratch, relying on *PySCF* [94–96][1] only as a backend for the quantum chemistry tasks such as mean-field calculations, molecular orbital initialization, and integrals generation. The representation and manipulation of quantum circuits are handled through *Qiskit* [236][2], which represents also the primary quantum computing interface. To ensure full flexibility and control, the framework also redefines some of the Qiskit datatypes, such as the fermionic operators and qubit operators, mainly oriented for allowing SparQ method to work optimally. We also included the possibility to automatically build Hamiltonian operators for both molecular systems and spin–lattices.

A key design principle of QuAQ is to be modular. Each algorithm can be executed either through standard *input* files, following what is a common convention of traditional quantum chemistry packages (like ORCA [178], Molcas [282], OpenMolcas [139, 140], QChem [283], etc), or directly inside a Python script for everywhone who require and want full control over the parameters and pipelining.

Among the various procedures, the Active Space Selection algorithm, *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS)* [281], is not yet publicly available. Conversely, the geometry optimization package, QA–q–GeomOpt, depends on our public codebase QuAQ, but the core algorithm is proprietary and currently under NDA for patenting purposes. QuAQ also includes standard VQE [53] and ADAPT-VQE [81] implementations. This optimization schemes, as we have presented in this thesis, are fully interfaced with our self-consistent field procedure, *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation Self–Consistent Field (WAHTOR–SCF)*, an extension of the previously theorized *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR)* algorithm [210, 212], also included in QuAQ.

---

[1] Any version above `2.0.0`.

[2] Version `1.4.3` and version `0.7.6` for *QiskitNature* [100].

Parallelization strategies are also well incorporated where beneficial. Intra-node parallelization relies on *joblib*, while *MPI* supports distributed multi-node workloads—features currently employed in WAHTOR–SCF and QA–q–GeomOpt. Furthermore, QA–q–GeomOpt provides optional quantum backends through *OpenFermion*, and supports GPU-accelerated simulations via *Cuda-Quantum* [101]. The integration of these features into the core QuAQ package is planned for future releases.

This chapter will recap briefly the structure of the QuAQ code base.

8.1 QuAQ stucture

QuAQ is split in three main folders: `algorithms`, `tools`, and `utils`.

8.1.1 Utils

Contains all the methods used to manipulate input and output, and conversion scripts. At top level, we have

- `input`: Main Python class used to read in the input file and collect all the information, parameters, and settings for the desired algorithms. The `input` class reads in a `fortran namelist` file, and collects, according to the different sections of the input, the arguments associated to each procedure implemented. The arguments are first passed through a round of consistency check. It is necessary to stress that, in any case, a Python script defining the overall pipeline that a user wants to follow is still required. The only difference between using an *input* file and the manual setup is the automatic generation of the arguments required by each method.

- `trust_region`: Implementation of the Trust-Region optimizer algorithm taking inspiration from `scipy` `trustregion`[3] and `trustregion_ncg` (Newton-Conjugate Gradient)[4].

Then, specific set of conversion functionalities are contained in the folder `converters`. Here, two different type of conversion scripts are collected: one dedicated to convert qubit operators and one for fermionic operators.

- `fermionic_operator_converters`: Contains the conversion for fermionic operators between QuAQ, Qiskit, and OpenFermion's `FermionOperator`. Given the different default ordering between Qiskit/QuAQ and OpenFermion, additional mode reordering functionalities are included.

- `qubit_operator_converters`: Contains the conversion for qubit operators between QuAQ 's `QubitOperator`, Qiskit's `SparsePauliOp`, and OpenFermion's `QubitOperator`.

On top of this, the file required to compile SparQ [163] method are contained in `efficient_toolbox`.

---

[3] Scipy `trustregion.py`

[4] Scipy `trustregion_ncg`

8.1.2 Tools

Contains all the system specific methods and classes, such as Hamiltonian generation, procedure to manipulate them, e.g. integrals trimming or unitary transformation application, and also quantum computing related dataclasses, such as qubit operators, fermionic operators, and ansatz. In particular, at top level we have defined:

- `system`: Main class of the whole code base contains the construction of the molecular or spin-lattice system. Focusing on the chemistry side, using PYSCF, it build the essential information for building the collection of arguments required to solve a single point calculations, or in general, provide information about the molecular system. If the resolution is a CAS calculation, then calls `orbital_reduction` to transform from full molecular space to the active space. It allows to restart from a previous mean-field guess, from a external PySCF mean-field calculation, and it is also the component that build the Iterative Natural Orbitals (INOs). The main field defined by the `system` are the one- and two-body integrals, already collected according to number of modes and spin–sector.

- `ansatz`: Wrapper class to Qiskit `QuantumCircuit` object. It contains already automated circuit construction routines in which rotation gates, entangling gates, and layer–specific correlator placement can be specified. As explained in Section 5.2.4, we also have routines for extending existing quantum circuit with arbitrary entangler maps placement or with the V-shape structure. We include the possibility to create circuit with the standard two-qubit gates contained in Qiskit and also custom gates such as *SO(4)*, *Givens* rotation (one - and two-qubit), *SU(4)*, and *O(4)* gates [79, 205, 208, 284]. For the specific excitation gates instead, we have a dedicated module named `excitation_operators`.

- `fermionic_operator`: Given a dictionary containing the one- and two-body integrals defines the associated fermionic operator, defined as `List[Tuple(list[Tuple(int, int)], complex)]`, where the first tuple are the fermionic operators composing a fermionic string and the associated complex coefficient.

- `qubit_operator`: It takes as input a fermionic operator in QuAQ format and it applies any of the desired Fermion–to–Qubit mapping either using existing procedure, e.g. Jordan–Wigner [48], or customs one using bonsai-mapping construction method by Miller *et al.* [51]. Internally, the custom qubit operator can be converted to Qiskit `SparsePauliOp` and also using Qiskit F2Q mapping.

- `orbital_reduction`: Tools for adapting integrals from being defined on the full molecular to the effective integrals, as defined in the procedure described by Equation 3.10 in Section 3.1.2.

- `hamiltonian_unitary_transformation`: Collection of methods for applying a unitary tranformation defined by orbita–rotation generators on a specified subset of orbitals, as a whole or grouped according to custom setup. This is a fundamental pieces for the all the WAHTOR variant, i.e. "vanilla" and self-consistent.

- `psiopsi`: Classes used to evaluate expectation values of an empty parametrized quantum circuit and an operator, given a set of parameters. It also contains the possibility of adding penalty terms, like used in Section 5.4, for a *constrained VQE* optimization. Additionally, it allows for individual Pauli string evaluation given an operator, similar to the procedure explained in Section 6.3. As backend, we also allows for Qiskit `Estimator`, which takes in input a Qiskit `QuantumCircuit`, an operator as `SparsePauliOp`, and evaluate the expectation value, allowing for shot based measurement.

- `excitation_operators`: Include all the existing implementation of single- and double-excitations used for either UCCSD or ADAPT-VQE ansätze. In particular, it contains: standard full fermionic excitations [81], Efficient fermionic excitations [73], single Pauli string excitation for Qubit-ADAPT [72], Qubit-Excitations for QEB-ADAPT [85], and single Pauli strings with Z-Chains [88, 222].
- `restrain_operators`: Two specific set of penalty operator definition for both Qiskit `SparsePauliOp` and OpenFermion `QubitOperator`. The two modules can defined spin projection operator, $\hat{S}_z$, spin-squared, $\hat{S}^2$, number of particles, $\hat{N}_e$, and energy shift, $(\hat{H} - w)^2$ for performing Folded-Spectrum VQE [285]. All of them can be used for targeting the first excited root for different multiplicity.

Then, we have a specific set of tools defined for the more complex algorithms, i.e. WAHTOR/WAHTOR–SCF and ADAPT-VQE. For the former, we defined

- `solvers`: Collections of solvers to be used in place of VQE such as QIDA, defiend in Section 5.1, and adaptive ansatz, as defined in Section 6.3.8.
- `parallel_derivatives`: Contains the `worker` used in the MPI/intra-node parallelized version of WAHTOR/WAHTOR–SCF. In particular, `fermionc_worker` and `pauli_worker` as defined in Appendix by 1.1 and 1.2, respectively.
- `integrals_trimming`: Simplified version of the effective integrals reductionm, in Equation 3.12 in Section 3.1.2, or by the active space projector operator $\mathcal{P}_{\text{act}}$, defined in Section 6.3, and used by the parallel workers, as defined by 1.2 and 1.1.

Finally, inside `adapt_tools` we collected all the methods associated to ADAPT-VQE. In particular, the modules are

- `pool_creation`: This modules include all the method for defining all the different operator pools, $\mathcal{P} = \{\hat{P}_i\}$, such as standard fermionic [81], Qubit-ADAPT [72], Qubit-Excitation based ADAPT [85], and Qubit with Z-chains [222]. We allows to have *generalized* excitations, where the number of particles and spin is not conserved, *restricted*, where only the number of particles is conserved, and *spin restricted*, where spin-flips are not allowed.
- `parallelization`: It contains all the used for the parallelized version of our ADAPT-VQE procedure. In particular, it contains method for evaluating explicitly the gradients, i.e. using the commutators of $\hat{H}$ and each opeator $\hat{P}_i \in \mathcal{P}$ as in Grimsley *et al.* [81] and using Roto–solve and Excitation–solve methods[286, 287].
- `generic`: Simple collection of methods used for seemingly connect all the different pools and excitation operators, like fermionic operator to Pauli string converters and TETRIS-ADAPT [82] utils.

8.1.3 Algorithms

This folder contains all the implemented algorithms and procedures. Each algorithm is defined as a Python `class` and it can take as argument an `input` object, read using `utils.input` class, or specific initialize arguments, which are going to have the priority over the first ones.

We start by listing the building algorithm used also by more convoluted procedure.

- `vqe`: Simple class containing the standard VQE procedure. In takes as input an `psiopsi` object defined on a specific operator of which, the groundstate is variationally searched, and the desired parametrized quantum circuit. It also allows for perform a constrained VQE optimization by providing a collection of *penalty operators*, $\{\hat{B}_i\}$, with corresponding multiplier λ_i and target value, b_i. The minimization is then performed on the combined operator $\hat{O}$ defined as
$$\hat{O} = \hat{H} + \sum_i \lambda_i (\hat{B}_i - b_i)^2. \tag{8.1}$$
- `layer_builder`: Method containing the procedure to extrapolate the QIDA layers from QMI matrices. It returns a collection of entangler maps, associated with each layers, and additional information such as rotations block, repetitions, and if required, a tag. It contains the reduction method defined in Section 5.2.2.

Looking at the more complex algorithms, inside this folder, we can find all the algorithm explained in the manuscript. In particular, `multi_layer_vqe`, together with the collection layer defined by `layer_builder`, it defines the Multi-QIDA procedure by performing the layer-wise VQE as defined in Section 5.2.5. Then, we have`adapt_vqe`, as explained both in Section 3.3.2 and Section 6.3.8, `wahtor` and `wahtor_MPI`, for both the standard, fully–parallelized, self–consistend, and adaptive version of WAHTOR procedure.

Conclusion

The work presented in this thesis explored the integration of quantum computing methodologies within the framework of electronic structure theory, with the ultimate goal of constructing an effective *Hybrid Quantum–Classical (HQC)* pipelne for chemical and industrial applications. The motivation comes from the growing recognition that, while fully fault-tolerant quantum computing remains beyond the reach of state-of-the-art technologies, the *Noisy Intermediate–Scale Quantum (NISQ)*-era already allows for a productive interplay between quantum and classical paradigms. Within this context, quantum processors and related algorithms are not viewed as replacements for existing computational chemistry methods, but rather as complementary tools capable of enhancing specific stages of the overall workflow. This vision guided the design and development of a suite of hybrid algorithms and protocols addressing four major components of the quantum chemistry pipeline: *active space selection*, *state preparation*, *correlation recovery*, and *geometry optimization*.

The first development, the *Atomic–Orbital and Entropy Guided Inference protocol for Space–Selection (AEGISS)*, focused on the initial and conceptually delicate stage of defining the active space in multi–reference quantum chemical calculations. This procedure introduced a semi-automated and reproducible approach that combines atomic orbital analysis with quantum information descriptors, most notably, single-orbital entropyes and mutual information measures, to guide the inclusion of orbitals in the active space. The method was applied to case studies relevant to *Photodynamic Therapy (PDT)*, offering insights into photosensitizer molecules whose accurate description requires a balanced treatment of static correlation. AEGISS thus represents a first step toward standardizing active space selection within hybrid workflows, alleviating one of the long-standing subjectivities of multi–configurational electronic structure theory.

The second contribution, the *Multi–threshold Quantum Information Driven Ansatz (Multi-QIDA)*, addressed the problem of *state preparation* in hybrid quantum algorithms such as the *Variational Quantum Eigensolver (VQE)*. By leveraging quantum information measures, Multi-QIDA provides a systematic, data-driven strategy for constructing ansätze that better capture the entanglement structure of the system while maintaining shallow circuit depth. This approach was tested on both molecular systems and spin-lattice models, demonstrating improved convergence and reduced resource requirements relative to generic hardware-efficient circuits. Multi-QIDA therefore acts as a warm-start or state-refinement procedure that links the theoretical elegance of quantum information concepts with the practical limitations of NISQ devices.

The third methodological block concerned the treatment of *dynamical correlation* through the extension of the *Wavefunction–Adapted Hamiltonian Through Orbital–Rotation (WAHTOR)* algorithm into its self-consistent form, *WAHTOR–SCF*. This hybrid framework allows for orbital relaxation and feedback between classical and quantum components, effectively bridging *Complete Active Space Configuration-Interaction (CASCI)* and *Complete Active Space Self–Consistent Field (CASSCF)* formulations in a quantum–classical environment. An adaptive version, *ADAPT–WAHTOR–SCF*, was further developed to enable dynamic generator selection and resource optimization. The method was benchmarked on both small molecular systems and a model *free-radical polymerization reaction*, highlighting its applicability to *industrial chemistry* scenarios. These developments shows how hybrid self-consistency loops can be incorporated efficiently, allowing quantum subroutines to contribute to the recovery of both static and dynamic correlation energies within scalable workflows.

The final technical contribution, the *Quasi-Adiabatic Quantum Geometry Optimization (QA–q–GeomOpt)* protocol, extended the HQC paradigm to the *nuclear degrees of freedom*, providing a proof-of-concept strategy for geometry refinement and structural optimization. By embedding quantum energy evaluations within a classical gradient-based search, the approach demonstrates how hybrid methods can be used to navigate *Potential Energy Surface (PES)* effectively, coupling quantum accuracy with classical efficiency. Its successful application to standard molecules and *DompeKey (DK)* molecular descriptors suggests potential use cases in *drug discovery* and materials design, where geometry optimization remains a computationally demanding step.

Taken together, these contributions outline a *unified and modular vision* for quantum-aided computational chemistry. Rather than treating quantum algorithms as isolated demonstrations, this thesis emphasized their integration as *functional components within a broader classical infrastructure*, supported by the in-house framework *Quantum @ L'Aquila (QuAQ)*. The results collectively validate the conceptual viability of an *end-to-end hybrid quantum chemistry pipeline*, demonstrating that hybridization can occur at multiple layers of the workflow, from orbital space definition to correlation treatment and structural refinement, without compromising chemical interpretability or computational efficiency.

In conclusion, this thesis represents a personal attempt to bridge two traditionally distinct worlds, quantum chemistry and quantum computation, through a consistent methodological framework that values both theoretical rigor and pragmatic adaptability. The work demonstrates that meaningful progress in this field does not arise from isolated algorithmic breakthroughs, but from the careful design of interfaces between methods, languages, and paradigms. Although the current hybrid approaches are still limited by hardware constraints and system scalability, the foundations laid here suggest that the integration of quantum processors within standard chemical workflows is not only feasible, but also intellectually and practically transformative. *It is my hope that this research contributes, even is a small part, to the broader effort of shaping a future where quantum computing becomes a standard and accessible tool for exploring the complexity of matter at its most fundamental level.*

List of Publications

Published

- **F. Tarocco**, D. Materia, L. Ratini and L. Guidoni. *Compact Multi-Threshold Quantum Information Driven Ansatz For Strongly Interactive Lattice Spin Models*, Journal of Physics A: Mathematical and Theoretical, (10.1088/1751-8121/adc4a1) 2025.
- H. Zehr, A. Baiardi, F. Tacchino, A. Gandon, L.E. Fischer, Y. Xu, F.P. DiFilippo, L. Guidoni, P.A.B. Haase, W.N. Talarico, M. Stella, **F. Tarocco**, A. Nykänen, A. Fitzpatrick, A. Miller, L. Thiessen, S. Knecht, E.-M. Borrelli, S. Maniscalco, F. Pavošević, I. Tavernelli, E. Maytin, V. Krishna.*Quantum Computing for Photosensitizer Design in Photodynamic Therapy*, Annual Review of Biomedical Data Science, (10.1146/annurev-biodatasci-103123-095644) 2025.

Under peer-review

- **F. Tarocco**, D. Materia, L. Ratini and L. Guidoni.*Multi-QIDA method for VQE state preparation in molecular systems*, (10.48550/arXiv.2508.11270) 2025.
- **F. Tarocco**, P. A. B. Haase, F. Pavošević, V. Krishna, L. Guidoni, S. Knecht, and M. Stella.*AEGISS – Atomic orbital and Entropy-based Guided Inference for Space Selection – A novel semi-automated active space selection workflow for quantum chemistry and quantum computing applications*, (10.48550/arXiv.2508.10671) 2025.

Pre-print

- L.U. Masci, **F.Tarocco**, D. Bonanni, M. Barberi, S. Rizzo, A. R. Beccari, L. Guidoni. *Toward High-throughput Virtual Screening via Quasi-Adiabatic Quantum Geometry Optimization Workflow*, (10.26434/chemrxiv.15001570/v1) 2026.

In production

- **F.Tarocco**, D.Materia, D. Chakravarti, L. Ratini, C. Capecci, L. Guidoni. *Enriching WAHTOR algorithm with Self-Consistent Orbital optimization: WAHTOR-SCF*, 2025.
- **F.Tarocco**, D. Chakravarti, L. Guidoni. *Exploration of Model Polymerization reaction with WATHOR-SCF*, 2025.

Bibliography

[1] E.R. Johnston, N. Harrigan, and M. Gimeno-Segovia. *Programming Quantum Computers: Essential Algorithms and Code Samples*. O'Reilly Media, Incorporated, 2019. ISBN: 9781492039686. URL: https://books.google.fi/books?id=LZY1vgEACAAJ.

[2] Phillip Kaye, Raymond Laflamme, and Michele Mosca. *An Introduction to Quantum Computing*. USA: Oxford University Press, Inc., 2007. ISBN: 0198570007.

[3] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[4] Dorit Aharonov. *A Simple Proof that Toffoli and Hadamard are Quantum Universal*. 2003. arXiv: quant-ph/0301040 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/0301040.

[5] He-Liang Huang, Dachao Wu, Daojin Fan, and Xiaobo Zhu. "Superconducting quantum computing: a review". In: *Science China Information Sciences* 63.8 (2020). ISSN: 1869-1919. DOI: 10.1007/s11432-020-2881-9. URL: http://dx.doi.org/10.1007/s11432-020-2881-9.

[6] Francesco Bernardini, Abhijit Chakraborty, and Carlos Ordóñez. *Quantum computing with trapped ions: a beginner's guide*. 2023. arXiv: 2303.16358 [quant-ph]. URL: https://arxiv.org/abs/2303.16358.

[7] Shannon P. Harvey. *Quantum Dots/Spin Qubits*. 2022. DOI: 10.1093/acrefore/9780190871994.013.83. URL: http://dx.doi.org/10.1093/acrefore/9780190871994.013.83.

[8] Paul Steinacker et al. "Industry-compatible silicon spin-qubit unit cells exceeding 99% fidelity". In: *Nature* (2025). DOI: 10.1038/s41586-025-09531-9. URL: https://doi.org/10.1038/s41586-025-09531-9.

[9] M. Saffman, T. G. Walker, and K. Mølmer. "Quantum information with Rydberg atoms". In: *Rev. Mod. Phys.* 82 (3 2010), pp. 2313–2363. DOI: 10.1103/RevModPhys.82.2313. URL: https://link.aps.org/doi/10.1103/RevModPhys.82.2313.

[10] H.-J. Briegel, T. Calarco, D. Jaksch, J. I. Cirac, and P. Zoller. "Quantum computing with neutral atoms". In: *Journal of Modern Optics* 47.2–3 (2000), 415–451. ISSN: 1362-3044. DOI: 10.1080/09500340008244052. URL: http://dx.doi.org/10.1080/09500340008244052.

[11] Dolev Bluvstein et al. "Logical quantum processor based on reconfigurable atom arrays". In: *Nature* 626.7997 (2023), 58–65. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06927-3. URL: http://dx.doi.org/10.1038/s41586-023-06927-3.

[12] Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. "High-fidelity parallel entangling gates on a neutral-atom quantum computer". In: *Nature* 622.7982 (2023), 268–272. ISSN: 1476-4687. DOI: `10.1038/s41586-023-06481-y`. URL: `http://dx.doi.org/10.1038/s41586-023-06481-y`.

[13] J. Eli Bourassa, Rafael N. Alexander, Michael Vasmer, Ashlesha Patil, Ilan Tzitrin, Takaya Matsuura, Daiqin Su, Ben Q. Baragiola, Saikat Guha, Guillaume Dauphinais, Krishna K. Sabapathy, Nicolas C. Menicucci, and Ish Dhand. "Blueprint for a Scalable Photonic Fault-Tolerant Quantum Computer". In: *Quantum* 5 (Feb. 2021), p. 392. ISSN: 2521-327X. DOI: `10.22331/q-2021-02-04-392`. URL: `http://dx.doi.org/10.22331/q-2021-02-04-392`.

[14] Lars S. Madsen, Fabian Laudenbach, Mohsen Falamarzi. Askarani, Fabien Rortais, Trevor Vincent, Jacob F. F. Bulmer, Filippo M. Miatto, Leonhard Neuhaus, Lukas G. Helt, Matthew J. Collins, Adriana E. Lita, Thomas Gerrits, Sae Woo Nam, Varun D. Vaidya, Matteo Menotti, Ish Dhand, Zachary Vernon, Nicolás Quesada, and Jonathan Lavoie. "Quantum computational advantage with a programmable photonic processor". In: *Nature* 606.7912 (2022), pp. 75–81. DOI: `10.1038/s41586-022-04725-x`. URL: `https://doi.org/10.1038/s41586-022-04725-x`.

[15] Han-Sen Zhong et al. "Quantum computational advantage using photons". In: *Science* 370.6523 (2020), 1460–1463. ISSN: 1095-9203. DOI: `10.1126/science.abe8770`. URL: `http://dx.doi.org/10.1126/science.abe8770`.

[16] Ville Lahtinen and Jiannis Pachos. "A Short Introduction to Topological Quantum Computation". In: *SciPost Physics* 3.3 (2017). ISSN: 2542-4653. DOI: `10.21468/scipostphys.3.3.021`. URL: `http://dx.doi.org/10.21468/SciPostPhys.3.3.021`.

[17] Fabian Hassler. *Topological quantum computing.* 2024. arXiv: `2410.13547 [quant-ph]`. URL: `https://arxiv.org/abs/2410.13547`.

[18] Philipp Hauke, Helmut G Katzgraber, Wolfgang Lechner, Hidetoshi Nishimori, and William D Oliver. "Perspectives of quantum annealing: methods and implementations". In: *Reports on Progress in Physics* 83.5 (2020), p. 054401. DOI: `10.1088/1361-6633/ab85b8`. URL: `https://doi.org/10.1088/1361-6633/ab85b8`.

[19] Atanu Rajak, Sei Suzuki, Amit Dutta, and Bikas K. Chakrabarti. "Quantum annealing: an overview". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 381.2241 (2022). ISSN: 1471-2962. DOI: `10.1098/rsta.2021.0417`. URL: `http://dx.doi.org/10.1098/rsta.2021.0417`.

[20] Daniel Gottesman. *Surviving as a quantum computer in a classical world.* textbook manuscript prepeprint, 2024. eprint: `https://www.cs.umd.edu/class/spring2024/cmsc858G/QECCbook-2024-ch1-15.pdf`. URL: `https://www.cs.umd.edu/class/spring2024/cmsc858G/QECCbook-2024-ch1-15.pdf`.

[21] Daniel Gottesman. *An Introduction to Quantum Error Correction and Fault-Tolerant Quantum Computation.* 2009. arXiv: `0904.2557 [quant-ph]`. URL: `https://arxiv.org/abs/0904.2557`.

[22] Daniel Gottesman. “Theory of fault-tolerant quantum computation”. In: *Physical Review A* 57.1 (Jan. 1998), 127–137. ISSN: 1094-1622. DOI: 10.1103/physreva.57.127. URL: http://dx.doi.org/10.1103/PhysRevA.57.127.

[23] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. “Error Mitigation for Short-Depth Quantum Circuits”. In: *Physical Review Letters* 119.18 (2017). ISSN: 1079-7114. DOI: 10.1103/physrevlett.119.180509. URL: http://dx.doi.org/10.1103/PhysRevLett.119.180509.

[24] Ying Li and Simon C. Benjamin. “Efficient Variational Quantum Simulator Incorporating Active Error Minimization”. In: *Physical Review X* 7.2 (2017). ISSN: 2160-3308. DOI: 10.1103/physrevx.7.021050. URL: http://dx.doi.org/10.1103/PhysRevX.7.021050.

[25] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. “Surface codes: Towards practical large-scale quantum computation”. In: *Physical Review A* 86.3 (2012). ISSN: 1094-1622. DOI: 10.1103/physreva.86.032324. URL: http://dx.doi.org/10.1103/PhysRevA.86.032324.

[26] Trygve Helgaker, Poul Jørgensen, and Jeppe Olsen. *Molecular Electronic-Structure Theory*. Wiley, 2000.

[27] Attila Szabo and Neil S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. First. Mineola: Dover Publications, Inc., 1996.

[28] Attila Szabo and Neil S. Ostlund. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover, 1996.

[29] C. David Sherrill and Henry F. Schaefer. “The Configuration Interaction Method: Advances in Highly Correlated Approaches”. In: Advances in Quantum Chemistry 34 (1999). Ed. by Per-Olov Löwdin, John R. Sabin, Michael C. Zerner, and Erkki Brändas, pp. 143–269. ISSN: 0065-3276. DOI: https://doi.org/10.1016/S0065-3276(08)60532-8. URL: https://www.sciencedirect.com/science/article/pii/S0065327608605328.

[30] Andreas Görling. “Density-functional theory beyond the Hohenberg-Kohn theorem”. In: *Phys. Rev. A* 59 (5 1999), pp. 3359–3374. DOI: 10.1103/PhysRevA.59.3359. URL: https://link.aps.org/doi/10.1103/PhysRevA.59.3359.

[31] Fernanda Bettanin, Luiz F. A. Ferrão, Max Jr. Pinheiro, Adélia J. A. Aquino, Hans Lischka, Francisco B. C. Machado, and Dana Nachtigallova. “Singlet La and Lb Bands for N-Acenes (N = 2-7): A CASSCF/CASPT2 Study”. In: *J. Chem. Theory Comput.* 13.9 (2017), pp. 4297–4306. DOI: 10.1021/acs.jctc.7b00302. eprint: https://doi.org/10.1021/acs.jctc.7b00302. URL: https://doi.org/10.1021/acs.jctc.7b00302.

[32] Marko Schreiber, Mario R. Silva-Junior, Stephan P. A. Sauer, and Walter Thiel. “Benchmarks for electronically excited states: CASPT2, CC2, CCSD, and CC3”. In: *J. Chem. Phys.* 128.13 (2008), p. 134110. ISSN: 0021-9606. DOI: 10.1063/1.2889385. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.2889385/13877643/134110_1_online.pdf. URL: https://doi.org/10.1063/1.2889385.

[33] Kerstin Andersson, Per-Åke Malmqvist, and Björn O. Roos. "Second-order perturbation theory with a complete active space self-consistent field reference function". In: *J. Phys. Chem.* 96.2 (1992), pp. 1218–1226. ISSN: 0021-9606. DOI: 10.1063/1.462209. eprint: https://pubs.aip.org/aip/jcp/article-pdf/96/2/1218/18996245/1218_1_online.pdf. URL: https://doi.org/10.1063/1.462209.

[34] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J.-P. Malrieu. "Introduction of n-electron valence states for multireference perturbation theory". In: *J. Chem. Pyhs.* 114.23 (2001), pp. 10252–10264. ISSN: 0021-9606. DOI: 10.1063/1.1361246. eprint: https://pubs.aip.org/aip/jcp/article-pdf/114/23/10252/19105543/10252_1_online.pdf. URL: https://doi.org/10.1063/1.1361246.

[35] Celestino Angeli, Renzo Cimiraglia, and Jean-Paul Malrieu. "N-electron valence state perturbation theory: a fast implementation of the strongly contracted variant". In: *Chem. Phys. Lett.* 350.3 (2001), pp. 297–305. ISSN: 0009-2614. DOI: https://doi.org/10.1016/S0009-2614(01)01303-3. URL: https://www.sciencedirect.com/science/article/pii/S0009261401013033.

[36] Celestino Angeli, Renzo Cimiraglia, and Jean-Paul Malrieu. "n-electron valence state perturbation theory: A spinless formulation and an efficient implementation of the strongly contracted and of the partially contracted variants". In: *J. Chem. Phys.* 117.20 (2002), pp. 9138–9153. ISSN: 0021-9606. DOI: 10.1063/1.1515317. eprint: https://pubs.aip.org/aip/jcp/article-pdf/117/20/9138/19244630/9138_1_online.pdf. URL: https://doi.org/10.1063/1.1515317.

[37] Steven R. White. "Density-matrix algorithms for quantum renormalization groups". In: *Phys. Rev. B* 48 (14 1993), pp. 10345–10356. DOI: 10.1103/PhysRevB.48.10345. URL: https://link.aps.org/doi/10.1103/PhysRevB.48.10345.

[38] Steven R. White. "Density matrix formulation for quantum renormalization groups". In: *Phys. Rev. Lett.* 69 (19 1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863. URL: https://link.aps.org/doi/10.1103/PhysRevLett.69.2863.

[39] Garnet Kin-Lic Chan and Sandeep Sharma. "The Density Matrix Renormalization Group in Quantum Chemistry". In: *Annual Review of Physical Chemistry* 62.Volume 62, 2011 (2011), pp. 465–481. ISSN: 1545-1593. DOI: https://doi.org/10.1146/annurev-physchem-032210-103338. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-physchem-032210-103338.

[40] Garnet Kin-Lic Chan and Martin Head-Gordon. "Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group". In: *J. Chem. Phys.* 116.11 (2002), pp. 4462–4476. ISSN: 0021-9606. DOI: 10.1063/1.1449459. eprint: https://pubs.aip.org/aip/jcp/article-pdf/116/11/4462/19222618/4462_1_online.pdf. URL: https://doi.org/10.1063/1.1449459.

[41] Ulrich Schollwöck. "The density-matrix renormalization group in the age of matrix product states". In: *Annals of Physics* 326.1 (2011). January 2011 Special Issue, pp. 96–192. ISSN: 0003-4916. DOI: https://doi.org/10.1016/j.aop.2010.09.012. URL: https://www.sciencedirect.com/science/article/pii/S0003491610001752.

[42] Rodney J. Bartlett and Monika Musiał. “Coupled-cluster theory in quantum chemistry”. In: *Rev. Mod. Phys.* 79 (1 2007), pp. 291–352. DOI: 10.1103/RevModPhys.79.291. URL: https://link.aps.org/doi/10.1103/RevModPhys.79.291.

[43] Rodney J Bartlett and Monika Musiał. “Coupled-Cluster Theory in Quantum Chemistry”. In: *Rev. Mod. Phys.* 79.1 (2007), p. 291.

[44] Krishnan Raghavachari, Gary W. Trucks, John A. Pople, and Martin Head-Gordon. “A fifth-order perturbation comparison of electron correlation theories”. In: *Chemical Physics Letters* 157.6 (1989), pp. 479–483. ISSN: 0009-2614. DOI: https://doi.org/10.1016/S0009-2614(89)87395-6. URL: https://www.sciencedirect.com/science/article/pii/S0009261489873956.

[45] P. Hohenberg and W. Kohn. “Inhomogeneous Electron Gas”. In: *Phys. Rev.* 136.3B (1964), B864–B871.

[46] W. Kohn and L. J. Sham. “Self-Consistent Equations Including Exchange and Correlation Effects”. In: *Phys. Rev.* 140.4A (1965), A1133–A1138.

[47] Robert G Parr and Yang Weitao. *Density-Functional Theory of Atoms and Molecules*. Oxford University Press, 1995. ISBN: 9780195092769. DOI: 10.1093/oso/9780195092769.001.0001. URL: https://doi.org/10.1093/oso/9780195092769.001.0001.

[48] Pascual Jordan and Eugene Wigner. “Über das Paulische Äquivalenzverbot”. In: *Zeitschrift für Physik* 47.9-10 (1928), pp. 631–651.

[49] Sergey Bravyi, Jay M. Gambetta, Antonio Mezzacapo, and Kristan Temme. *Tapering off qubits to simulate fermionic Hamiltonians*. 2017. arXiv: 1701.08213 [quant-ph]. URL: https://arxiv.org/abs/1701.08213.

[50] Sergey B. Bravyi and Alexei Yu. Kitaev. “Fermionic Quantum Computation”. In: *Ann. Phys.* 298.1 (2002), pp. 210–226. ISSN: 0003-4916. DOI: https://doi.org/10.1006/aphy.2002.6254. URL: https://www.sciencedirect.com/science/article/pii/S0003491602962548.

[51] A. Miller, Z. Zimborás, S. Knecht, S. Maniscalco, and G. García-Pérez. “Bonsai Algorithm: Grow Your Own Fermion-to-Qubit Mappings”. In: *PRX Quantum* 4.3 (2023). DOI: 10.1103/prxquantum.4.030314.

[52] A. Miller, A. Glos, and Z. Zimborás. *Treespilation: Architecture- and State-Optimised Fermion-to-Qubit Mappings*. 2024. DOI: 10.48550/ARXIV.2403.03992. URL: https://arxiv.org/abs/2403.03992.

[53] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. “A variational eigenvalue solver on a photonic quantum processor”. In: *Nature Communications* 5.1 (2014), p. 4213. ISSN: 2041-1723. DOI: 10.1038/ncomms5213. URL: https://doi.org/10.1038/ncomms5213.

[54] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. “The theory of variational hybrid quantum-classical algorithms”. In: *New Journal of Physics* 18.2 (2016), p. 023023. ISSN: 1367-2630. DOI: 10.1088/1367-2630/18/2/023023. URL: http://dx.doi.org/10.1088/1367-2630/18/2/023023.

[55] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. In: *Nature* 549.7671 (2017), 242–246. ISSN: 1476-4687. DOI: `10.1038/nature23879`. URL: `http://dx.doi.org/10.1038/nature23879`.

[56] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (2021), 625–644. ISSN: 2522-5820. DOI: `10.1038/s42254-021-00348-9`. URL: `http://dx.doi.org/10.1038/s42254-021-00348-9`.

[57] Dmitry A. Fedorov, Bo Peng, Niranjan Govind, and Yuri Alexeev. *VQE Method: A Short Survey and Recent Developments*. 2021. arXiv: `2103.08505 [quant-ph]`.

[58] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. “Noisy intermediate-scale quantum algorithms”. In: *Reviews of Modern Physics* 94.1 (2022). ISSN: 1539-0756. DOI: `10.1103/revmodphys.94.015004`. URL: `http://dx.doi.org/10.1103/RevModPhys.94.015004`.

[59] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H Booth, et al. “The Variational Quantum Eigensolver: A Review of Methods and Best Practices”. In: *Phys. Rep.* 986 (2022), pp. 1–128.

[60] Walter Ritz. “Über eine neue Methode zur Lösung gewisser Variationsprobleme der mathematischen Physik.” ger. In: *Journal für die reine und angewandte Mathematik* 135 (1909), pp. 1–61. URL: `http://eudml.org/doc/149295`.

[61] Ophelia Crawford, Barnaby van Straaten, Daochen Wang, Thomas Parks, Earl Campbell, and Stephen Brierley. “Efficient quantum measurement of Pauli operators in the presence of finite sampling error”. In: *Quantum* 5 (2021), p. 385. ISSN: 2521-327X. DOI: `10.22331/q-2021-01-20-385`. URL: `http://dx.doi.org/10.22331/q-2021-01-20-385`.

[62] Abhinav Anand and Kenneth R Brown. “Leveraging commuting groups for an efficient variational Hamiltonian ansatz”. In: *Quantum Science and Technology* 10.4 (2025), p. 045009. ISSN: 2058-9565. DOI: `10.1088/2058-9565/adf507`. URL: `http://dx.doi.org/10.1088/2058-9565/adf507`.

[63] Laurin E. Fischer, Timothée Dao, Ivano Tavernelli, and Francesco Tacchino. “Dual-frame optimization for informationally complete quantum measurements”. In: *Physical Review A* 109.6 (2024). ISSN: 2469-9934. DOI: `10.1103/physreva.109.062415`. URL: `http://dx.doi.org/10.1103/PhysRevA.109.062415`.

[64] Roger Fletcher. *Practical Methods of Optimization*. Second. New York, NY, USA: John Wiley & Sons, 1987.

[65] M. J. D. Powell. “A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation”. In: 1994. URL: `https://api.semanticscholar.org/CorpusID:118045691`.

[66] James C. Spall. “A Stochastic Approximation Technique for Generating Maximum Likelihood Parameter Estimates”. In: *1987 American Control Conference* (1987), pp. 1161–1167. URL: https://api.semanticscholar.org/CorpusID:19559269.

[67] Martin Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. *A Review of Barren Plateaus in Variational Quantum Computing*. 2024. arXiv: 2405.00781 [quant-ph].

[68] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. In: *Nature* 549.7671 (2017), 242–246. ISSN: 1476-4687. DOI: 10.1038/nature23879. URL: http://dx.doi.org/10.1038/nature23879.

[69] M. Ganzhorn, D.J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp. “Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer”. In: *Phys. Rev. Appl.* 11 (4 2019), p. 044092. DOI: 10.1103/PhysRevApplied.11.044092. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.11.044092.

[70] Arthur G. Rattew, Shaohan Hu, Marco Pistoia, Richard Chen, and Steve Wood. *A Domain-agnostic, Noise-resistant, Hardware-efficient Evolutionary Variational Quantum Eigensolver*. 2020. arXiv: 1910.09694 [quant-ph].

[71] Nikolay V. Tkachenko, James Sud, Yu Zhang, Sergei Tretiak, Petr M. Anisimov, Andrew T. Arrasmith, Patrick J. Coles, Lukasz Cincio, and Pavel A. Dub. “Correlation-Informed Permutation of Qubits for Reducing Ansatz Depth in the Variational Quantum Eigensolver”. In: *PRX Quantum* 2 (2 2021), p. 020337. DOI: 10.1103/PRXQuantum.2.020337. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.020337.

[72] Ho Lun Tang, V.O. Shkolnikov, George S. Barron, Harper R. Grimsley, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. “Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor”. In: *PRX Quantum* 2 (2 2021), p. 020310. DOI: 10.1103/PRXQuantum.2.020310. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.020310.

[73] Yordan S. Yordanov, David R. M. Arvidsson-Shukur, and Crispin H. W. Barnes. “Efficient quantum circuits for quantum computational chemistry”. In: *Physical Review A* 102.6 (2020). ISSN: 2469-9934. DOI: 10.1103/physreva.102.062612. URL: http://dx.doi.org/10.1103/PhysRevA.102.062612.

[74] Mark R. Hoffmann and Jack Simons. “A unitary multiconfigurational coupled-cluster method: Theory and applications”. In: *The Journal of Chemical Physics* 88.2 (1988), pp. 993–1002. ISSN: 0021-9606. DOI: 10.1063/1.454125. eprint: https://pubs.aip.org/aip/jcp/article-pdf/88/2/993/11189139/993_1_online.pdf. URL: https://doi.org/10.1063/1.454125.

[75] Bridgette Cooper and Peter J. Knowles. “Benchmark studies of variational, unitary and extended coupled cluster methods”. In: *The Journal of Chemical Physics* 133.23 (2010), p. 234102. ISSN: 0021-9606. DOI: 10.1063/1.3520564. eprint: https://pubs.aip.org/aip/

jcp/article-pdf/doi/10.1063/1.3520564/15434716/234102_1_online.pdf. URL: https://doi.org/10.1063/1.3520564.

[76] Francesco A. Evangelista. "Alternative single-reference coupled cluster approaches for multireference problems: The simpler, the better". In: *The Journal of Chemical Physics* 134.22 (2011), p. 224102. ISSN: 0021-9606. DOI: 10.1063/1.3598471. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.3598471/15437394/224102_1_online.pdf. URL: https://doi.org/10.1063/1.3598471.

[77] Jonathan Romero, Ryan Babbush, Jarrod R. McClean, Cornelius Hempel, Peter J. Love, and Alán Aspuru-Guzik. *Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz.* 2018. DOI: 10.1088/2058-9565/aad3e4. arXiv: 1701.02691 [quant-ph].

[78] Ilias Magoulas and Francesco A. Evangelista. "Unitary Coupled Cluster: Seizing the Quantum Moment". In: *The Journal of Physical Chemistry A* 127 (31 2023), pp. 6567–6576. ISSN: 1089-5639. DOI: 10.1021/acs.jpca.3c02781.

[79] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, Stefan Filipp, and Ivano Tavernelli. "Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions". In: *Phys. Rev. A* 98 (2 2018), p. 022322. DOI: 10.1103/PhysRevA.98.022322. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.022322.

[80] Joonho Lee, William J. Huggins, Martin Head-Gordon, and K. Birgitta Whaley. "Generalized Unitary Coupled Cluster Wave functions for Quantum Computation". In: *Journal of Chemical Theory and Computation* 15 (1 2019), pp. 311–324. ISSN: 15499626. DOI: 10.1021/acs.jctc.8b01004.

[81] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. "An adaptive variational algorithm for exact molecular simulations on a quantum computer". In: *Nature Communications* 10.1 (2019), p. 3007. ISSN: 2041-1723. DOI: 10.1038/s41467-019-10988-2. URL: https://doi.org/10.1038/s41467-019-10988-2.

[82] Panagiotis G. Anastasiou, Yanzhu Chen, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. "TETRIS-ADAPT-VQE: An adaptive algorithm that yields shallower, denser circuit Ansätze". In: *Physical Review Research* 6.1 (2024). ISSN: 2643-1564. DOI: 10.1103/physrevresearch.6.013254. URL: http://dx.doi.org/10.1103/PhysRevResearch.6.013254.

[83] César Feniou, Muhammad Hassan, Diata Traoré, Emmanuel Giner, Yvon Maday, and Jean-Philip Piquemal. "Overlap-ADAPT-VQE: practical quantum chemistry on quantum computers via overlap-guided compact Ansätze". In: *Communications Physics* 6.1 (2023). ISSN: 2399-3650. DOI: 10.1038/s42005-023-01312-y. URL: http://dx.doi.org/10.1038/s42005-023-01312-y.

[84] Mafalda Ramôa, Panagiotis G. Anastasiou, Luis Paulo Santos, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. "Reducing the resources required by ADAPT-VQE using coupled exchange operators and improved subroutines". In: *npj Quantum Information* 11.1 (2025). ISSN: 2056-6387. DOI: 10.1038/s41534-025-01039-4. URL: http://dx.doi.org/10.1038/s41534-025-01039-4.

[85] Yordan S. Yordanov, V. Armaos, Crispin H. W. Barnes, and David R. M. Arvidsson-Shukur. "Qubit-excitation-based adaptive variational quantum eigensolver". In: *Communications Physics* 4.1 (2021). ISSN: 2399-3650. DOI: 10.1038/s42005-021-00730-0. URL: http://dx.doi.org/10.1038/s42005-021-00730-0.

[86] Zi-Jian Zhang, Thi Ha Kyaw, Jakob S Kottmann, Matthias Degroote, and Alán Aspuru-Guzik. "Mutual information-assisted adaptive variational quantum eigensolver". In: *Quantum Science and Technology* 6.3 (2021), p. 035001. ISSN: 2058-9565. DOI: 10.1088/2058-9565/abdca4. URL: http://dx.doi.org/10.1088/2058-9565/abdca4.

[87] J. Wayne Mullinax, Panagiotis G. Anastasiou, Jeffrey Larson, Sophia E. Economou, and Norm M. Tubman. "Classical Preoptimization Approach for ADAPT-VQE: Maximizing the Potential of High-Performance Computing Resources to Improve Quantum Simulation of Chemical Applications". In: *Journal of Chemical Theory and Computation* 21.8 (2025), 4006–4015. ISSN: 1549-9626. DOI: 10.1021/acs.jctc.5c00150. URL: http://dx.doi.org/10.1021/acs.jctc.5c00150.

[88] Aaron Fitzpatrick, N. Walter Talarico, Roberto Di Remigio Eikås, and Stefan Knecht. *Quantum-centric strong and dynamical electron correlation: A resource-efficient second-order N-electron valence perturbation theory formulation for near-term quantum devices.* 2024. arXiv: 2405.15422 [quant-ph]. URL: https://arxiv.org/abs/2405.15422.

[89] Wen Wei Ho and Timothy H. Hsieh. "Efficient variational simulation of non-trivial quantum states". In: *SciPost Phys.* 6 (2019), p. 029. DOI: 10.21468/SciPostPhys.6.3.029. URL: https://scipost.org/10.21468/SciPostPhys.6.3.029.

[90] Roeland Wiersema, Cunlu Zhou, Yvette de Sereville, Juan Felipe Carrasquilla, Yong Baek Kim, and Henry Yuen. "Exploring Entanglement and Optimization within the Hamiltonian Variational Ansatz". In: *PRX Quantum* 1 (2 2020), p. 020319. DOI: 10.1103/PRXQuantum.1.020319. URL: https://link.aps.org/doi/10.1103/PRXQuantum.1.020319.

[91] Chris Cade, Lana Mineh, Ashley Montanaro, and Stasja Stanisic. "Strategies for solving the Fermi-Hubbard model on near-term quantum computers". In: *Phys. Rev. B* 102 (23 2020), p. 235122. DOI: 10.1103/PhysRevB.102.235122. URL: https://link.aps.org/doi/10.1103/PhysRevB.102.235122.

[92] Susan Monro, Katsuya L. Colon, Huimin Yin, John III Roque, Prathyusha Konda, Shashi Gujar, Randolph P. Thummel, Lothar Lilge, Colin G. Cameron, and Sherri A. McFarland. "Transition Metal Complexes and Photodynamic Therapy from a Tumor-Centered Approach: Challenges, Opportunities, and Highlights from the Development of TLD1433". In: *Chem. Rev.* 119.2 (2019), pp. 797–828. DOI: 10.1021/acs.chemrev.8b00211. eprint: https://doi.org/10.1021/acs.chemrev.8b00211. URL: https://doi.org/10.1021/acs.chemrev.8b00211.

[93] Candida Manelfi, Valerio Tazzari, Filippo Lunghini, Carmen Cerchia, Anna Fava, Alessandro Pedretti, Pieter F. W. Stouten, Giulio Vistoli, and Andrea Rosario Beccari. "“DompeKeys”: a set of novel substructure-based descriptors for efficient chemical space mapping, development and structural interpretation of machine learning models, and indexing of large databases". In: *Journal of Cheminformatics* 16.1 (2024), p. 21.

[94] Qiming Sun et al. "Recent developments in the PySCF program package". In: *J. Phys. Chem.* 153.2 (2020), p. 024109. ISSN: 0021-9606. DOI: 10.1063/5.0006074. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0006074/16722275/024109_1_online.pdf. URL: https://doi.org/10.1063/5.0006074.

[95] Qiming Sun, Timothy C Berkelbach, Nick S Blunt, George H Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D McClain, Elvira R Sayfutyarova, Sandeep Sharma, et al. "PySCF: the Python-based simulations of chemistry framework". In: *WIREs Comput. Mol. Sci.* 8.1 (2018), e1340.

[96] Qiming Sun, Xing Zhang, Samragni Banerjee, Peng Bao, Marc Barbry, Nick S Blunt, Nikolay A Bogdanov, George H Booth, Jia Chen, Zhi-Hao Cui, et al. "Recent developments in the PySCF program package". In: *J. Chem. Phys.* 153.2 (2020).

[97] Juan Angel de Gracia Triviño, Iulia Emilia Brumboiu, David Carrasco-Busturia, Xin Li, Chenxi Li, Mathieu Linares, Valentin Lindfeld, Young Min Rhee, Julia Rune, Bastiaan van Hoorn, Patrick Norman, and Mårten S. G. Ahlquist. "VeloxChem Quantum–Classical Interoperability for Modeling of Complex Molecular Systems". In: *The Journal of Physical Chemistry A* 129.32 (2025). PMID: 40754762, pp. 7575–7587. DOI: 10.1021/acs.jpca.5c03187. eprint: https://doi.org/10.1021/acs.jpca.5c03187. URL: https://doi.org/10.1021/acs.jpca.5c03187.

[98] Zilvinas Rinkevicius, Xin Li, Olav Vahtras, Karan Ahmadzadeh, Manuel Brand, Magnus Ringholm, Nanna Holmgaard List, Maximilian Scheurer, Mikael Scott, Andreas Dreuw, and Patrick Norman. "VeloxChem: A Python-driven density-functional theory program for spectroscopy simulations in high-performance computing environments". In: *WIREs Computational Molecular Science* 10.5 (2020), e1457. DOI: https://doi.org/10.1002/wcms.1457. eprint: https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1457. URL: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1457.

[99] Xin Li, Mathieu Linares, and Patrick Norman. "VeloxChem: GPU-Accelerated Fock Matrix Construction Enabling Complex Polarization Propagator Simulations of Circular Dichroism Spectra of G-Quadruplexes". In: *The Journal of Physical Chemistry A* 129.2 (2025). PMID: 39739694, pp. 633–642. DOI: 10.1021/acs.jpca.4c07510. eprint: https://doi.org/10.1021/acs.jpca.4c07510. URL: https://doi.org/10.1021/acs.jpca.4c07510.

[100] The Qiskit Nature developers and contributors. *Qiskit Nature 0.6.0.* Version 0.6.0. 2023. DOI: 10.5281/zenodo.7828768. URL: https://doi.org/10.5281/zenodo.7828768.

[101] Jin-Sung Kim, Alex McCaskey, Bettina Heim, Manish Modani, Sam Stanwyck, and Timothy Costa. "CUDA Quantum: The Platform for Integrated Quantum-Classical Computing". In: *2023 60th ACM/IEEE Design Automation Conference (DAC)*. 2023, pp. 1–4. DOI: 10.1109/DAC56929.2023.10247886.

[102] Priyanka Makkar and Narendra Nath Ghosh. "A review on the use of DFT for the prediction of the properties of nanomaterials". In: *RSC Adv.* 11 (45 2021), pp. 27897–27924. DOI: 10.1039/D1RA04876G. URL: http://dx.doi.org/10.1039/D1RA04876G.

[103] Kurt Lejaeghere et al. "Reproducibility in density functional theory calculations of solids". In: *Science* 351.6280 (2016), aad3000. DOI: 10.1126/science.aad3000. URL: https://www.science.org/doi/abs/10.1126/science.aad3000.

[104] Kieron Burke. "Perspective on Density Functional Theory". In: *J. Chem. Phys.* 136.15 (2012), p. 150901.

[105] Hans Lischka, Dana Nachtigallová, Adélia J. A. Aquino, Péter G. Szalay, Felix Plasser, Francisco B. C. Machado, and Mario Barbatti. "Multireference Approaches for Excited States of Molecules". In: *Chem. Rev.* 118.15 (2018), pp. 7293–7361. DOI: 10.1021/acs.chemrev.8b00244. URL: https://doi.org/10.1021/acs.chemrev.8b00244.

[106] Péter G. Szalay, Thomas Müller, Gergely Gidofalvi, Hans Lischka, and Ron Shepard. "Multiconfiguration Self-Consistent Field and Multireference Configuration Interaction Methods and Applications". In: *Chem. Rev.* 112.1 (2011), pp. 108–181. DOI: 10.1021/cr200137a.

[107] T Daniel Crawford and Henry F Schaefer. "An Introduction to Coupled Cluster Theory for Computational Chemists". In: *Rev. Comput. Chem.* 14 (2000), pp. 33–136.

[108] Dmitry I Lyakh, Monika Musiał, Victor F Lotrich, and Rodney J Bartlett. "Multireference Nature of Chemistry: The Coupled-Cluster View". In: *Chem. Rev.* 112.1 (2012), pp. 182–243.

[109] Björn O. Roos, Peter R. Taylor, and Per E.M. Sigbahn. "A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach". In: *Chem. Phys.* 48.2 (1980), pp. 157–173. ISSN: 0301-0104. DOI: https://doi.org/10.1016/0301-0104(80)80045-0. URL: https://www.sciencedirect.com/science/article/pii/0301010480800450.

[110] Kerstin. Andersson, Per-Åke. Malmqvist, Bjöern O. Roos, Andrzej J. Sadlej, and Krzysztof. Wolinski. "Second-order perturbation theory with a CASSCF reference function". In: *J. Phys. Chem.* 94.14 (1990), pp. 5483–5488. DOI: 10.1021/j100377a012. eprint: https://doi.org/10.1021/j100377a012. URL: https://doi.org/10.1021/j100377a012.

[111] Konstantinos D Vogiatzis, Dongxia Ma, Jeppe Olsen, Laura Gagliardi, and Wibe A De Jong. "Pushing configuration-interaction to the limit: Towards massively parallel MCSCF calculations". In: *J. Chem. Phys.* 147.18 (2017).

[112] S. Knecht, E. D. Hedegaard, S. Keller, A. Kovyrshin, Y. Ma, A. Muolo, C. J. Stein, and M. Reiher. "New approaches for ab initio calculations of molecules with strong electron correlation". In: *Chimia* 70 (2016), pp. 244–251. DOI: 10.2533/chimia.2016.244.

[113] A. Baiardi and M. Reiher. "The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges". In: *J. Chem. Phys.* 152 (2020), p. 040903. DOI: 10.1063/1.5129672.

[114] George H Booth, Alex J W Thom, and Ali Alavi. "Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space". In: *J. Chem. Phys.* 131 (2009), pp. 054106–11. DOI: 10.1063/1.3193710.

[115] Junhao Li, Matthew Otten, Adam A. Holmes, Sandeep Sharma, and C. J. Umrigar. “Fast semistochastic heat-bath configuration interaction”. In: *J. Chem. Phys.* 149.21 (2018), p. 214110. DOI: `10.1063/1.5055390`.

[116] Norm M. Tubman, Joonho Lee, Tyler Y. Takeshita, Martin Head-Gordon, and K. Birgitta Whaley. “A deterministic alternative to the full configuration interaction quantum Monte Carlo method”. In: *J. Chem. Phys.* 145 (2016), p. 044112. DOI: `10.1063/1.4955109`.

[117] V. G. Chilkuri and F. Neese. “Comparison of Many-Particle Representations for Selected-CI I: A Tree Based Approach”. In: *J. Comp. Chem.* 42 (2021), pp. 982–1005. DOI: `10.1002/jcc.26510`.

[118] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. “Quantum Chemistry in the Age of Quantum Computing”. In: *Chem. Rev.* 119.19 (2019), pp. 10856–10915.

[119] Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. “Quantum Algorithms for Quantum Chemistry and Quantum Materials Science”. In: *Chem. Rev.* 120.22 (2020), pp. 12685–12717.

[120] Hongbin Liu, Guang Hao Low, Damian S. Steiger, Thomas Häner, Markus Reiher, and Matthias Troyer. “Prospects of quantum computing for molecular sciences”. In: *Mater. Theory* 6.1 (2022), p. 11. ISSN: 2509-8012. DOI: `10.1186/s41313-021-00039-z`. URL: `https://doi.org/10.1186/s41313-021-00039-z`.

[121] Yudong Cao, Jonathan Romero, Jonathan P. Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D. Kivlichan, Tim Menke, Borja Peropadre, Nicolas P. D. Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik. “Quantum Chemistry in the Age of Quantum Computing”. In: *Chem. Rev.* 119.19 (2019). PMID: 31469277, pp. 10856–10915. DOI: `10.1021/acs.chemrev.8b00803`. eprint: `https://doi.org/10.1021/acs.chemrev.8b00803`. URL: `https://doi.org/10.1021/acs.chemrev.8b00803`.

[122] Jared D. Weidman, Manas Sajjan, Camille Mikolas, Zachary J. Stewart, Johannes Pollanen, Sabre Kais, and Angela K. Wilson. “Quantum computing and chemistry”. In: *Cell Reports Physical Science* 5 (2024), p. 102105. DOI: `10.1016/j.xcrp.2024.102105`. URL: `https://doi.org/10.1016/j.xcrp.2024.102105`.

[123] Y. Alexeev et al. *A Perspective on Quantum Computing Applications in Quantum Chemistry using 25–100 Logical Qubits*. arXiv: 2506.19337. 2025. URL: `https://arxiv.org/abs/2506.19337`.

[124] Philipp Schleich, Luis Mantilla Calderón, Chong Sun, Mohsen Bagherimehrab, Abdulrahman Aldossary, Jakob S. Kottmann, and Alán Aspuru-Guzik. *Quantum Computing for Quantum Chemistry*. Washington, DC, USA: American Chemical Society, 2025. DOI: `10.1021/acsinfocus.7e9012`. eprint: `https://pubs.acs.org/doi/pdf/10.1021/acsinfocus.7e9012`. URL: `https://pubs.acs.org/doi/abs/10.1021/acsinfocus.7e9012`.

[125] David M. Rogers, Claire Wells, Melanie Joseph, Vanessa J. Boddington, and Joseph J.W. McDouall. “On the choice of active space orbitals in MCSCF calculations”. In: *J. Mol. Struct. THEOCHEM* 434.1 (1998), pp. 239–245. ISSN: 0166-1280. DOI: `https://doi.org/10.1016/S0166-1280(98)00123-7`. URL: `https://www.sciencedirect.com/science/article/pii/S0166128098001237`.

[126] Zsuzsanna Tóth and Peter Pulay. “Comparison of Methods for Active Orbital Selection in Multiconfigurational Calculations”. In: *J. Chem. Theory Comput.* 16.12 (2020). PMID: 33170653, pp. 7328–7341. DOI: `10.1021/acs.jctc.0c00123`. eprint: `https://doi.org/10.1021/acs.jctc.0c00123`. URL: `https://doi.org/10.1021/acs.jctc.0c00123`.

[127] Abhishek Khedkar and Michael Roemelt. “Active Space Selection Based on Natural Orbital Occupation Numbers from n-Electron Valence Perturbation Theory”. In: *J. Chem. Theory Comput.* 15.6 (2019). PMID: 31059643, pp. 3522–3536. DOI: `10.1021/acs.jctc.8b01293`. eprint: `https://doi.org/10.1021/acs.jctc.8b01293`. URL: `https://doi.org/10.1021/acs.jctc.8b01293`.

[128] Abhishek Khedkar and Michael Roemelt. “Extending the ASS1ST Active Space Selection Scheme to Large Molecules and Excited States”. In: *J. Chem. Theory Comput.* 16.8 (2020). PMID: 32644789, pp. 4993–5005. DOI: `10.1021/acs.jctc.0c00332`. eprint: `https://doi.org/10.1021/acs.jctc.0c00332`. URL: `https://doi.org/10.1021/acs.jctc.0c00332`.

[129] Elvira R. Sayfutyarova, Qiming Sun, Garnet Kin-Lic Chan, and Gerald Knizia. “Automated Construction of Molecular Active Spaces from Atomic Valence Orbitals”. In: *J. Chem. Theory Comput.* 13.9 (2017), 4063–4078. ISSN: 1549-9626. DOI: `10.1021/acs.jctc.7b00128`. URL: `http://dx.doi.org/10.1021/acs.jctc.7b00128`.

[130] Christopher J. Stein and Markus Reiher. “Automated Selection of Active Orbital Spaces”. In: *J. Chem. Theory Comput.* 12.4 (2016), pp. 1760–1771. ISSN: 1549-9618. DOI: `10.1021/acs.jctc.6b00156`. URL: `https://doi.org/10.1021/acs.jctc.6b00156`.

[131] Christopher J Stein and Markus Reiher. “Automated Identification of Relevant Frontier Orbitals for Chemical Compounds and Processes”. en. In: *Chimia (Aarau)* 71.4 (2017), pp. 170–176.

[132] Christopher J. Stein and Markus Reiher. “autoCAS: A Program for Fully Automated Multiconfigurational Calculations”. In: *J. Comp. Chem.* 40.25 (2019), pp. 2216–2226. DOI: `https://doi.org/10.1002/jcc.25869`. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.25869`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.25869`.

[133] Björn O Roos. “MULTICONFIGURATIONAL QUANTUM CHEMISTRY FOR GROUND AND EXCITED STATES”. In: *Radiation Induced Molecular Phenomena in Nucleic Acids*. Ed. by M. K. Shukla and J. Leszczynski. Springer, 2008, pp. 125–156.

[134] H. Hu, S. Upadhyay, L. Lu, A. J. Jenkins, T. Zhang, A. Shayit, S. Knecht, and X. Li. “Small Tensor Product Distributed Active Space (STP-DAS) Framework for Relativistic and Nonrelativistic Multiconfiguration Calculations: Scaling from 10^9 on a Laptop to $10^1$2 Determinants on a Supercomputer”. In: *Chem. Phys. Rev.* 5 (2024), p. 041404.

[135] H. Gao, S. Imamura, A. Kasagi, and E. Yoshida. “Distributed Implementation of Full Configuration Interaction for One Trillion Determinants”. In: *J. Chem. Theory Comput.* 20 (2024), pp. 1185–1192. DOI: `10.1021/acs.jctc.3c01190`.

[136] A. Shayit, C. Liao, S. Upadhyay, H. Hu, T. Zhang, A. E. DePrince, C. Yang, and X. Li. *Breaking the Quadrillion Determinant Barrier in Numerically Exact Configuration Interaction.* arXiv: 2505.20375. 2025. URL: `https://arxiv.org/abs/2505.20375`.

[137] U. Schollwöck. “The density-matrix renormalization group”. In: *Rev. Mod. Phys.* 77 (1 2005), pp. 259–315. DOI: `10.1103/RevModPhys.77.259`. URL: `https://link.aps.org/doi/10.1103/RevModPhys.77.259`.

[138] S. Keller and M. Reiher. “Determining Factors for the Accuracy of DMRG in Chemistry”. In: *Chimia* 68 (2014), pp. 200–203.

[139] Ignacio Fdez. Galván et al. “OpenMolcas: From Source Code to Insight”. In: *J. Chem. Theory Comput.* 15 (2019), pp. 5925–5964.

[140] Giovanni Li Manni et al. “The OpenMolcas Web: A Community-Driven Approach to Advancing Computational Chemistry”. In: *J. Chem. Theory Comput.* 19.20 (2023), pp. 6933–6991. DOI: `10.1021/acs.jctc.3c00182`. eprint: `https://doi.org/10.1021/acs.jctc.3c00182`. URL: `https://doi.org/10.1021/acs.jctc.3c00182`.

[141] T. Dresselhaus, J. Neugebauer, S. Knecht, S. Keller, Y. Ma, and M. Reiher. “Self-consistent embedding of density-matrix renormalization group wavefunctions in a density functional environment”. In: *J. Chem. Phys.* 142 (2015), p. 044111.

[142] Y. Ma, S. Knecht, S. Keller, and M. Reiher. “Second-order self-consistent-field density-matrix renormalization group”. In: *J. Chem. Theory Comput.* 13 (2017), pp. 2533–2549.

[143] Huanchen Zhai and Garnet Kin-Lic Chan. “Low communication high performance ab initio density matrix renormalization group algorithms”. In: *J. Chem. Phys.* 154.22 (2021), p. 224116. ISSN: 0021-9606. DOI: `10.1063/5.0050902`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0050902/14003774/224116\_1\_online.pdf`. URL: `https://doi.org/10.1063/5.0050902`.

[144] Yihe Xu, Yifan Cheng, Yinxuan Song, and Haibo Ma. “New Density Matrix Renormalization Group Approaches for Strongly Correlated Systems Coupled with Large Environments”. In: *J. Chem. Theory Comput.* 19.15 (2023). PMID: 37471519, pp. 4781–4795. DOI: `10.1021/acs.jctc.2c01316`. eprint: `https://doi.org/10.1021/acs.jctc.2c01316`. URL: `https://doi.org/10.1021/acs.jctc.2c01316`.

[145] Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M. Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, Michael Kastoryano, Ryan Babbush, John Preskill, David R. Reichman, Earl T. Campbell, Edward F. Valeev, Lin Lin, and Garnet Kin-Lic Chan. “Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry”. In: *Nat. Commun.* 14.1 (2023), p. 1952. ISSN: 2041-1723. DOI: `10.1038/s41467-023-37587-6`. URL: `https://doi.org/10.1038/s41467-023-37587-6`.

[146] Sandeep Sharma, Kantharuban Sivalingam, Frank Neese, and Garnet Kin-Lic Chan. "Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics". In: *Nat. Chem.* 6.10 (2014), pp. 927–933. ISSN: 1755-4349. DOI: 10.1038/nchem.2041. URL: https://doi.org/10.1038/nchem.2041.

[147] Yuki Kurashige, Garnet Kin-Lic Chan, and Takeshi Yanai. "Entangled quantum electronic wavefunctions of the Mn_4CaO_5 cluster in photosystem II". en. In: *Nat. Chem.* 5.8 (2013), pp. 660–666.

[148] S. Knecht, Ö. Legeza, and M. Reiher. "Four-component density matrix renormalization group." In: *J. Chem. Phys.* 140 (2014), p. 041101.

[149] S. Battaglia, S. Keller, and S. Knecht. "An efficient relativistic density-matrix renormalization group implementation in a matrix-product operator formulation". In: *J. Chem. Theory Comput.* 14 (2018), pp. 2353–2369.

[150] Chad E. Hoyer, Hang Hu, Lixin Lu, Stefan Knecht, and Xiaosong Li. "Relativistic Kramers-Unrestricted Exact-Two-Component Density Matrix Renormalization Group". In: *J. Phys. Chem. A* 126 (2022), pp. 5011–5020. DOI: 10.1021/acs.jpca.2c02150.

[151] Chad E. Hoyer, Lixin Lu, Hang Hu, Kirill D. Shumilov, Shichao Sun, Stefan Knecht, and Xiaosong Li. "Correlated Dirac–Coulomb–Breit multiconfigurational self-consistent-field methods". In: *J. Chem. Phys.* 158 (2023), p. 044101. DOI: 10.1063/5.0133741.

[152] Per-Olov Löwdin. "Quantum Theory of Many-Particle Systems. I. Physical Interpretations by Means of Density Matrices, Natural Spin-Orbitals, and Convergence Problems in the Method of Configurational Interaction". In: *Phys. Rev.* 97 (1955), pp. 1474–1489.

[153] Jie J. Bao, Sijia S. Dong, Laura Gagliardi, and Donald G. Truhlar. "Automatic Selection of an Active Space for Calculating Electronic Excitation Spectra by MS-CASPT2 or MC-PDFT". In: *J. Chem. Theory Comput.* 14.4 (2018). PMID: 29486125, pp. 2017–2025. DOI: 10.1021/acs.jctc.8b00032. eprint: https://doi.org/10.1021/acs.jctc.8b00032. URL: https://doi.org/10.1021/acs.jctc.8b00032.

[154] Giovanni Li Manni, Rebecca K Carlson, Sijie Luo, Dongxia Ma, Jeppe Olsen, Donald G Truhlar, and Laura Gagliardi. "Multiconfiguration Pair-Density Functional Theory". en. In: *J. Chem. Theory Comput.* 10.9 (2014), pp. 3669–3680.

[155] Giovanni Li Manni, Rebecca K Carlson, Sijie Luo, Dongxia Ma, Jeppe Olsen, Donald G Truhlar, and Laura Gagliardi. "Correction to multiconfiguration pair-density functional theory". en. In: *J. Chem. Theory Comput.* 12.1 (2016), p. 458.

[156] Moritz Bensberg and Markus Reiher. "Corresponding Active Orbital Spaces along Chemical Reaction Paths". In: *J. Phys. Chem. Lett* 14.8 (2023). PMID: 36802629, pp. 2112–2118. DOI: 10.1021/acs.jpclett.2c03905. eprint: https://doi.org/10.1021/acs.jpclett.2c03905. URL: https://doi.org/10.1021/acs.jpclett.2c03905.

[157] K. Boguslawski and P. Tecmer. "Orbital Entanglement in Quantum Chemistry". In: *Int. J. Quantum Chem.* 115 (2015), pp. 1289–1295.

[158] Matthew Otten, Matthew R. Hermes, Riddhish Pandharkar, Yuri Alexeev, Stephen K. Gray, and Laura Gagliardi. "Localized Quantum Chemistry on Quantum Computers". In: *J. Chem. Theory Comput.* 18.12 (2022). PMID: 36346785, pp. 7205–7217. DOI: `10.1021/acs.jctc.2c00388`. eprint: `https://doi.org/10.1021/acs.jctc.2c00388`. URL: `https://doi.org/10.1021/acs.jctc.2c00388`.

[159] Carlos Sosa, Jan Geertsen, Gary W. Trucks, Rodney J. Bartlett, and James A. Franz. "Selection of the reduced virtual space for correlated calculations. An application to the energy and dipole moment of H2O". In: *Chem. Phys. Lett.* 159.2 (1989), pp. 148–154. ISSN: 0009-2614. DOI: `https://doi.org/10.1016/0009-2614(89)87399-3`. URL: `https://www.sciencedirect.com/science/article/pii/0009261489873993`.

[160] Andrew G Taube and Rodney J Bartlett. "Frozen natural orbitals: Systematic basis set truncation for coupled-cluster theory". In: *Collect. Czechoslov. Chem. Commun* 70.6 (2005), pp. 837–850. DOI: `https://doi.org/10.1135/cccc20050837`. URL: `http://cccc.uochb.cas.cz/70/6/0837/`.

[161] Andrew G. Taube and Rodney J. Bartlett. "Frozen natural orbital coupled-cluster theory: Forces and application to decomposition of nitroethane". In: *J. Chem. Phys.* 128.16 (2008), p. 164101. ISSN: 0021-9606. DOI: `10.1063/1.2902285`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.2902285/15413526/164101\_1\_online.pdf`. URL: `https://doi.org/10.1063/1.2902285`.

[162] Prakash Verma, Lee Huntington, Marc P. Coons, Yukio Kawashima, Takeshi Yamazaki, and Arman Zaribafiyan. "Scaling up electronic structure calculations on quantum computers: The frozen natural orbital based method of increments". In: *J. Chem. Phys.* 155.3 (2021), p. 034110. ISSN: 0021-9606. DOI: `10.1063/5.0054647`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0054647/13511272/034110\_1\_online.pdf`. URL: `https://doi.org/10.1063/5.0054647`.

[163] Davide Materia, Leonardo Ratini, and Leonardo Guidoni. *Quantum information theory on sparse wavefunctions and applications for Quantum Chemistry.* 2024. arXiv: `2408.02631 [quant-ph]`. URL: `https://arxiv.org/abs/2408.02631`.

[164] W. Kohn and L. J. Sham. "Self-Consistent Equations Including Exchange and Correlation Effects". In: *Phys. Rev.* 140 (4A 1965), A1133–A1138. DOI: `10.1103/PhysRev.140.A1133`. URL: `https://link.aps.org/doi/10.1103/PhysRev.140.A1133`.

[165] Lara Martinez-Fernandez and Antonio Frances-Monerris. "Chapter 12 - DNA photostability". In: *Comput. Theor. Chem.* Ed. by Cristina Garcia-Iriepa and Marco Marazzi. Elsevier, 2023, pp. 311–336. ISBN: 978-0-323-91738-4. DOI: `https://doi.org/10.1016/B978-0-323-91738-4.00001-4`. URL: `https://www.sciencedirect.com/science/article/pii/B9780323917384000014`.

[166] José H. Correia, José A. Rodrigues, Sara Pimenta, Tao Dong, and Zhaochu Yang. "Photodynamic Therapy Review: Principles, Photosensitizers, Applications, and Future Directions". In: *Pharmaceutics* 13.9 (2021). ISSN: 1999-4923. DOI: `10.3390/pharmaceutics13091332`. URL: `https://www.mdpi.com/1999-4923/13/9/1332`.

[167] Spiridoula Matsika and Anna I. Krylov. "Introduction: Theoretical Modeling of Excited State Processes". In: *Chem. Rev.* 118.15 (2018), pp. 6925–6926. DOI: `10.1021/acs.chemrev.8b00436`. eprint: `https://doi.org/10.1021/acs.chemrev.8b00436`. URL: `https://doi.org/10.1021/acs.chemrev.8b00436`.

[168] Daniel Escudero and Leticia González. "RASPT2/RASSCF vs Range-Separated/Hybrid DFT Methods: Assessing the Excited States of a Ru(II)bipyridyl Complex". In: *J. Chem. Theory Comput.* 8.1 (2012). PMID: 26592882, pp. 203–213. DOI: `10.1021/ct200640q`. eprint: `https://doi.org/10.1021/ct200640q`. URL: `https://doi.org/10.1021/ct200640q`.

[169] Andrew J. Atkins, Francesco Talotta, Leon Freitag, Martial Boggio-Pasqua, and Leticia González. "Assessing Excited State Energy Gaps with Time-Dependent Density Functional Theory on Ru(II) Complexes". In: *J. Chem. Theory Comput.* 13.9 (2017). PMID: 28787162, pp. 4123–4145. DOI: `10.1021/acs.jctc.7b00379`. eprint: `https://doi.org/10.1021/acs.jctc.7b00379`. URL: `https://doi.org/10.1021/acs.jctc.7b0037`.

[170] Marta Erminia Alberto, Jenny Pirillo, Nino Russo, and Carlo Adamo. "Theoretical Exploration of Type I/Type II Dual Photoreactivity of Promising Ru(II) Dyads for PDT Approach". In: *Inorganic Chemistry* 55.21 (2016), pp. 11185–11192. ISSN: 0020-1669. DOI: `10.1021/acs.inorgchem.6b01782`. URL: `https://doi.org/10.1021/acs.inorgchem.6b01782`.

[171] Casida Mark E. "Time-Dependent Density Functional Response Theory for Molecules". In: *Recent Advances in Density Functional Methods*. 1984, pp. 155–192. DOI: `10.1142/9789812830586_0005`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/9789812830586_0005`. URL: `https://www.worldscientific.com/doi/abs/10.1142/9789812830586_0005`.

[172] Mark E. Casida. "Time-dependent density-functional theory for molecules and molecular solids". In: *J. Mol. Struct.: THEOCHEM* 914.1 (2009). Time-dependent density-functional theory for molecules and molecular solids, pp. 3–18. ISSN: 0166-1280. DOI: `https://doi.org/10.1016/j.theochem.2009.08.018`. URL: `https://www.sciencedirect.com/science/article/pii/S0166128009005363`.

[173] Erich Runge and E. K. U. Gross. "Density-Functional Theory for Time-Dependent Systems". In: *Phys. Rev. Lett.* 52 (12 1984), pp. 997–1000. DOI: `10.1103/PhysRevLett.52.997`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.52.997`.

[174] Axel D. Becke. "Density-functional thermochemistry. III. The role of exact exchange". In: *J. Chem. Phys.* 98.7 (1993), pp. 5648–5652. ISSN: 0021-9606. DOI: `10.1063/1.464913`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/98/7/5648/19277469/5648\_1\_online.pdf`. URL: `https://doi.org/10.1063/1.464913`.

[175] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch. "Ab Initio Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields". In: *J. Phys. Chem* 98.45 (1994), pp. 11623–11627. DOI: `10.1021/j100096a001`. eprint: `https://doi.org/10.1021/j100096a001`. URL: `https://doi.org/10.1021/j100096a001`.

[176] Florian Weigend and Reinhart Ahlrichs. "Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy". In: *Phys. Chem. Chem. Phys.* 7 (18 2005), pp. 3297–3305. DOI: `10.1039/B508541A`. URL: `http://dx.doi.org/10.1039/B508541A`.

[177] D Andrae, U Häußermann, M Dolg, H Stoll, and H Preuß. "Energy-adjustedab initio pseudopotentials for the second and third row transition elements". In: *Theor. Chim. Acta* 77.2 (1990), pp. 123–141.

[178] F. Neese. "The ORCA program system". In: *WIRES Comput. Molec. Sci.* 2.1 (2012), pp. 73–78. DOI: `10.1002/wcms.81`.

[179] Kristine Pierloot, Birgit Dumez, Per-Olof Widmark, and Björn O Roos. "Density matrix averaged atomic natural orbital (ANO) basis sets for correlated molecular wave functions". In: *Theor. Chim. Acta* 90.2 (1995), pp. 87–114.

[180] Björn O Roos, Roland Lindh, Per-Åke Malmqvist, Valera Veryazov, and Per-Olof Widmark. "New relativistic ANO basis sets for transition metal atoms". en. In: *J. Phys. Chem. A* 109.29 (2005), pp. 6575–6579.

[181] Martin Feyereisen, George Fitzgerald, and Andrew Komornicki. "Use of approximate integrals in ab initio theory. An application in MP2 energy calculations". In: *Chem. Phys. Lett.* 208.5 (1993), pp. 359–363. ISSN: 0009-2614. DOI: `https://doi.org/10.1016/0009-2614(93)87156-W`. URL: `https://www.sciencedirect.com/science/article/pii/000926149387156W`.

[182] Marcus D. Hanwell, Donald E. Curtis, David C. Lonie, Tim Vandermeersch, Eva Zurek, and Geoffrey R. Hutchison. "Avogadro: an advanced semantic chemical editor, visualization, and analysis platform". In: *J. Cheminform* 4.1 (2012), p. 17. ISSN: 1758-2946. DOI: `10.1186/1758-2946-4-17`. URL: `https://doi.org/10.1186/1758-2946-4-17`.

[183] A. D. Becke. "Density-functional exchange-energy approximation with correct asymptotic behavior". In: *Phys. Rev. A* 38 (6 1988), pp. 3098–3100. DOI: `10.1103/PhysRevA.38.3098`. URL: `https://link.aps.org/doi/10.1103/PhysRevA.38.3098`.

[184] John P. Perdew. "Density-functional approximation for the correlation energy of the inhomogeneous electron gas". In: *Phys. Rev. B* 33 (12 1986), pp. 8822–8824. DOI: `10.1103/PhysRevB.33.8822`. URL: `https://link.aps.org/doi/10.1103/PhysRevB.33.8822`.

[185] M. J. Frisch et al. *Gaussian~16 Revision C.01.* Gaussian Inc. Wallingford CT. 2016.

[186] Matthias Ernzerhof and Gustavo E. Scuseria. "Assessment of the Perdew–Burke–Ernzerhof exchange-correlation functional". In: *J. Chem. Phys.* 110.11 (1999), pp. 5029–5036. ISSN: 0021-9606. DOI: `10.1063/1.478401`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/110/11/5029/19111374/5029\_1\_online.pdf`. URL: `https://doi.org/10.1063/1.478401`.

[187] Carlo Adamo and Vincenzo Barone. "Toward reliable density functional methods without adjustable parameters: The PBE0 model". In: *J. Chem. Phys.* 110.13 (1999), pp. 6158–6170. ISSN: 0021-9606. DOI: `10.1063/1.478522`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/110/13/6158/19068890/6158\_1\_online.pdf`. URL: `https://doi.org/10.1063/1.478522`.

[188] Thanh N. Truong and Eugene V. Stefanovich. “A new method for incorporating solvent effect into the classical, ab initio molecular orbital and density functional theory frameworks for arbitrary shape cavity”. In: *Chem. Phys. Lett.* 240.4 (1995), pp. 253–260. ISSN: 0009-2614. DOI: https://doi.org/10.1016/0009-2614(95)00541-B. URL: https://www.sciencedirect.com/science/article/pii/000926149500541B.

[189] Vincenzo Barone and Maurizio Cossi. “Quantum Calculation of Molecular Energies and Energy Gradients in Solution by a Conductor Solvent Model”. In: *J. Phys. Chem. A.* 102.11 (1998), pp. 1995–2001. ISSN: 1089-5639. DOI: 10.1021/jp9716997. URL: https://doi.org/10.1021/jp9716997.

[190] Dmitrij Rappoport and Filipp Furche. “Property-optimized Gaussian basis sets for molecular response calculations”. In: *J. Chem. Phys.* 133 (2010), p. 134105. DOI: 10.1063/1.3484283.

[191] Pavlo Golub, Andrej Antalik, Libor Veis, and Jiri Brabec. *Automatic selection of active spaces for strongly correlated systems using machine learning algorithms*. 2020. arXiv: 2011.14715 [physics.chem-ph]. URL: https://arxiv.org/abs/2011.14715.

[192] Lexin Ding, Stefan Knecht, Zoltán Zimborás, and Christian Schilling. “Quantum correlations in molecules: from quantum resourcing to chemical bonding”. In: *Quantum Sci. Technol.* 8.1 (2022), p. 015015. DOI: 10.1088/2058-9565/aca4ee. URL: https://dx.doi.org/10.1088/2058-9565/aca4ee.

[193] Francesco A. Evangelista. *Mutual Correlation*. 2025. arXiv: 2506.07344 [physics.chem-ph]. URL: https://arxiv.org/abs/2506.07344.

[194] HQS Quantum Simulations. *Active-Space Finder*. https://github.com/HQSquantumsimulations/ActiveSpaceFinder. 2025.

[195] Luning Zhao et al. *Quantum-Classical Auxiliary Field Quantum Monte Carlo with Matchgate Shadows on Trapped Ion Quantum Computers*. 2025. arXiv: 2506.22408 [quant-ph]. URL: https://arxiv.org/abs/2506.22408.

[196] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, and Jonathan Tennyson. “The Variational Quantum Eigensolver: A review of methods and best practices”. In: *Phys. Rep.* 986 (2022). The Variational Quantum Eigensolver: a review of methods and best practices, pp. 1–128. ISSN: 0370-1573. DOI: https://doi.org/10.1016/j.physrep.2022.08.003. URL: https://www.sciencedirect.com/science/article/pii/S0370157322003118.

[197] A. Yu. Kitaev. *Quantum measurements and the Abelian Stabilizer Problem*. 1995. arXiv: quant-ph/9511026 [quant-ph]. URL: https://arxiv.org/abs/quant-ph/9511026.

[198] Francesco Benfenati, Guglielmo Mazzola, Chiara Capecci, Panagiotis Kl. Barkoutsos, Pauline J. Ollitrault, Ivano Tavernelli, and Leonardo Guidoni. *Improved Accuracy on Noisy Devices by Nonunitary Variational Quantum Eigensolver for Chemistry Applications*. 2021. DOI: 10.1021/acs.jctc.1c00091. URL: https://doi.org/10.1021/acs.jctc.1c00091.

[199] M. Ganzhorn, D.J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp. "Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer". In: *Physical Review Applied* 11.4 (2019). ISSN: 2331-7019. DOI: `10.1103/physrevapplied.11.044092`. URL: `http://dx.doi.org/10.1103/PhysRevApplied.11.044092`.

[200] Davide Materia, Leonardo Ratini, Celestino Angeli, and Leonardo Guidoni. *Quantum Information Driven Ansatz (QIDA): Shallow-Depth Empirical Quantum Circuits from Quantum Chemistry*. 2024. DOI: `10.1021/acs.jpca.4c03756`. URL: `https://doi.org/10.1021/acs.jpca.4c03756`.

[201] Keita Kanno, Masaya Kohda, Ryosuke Imai, Sho Koh, Kosuke Mitarai, Wataru Mizukami, and Yuya O. Nakagawa. *Quantum-Selected Configuration Interaction: classical diagonalization of Hamiltonians in subspaces selected by quantum computers*. 2023. arXiv: `2302.11320 [quant-ph]`. URL: `https://arxiv.org/abs/2302.11320`.

[202] Ethan N. Epperly, Lin Lin, and Yuji Nakatsukasa. "A Theory of Quantum Subspace Diagonalization". In: *SIAM Journal on Matrix Analysis and Applications* 43.3 (2022), 1263–1290. ISSN: 1095-7162. DOI: `10.1137/21m145954x`. URL: `http://dx.doi.org/10.1137/21M145954X`.

[203] Javier Robledo-Moreno, Mario Motta, Holger Haas, Ali Javadi-Abhari, Petar Jurcevic, William Kirby, Simon Martiel, Kunal Sharma, Sandeep Sharma, Tomonori Shirakawa, Iskandar Sitdikov, Rong-Yang Sun, Kevin J. Sung, Maika Takita, Minh C. Tran, Seiji Yunoki, and Antonio Mezzacapo. *Chemistry Beyond Exact Solutions on a Quantum-Centric Supercomputer*. 2024. arXiv: `2405.05068 [quant-ph]`. URL: `https://arxiv.org/abs/2405.05068`.

[204] Ludwig Nützel, Alexander Gresch, Lukas Hehn, Lucas Marti, Robert Freund, Alex Steiner, Christian D Marciniak, Timo Eckstein, Nina Stockinger, Stefan Wolf, Thomas Monz, Michael Kühn, and Michael J Hartmann. "Solving an industrially relevant quantum chemistry problem on quantum hardware". In: *Quantum Science and Technology* 10.1 (2025), p. 015066. DOI: `10.1088/2058-9565/ad9ed3`. URL: `https://dx.doi.org/10.1088/2058-9565/ad9ed3`.

[205] Gian-Luca R Anselmetti, David Wierichs, Christian Gogolin, and Robert M Parrish. "Local, expressive, quantum-number-preserving VQE ansätze for fermionic systems". In: *New Journal of Physics* 23.11 (2021), p. 113010. DOI: `10.1088/1367-2630/ac2cb3`. URL: `https://dx.doi.org/10.1088/1367-2630/ac2cb3`.

[206] Fabio Tarocco, Davide Materia, Leonardo Ratini, and Leonardo Guidoni. "Compact multi-threshold quantum information driven ansatz for strongly interactive lattice spin models". In: *Journal of Physics A: Mathematical and Theoretical* 58.16 (2025), p. 165302. DOI: `10.1088/1751-8121/adc4a1`. URL: `https://doi.org/10.1088/1751-8121/adc4a1`.

[207] Fabio Tarocco, Davide Materia, Leonardo Ratini, and Leonardo Guidoni. *Multi-QIDA method for VQE state preparation in molecular systems*. 2025. arXiv: `2508.11270 [quant-ph]`. URL: `https://arxiv.org/abs/2508.11270`.

[208] Farrokh Vatan and Colin Williams. "Optimal quantum circuits for general two-qubit gates". In: *Phys. Rev. A* 69 (3 2004), p. 032315. DOI: `10.1103/PhysRevA.69.032315`. URL: `https://link.aps.org/doi/10.1103/PhysRevA.69.032315`.

[209] Per-Olov Löwdin and Harrison Shull. “Natural Orbitals in the Quantum Theory of Two-Electron Systems”. In: *Phys. Rev.* 101 (6 1956), pp. 1730–1739. DOI: `10.1103/PhysRev.101.1730`. URL: `https://link.aps.org/doi/10.1103/PhysRev.101.1730`.

[210] Leonardo Ratini, Chiara Capecci, Francesco Benfenati, and Leonardo Guidoni. “Wave Function Adapted Hamiltonians for Quantum Computing”. In: *Journal of Chemical Theory and Computation* 18.2 (2022). PMID: 35041784, pp. 899–909. DOI: `10.1021/acs.jctc.1c01170`. eprint: `https://doi.org/10.1021/acs.jctc.1c01170`. URL: `https://doi.org/10.1021/acs.jctc.1c01170`.

[211] Jawed A. Jafri and Jerry L. Whitten. “Iterative natural orbitals for configuration interaction using perturbation theory”. In: *Theoretica Chimica Acta* 44 (3 1977), pp. 305–313. ISSN: 0040-5744. DOI: `10.1007/BF00551172`.

[212] Leonardo Ratini, Chiara Capecci, and Leonardo Guidoni. “Optimization strategies in WAHTOR algorithm for quantum computing empirical ansatz: a comparative study”. In: *Electronic Structure* 5.4 (2023), p. 045006. DOI: `10.1088/2516-1075/ad018e`. URL: `https://dx.doi.org/10.1088/2516-1075/ad018e`.

[213] Leonardo Ratini, Chiara Capecci, and Leonardo Guidoni. “Natural Orbitals and Sparsity of Quantum Mutual Information”. In: *Journal of Chemical Theory and Computation* 20.9 (2024), 3535–3542. ISSN: 1549-9626. DOI: `10.1021/acs.jctc.3c01325`. URL: `http://dx.doi.org/10.1021/acs.jctc.3c01325`.

[214] Tyler Takeshita, Nicholas C. Rubin, Zhang Jiang, Eunseok Lee, Ryan Babbush, and Jarrod R. McClean. “Increasing the Representation Accuracy of Quantum Simulations of Chemistry without Extra Quantum Resources”. In: *Phys. Rev. X* 10 (1 2020), p. 011004. DOI: `10.1103/PhysRevX.10.011004`. URL: `https://link.aps.org/doi/10.1103/PhysRevX.10.011004`.

[215] Jules Tilly, P. V. Sriluckshmy, Akashkumar Patel, Enrico Fontana, Ivan Rungger, Edward Grant, Robert Anderson, Jonathan Tennyson, and George H. Booth. “Reduced density matrix sampling: Self-consistent embedding and multiscale electronic structure on current generation quantum computers”. In: *Phys. Rev. Res.* 3 (3 2021), p. 033230. DOI: `10.1103/PhysRevResearch.3.033230`. URL: `https://link.aps.org/doi/10.1103/PhysRevResearch.3.033230`.

[216] Igor O. Sokolov, Panagiotis Kl. Barkoutsos, Pauline J. Ollitrault, Donny Greenberg, Julia Rice, Marco Pistoia, and Ivano Tavernelli. “Quantum orbital-optimized unitary coupled cluster methods in the strongly correlated regime: Can quantum algorithms outperform their classical equivalents?” In: *The Journal of Chemical Physics* 152.12 (2020), p. 124107. ISSN: 0021-9606. DOI: `10.1063/1.5141835`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.5141835/15573242/124107_1_online.pdf`. URL: `https://doi.org/10.1063/1.5141835`.

[217] Wataru Mizukami, Kosuke Mitarai, Yuya O. Nakagawa, Takahiro Yamamoto, Tennin Yan, and Yu-ya Ohnishi. “Orbital optimized unitary coupled cluster theory for quantum computer”. In: *Phys. Rev. Res.* 2 (3 2020), p. 033421. DOI: `10.1103/PhysRevResearch.2.033421`. URL: `https://link.aps.org/doi/10.1103/PhysRevResearch.2.033421`.

[218] Keita Omiya, Yuya O. Nakagawa, Sho Koh, Wataru Mizukami, Qi Gao, and Takao Kobayashi. “Analytical Energy Gradient for State-Averaged Orbital-Optimized Variational Quantum Eigensolvers and Its Application to a Photochemical Reaction”. In: *Journal of Chemical Theory and Computation* 18.2 (2022). PMID: 35060747, pp. 741–748. DOI: `10.1021/acs.jctc.1c00877`. eprint: `https://doi.org/10.1021/acs.jctc.1c00877`. URL: `https://doi.org/10.1021/acs.jctc.1c00877`.

[219] Juan Angel de Gracia Triviño, Mickael G. Delcey, and Göran Wendin. “Complete Active Space Methods for NISQ Devices: The Importance of Canonical Orbital Optimization for Accuracy and Noise Resilience”. In: *Journal of Chemical Theory and Computation* 19.10 (2023). PMID: 37103120, pp. 2863–2872. DOI: `10.1021/acs.jctc.3c00123`. eprint: `https://doi.org/10.1021/acs.jctc.3c00123`. URL: `https://doi.org/10.1021/acs.jctc.3c00123`.

[220] Joel Bierman, Yingzhou Li, and Jianfeng Lu. “Improving the Accuracy of Variational Quantum Eigensolvers with Fewer Qubits Using Orbital Optimization”. In: *Journal of Chemical Theory and Computation* 19.3 (2023). PMID: 36696487, pp. 790–798. DOI: `10.1021/acs.jctc.2c00895`. eprint: `https://doi.org/10.1021/acs.jctc.2c00895`. URL: `https://doi.org/10.1021/acs.jctc.2c00895`.

[221] Shigeki Gocho, Hajime Nakamura, Shu Kanno, Qi Gao, Takao Kobayashi, Taichi Inagaki, and Miho Hatanaka. *Excited state calculations using variational quantum eigensolver with spin-restricted ansätze and automatically-adjusted constraints.* 2022. arXiv: `2110.14448 [quant-ph]`. URL: `https://arxiv.org/abs/2110.14448`.

[222] Aaron Fitzpatrick, Anton Nykänen, N. Walter Talarico, Alessandro Lunghi, Sabrina Maniscalco, Guillermo García-Pérez, and Stefan Knecht. *A self-consistent field approach for the variational quantum eigensolver: orbital optimization goes adaptive.* 2022. arXiv: `2212.11405 [quant-ph]`.

[223] Saad Yalouz, Bruno Senjean, Jakob Günther, Francesco Buda, Thomas E O’Brien, and Lucas Visscher. “A state-averaged orbital-optimized hybrid quantum–classical algorithm for a democratic description of ground and excited states”. In: *Quantum Sci. Technol.* 6.2 (2021), p. 024004.

[224] Fabio Tarocco, Davide Materia, Leonardo Ratini, Chiara Capecci, and Leonardo Guidoni. *Quantum @ L’Aquila*. `https://gitlab.com/leonardoguidoni/quaq`. 2025.

[225] 1930-2007 Cotton F. Albert (Frank Albert). *Chemical applications of group theory.* New York (State), United States: New York : Wiley, 1990. ISBN: 0471510947; 9780471510949. URL: `https://searchworks.stanford.edu/view/1362247`.

[226] W. J. Hehre, R. Ditchfield, and J. A. Pople. “Self—Consistent Molecular Orbital Methods. XII. Further Extensions of Gaussian—Type Basis Sets for Use in Molecular Orbital Studies of Organic Molecules”. In: *The Journal of Chemical Physics* 56.5 (1972), pp. 2257–2261. ISSN: 0021-9606. DOI: `10.1063/1.1677527`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/56/5/2257/18878569/2257_1_online.pdf`. URL: `https://doi.org/10.1063/1.1677527`.

[227] Jr. Bauschlicher Charles W. and Peter R. Taylor. "Benchmark full configuration-interaction calculations on H2O, F, and F-". In: *The Journal of Chemical Physics* 85.5 (1986), pp. 2779–2783. ISSN: 0021-9606. DOI: 10.1063/1.451034. eprint: https://pubs.aip.org/aip/jcp/article-pdf/85/5/2779/18961316/2779_1_online.pdf. URL: https://doi.org/10.1063/1.451034.

[228] Péter G. Szalay and Rodney J. Bartlett. "Approximately extensive modifications of the multireference configuration interaction method: A theoretical and practical analysis". In: *The Journal of Chemical Physics* 103.9 (1995), pp. 3600–3612. ISSN: 0021-9606. DOI: 10.1063/1.470243. eprint: https://pubs.aip.org/aip/jcp/article-pdf/103/9/3600/19296194/3600_1_online.pdf. URL: https://doi.org/10.1063/1.470243.

[229] Mihály Kállay, Péter G. Szalay, and Péter R. Surján. "A general state-selective multireference coupled-cluster algorithm". In: *The Journal of Chemical Physics* 117.3 (2002), pp. 980–990. ISSN: 0021-9606. DOI: 10.1063/1.1483856. eprint: https://pubs.aip.org/aip/jcp/article-pdf/117/3/980/19192251/980_1_online.pdf. URL: https://doi.org/10.1063/1.1483856.

[230] Takeshi Yanai and Garnet Kin-Lic Chan. "Canonical transformation theory from extended normal ordering". In: *The Journal of Chemical Physics* 127.10 (2007), p. 104107. ISSN: 0021-9606. DOI: 10.1063/1.2761870. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.2761870/14913671/104107_1_online.pdf. URL: https://doi.org/10.1063/1.2761870.

[231] Sanghamitra Das, Debashis Mukherjee, and Mihály Kállay. "Full implementation and benchmark studies of Mukherjee's state-specific multireference coupled-cluster ansatz". In: *The Journal of Chemical Physics* 132.7 (2010), p. 074103. ISSN: 0021-9606. DOI: 10.1063/1.3310288. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.3310288/15775079/074103_1_online.pdf. URL: https://doi.org/10.1063/1.3310288.

[232] Francesco A. Evangelista and Jürgen Gauss. "An orbital-invariant internally contracted multireference coupled cluster approach". In: *The Journal of Chemical Physics* 134.11 (2011), p. 114102. ISSN: 0021-9606. DOI: 10.1063/1.3559149. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.3559149/13525085/114102_1_online.pdf. URL: https://doi.org/10.1063/1.3559149.

[233] Ilya G. Ryabinkin, Tzu-Ching Yen, Scott N. Genin, and Artur F. Izmaylov. "Qubit Coupled Cluster Method: A Systematic Approach to Quantum Chemistry on a Quantum Computer". In: *Journal of Chemical Theory and Computation* 14.12 (2018). PMID: 30427679, pp. 6317–6326. DOI: 10.1021/acs.jctc.8b00932. eprint: https://doi.org/10.1021/acs.jctc.8b00932. URL: https://doi.org/10.1021/acs.jctc.8b00932.

[234] Emmanuel Giner, Celestino Angeli, Yann Garniron, Anthony Scemama, and Jean-Paul Malrieu. "A Jeziorski-Monkhorst fully uncontracted multi-reference perturbative treatment. I. Principles, second-order versions, and tests on ground state potential energy curves". In: *The Journal of Chemical Physics* 146.22 (2017), p. 224108. ISSN: 0021-9606. DOI: 10.1063/1.4984616. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.4984616/15526578/224108_1_online.pdf. URL: https://doi.org/10.1063/1.4984616.

[235] Anna Engels-Putzka and Michael Hanrath. “Dissociating N2: a multi-reference coupled cluster study on the potential energy surfaces of ground and excited states”. In: *Molecular Physics* 107.2 (2009), pp. 143–155. DOI: `10.1080/00268970902724922`. eprint: `https://doi.org/10.1080/00268970902724922`. URL: `https://doi.org/10.1080/00268970902724922`.

[236] Qiskit contributors. *Qiskit: An Open-source Framework for Quantum Computing*. 2023. DOI: `10.5281/zenodo.2573505`.

[237] J J Boor. *Ziegler-Natta Catalysts Polymerizations*. Academic Press, 2012.

[238] G Odian. *Principles of Polymerization*. Wiley, 2004.

[239] Erik P. Bierwagen, John E. Bercaw, and W. A. III Goddard. “Theoretical Studies of Ziegler-Natta Catalysis: Structural Variations and Tacticity Control”. In: *Journal of the American Chemical Society* 116.4 (1994), pp. 1481–1489. DOI: `10.1021/ja00083a037`. eprint: `https://doi.org/10.1021/ja00083a037`. URL: `https://doi.org/10.1021/ja00083a037`.

[240] P Cossee. “Ziegler-Natta catalysis I. Mechanism of polymerization of

$$\alpha$$

-olefins with Ziegler-Natta catalysts”. In: *Journal of Catalysis* 3.1 (1964), pp. 80–88. ISSN: 0021-9517. DOI: `https://doi.org/10.1016/0021-9517(64)90095-8`. URL: `https://www.sciencedirect.com/science/article/pii/0021951764900958`.

[241] E.J Arlman. “Ziegler-Natta catalysis II. Surface structure of layer-lattice transition metal chlorides”. In: *Journal of Catalysis* 3.1 (1964), pp. 89–98. ISSN: 0021-9517. DOI: `https://doi.org/10.1016/0021-9517(64)90096-X`. URL: `https://www.sciencedirect.com/science/article/pii/002195176490096X`.

[242] E.J Arlman and P Cossee. “Ziegler-Natta catalysis III. Stereospecific polymerization of propene with the catalyst system TiCl3-AlEt3”. In: *Journal of Catalysis* 3.1 (1964), pp. 99–104. ISSN: 0021-9517. DOI: `https://doi.org/10.1016/0021-9517(64)90097-1`. URL: `https://www.sciencedirect.com/science/article/pii/0021951764900971`.

[243] Shannon S. Stahl. “Organotransition Metal Chemistry: From Bonding to Catalysis”. In: *Journal of the American Chemical Society* 132.24 (2010), pp. 8524–8525. DOI: `10.1021/ja103695e`. eprint: `https://doi.org/10.1021/ja103695e`. URL: `https://doi.org/10.1021/ja103695e`.

[244] Robert O. Symcox and Paul. Ehrlich. “The Free Radical, High Pressure Polymerization of Ethylenes. The Effects of Initiator Concentration, Monomer Concentration and Pressure on the Polymerization Rate”. In: *Journal of the American Chemical Society* 84.4 (1962), pp. 531–536. DOI: `10.1021/ja00863a006`. eprint: `https://doi.org/10.1021/ja00863a006`. URL: `https://doi.org/10.1021/ja00863a006`.

[245] M. Asteasuain, S. Pereda, M. H. Lacunza, P. E. Ugrin, and A. Brandolin. “Industrial high pressure ethylene polymerization initiated by peroxide mixtures: A reduced mathematical model for parameter adjustment”. In: *Polymer Engineering & Science* 41.5 (2001), pp. 711–726. DOI: `https://doi.org/10.1002/pen.10767`. eprint: `https://4spepublications.`

onlinelibrary.wiley.com/doi/pdf/10.1002/pen.10767. URL: https://4spepublications.onlinelibrary.wiley.com/doi/abs/10.1002/pen.10767.

[246] Horst Weiss, Michael Ehrig, and Reinhart Ahlrichs. "Ethylene insertion in the homogeneous Ziegler-Natta catalysis: an ab initio investigation on a correlated level". en. In: *Journal of the American Chemical Society* 116.11 (June 1994), pp. 4919–4928. ISSN: 0002-7863, 1520-5126. DOI: 10.1021/ja00090a042. URL: https://pubs.acs.org/doi/abs/10.1021/ja00090a042 (visited on 05/14/2024).

[247] Nahal Nassabeh, Mark Tran, and Patrick E. Fleming. "Dissociation of the Ethyl Radical: An Exercise in Computational Chemistry". In: *Journal of Chemical Education* 91.8 (2014), pp. 1248–1253. DOI: 10.1021/ed4007748. eprint: https://doi.org/10.1021/ed4007748. URL: https://doi.org/10.1021/ed4007748.

[248] P. M. Goldfeder, V. A. Volpert, V. M. Ilyashenko, A. M. Khan, J. A. Pojman, and S. E. Solovyov. "Mathematical Modeling of Free-Radical Polymerization Fronts". In: *The Journal of Physical Chemistry B* 101.18 (1997), pp. 3474–3482. DOI: 10.1021/jp962150v. eprint: https://doi.org/10.1021/jp962150v. URL: https://doi.org/10.1021/jp962150v.

[249] Hans Hellmann. "Zur Einführung in die Quantenchemie". In: *Deuticke, Leipzig* (1937).

[250] R. P. Feynman. "Forces in Molecules". In: *Phys. Rev.* 56 (4 1939), pp. 340–343. DOI: 10.1103/PhysRev.56.340. URL: https://link.aps.org/doi/10.1103/PhysRev.56.340.

[251] P. Pulay. "Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules". In: *Molecular Physics* 17.2 (1969), pp. 197–204. DOI: 10.1080/00268976900100941. eprint: https://doi.org/10.1080/00268976900100941. URL: https://doi.org/10.1080/00268976900100941.

[252] P. PULAY. "Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. I. Theory". In: *Molecular Physics* 100.1 (2002), pp. 57–62. DOI: 10.1080/00268970110088884. eprint: https://doi.org/10.1080/00268970110088884. URL: https://doi.org/10.1080/00268970110088884.

[253] Robert M. Parrish, Edward G. Hohenstein, Jarrod R. McClean, and Todd J. Martínez. "Hybrid Quantum/Classical Derivative Theory: Analytical Gradients and Excited-State Dynamics for the Multistate Contracted Variational Quantum Eigensolver". In: *arXiv preprint arXiv:1906.08728* (2019).

[254] Thomas E. O'Brien, Michael Streif, Nicholas C. Rubin, Raffaele Santagati, Yuan Su, William J. Huggins, Joshua J. Goings, Nikolaj Moll, Elica Kyoseva, Matthias Degroote, Christofer S. Tautermann, Joonho Lee, Dominic W. Berry, Nathan Wiebe, and Ryan Babbush. "Efficient quantum computation of molecular forces and other energy gradients". In: *Phys. Rev. Res.* 4 (4 2022), p. 043210. DOI: 10.1103/PhysRevResearch.4.043210. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.4.043210.

[255] Alain Delgado, Juan Miguel Arrazola, Soran Jahangiri, Zeyue Niu, Josh Izaac, Chase Roberts, and Nathan Killoran. "Variational quantum algorithm for molecular geometry optimization". In: *Phys. Rev. A* 104 (5 2021), p. 052402. DOI: 10.1103/PhysRevA.104.052402. URL: https://link.aps.org/doi/10.1103/PhysRevA.104.052402.

[256] Emiel Koridon, Joana Fraxanet, Alexandre Dauphin, Lucas Visscher, Thomas E. O'Brien, and Stefano Polla. "A hybrid quantum algorithm to detect conical intersections". In: *Quantum* 8 (2024), p. 1259. ISSN: 2521-327X. DOI: 10.22331/q-2024-02-20-1259. URL: https://doi.org/10.22331/q-2024-02-20-1259.

[257] Nonia Vaquero-Sabater, Abel Carreras, Román Orús, Nicholas J. Mayhall, and David Casanova. "Physically Motivated Improvements of Variational Quantum Eigensolvers". In: *Journal of Chemical Theory and Computation* 20.12 (2024). PMID: 38853416, pp. 5133–5144. DOI: 10.1021/acs.jctc.4c00329. eprint: https://doi.org/10.1021/acs.jctc.4c00329. URL: https://doi.org/10.1021/acs.jctc.4c00329.

[258] Hugh G. A. Burton, Daniel Marti-Dafcik, David P. Tew, and David J. Wales. "Exact electronic states with shallow quantum circuits from global optimisation". In: *npj Quantum Information* 9.1 (2023), p. 75. DOI: 10.1038/s41534-023-00744-2. URL: https://doi.org/10.1038/s41534-023-00744-2.

[259] Weitang Li, Zigeng Huang, Changsu Cao, Yifei Huang, Zhigang Shuai, Xiaoming Sun, Jinzhao Sun, Xiao Yuan, and Dingshun Lv. "Toward practical quantum embedding simulation of realistic chemical systems on near-term quantum computers". In: *Chem. Sci.* 13 (31 2022), pp. 8953–8962. DOI: 10.1039/D2SC01492K. URL: http://dx.doi.org/10.1039/D2SC01492K.

[260] Yajie Hao, Qiming Ding, Xiaoting Wang, and Xiao Yuan. "Large-scale Efficient Molecule Geometry Optimization with Hybrid Quantum-Classical Computing". In: (2025). arXiv: 2509.07460 [quant-ph]. URL: https://arxiv.org/abs/2509.07460.

[261] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. "The theory of variational hybrid quantum-classical algorithms". In: *New Journal of Physics* 18.2 (2016), p. 023023. ISSN: 1367-2630. DOI: 10.1088/1367-2630/18/2/023023. URL: http://dx.doi.org/10.1088/1367-2630/18/2/023023.

[262] Mickaël G. Delcey. "MultiPsi: A python-driven MCSCF program for photochemistry and spectroscopy simulations on modern HPC environments". In: *WIREs Computational Molecular Science* 13.6 (2023), e1675. DOI: https://doi.org/10.1002/wcms.1675. eprint: https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1675. URL: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1675.

[263] J. Schirmer, L. S. Cederbaum, and O. Walter. "New approach to the one-particle Green's function for finite Fermi systems". In: *Phys. Rev. A* 28 (3 1983), pp. 1237–1259. DOI: 10.1103/PhysRevA.28.1237. URL: https://link.aps.org/doi/10.1103/PhysRevA.28.1237.

[264] Dirk R. Rehn, Zilvinas Rinkevicius, Michael F. Herbst, Xin Li, Maximilian Scheurer, Manuel Brand, Adrian L. Dempwolff, Iulia E. Brumboiu, Thomas Fransson, Andreas Dreuw, and Patrick Norman. "Gator: A Python-driven program for spectroscopy simulations using correlated wave functions". In: *WIREs Computational Molecular Science* 11.6 (2021), e1528. DOI: https://doi.org/10.1002/wcms.1528. eprint: https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1528. URL: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1528.

[265] H. van de Waterbeemd, R. E. Carter, G. Grassy, H. Kubinyi, Y. C. Martin, M. S. Tute, and P. Willett. "Glossary of terms used in computational drug design (IUPAC Recommendations 1997)". In: *Pure and Applied Chemistry* 69.5 (1997), pp. 1137–1152. DOI: `doi:10.1351/pac199769051137`. URL: `https://doi.org/10.1351/pac199769051137`.

[266] Stephen Heller, Alan McNaught, Stephen Stein, Dmitrii Tchekhovskoi, and Igor Pletnev. "InChI - the worldwide chemical structure identifier standard". en. In: *J. Cheminform.* 5.1 (2013), p. 7.

[267] Ingo Muegge and Prasenjit Mukherjee. "An overview of molecular fingerprint similarity search in virtual screening". en. In: *Expert Opin. Drug Discov.* 11.2 (2016), pp. 137–148.

[268] Wendy A. Warr. "Representation of chemical structures". In: *WIREs Computational Molecular Science* 1.4 (2011), pp. 557–579. DOI: `https://doi.org/10.1002/wcms.36`. eprint: `https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.36`. URL: `https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.36`.

[269] Robert W. Layer. "The Chemistry of Imines." In: *Chemical Reviews* 63.5 (1963), pp. 489–510. DOI: `10.1021/cr60225a003`. eprint: `https://doi.org/10.1021/cr60225a003`. URL: `https://doi.org/10.1021/cr60225a003`.

[270] Matthew E. Belowich and J. Fraser Stoddart. "Dynamic imine chemistry". In: *Chem. Soc. Rev.* 41 (6 2012), pp. 2003–2024. DOI: `10.1039/C2CS15305J`. URL: `http://dx.doi.org/10.1039/C2CS15305J`.

[271] Maria Ciaccia and Stefano Di Stefano. "Mechanisms of imine exchange reactions in organic solvents". In: *Org. Biomol. Chem.* 13 (3 2015), pp. 646–654. DOI: `10.1039/C4OB02110J`. URL: `http://dx.doi.org/10.1039/C4OB02110J`.

[272] Jr. Pearson Richard and Frank J. Lovas. "Microwave spectrum and molecular structure of methylenimine (CH2NH)". In: *The Journal of Chemical Physics* 66.9 (May 1977), pp. 4149–4156. ISSN: 0021-9606. DOI: `10.1063/1.434490`. eprint: `https://pubs.aip.org/aip/jcp/article-pdf/66/9/4149/18905919/4149_1_online.pdf`. URL: `https://doi.org/10.1063/1.434490`.

[273] V. Vuitton, R. V. Yelle, and V. G. Anicich. "The Nitrogen Chemistry of Titan's Upper Atmosphere Revealed". In: *The Astrophysical Journal* 647.2 (2006), p. L175. DOI: `10.1086/507467`. URL: `https://doi.org/10.1086/507467`.

[274] Christian Gampe and Vishal A Verma. "Curse or cure? A perspective on the developability of aldehydes as active pharmaceutical ingredients". en. In: *J. Med. Chem.* 63.23 (2020), pp. 14357–14381.

[275] Sriram Mahesh, Kuei-Chien Tang, and Monika Raj. "Amide bond activation of biological molecules". en. In: *Molecules* 23.10 (2018), p. 2615.

[276] Guangrong Meng, Jin Zhang, and Michal Szostak. "Acyclic Twisted Amides". In: *Chemical Reviews* 121.20 (2021). PMID: 34406005, pp. 12746–12783. DOI: `10.1021/acs.chemrev.1c00225`. eprint: `https://doi.org/10.1021/acs.chemrev.1c00225`. URL: `https://doi.org/10.1021/acs.chemrev.1c00225`.

[277] "amides". In: (2025). DOI: doi:10.1351/goldbook.A00266. URL: https://doi.org/10.1351/goldbook.A00266.

[278] Peter W. Seavill and Jonathan D. Wilden. "The preparation and applications of amides using electrosynthesis". In: *Green Chem.* 22 (22 2020), pp. 7737–7759. DOI: 10.1039/D0GC02976A. URL: http://dx.doi.org/10.1039/D0GC02976A.

[279] Hosea Cheung, Robin S. Tanke, and G. Paul Torrence. "Acetic Acid". In: *Ullmann's Encyclopedia of Industrial Chemistry.* John Wiley and Sons, Ltd, 2011. ISBN: 9783527306732. DOI: https://doi.org/10.1002/14356007.a01_045.pub2. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/14356007.a01_045.pub2. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/14356007.a01_045.pub2.

[280] Javier Robledo-Moreno, Mario Motta, Holger Haas, Ali Javadi-Abhari, Petar Jurcevic, William Kirby, Simon Martiel, Kunal Sharma, Sandeep Sharma, Tomonori Shirakawa, Iskandar Sitdikov, Rong-Yang Sun, Kevin J. Sung, Maika Takita, Minh C. Tran, Seiji Yunoki, and Antonio Mezzacapo. "Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer". In: *Science Advances* 11.25 (2025). ISSN: 2375-2548. DOI: 10.1126/sciadv.adu9991. URL: http://dx.doi.org/10.1126/sciadv.adu9991.

[281] Fabio Tarocco, Pi A. B. Haase, Fabijan Pavošević, Vijay Krishna, Leonardo Guidoni, Stefan Knecht, and Martina Stella. *AEGISS – Atomic orbital and Entropy-based Guided Inference for Space Selection – A novel semi-automated active space selection workflow for quantum chemistry and quantum computing applications.* 2025. arXiv: 2508.10671 [physics.chem-ph]. URL: https://arxiv.org/abs/2508.10671.

[282] Francesco Aquilante et al. "Molcas 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table". In: *Journal of Computational Chemistry* 37.5 (2016), pp. 506–541. DOI: https://doi.org/10.1002/jcc.24221. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.24221. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.24221.

[283] Yihan Shao et al. "Advances in molecular quantum chemistry contained in the Q-Chem 4 program package". In: *Molecular Physics* 113.2 (2015), pp. 184–215. DOI: 10.1080/00268976.2014.952696. eprint: https://doi.org/10.1080/00268976.2014.952696. URL: https://doi.org/10.1080/00268976.2014.952696.

[284] Juan Miguel Arrazola, Olivia Di Matteo, Nicolás Quesada, Soran Jahangiri, Alain Delgado, and Nathan Killoran. "Universal quantum circuits for quantum chemistry". In: *Quantum* 6 (June 2022), p. 742. ISSN: 2521-327X. DOI: 10.22331/q-2022-06-20-742. URL: http://dx.doi.org/10.22331/q-2022-06-20-742.

[285] Lila Cadi Tazi and Alex J. W. Thom. "Folded Spectrum VQE: A Quantum Computing Method for the Calculation of Molecular Excited States". In: *Journal of Chemical Theory and Computation* 20.6 (2024), 2491–2504. ISSN: 1549-9626. DOI: 10.1021/acs.jctc.3c01378. URL: http://dx.doi.org/10.1021/acs.jctc.3c01378.

[286] Mateusz Ostaszewski, Edward Grant, and Marcello Benedetti. "Structure optimization for parameterized quantum circuits". In: *Quantum* 5 (2021), p. 391. ISSN: 2521-327X. DOI: 10.22331/q-2021-01-28-391. URL: https://doi.org/10.22331/q-2021-01-28-391.

[287] César Feniou, Muhammad Hassan, Baptiste Claudon, Axel Courtat, Olivier Adjoua, Yvon Maday, and Jean-Philip Piquemal. “Greedy gradient-free adaptive variational quantum algorithms on a noisy intermediate scale quantum computer”. In: *Scientific Reports* 15.1 (May 2025). ISSN: 2045-2322. DOI: `10.1038/s41598-025-99962-1`. URL: `http://dx.doi.org/10.1038/s41598-025-99962-1`.

[288] Koichi Momma and Fujio Izumi. “*VESTA*: a three-dimensional visualization system for electronic and structural analysis”. In: *J. Appl. Crystallogr.* 41.3 (2008), pp. 653–658. DOI: `10.1107/S0021889808012016`. URL: `https://doi.org/10.1107/S0021889808012016`.

[289] Daniel Claudino. *PsiEmb*. `https://github.com/danclaudino/PsiEmbed`. 2020.

[290] Huanchen Zhai, Henrik R Larsson, Seunghoon Lee, Zhi-Hao Cui, Tianyu Zhu, Chong Sun, Linqing Peng, Ruojing Peng, Ke Liao, Johannes Tölle, et al. “Block2: A Comprehensive Open Source Framework to Develop and Apply State-Of-The-Art DMRG Algorithms in Electronic Structure and Beyond”. In: *J. Chem. Phys.* 159.23 (2023).

Appendix A

Chapter 02 Appendix: Active Space Selection

1 Computational Details

Pre-processing

AEGISS has been implemented in Python. Following the workflow, the atom-labeling step is software/package agnostic, and it is done with the help of any molecular visualization tool (e.g., Avogadro[182], VESTA[288], etc.), by adding the relative cluster identification, $\mathcal{C} \in \{\mathcal{R}\}$ to each associated atom. For the initial guess (and eventually needed pre-processing), the calculation of the starting MOs ($[\mathbf{C}]_{\mu i}$ or $[\mathbf{C}^{\mathbf{A}}]_{\mu i}$) is computed using PySCF [94]. We choose PySCF because of its ease of use, integrability with other Python packages, availability of Δ-SCF method, embedding[289], and direct compatibility with the Python implementation of the DMRG algorithm.

DMRG-setup

All the DMRG calculations and reference wavefunctions have been computed using Block2 [290] with its Python DMRG implementation. The computation of the single orbital entropies has been defined from Block2 functionalities, in particular, with the method `DMRGDriver.get_orbital_entropies()` from the resulting MPS (i.e., approximated ground state at a given bond-dimension). For all the systems, the approximated calculation has been obtained from a DMRG run with fixed bond-dimension, `m`, number of sweeps `n_sweeps = 5`, `SymmetryTypes.SZ`, and a single root, `n_roots = 1`. At this step, it is also possible to compute the reference calculation on multiple roots, and then maintain the union of the MOs pre-selected according to Single orbital entropies for each root.

AO projection

For the AO projection, the selection of the atomic orbitals in each group $D \in \mathcal{D}$ is done exploiting the function `search_ao_labels(`D`)`. This function allows only the indices of the desired AO contributions in the full expansion. The overlap matrix , $[\mathbf{S}]_{\nu\mu}$, and the overlap projector, $[\mathbf{S}^{D}]_{\eta\mu}$, are obtained using built-in PySCF functions. In particular, the non-orthogonal mixed-basis overlap matrix $[\mathbf{S}]_{\nu\mu}$ is defined using `Mole.intor('int1e_cross')`, which returns the *cross-integral* between two different basis MO coefficient matrices. The AO contributions are associated to one or more

AO-labels. The AO-labels can be either single AO contributions or groups of them. A single label is composed by "X_c AO_s", where X is the atom, c is the cluster identifier, and AO_s is the atomic orbital contribution.

Reference calculation and final assessment

After selecting the MOs to be included in the final active spaces, the reference calculations for the two smaller molecules have been obtained from PySCF state-averaged CASSCF. For TLD-1433, the reference calculation is instead obtained from DMRG-CI, using the Block2 DMRG implementation. For the computational setup for the DMRG calculations, and since all systems exhibit a closed-shell singlet ground state, the highest symmetry can be used i.e. `SymmetryTypes.SU2`, with bond-dimension `m = 1000`, `n_sweeps = 100`, and a noise schedule set according to `noises = [1e-4] * n_sweeps //2 + [1e-5] * n_sweeps // 2 + [0]`.

2 Trans-Cl Benchmarks

2.1 TDDFT

The performance of several TD-DFT flavors has been assessed as we need a general method which allows describing the different transitions contributing to the UV-vis spectrum. While it is known that solvent effects are found to be mandatory to obtain spectroscopic accuracy, especially in the case of MLCT states, for the investigation of this prototype complex we only performed gas phase calculations to the purpose of benchmarking TDDFT calculations against CASSCF calculations based on the active spaces we found. The results of the TD-DFT calculations are collected in Table A.1 for the singlet states, with also the oscillator strengths, and in Table A.2 for the triplet states.

As a general trend, we see our calculations being in good agreement with the reference [168]. As a result, we expect our findings to be aligned with the description that MC transitions are rather robust to any of the functionals tested while MLCT states are only well described with functionals bearing intermediate amounts of exact exchange, i.e. PBE0 and B3LYP, possibly in combination with solvent effects. IL states are also best described with these functionals.

	PBE		PBE0		B3LYP		CAM-B3LYP	
State	ΔE	f	ΔE	f	ΔE	f	ΔE	f
S_1	1.395	*0.0003*	2.471	*0.0003*	2.282	*0.0003*	3.303	*0.0002*
S_2	1.56	***0.0122***	2.58	***0.0141***	2.40	***0.0132***	3.38	*0.0000*
S_3	2.164	*0.0028*	3.301	*0.0000*	3.225	*0.0025*	3.41	***0.0191***
S_4	2.265	*0.0008*	3.377	*0.0031*	3.252	*0.0000*	3.472	*0.0038*
S_5	2.375	*0.0000*	3.466	*0.0024*	3.284	*0.0002*	4.131	*0.0041*
S_6	2.386	*0.0000*	3.529	*0.0007*	3.306	*0.0012*	4.292	*0.0000*
S_7	2.394	*0.0105*	3.538	*0.0000*	3.139	*0.0000*	4.41	*0.0007*
S_8	2.469	*0.0044*	3.547	*0.0014*	3.347	*0.0028*	4.461	*0.0000*
S_9	3.078	*0.0000*	3.683	*0.0084*	3.439	*0.0083*	4.49	*0.0008*
S_{10}	3.137	*0.0016*	3.750	*0.0025*	3.506	*0.0027*	4.582	*0.0021*
S_{11}	3.149	*0.0005*	3.887	*0.0048*	3.791	*0.0045*	4.673	*0.0079*

Table A.1: Electronic transition energies (in eV) and oscillator strength of the first 11 singlets states from TD-DFT calculation performed on Trans-Cl complex in GAS PHASE. Functionals: PBE, PBE0, B3LYP, and CAM-B3LYP. Basis-set used 6-311G* for all the atoms except for Ruthenium. On Ruthenium def2-TZVP basis has been used with quasi-relativistic Stuttgard-Dresden pseudopotential.

State	PBE	PBE0	B3LYP	CAM-B3LYP
T_1	1.368	2.439	2.254	2.847
T_2	1.453	2.509	2.324	2.981
T_3	2.144	2.753	2.774	3.212
T_4	2.233	2.875	2.888	3.327
T_5	2.347	3.271	3.207	3.429
T_6	2.353	3.447	3.271	3.633
T_7	2.364	3.458	3.277	3.832
T_8	2.439	3.495	3.278	4.207
T_9	2.636	3.503	3.285	4.231
T_{10}	2.741	3.512	3.409	4.380
T_{11}	3.034	3.651	3.440	4.394

Table A.2: Electronic transition energies (in eV) of the first 11 triplet states from TD-DFT calculation performed on Trans-Cl complex in GAS PHASE. Functionals: PBE, PBE0, B3LYP and CAM-B3LYP. Basis-set used 6-311G* for all the atoms except for Ruthenium. On Ruthenium def2-TZVP basis has been used with quasi-relativistic Stuttgard-Dresden pseudopotential.

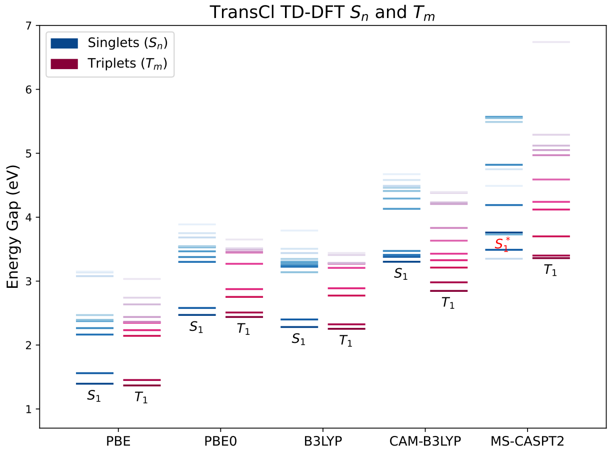

Figure A.1: Diagram of S_n and T_m vertical excitation states for Trans-Cl complex obtained from TD-DFT calculation with different functionals (PBE, PBE0, B3LYP, and CAM-B3LYP) and reference benchmark MS-CASPT2. Singlet states computed by MS-CASPT2 are subject to a state reordering.

Finally, as highlighted in Figure A.1, it is clear that TDDFT alone provides a strongly functional-dependent description of the excitation spectra and great caution should be used when employing this tool for predicting UV-vis spectra.

3 Ruby-3 Benchmarks

3.1 TDDFT

For the TD-DFT calculations, the exchange-correlation functionals used in this comparative study are selected from the reference paper [169]. The XC functionals selected are BP86, B3LYP, and CAM-B3LYP. For these calculations, we employed the same basis set as for the optimizations. No TDA has been used in this investigation as it has been shown in previous work not to have a meaningful effect on the results. Also in this case, while solvent effects might be necessary for thorough comparison with experiments, this is out of the scope of the present investigation as this molecule is not our main target system. As a result, we will present gas-phase calculations only for Ruby3.

The calculated TD-DFT with no TDA vertical excitation energies in gas-phase with BP86, B3LYP and CAM-B3LYP are shown in Table A.3 and Table A.4.

We find our calculations to show excellent agreement with the work of Gonzales et. al, where the main findings are that BP86 severely underestimates the excitation energies in comparison to MS-CASPT2. The inclusion of HF exact exchange in B3LYP seems to reduce the underestimation of the excitation energies, with the S_1 being underestimated by 0.11 eV and the T_2 by 0.38 eV in the MS-CASPT2 case. Furthermore, it is worth mentioning that MS-CASPT2 predicts an inverted order of T_1 being higher in energy than S_1. In all cases the triplets are further underestimated than the singlet excited states but ordering is always correct with TD-DFT. Cam-B3LYP results, that we have also carried out in this work, seem to overestimate CASSCF results from all functionals. We will compare our TDDFT results with CASSCF results we are going to present herein.

	BP86		B3LYP		CAM-B3LYP	
State	ΔE	f	ΔE	f	ΔE	f
S_1	1.851	*0.0018*	2.508	*0.0013*	3.135	*0.0001*
S_2	2.015	*0.0001*	2.522	*0.0001*	3.136	*0.0001*
S_3	2.015	*0.0001*	2.523	*0.0001*	3.207	*0.0006*
S_4	2.203	*0.0179*	2.697	*0.0001*	3.298	*0.0000*
S_5	2.206	*0.0179*	2.720	*0.0102*	3.363	*0.0021*
S_6	2.275	*0.0002*	2.723	*0.0100*	3.366	*0.0017*
S_7	2.497	*0.0732*	2.903	*0.1332*	3.429	*0.1737*
S_8	2.498	***0.0734***	2.903	***0.1335***	3.430	***0.1740***
S_9	2.502	*0.0097*	3.173	*0.0000*	3.773	*0.0000*

Table A.3: Electronic transition energies (in eV) and oscillator strength of the first 9 singlet states from TD-DFT calculation performed on Ruby-3 complex in GAS PHASE. Functionals: BP86, B3LYP and CAM-B3LYP. Basis-set used def2-TZVP for all the atoms. On Ruthenium has been used a quasi-relativistic Stuttgard-Dresden pseudopotential for ECP.

State	BP86	B3LYP	CAM-B3LYP
T_1	1.790	2.367	2.867
T_2	1.885	2.368	2.868
T_3	1.885	2.410	2.920
T_4	2.066	2.430	2.940
T_5	2.088	2.518	3.021
T_6	2.089	2.519	3.023
T_7	2.151	2.625	3.227
T_8	2.152	2.627	3.228
T_9	2.243	2.659	3.256

Table A.4: Electronic transition energies (in eV) of the first 9 triplets states from TD-DFT calculation performed on Ruby-3 complex in GAS PHASE. Functionals: BP86, B3LYP and CAM-B3LYP. Basis-set used def2-TZVP for all the atoms. On Ruthenium has been used a quasi-relativistic Stuttgard-Dresden pseudopotential for ECP.

3.2 SA-CASSCF

The complete tables for the CASSCF calculations ran using the selected actives spaces are reported below.

4 TLD1433 Benchmarks

In our study, we optimized the ground-state geometry for the conformation with alternated S-Ring and opposite H-S-Position, namely TLD1433-3(G), obtaining a new local minima of the geometry in

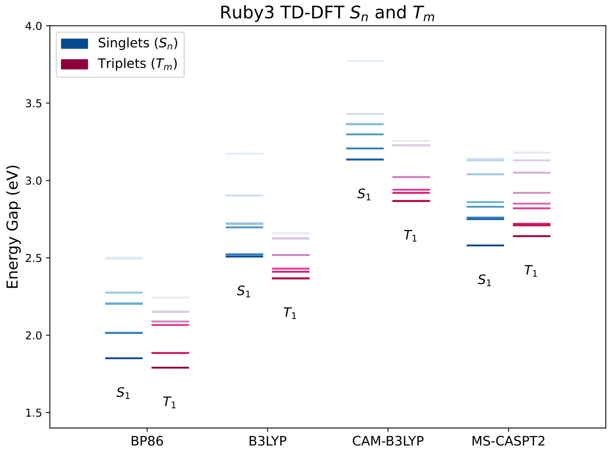

Figure A.2: Diagram of S_n and T_m vertical excitation states for Ruby-3 complex obtained from TDDFT calculation with different functionals (BP86, B3LYP, and CAM-B3LYP) and MS-CASPT2 used in the reference paper as benchmark.

State	AS1 ΔE	AS2 ΔE	AS3 ΔE
S_1	2.4498	2.1372	1.8233
S_2	2.4520	2.4521	2.1381
S_3	2.4521	2.4581	2.3243
S_4	2.8565	2.5375	2.3968
S_5	2.8567	2.6266	2.4458
S_6	2.8726	2.7949	2.6824
S_7	2.9702	2.8498	2.7117
S_8	2.9703	2.882	2.7189

Table A.5: Singlet excited state transition energies (in eV) obtained from SA-CASSCF calculations on the Ruby-3 complex.

State	AS1 ΔE	AS2 ΔE	AS3 ΔE
T_1	2.363	2.0699	1.7686
T_2	2.3632	2.3417	1.8357
T_3	2.4018	2.36	2.0833
T_4	2.6707	2.4091	2.2579
T_5	2.7658	2.4657	2.2984
T_6	2.7659	2.6137	2.5516
T_7	2.8212	2.7224	2.5609
T_8	2.8212	2.7591	2.6643
T_9	2.8327	2.8208	2.6946

Table A.6: Triplet excited state transition energies (in eV) obtained from SA-CASSCF calculations on the Ruby-3 complex.

both gas and solvent. Our structure, TLD1433-3(G), is lower in energy with respect to the optimal structure of the reference paper, TLD1433-4(G), of 0.042 mHa in gas phase and 1.84 mHa in H_2O.

4.1 TDDFT

In reference [170], a series of preliminary calculations using various exchange-correlation density functionals on smaller Ru compounds in methanol, with available experimental data, showed that PBE0 accurately reproduces the geometrical parameters. Therefore, we selected it as the most suitable XC functional for ground-state molecular optimizations. However, the computed absorption spectra indicate that M06 significantly outperforms the other XC functionals, particularly at longer wavelengths. The excellent performance of M06 has also been noted in previous studies on transition-metal complexes. To evaluate the reliability of our data, we focused on comparing our results with those of Alberto et al[170]. In their study, they presented M06 DFT results for TDL1433, and we find that our results are in excellent agreement with theirs. However, in this work, in particular to gain a deeper understanding of the excitations of this molecular system as well as to be able to accurately select the active space for our quantum circuit simulations, we decided to go beyond existing literature calculations. Specifically, we explore the effect of different exchange-correlation functionals on the excitations, including M062x (27% exact HF exchange), which has twice the exact HF exchange of M06 (15% exact HF exchange).

In previous systems, we have consistently tested and reported results both with and without the TD approximation. For this particular case, we have also run calculations with both options. However, to simplify the discussion and given the limited experimental data available to definitively confirm the accuracy of one approach over the other, we find that using the TD approximation is not only acceptable but also offers the advantage of faster computation at this stage. This choice is based on the practical consideration that the TD approximation simplifies the calculations without introducing significant discrepancies, allowing us to make reasonable progress while acknowledging the limitations posed by the lack of sufficient experimental validation. We also provide a table to compare the results for both approaches. Still, from now on, we will primarily present results based on the TDA-based approximation. Furthermore, the main focus and results of current benchmark investigation is on calculations with implicit solvent in water.

	PBE0		B3LYP		CAM-B3LYP		M06	
State	ΔE	f	ΔE	f	ΔE	f	ΔE	f
S_1	2.625	*0.0008*	2.473	*0.0011*	3.133	***1.9516***	2.387	*0.0008*
S_2	2.709	*0.0017*	2.542	*0.0013*	3.134	*0.3315*	2.462	*0.0016*
S_3	2.718	*0.0334*	2.563	*0.0597*	3.215	*0.0141*	2.467	*0.0006*
S_4	2.733	***0.9679***	2.569	***0.4472***	3.249	*0.0014*	2.569	*0.0042*
S_5	2.798	*0.0080*	2.658	*0.0102*	3.300	*0.0011*	2.620	***0.4991***
S_6	2.873	*0.0093*	2.717	*0.0073*	3.352	*0.0299*	2.639	*0.0191*
S_7	2.886	*0.1651*	2.720	*0.0122*	3.388	*0.0121*	2.705	*0.4908*
S_8	2.921	*1.2190*	2.764	*1.2701*	3.417	*0.0007*	2.780	*0.1848*
S_9	2.995	*0.1885*	2.816	*0.4539*	3.429	*0.2406*	2.868	*0.8037*
S_{10}	2.998	*0.0298*	2.831	*0.1042*	3.674	*0.0461*	2.895	*0.3348*
S_{11}	3.066	*0.0027*	2.874	*0.0500*	3.805	*0.0032*	2.941	*0.1182*
S_{12}	3.138	*0.0175*	2.907	*0.0995*	3.956	*0.0568*	3.076	*0.0023*
S_{13}	3.164	*0.0059*	2.972	*0.0062*	4.009	*0.0030*	3.112	*0.0031*
S_{14}	3.288	*0.0532*	3.081	*0.0932*	4.071	*0.0777*	3.153	*0.0142*
S_{15}	3.327	*0.0137*	3.134	*0.0151*	4.101	*0.0003*	3.161	*0.0086*

Table A.7: Electronic transition energies (in eV) and oscillator strength of the first 15 singlet states from TD-DFT calculation performed on TLD1433 complex in implicit solvent WATER with TDA approximation. Functionals: PBE0, B3LYP, CAM-B3LYP, and M06. Basis-set used def2-TZVP for all the atom and quasi-relativistic Stuttgard-Dresden pseudopotential on Ruthenium.

We report electronic transition energies and oscillator strength for singlet and triplet states for calculations in aqueous solution are reported in Tables A.7, A.8 and A.9, respectively.

Let us start from inspecting the character of the first bright singlet states (higher oscillator strengths) for all the functionals. We find that not all functionals predict the first singlet state to be the first bright state but all functionals, except M06, show that the first bright singlet has a stronger HOMO-LUMO character (cf Figure A.4). As shown in Figure A.3, DFT consistently predicts HOMO to be solely localised on the tail while LUMO has both contributions from the ligands around the metal and smaller contribution on the tail, suggesting this excitation can be

State	M062x W	*f*	M062x G	*f*
S_1	3.182	*2.3199*	2.278	*0.0442*
S_2	3.381	*0.0011*	2.455	*0.0005*
S_3	3.442	*0.0021*	2.517	*0.0125*
S_4	3.513	*0.0058*	2.663	*0.3428*
S_5	3.521	*0.0105*	3.048	*0.6417*
S_6	3.570	*0.0132*	3.313	*0.1432*
S_7	3.613	*0.0077*	3.372	*0.0043*
S_8	3.621	*0.0023*	3.428	*0.0013*
S_9	3.657	*0.0124*	3.445	*0.0006*
S_{10}	3.686	*0.0152*	3.544	*0.0019*
S_{11}	3.706	*0.1835*	3.588	*0.0026*
S_{12}	3.712	*0.0137*	3.600	*0.0002*
S_{13}	3.795	*0.0916*	3.610	*0.0017*
S_{14}	3.798	*0.0052*	3.652	*0.0009*
S_{15}	4.006	*0.0426*	3.661	*0.0017*

Table A.8: Electronic transition energies (in eV) of the first 15 singlet states from TD-DFT calculation performed on TLD1433 complex in implicit solvent WATER (W) and GAS (G) phase with TDA approximation. Functionals M062x. Basis-set used def2-TZVP for all the atom and quasi-relativistic Stuttgard-Dresden pseudopotential on Ruthenium.

State	PBE0	B3LYP	CAM-B3LYP	M06	M062x
T_1	2.007	1.974	2.141	2.059	2.349
T_2	2.464	2.344	2.820	2.274	3.037
T_3	2.465	2.361	2.850	2.339	3.092
T_4	2.549	2.419	2.911	2.345	3.114
T_5	2.580	2.435	2.965	2.369	3.151
T_6	2.617	2.492	3.011	2.439	3.243
T_7	2.651	2.533	3.033	2.463	3.307
T_8	2.694	2.569	3.076	2.509	3.318
T_9	2.722	2.583	3.192	2.577	3.344
T_{10}	2.749	2.645	3.269	2.588	3.355
T_{11}	2.802	2.653	3.280	2.690	3.375
T_{12}	2.818	2.676	3.327	2.720	3.431
T_{13}	2.914	2.796	3.356	2.884	3.465
T_{14}	3.047	2.832	3.539	2.896	3.498
T_{15}	3.110	2.868	3.546	3.071	3.544

Table A.9: Electronic transition energies (in eV) of the first 15 triplet states from TD-DFT calculation performed on TLD1433 complex in implicit solvent WATER with TDA approximation. Functionals: PBE0, B3LYP, CAM-B3LYP, and M062x. Basis-set used def2-TZVP for all the atom and quasi-relativistic Stuttgard-Dresden pseudopotential on Ruthenium.

classified as ligand to ligand. To gain a deeper understanding of the character of the excitation we look into the natural transition orbitals (NTOs) for the first bright state. In Figure A.4 we report the NTOs for the bright transition singlet and most relevant contributions (for simplicity we are only showing M06, CAM-B3LYP and M062x, but B3LYP shows same behaviour as CAM-B3LYP and M062x). We observe that, focusing only on M06 (see also Ref. [170]), gives a limited picture of this excitation. Indeed, CAM-B3LYP, M062x, and B3LYP predict a higher contribution of the tail ligand and a lower one coming from metal and bipyridines.

Concerning the investigation of the triplet excited-state manifold, we decided that for the purpose and relevance of the current work, only the first triplet state, T_1, is investigated. From the analysis of the excitation character carried out on various functionals we find that T1 is mostly a tail-to-tail excitation with contributions from the bridging bipyridine.

In the cases of CAM-B3LYP and M062x we observe that the effect of long-range interactions and higher exact exchange pushed up the bright state, which in these cases is also the S_1. In general, it is important to point out that, although more computationally demanding, for a heavy metal like ruthenium, these types of functionals can better handle the metal's electron distribution and strong correlation effects, leading to improved accuracy in describing bond strengths, electronic structure, and reactivity in such large complexes.

For simplicity and because of a satisfactory agreement with the limited experimental data available, we will mainly refer to CAM-B3LYP as our reference exchange correlation functional in this benchmark work.

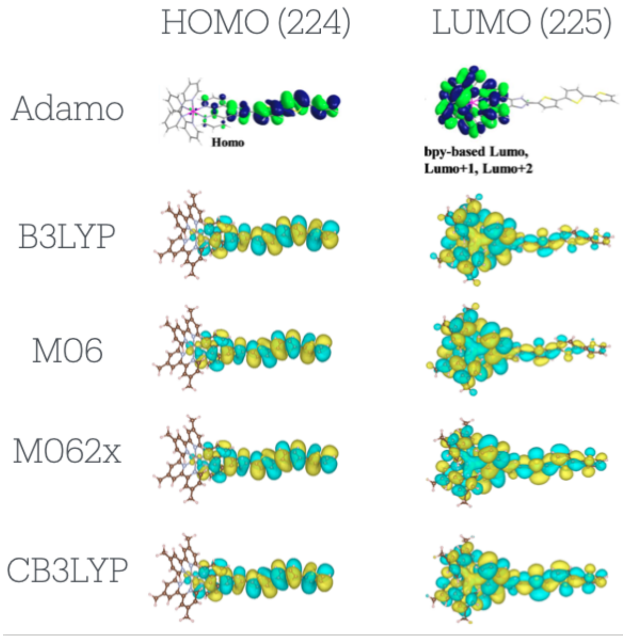

Figure A.3: HOMO and LUMO orbitals as computed by DFT and compared to Ref. [170].

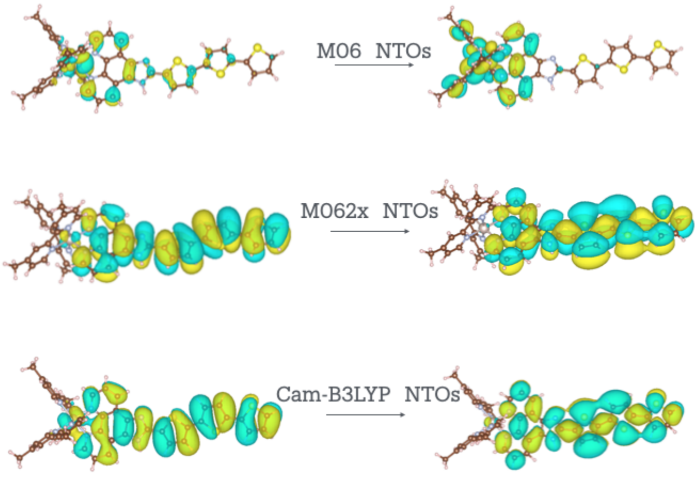

Figure A.4: Natural Transition Orbitals for S_1 computed with PBE0, M06, CAM-B3LYP and M062x.

Appendix B

Chapter 03 Appendix: Multi-QIDA results

1 Spin Systems

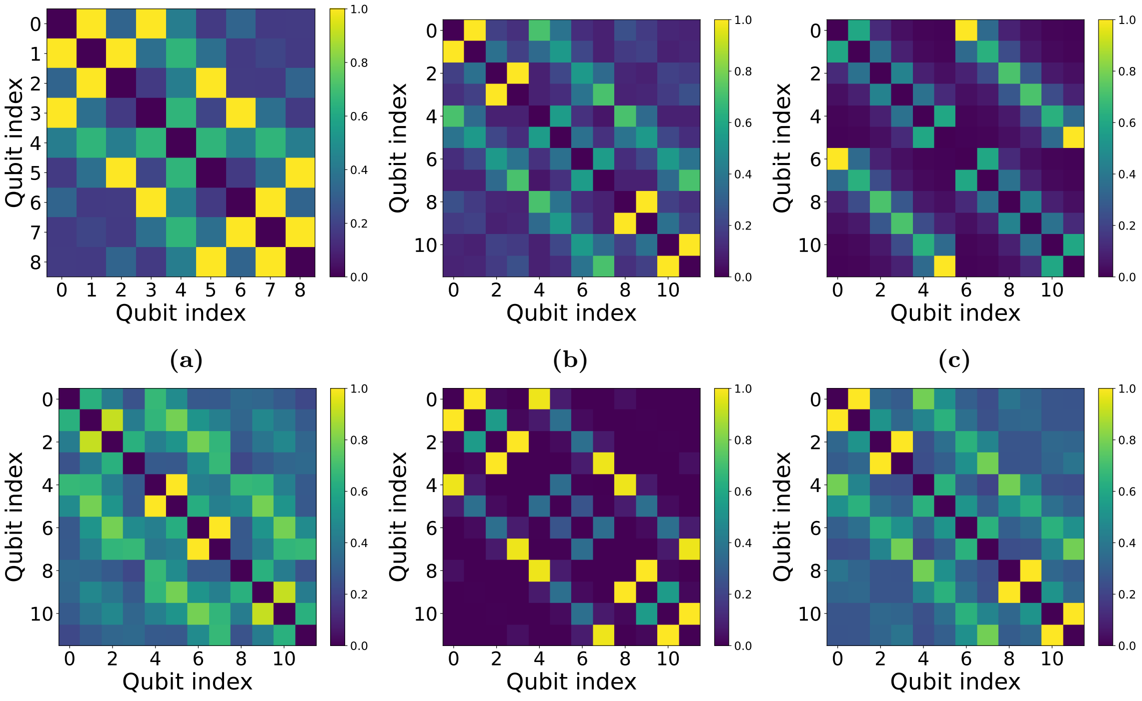

(a) (b) (c) (d) (e) (f)

Figure B.1: QMI matrices obtained by the tested systems. Maps (a)-(b)-(c) refer to isotropic Heisenberg systems, (d) to isotropic Heisenberg with external magnetic field $h = 2$, (e)-(f) to anisotropic Heisenberg with $\Delta = 1/10$ and $\Delta = 2/3$, respectively.

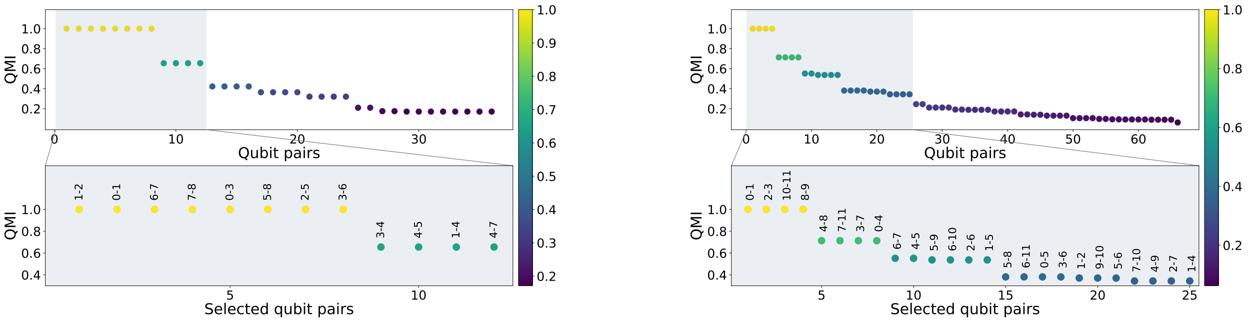

(a) 3×3 isotropic Heisenberg Hamiltonian QMI value pairing selection.

(b) 3×4 isotropic Heisenberg Hamiltonian QMI value pairing selection.

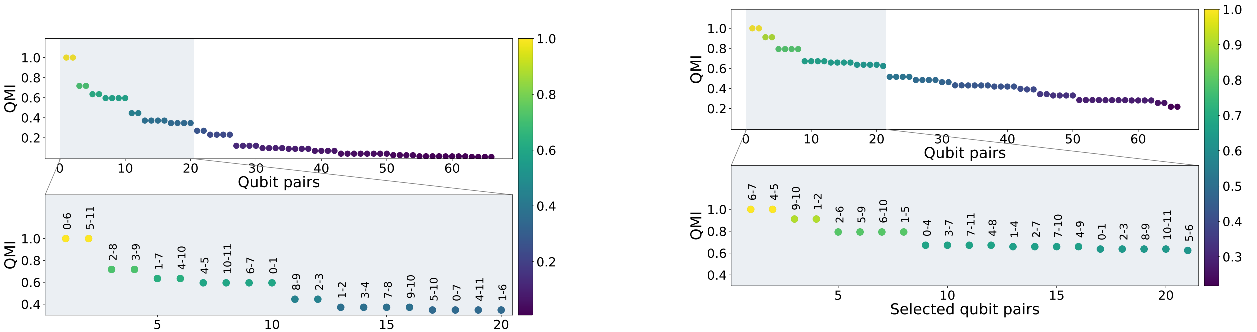

(c) 2×6 isotropic Heisenberg Hamiltonian QMI pairings selection.

(d) 3×4 isotropic Heisenberg Hamiltonian with external magnetic field $h = 2$ QMI entangling pairings selection.

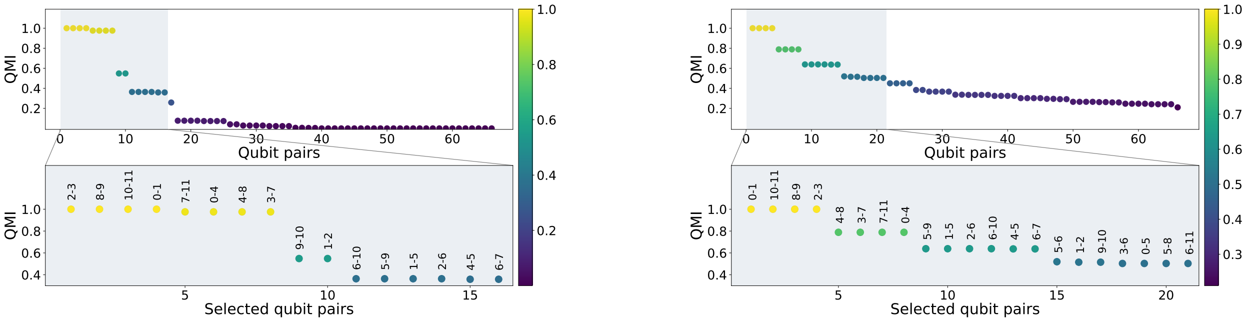

(e) Anisotropic 3×4 Heisenberg Hamiltonian with $\Delta = 1/10$ QMI value pairing selection.

(f) Anisotropic 3×4 Heisenberg Hamiltonian with $\Delta = 2/3$ QMI value pairing selection.

Figure B.2: Qubit–Pairs selection plot. For each image, the higher plot shows all the qubit-pairings in the system. The lower plot is a close-up on the couplers that have been selected by the first phase of the Multi-QIDA layers-building procedure. The qubit-pairs obtained from Algorithm 1 are shown in Table B.1.

System	*Layer 1*	*Layer 2*	*Layer 2*	*Layer 4*
3×3	$[0,1],[0,3],[1,2],[2,5],$ $[5,8],[7,8],[6,7],[3,6]$	$[1,4],[3,4],[4,5],[4,7]$	-	-
2×6	$[0,6],[1,7],[2,8],[3,9],$ $[4,10],[5,11]$(*)	$[0,1],[4,5],[6,7],[10,11]$	$[2,3],[8,9]$	$[1,2],[3,4],[7,8],[9,10]$
3×4	$[0,1],[2,3],[8,9],[10,11]$	$[0,4],[3,7],[4,8],[7,11]$	$[1,5],[2,6],[4,5],[5,9],$ $[6,7],[6,10]$	$[1,2],[5,6],[9,10]$
$3 \times 4, h = 2$	$[1,2],[4,5],[6,7],[9,10]$	$[1,5],[2,6],[5,9],[6,10]$	$[0,1],[0,4],[2,3],[3,7],$ $[4,8],[7,11],[8,9],[10,11]$	-
$3 \times 4, \Delta = 1/10$	$[0,1],[0,4],[2,3],[3,7],$ $[4,8],[7,11],[8,9],[10,11]$	$[1,2],[9,10]$	$[1,5],[2,6],[4,5],[5,9],$ $[6,7],[6,10]$	-
$3 \times 4, \Delta = 2/3$	$[0,1],[2,3],[8,9],[10,11]$	$[0,4],[3,7],[4,8],[7,11]$	$[1,5],[2,6],[4,5],[6,7],$ $[5,9],[6,10]$	$[1,2],[5,6],[9,10]$

Table B.1: Entangling map for each Multi-QIDA layer divided by system configurations. (*) in this configuration, the grouping optimization is used allowing the contraction of 3 layers into a unique layer. The selection has been done using a finesse ratio set to $\mu = 0.1$. Upon these layers, each Multi-QIDA configuration is completed with an additional ladder layer. In $(QIDA)^{CX}$ configuration, every layer after the first one is applied in V-shape i.e. only the first half of the full additional layers is shown.

	3×3	2×6	3×4	3×4 $h = 2$	3×4 $\Delta = 2/3$	3×4 $\Delta = 1/10$
$(L)_4^{CX}$	32	44	44	44	44	44
$(L)_5^{CX}$	40	55	55	55	55	55
$(L)_6^{CX}$	–	–	66	–	66	–
$(QIDA)^{CX}$	32	48	52	50	52	46
$(QIDA)^{SO4}$	40	54	56	54	56	54

Table B.2: Number of CNOTs for every system and configuration.$(QIDA)^{CX}$ and $(QIDA)^{SO4}$ represent QIDA configuration with CNOTs and $\mathbf{SO}(4)$ gates,respectively, while $(L)_d^{CX}$ describe ladder ansatz with depth d and CNOT as entangling gates.

	3×3	2×6	3×4	3×4 $h = 2$	3×4 $\Delta = 2/3$	3×4 $\Delta = 1/10$
$(L)_4^{CX}$	211	393	338	506	426	480
$(L)_5^{CX}$	314	456	427	730	623	791
$(L)_6^{CX}$	-	-	538	-	666	-
$(QIDA)^{CX}$	274	799	825	682	1082	1578
$(QIDA)^{SO4}$	638	1011	1073	1592	1161	1115

Table B.3: Average number of iterations required for convergence.

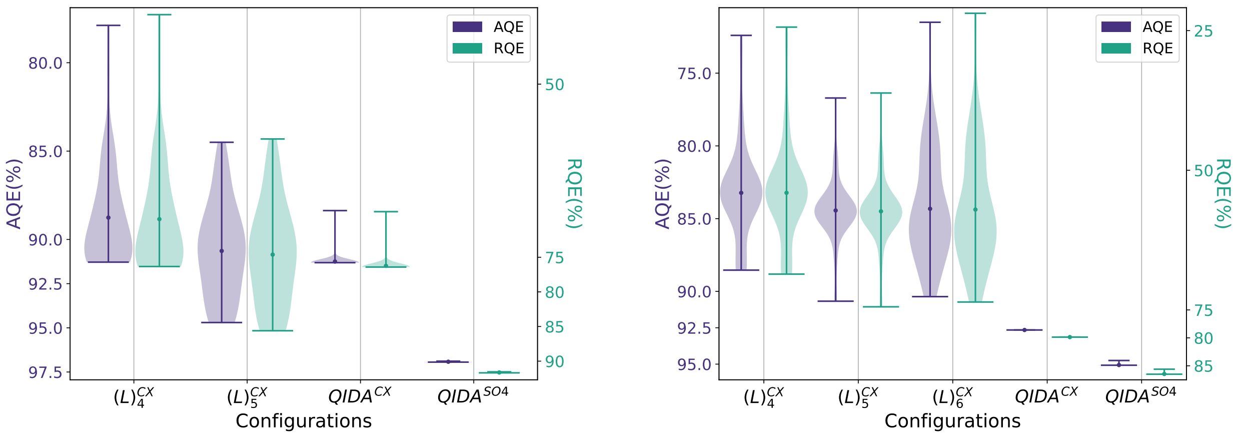

(a) 3×3 Isotropic Heisenberg model system results. HEA: 32,40. Multi-QIDA: 32,40.

(b) 3×4 Isotropic Heisenberg model system results. HEA: 44,55,66. Multi-QIDA: 52,56.

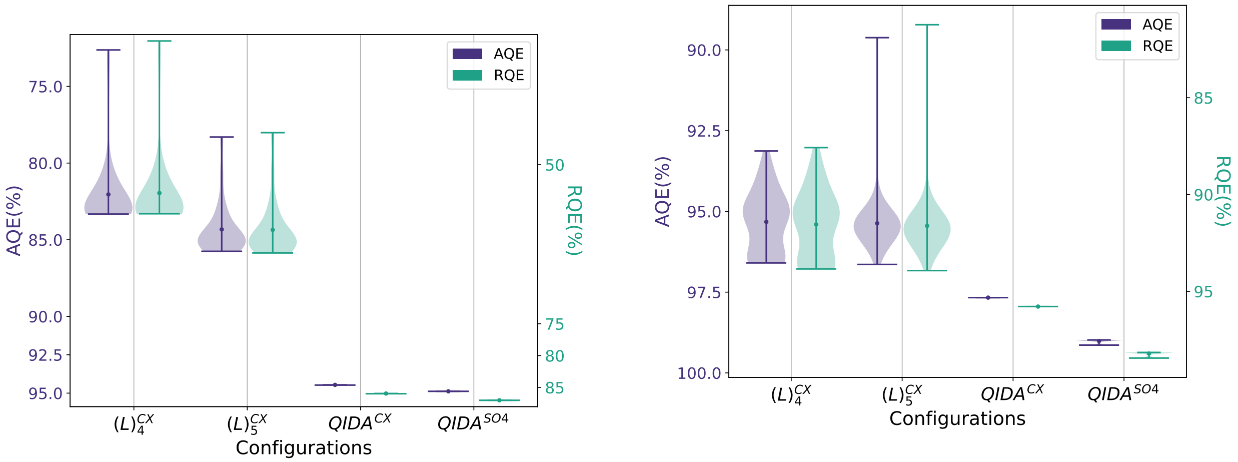

(c) 2×6 Isotropic Heisenberg model system results. HEA: 44,55. Multi-QIDA: 48,54.

(d) 3×3 Isotropic Heisenberg model with External magnetic field $h = 2$ results. HEA: 44,55. Multi-QIDA: 50,54.

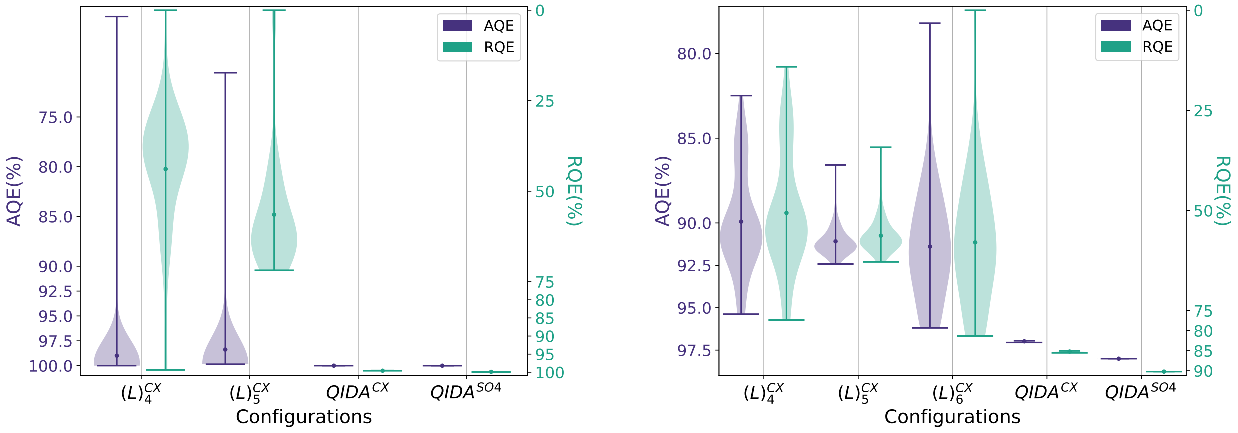

(e) 3×4 Anisotropic Heisenberg model with $\Delta = 1/10$ results.
HEA: 44,55. Multi-QIDA: 46,54.

(f) 3×4 Anisotropic Heisenberg model with $\Delta = 2/3$ results.
HEA: 44,55,66. Multi-QIDA: 52,56.

Figure B.3: Comparison between *AQE* and *RQE* for different systems and ansatz configuration. The number of *CNOT*s are showed in the caption for both hardware-efficient ansätze (HEA) and for multi-QIDA ansätze. (a)-(b)-(c) refer to isotropic Heisenberg systems, (d) to isotropic Heisenberg with external magnetic field $h = 2$, (e)-(f) to anisotropic Heisenberg with $\Delta = 1/10$ and $\Delta = 2/3$, respectively.

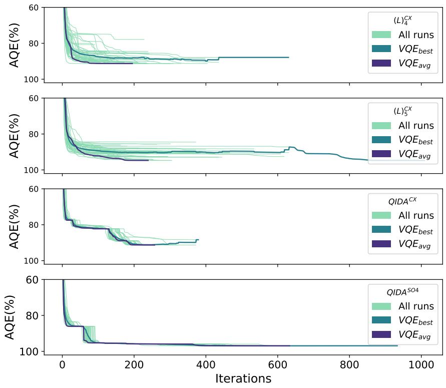

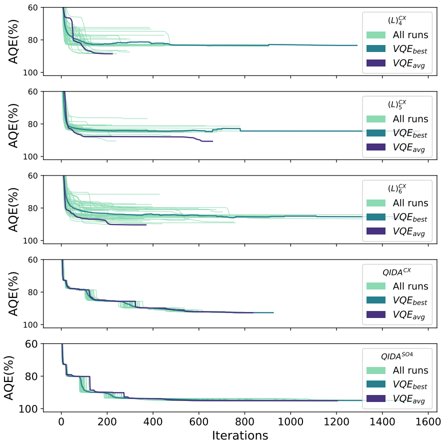

(a) 3 × 3 isotropic Heisenberg Hamiltonian. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

(b) 3 × 4 isotropic Heisenberg Hamiltonian. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $(L)_6^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

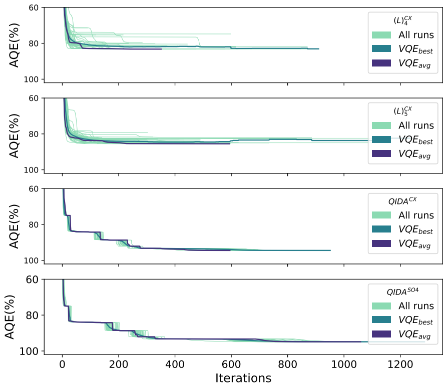

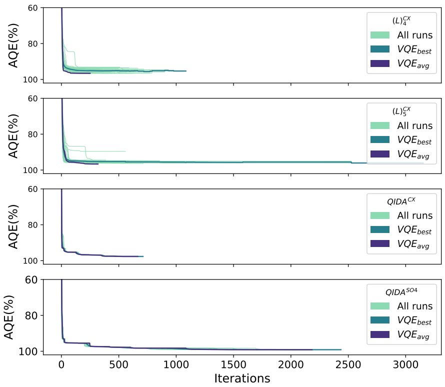

(c) 2 × 6 isotropic Heisenberg Hamiltonian. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

(d) 3 × 4 isotropic Heisenberg Hamiltonian with an external magnetic field $h = 2$. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

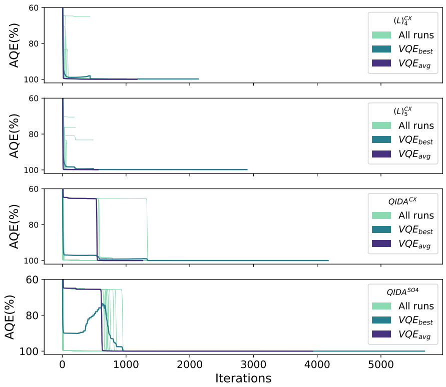

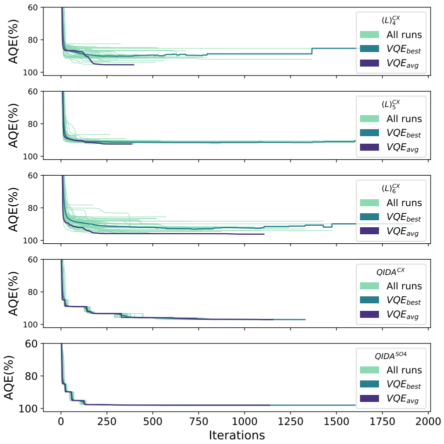

(e) 3 × 4 anisotropic Heisenberg Hamiltonian with $\Delta = 1/10$. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

(f) 3 × 4 anisotropic Heisenberg Hamiltonian with $\Delta = 2/3$. From the top to the bottom: $(L)_4^{CX}$, $(L)_5^{CX}$, $(L)_6^{CX}$, $QIDA^{CX}$, and $QIDA^{SO4}$

Figure B.4: Comparison between different systems of the optimization trajectories. The number of iterations corresponds to the number of evaluations required for the optimizator to converge.

2 Molecular system results

Mol.	Coordinates(Å)	Basis	Qubits
H_2O	H 0.757 0.586 0.0 H -0.757 0.586 0.0 O 0.0 0.0 1.595	INOs RCISD (STO-3G)	12
BeH_2	Be 0.0 0.0 1.334 H 0.0 0.0 0.0 H 0.0 0.0 2.668	INOs RCISD (STO-3G)	12
NH_3	N 0.0 0.0 0.1211 H 0.0 0.9306 -0.2826 H 0.8059 -0.4653 -0.2826 H -0.8059 -0.4653 -0.2826	INOs RCISD (STO-3G)	14
H_2O CAS(4,4)	H 0.847 0.0 0.0 H -0.298 0.0 0.793 O 0.0 0.0 0.0	6-31G	8
N_2 CAS(6,6)	N 0.0 0.0 -0.5488 N 0.0 0.0 0.5488	cc-pVTZ	12

Table B.4: Molecular systems under analysis. The number of qubits is computed as $2{*}M_{act}$, where M_{act} are the active orbitals. For the first three system, the basis-set in brackets is the initial one on which INOs are constructed.

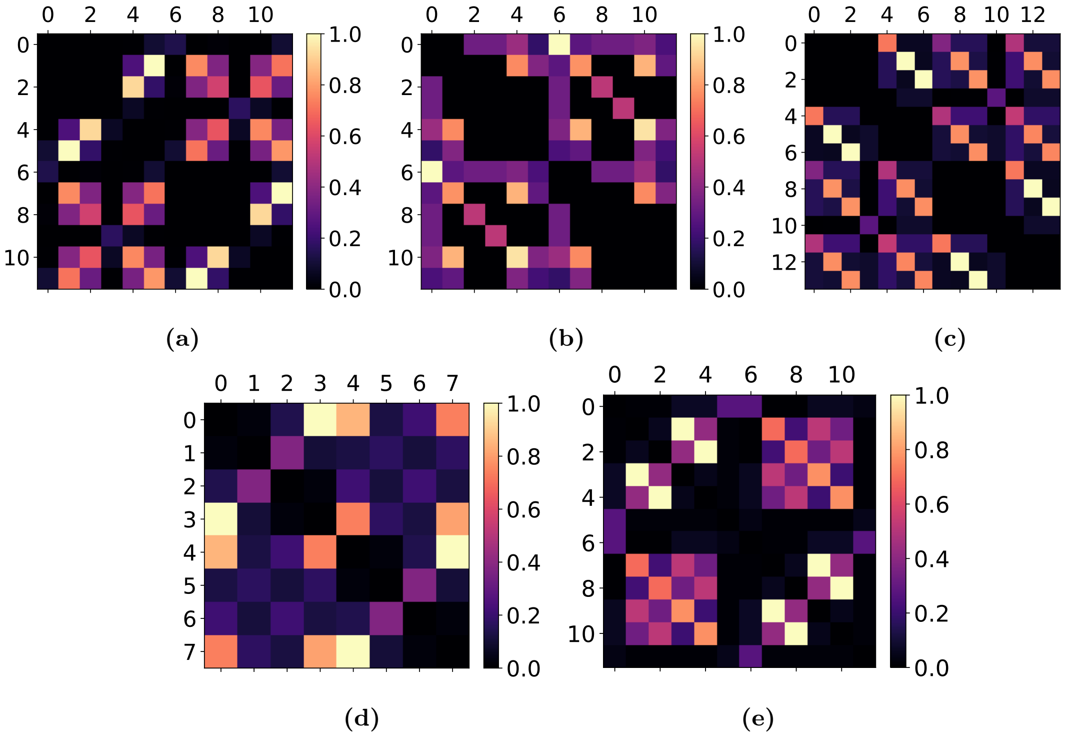

(a) (b) (c)

(d) (e)

Figure B.5: QMI matrices obtained by the tested systems: **(a)** H_2O INOs/12 qubits. **(b)** BeH_2 INOs/12 qubits. **(c)** NH_3 INOs/14 qubits. **(d)** H_2O CAS(4,4). **(d)** N_2 CAS(6,6). All obtained with SparQ at R-CISD level.

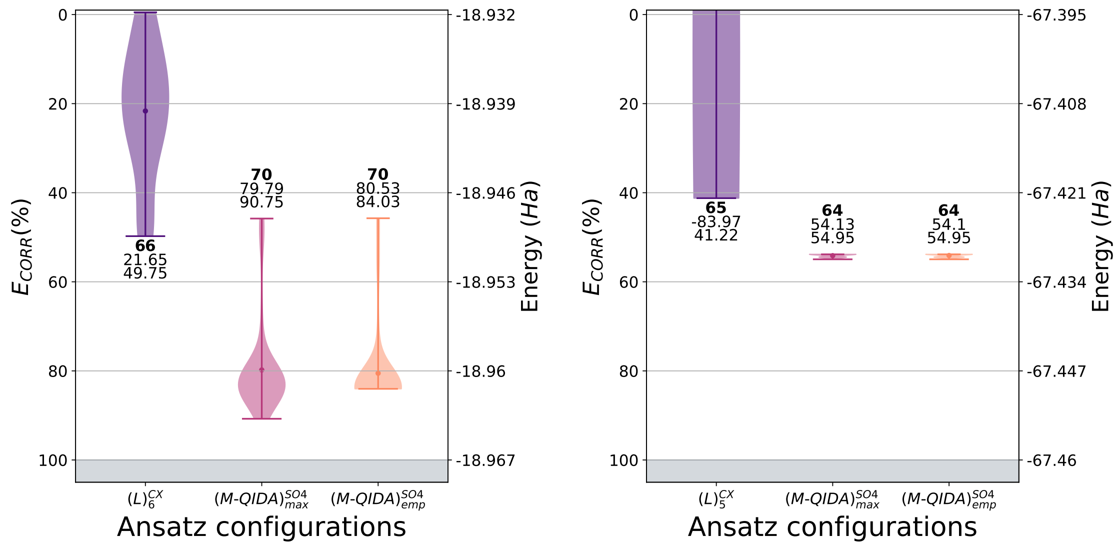

(a) BeH_2 INOs system. **(b)** NH_3 INOs system.

Figure B.6: Convergence percentage correlation energy/absolute energy HEA against Multi-QIDA configurations for INOs molecular systems.

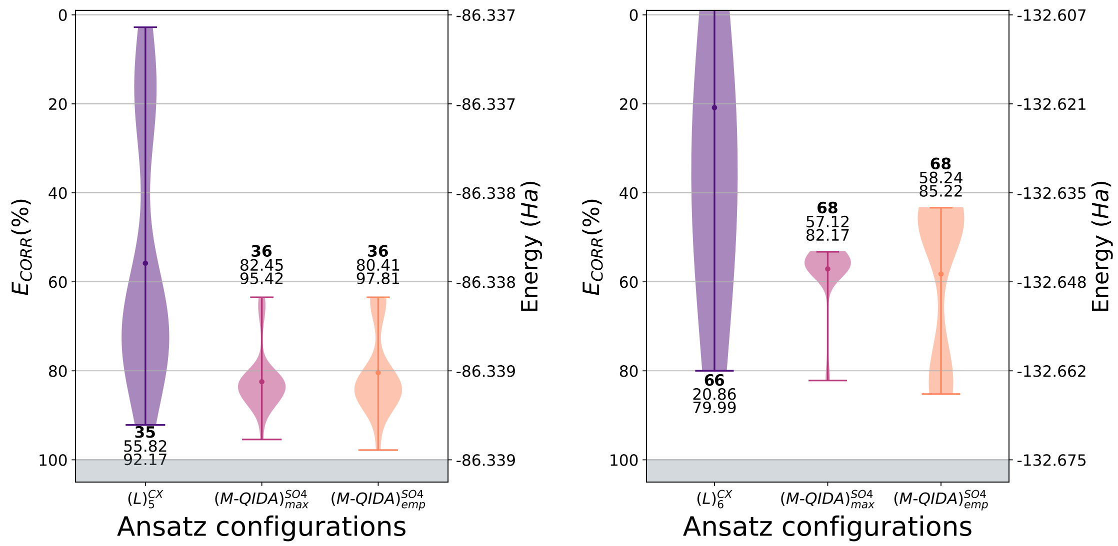

(a) H_2O 6-31G CAS(4,4) system. **(b)** N_2 cc-PVTZ CAS(6,6) system.

Figure B.7: Convergence percentage correlation energy/absolute energy HEA against Multi-QIDA configurations for CASCI/Active region molecular systems.

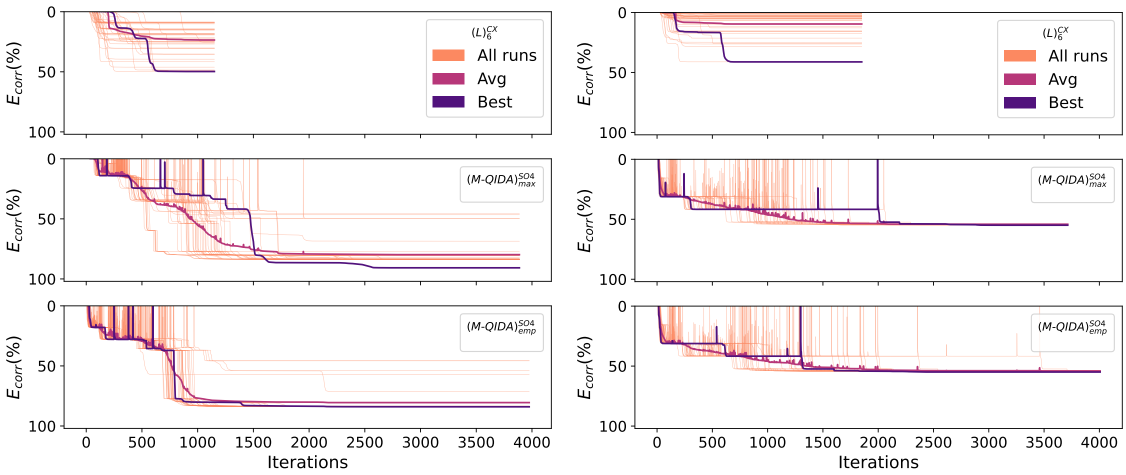

(a) BeH_2 INOs system convergence trajectories **(b)** NH_3 INOs system convergence trajectories

Figure B.8: Optimization trajectories for each of the INOs system of 50 VQE for HEA ladder-fashion circuit against Multi-QIDA circuits.

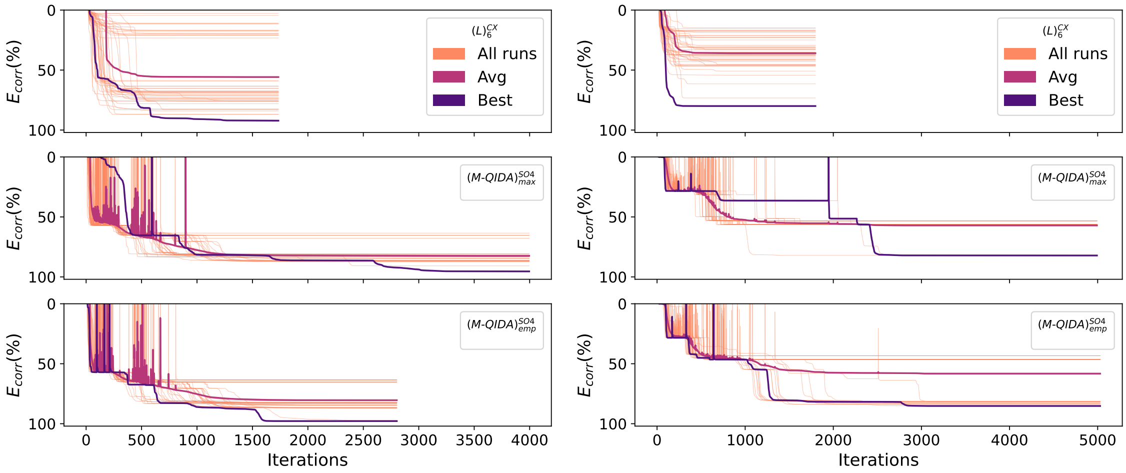

(a) H_2O 6-31G CAS(4,4) system convergence trajectories. **(b)** N_2 cc-PVDZ CAS(6,6) system convergence trajectories.

Figure B.9: Optimization trajectories for each of the CASCI/Active Region system of 50 VQE for HEA ladder-fashion circuit against Multi-QIDA circuits.

	BeH_2			H_2O			NH_3		
	$(L)_6^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$
$\mathcal{F}_{avg}[\%]$	98.89579	99.60117	99.61037	89.08740	99.86803	99.86958	92.23381	99.19447	99.19421
$\hat{S}_{z\,avg}$	-0.00002	0	0	0.10343	0	0	0.05542	0	0
$\hat{S}^2_{avg}$	0.00261	0.00111	0.00062	0.22269	0.00128	0.00115	0.13443	0.00003	0.00004
$\hat{N}_{avg}$	4.00001	4.00000	3.99999	7.97648	7.99998	7.99998	7.98916	8.00000	8.00000
$\mathcal{F}_{best}[\%]$	99.28972	99.80993	99.65462	99.77667	99.88329	99.88125	98.98914	99.20485	99.20485
$\hat{S}_{z\,best}$	0	0	0.00001	0	0	0.00001	0	0	0
$\hat{S}^2_{best}$	0.00016	0.00026	0.00002	0.00001	0.00015	0.00042	0.00002	0	0
$\hat{N}_{best}$	4.00000	4.00010	3.99999	8.00000	8.00000	7.99998	7.99998	8.00000	8.00000

Table B.5: Properties of BeH_2, H_2O, and NH_3 INOs system.

	H_2O $CAS(4,4)$			N_2 $CAS(6,6)$		
	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$	$(L)_5^{cx}$	$QIDA_{max}$	$QIDA_{emp}$
$\mathcal{F}_{avg}[\%]$	$99,95724$	$99,98186$	$99,97978$	$94,16438$	$98,77033$	$98,70980$
$\hat{S}_{z\,avg}$	0	0	0	$-0,00007$	0	0
$\hat{S}^2_{avg}$	$0,00025$	$0,00035$	$0,00030$	$0,05629$	$0,00049$	$0,04372$
$\hat{N}_{avg}$	$4,00000$	$4,00000$	$4,00000$	$6,04056$	$6,00000$	$6,00000$
$\mathcal{F}_{best}[\%]$	$99,99057$	$99,99390$	$99,9967$	$99,5714$	$99,55640$	$99,67421$
$\hat{S}_{z\,best}$	0	0	0	$-0,0002$	0	0
$\hat{S}^2_{best}$	$0,00018$	$0,00010$	$0,00004$	$0,00192$	$0,00766$	0
$\hat{N}_{best}$	$4,00000$	$4,00000$	$4,00000$	$5,99956$	$6,00000$	$6,00000$

Table B.6: Properties of H_2O 6-31G CAS(4,4) and N_2 cc-PVTZ CAS(6,6) system.

Appendix C

Chapter 06 Appendix: Hybrid SCF

1 Derivatives Evaluators

In this section of the Appendix, we report the two type of `joblib` workers method for evaluating the first- and second-order derivatives in WAHTOR–SCF and different variations. In order, *fermionc_worker* and `pauli_worker`.

1.1 Fermionic Derivative Worker

```python
def fermionic_worker(args,
                psi: np.ndarray,
                num_qubits: int,
                masks: dict,
                core: list | np.ndarray,
                virtual: list | np.ndarray,
                memmap: dict | None = None,
                h: dict[tuple[str, str], np.ndarray] | None = None,
                second_order: bool | None = False) -> tuple[int, float, int|None, float | None]:

jac_i, hess_ij = None, None
orbs_g_i, t_i, i_idx, orbs_g_j, t_j, j_idx= args
if memmap is not None: #MEMMAP BRANCH
    h = {}
    for key, info in memmap.items():
        arr = np.memmap(info[’filename’], dtype=info[’dtype’], mode=’r+’, shape=info[’shape’])
        h[key] = np.array(arr, dtype=np.complex64)

    eff = list(wrapper_single_restricted_dh(T=t_i,group_idx=orbs_g_i,h=h))

    for k in list(h.keys()):
        del h[k]
    h.clear()
    del h

    h_first = {}
    for k, block in eff:
```

```python
            h_first[k] = block
        del eff
        jac_i = fermionic_dh_evaluation(h_first, psi, masks, num_qubits, core, virtual)

        if second_order:
            eff = list(wrapper_single_restricted_dh(T=t_j,group_idx=orbs_g_j,h=h_first))
            for k in list(h_first.keys()):
                del h_first[k]
            h_first.clear()
            del h_first
            h_second = {}
            for k, block in eff:
                h_second[k] = block

            hess_ij = fermionic_dh_evaluation(h_second, psi, masks, num_qubits, core, virtual)
            for k in list(h_second.keys()):
                del h_second[k]
            h_second.clear()
            del h_second
            return (i_idx, jac_i, j_idx, hess_ij)
        else:
            return (i_idx, jac_i, None, None)

def fermionic_dh_evaluation(dh: dict,
                       psi: np.ndarray,
                       masks: dict,
                       num_qubits: int,
                       core: list | np.ndarray,
                       virtual: list | np.ndarray ) -> float:

    if not (bool(core) or bool(virtual)):
        xc = 0
    else:
        xc, dh = local_cut_integrals(dh, masks, core, virtual)
    fermop = FermionicOperator(h=dh).fermionic_op
    del dh
    if len(fermop) == 0:
        del fermop
        return np.real(xc)
    q_temp = QubitOperator(data_op=fermop, qiskit=True, map_method = 'jw', num_qubits=num_qubits).qubit_op.simplify()
    T = q_temp.to_matrix(sparse=True, force_serial=True)
    exp = np.real(psi.conj().T @ T @ psi) + xc
    del T, q_temp, fermop
    return exp
```

Snippet C.1: `memmap` branch of `fermionic_worker` method. Here, the method `local_cut_integrals` perform the reduction from full molecular space to the effective operator, as defined in Section 3.1.2

1.2 Pauli Recycling Derivative Worker

```python
def pauli_worker(args,
                cached_h: np.ndarray,
                num_qubits: int,
                masks: dict,
                core: list | np.ndarray,
                virtual: list | np.ndarray,
                h: dict[tuple[str, str], np.ndarray] | None = None,
                memmap: dict | None = None,
                second_order: bool | None = False) -> tuple[int, float, int| None, float | None]:

jac_i, hess_ij = None, None
orbs_g_i, t_i, i_idx, orbs_g_j, t_j, j_idx = args
if memmap is not None:
    h = {}
    for key, info in memmap.items():
        arr = np.memmap(info['filename'], dtype=info['dtype'], mode='r+', shape=info['shape'])
        h[key] = np.array(arr, dtype=np.complex64)

    eff = list(wrapper_single_restricted_dh(T=t_i,group_idx=orbs_g_i,h=h))
    if memmap is not None:
        for k in list(h.keys()):
            del h[k]
        h.clear()
        del h
    h_first = {}
    for k, block in eff:
        h_first[k] = block
    del eff
    jac_i = pauli_dh_evaluation(h_first, cached_h, masks, num_qubits, core, virtual)

    if second_order:
        eff = list(wrapper_single_restricted_dh(T=t_j,group_idx=orbs_g_j,h=h_first))
        for k in list(h_first.keys()):
            del h_first[k]
        h_first.clear()
        del h_first
        h_second = {}
        for k, block in eff:
            h_second[k] = block

        hess_ij = pauli_dh_evaluation(h_second, cached_h, masks, num_qubits, core, virtual)
        for k in list(h_second.keys()):
            del h_second[k]
        h_second.clear()
        del h_second
    gc.collect()
```

```
        return (i_idx, jac_i, j_idx, hess_ij)

def pauli_dh_evaluation(dh: dict,
                        cached_h: np.ndarray,
                        masks: dict,
                        num_qubits: int,
                        core: list | np.ndarray,
                        virtual: list | np.ndarray ) -> float:
    if not (bool(core) or bool(virtual)):
        xc = 0
    else:
        xc, xh = local_cut_integrals(dh, masks, core, virtual)
    fermop = []
    for op in quick_FermionicOp(h=xh):
        fermop.append(op)
    if len(fermop) == 0:
        for k in list(xh.keys()):
            del xh[k]
        del dh, fermop, xh
        return np.real(xc)
    for k in list(xh.keys()):
        del xh[k]

    q_temp = QubitOperator(data_op=fermop, qiskit=True, map_method = 'jw', num_qubits=num_qubits).qubit_op.simplify()
    arr = get_expectation_vector_numba(cache=cached_h, query_paulis=q_temp._pauli_list, query_coeffs=q_temp._coeffs)
    exp = np.real(arr) + xc
    del arr, dh, fermop, q_temp, xh
    gc.collect()
    return exp
```

Snippet C.2: `memmap` branch of `pauli_worker` method. Here, the method `local_cut_integrals` perform the reduction from full molecular space to the effective operator, as defined in Section 3.1.2

2 HEA–WAHTOR–SCF result collection

2.1 Statistic of the VQEs for each PEC and distance

2.2 Potential Energies Curve plots

3 Resource requirements for HEA–WAHTOR–SCF and ADAPT–WAHTOR–SCF

4 WAHTOR–SCF for Industrial chemistry

4.1 Model reaction geometries

Distance[Å]	Gate	#Params	#CNOTs	Best VQE	STD
1.68	Cx	20	12	99.99998	0.00244
1.68	SO4	48	12	99.99475	0.00343
2.65	Cx	20	12	99.99896	7.793e-08
2.65	SO4	48	12	99.99888	0.0048
3.10	Cx	20	12	99.99985	1.354e-09
3.10	SO4	48	12	99.99975	0.00249
3.65	Cx	20	12	99.99923	0.00210
3.65	SO4	48	12	99.99990	0.0032
5.30	Cx	20	12	99.99999	7.417e-09
5.30	SO4	48	12	99.99999	1.158e-09

Table C.1: Comparison between Cx and SO4 ladder ansatz for WAHTOR-SCF calculation. 50 VQE runs for each geometry. LiH, 6-31G basis, CAS(2,2)

Distance[Å]	Gate	#Params	#CNOTs	Best VQE	STD
1.84	Cx	72	56	97.62223	0.01379
1.84	SO4	208	56	99.99926	0.00838
2.10	Cx	72	56	99.13215	0.01577
2.10	SO4	208	56	99.99149	0.01420
3.00	Cx	72	56	99.95300	0.03015
3.00	SO4	208	56	99.99696	0.00917
4.30	Cx	72	56	97.72961	0.01120
4.30	SO4	208	56	99.58357	0.01086
6.00	Cx	72	56	99.90597	0.00028
6.00	SO4	208	56	99.90626	0.00033

Table C.2: Comparison between Cx and SO4 ladder ansatz for WAHTOR-SCF calculation. 50 VQE runs for each geometry. H_2O, 6-31G basis, CAS(4,4)

Distance[Å]	Gate	#Params	#CNOTs	Best VQE	STD
1.34	Cx	72	56	99.99675	0.00086
1.34	SO4	208	56	99.286326	0.00121
1.46	Cx	72	56	99.89205	0.00319
1.46	SO4	208	56	99.41380	0.00108
2.00	Cx	72	56	99.97348	0.02985
2.00	SO4	208	56	99.99921	0.02349
2.75	Cx	72	56	98.97621	0.02436
2.75	SO4	208	56	99.46493	0.02699
4.00	Cx	72	56	99.78867	0.01101
4.00	SO4	208	56	99.78753	0.01390

Table C.3: Comparison between Cx and SO4 ladder ansatz for WAHTOR-SCF calculation. 50 VQE runs for each geometry. C_2H_4, cc-pvdz basis, CAS(4,4)

Distance[Å]	Gate	#Params	#CNOTs	Best VQE	STD
1.12	Cx	156	132	82.23768	0.10004
1.12	SO4	480	132	96.50521	0.03709
1.20	Cx	156	132	99.87838	0.07562
1.20	SO4	480	132	96.45088	0.04497
1.60	Cx	156	132	83.25998	0.10314
1.60	SO4	480	132	95.85233	0.10349
2.00	Cx	156	132	88.34794	0.09953
2.00	SO4	480	132	83.49238	0.08159
3.50	Cx	156	132	99.98126	0.13224
3.50	SO4	480	132	87.60787	0.02990

Table C.4: Comparison between Cx and SO4 ladder ansatz for WAHTOR-SCF calculation. 50 VQE runs for each geometry. N_2, cc-pvdz basis, CAS(6,6)

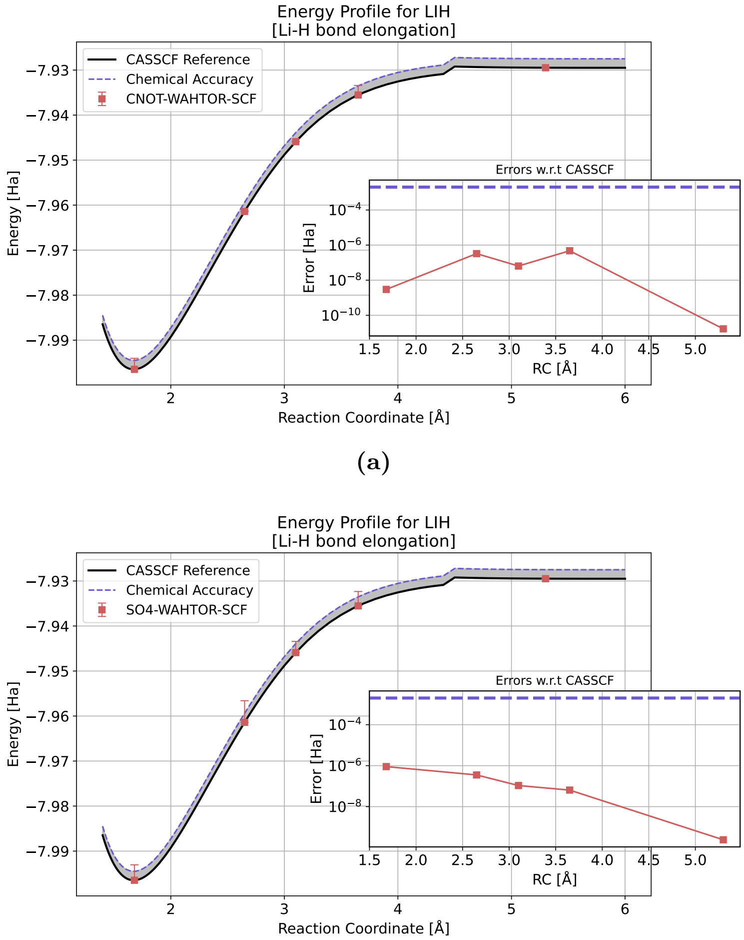

(a)

(b)

Figure C.1: Simulation results for HEA–WAHTOR–SCF with CNOTS **(a)** and *SO(4)* **(a)**. Here, we report in the main panel of each image the comparison against CASSCF (black line), the dashed-line corresponds to chemical accuracy (2mHa error), the squares are best WAHTOR–SCF for each coodinat, and corresponding error bars. In the smaller panel instead, we report only the error of the best WAHTOR–SCF run compared to CASSCF results.

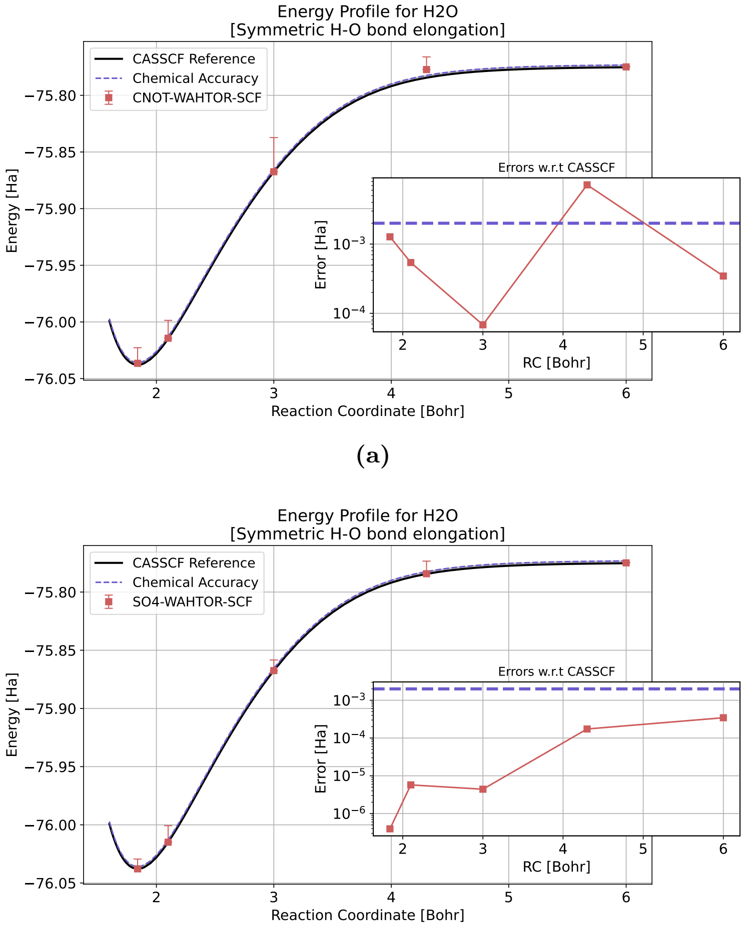

(a)

(b)

Figure C.2: Simulation results for HEA–WAHTOR–SCF with CNOTS **(a)** and *SO(4)* **(b)**. Here, we report in the main panel of each image the comparison against CASSCF (black line), the dashed-line corresponds to chemical accuracy (2mHa error), the squares are best WAHTOR–SCF for each coodinat, and corresponding error bars. In the smaller panel instead, we report only the error of the best WAHTOR–SCF run compared to CASSCF results.

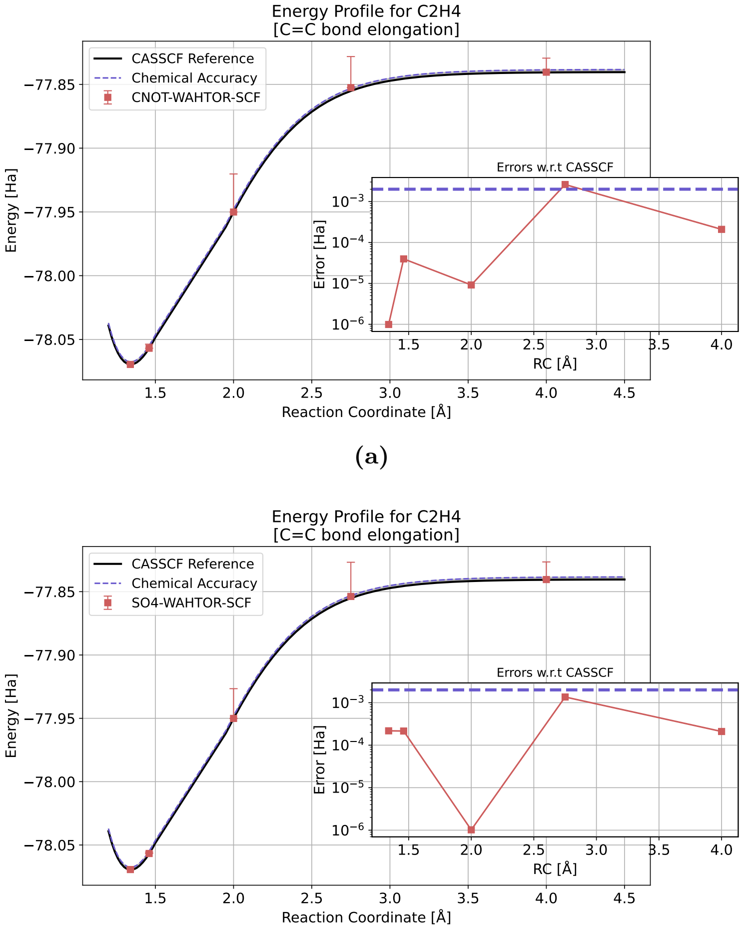

(a)

(b)

Figure C.3: Simulation results for HEA–WAHTOR–SCF with CNOTS **(a)** and *SO(4)* **(b)**. Here, we report in the main panel of each image the comparison against CASSCF (black line), the dashed-line corresponds to chemical accuracy (2mHa error), the squares are best WAHTOR–SCF for each coodinat, and corresponding error bars. In the smaller panel instead, we report only the error of the best WAHTOR–SCF run compared to CASSCF results.

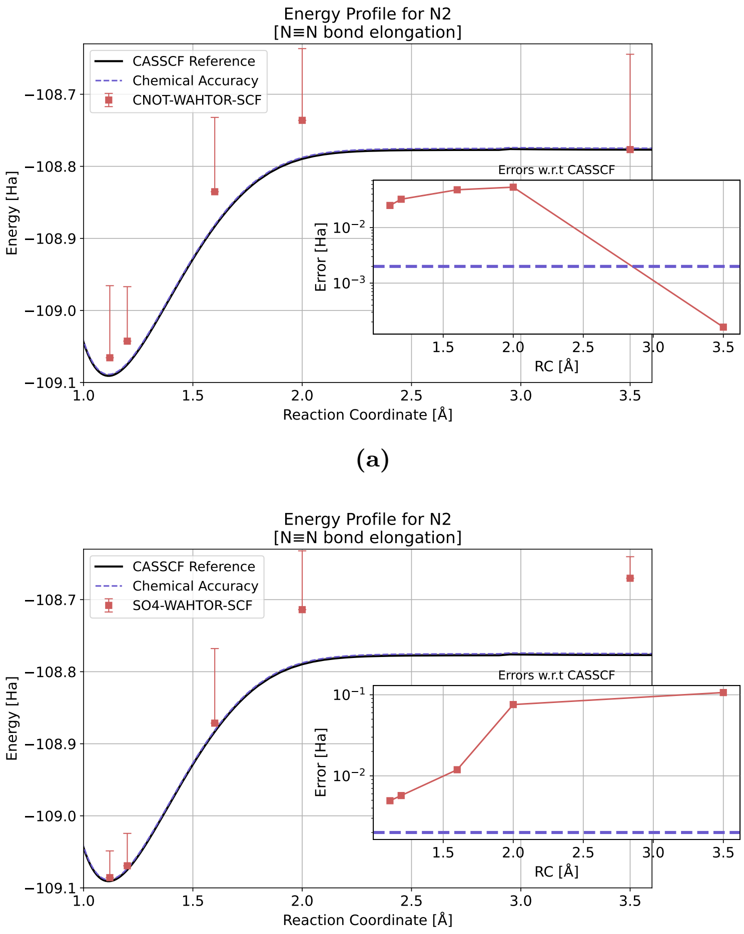

(a)

(b)

Figure C.4: Simulation results for HEA–WAHTOR–SCF with CNOTS **(a)** and *SO(4)* **(b)**. Here, we report in the main panel of each image the comparison against CASSCF (black line), the dashed-line corresponds to chemical accuracy (2mHa error), the squares are best WAHTOR–SCF for each coodinat, and corresponding error bars. In the smaller panel instead, we report only the error of the best WAHTOR–SCF run compared to CASSCF results.

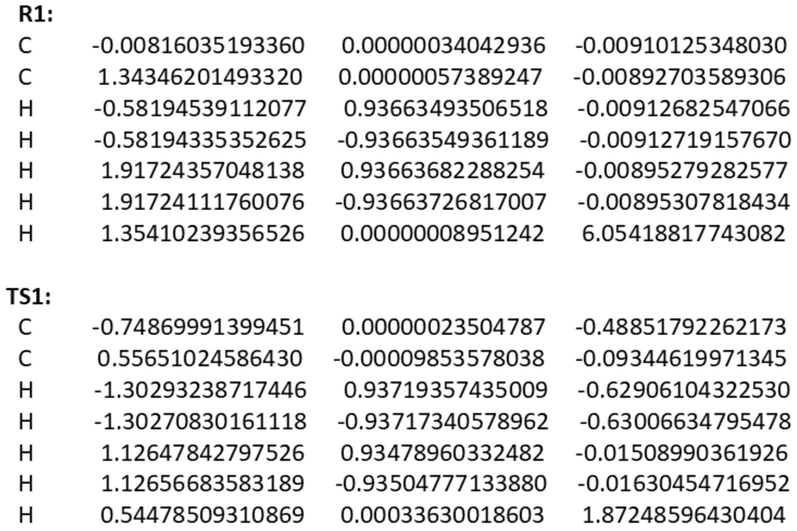

R1:

C	-0.00816035193360	0.00000034042936	-0.00910125348030
C	1.34346201493320	0.00000057389247	-0.00892703589306
H	-0.58194539112077	0.93663493506518	-0.00912682547066
H	-0.58194335352625	-0.93663549361189	-0.00912719157670
H	1.91724357048138	0.93663682288254	-0.00895279282577
H	1.91724111760076	-0.93663726817007	-0.00895307818434
H	1.35410239356526	0.00000008951242	6.05418817743082

TS1:

C	-0.74869991399451	0.00000023504787	-0.48851792262173
C	0.55651024586430	-0.00009853578038	-0.09344619971345
H	-1.30293238717446	0.93719357435009	-0.62906104322530
H	-1.30270830161118	-0.93717340578962	-0.63006634795478
H	1.12647842797526	0.93478960332482	-0.01508990361926
H	1.12656683583189	-0.93504777133880	-0.01630454716952
H	0.54478509310869	0.00033630018603	1.87248596430404

Figure C.5: Optimized CCSD(T) geometries of the reactant and transition state in the initiation step using cc-pVDZ basis.

Molecule (Basis, Active Space)	Geometry	CNOTs	Parameters
LiH (6-31G, CAS(2,2))	1.68 Å	10	3
	2.65 Å	10	3
	3.10 Å	10	3
	3.65 Å	10	3
	5.30 Å	14	3
H_2O (6-31G, CAS(4,4))	1.84 Bohr	118	25
	2.10 Bohr	132	28
	3.00 Bohr	96	16
	4.30 Bohr	104	18
C_2H_4 (cc-pVDZ, CAS(4,4))	1.34 Å	78	13
	1.46 Å	114	19
	2.00 Å	108	18
	2.75 Å	108	18
N_2 (cc-pVDZ, CAS(6,6))	1.12 Å	486	81
	1.20 Å	258	43
	1.60 Å	432	72
	2.00 Å	570	95

Table C.5: CNOT count and number of parameters required by the ADAPT-WAHTOR-SCF ansatz to reach chemical accuracy for representative geometries of each molecular system.

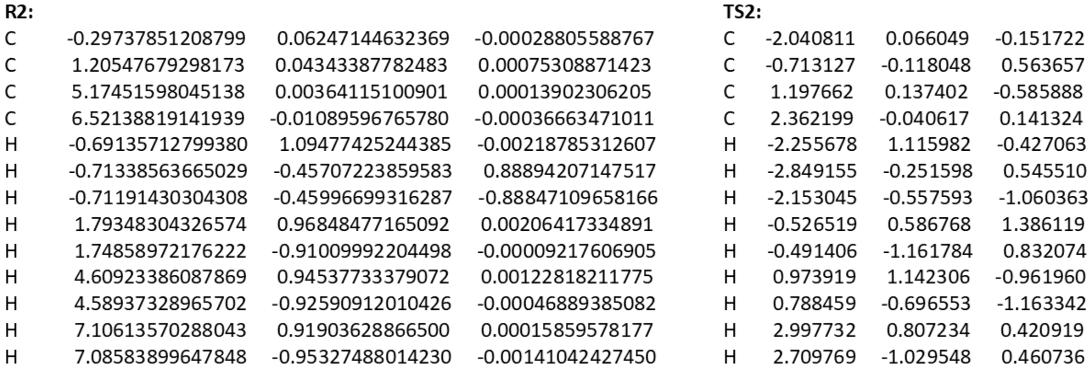

R2:

C	-0.29737851208799	0.06247144632369	-0.00028805588767
C	1.20547679298173	0.04343387782483	0.00075308871423
C	5.17451598045138	0.00364115100901	0.00013902306205
C	6.52138819141939	-0.01089596765780	-0.00036663471011
H	-0.69135712799380	1.09477425244385	-0.00218785312607
H	-0.71338563665029	-0.45707223859583	0.88894207147517
H	-0.71191430304308	-0.45996699316287	-0.88847109658166
H	1.79348304326574	0.96848477165092	0.00206417334891
H	1.74858972176222	-0.91009992204498	-0.00009217606905
H	4.60923386087869	0.94537733379072	0.00122818211775
H	4.58937328965702	-0.92590912010426	-0.00046889385082
H	7.10613570288043	0.91903628866500	0.00015859578177
H	7.08583899647848	-0.95327488014230	-0.00141042427450

TS2:

C	-2.040811	0.066049	-0.151722
C	-0.713127	-0.118048	0.563657
C	1.197662	0.137402	-0.585888
C	2.362199	-0.040617	0.141324
H	-2.255678	1.115982	-0.427063
H	-2.849155	-0.251598	0.545510
H	-2.153045	-0.557593	-1.060363
H	-0.526519	0.586768	1.386119
H	-0.491406	-1.161784	0.832074
H	0.973919	1.142306	-0.961960
H	0.788459	-0.696553	-1.163342
H	2.997732	0.807234	0.420919
H	2.709769	-1.029548	0.460736

Figure C.6: Optimized CCSD(T) geometries of the reactant and transition state in the propagation step using cc-pVDZ basis.

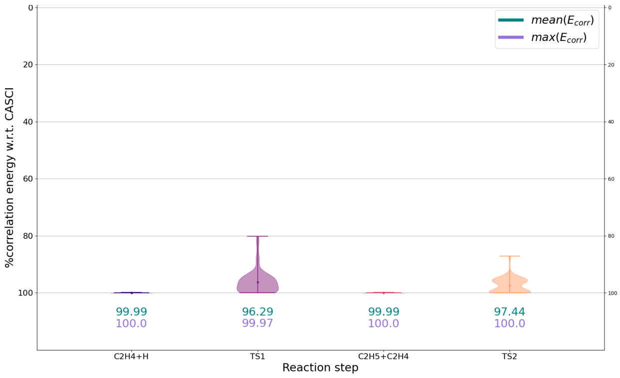

Figure C.7: Preliminary Single points calculations at CASCI-theory level for all the structure of interest, reactants and transition state for both the initiation and propagation step. Product structures are avoided given the single–determinant solution. Each violing plot represent 50 VQEs with ladder shaped ansatz and *SO(4)* gates.

Appendix D

Chapter 06 Appendix: Quantum Geometry Optimization

1 Initial DompeKeys structures

```
# Methylene Imine
5

C  0.0580   0.5852   0.0000
N  0.0580   -0.6875  0.0000
H  -0.8313  1.2172   0.0000
H  1.0094   1.1151   0.0000
H  -0.9323  -1.0310  0.0000

# Acetaldehyde
7

C  0.0000    0.4758    0.0000
C  -0.9440   -0.7365   0.0000
O  1.2147    0.3988    0.0000
H  -0.5153   1.4524    0.0000
H  -0.3753   -1.6605   0.0000
H  -1.5815   -0.7090   0.8799
H  -1.5815   -0.7090   -0.8799

# Acetamide
9

C  1.3915    -0.3288   -0.0017
C  -0.0660   0.1729    -0.0014
N  -1.0418   -0.8872   -0.1078
O  -0.3914   1.3461    0.0184
H  2.0697    0.5177    -0.0135
H  1.5861    -0.9266   0.8846
H  1.5692    -0.9450   -0.8783
H  -1.9716   -0.5598   0.1821
```

```
H   -0.7820    -1.7086    0.4510
```

Snippet D.1: Initial geometries of the three analyzed DompeKeys optimized with Restricted-HF and minimal basis set STO-3g.

Tesi redatta con il cofinanziamento con risorse dell’Unione europea-*NextGeneration EU*
Piano Nazionale di Ripresa e Resilienza Missione 4 – Componente 1 –
Investimento 4.1. del PNRR – “Ricerca PNRR” – CUP E11I22000150001

Finanziato
dall'Unione europea
NextGenerationEU

Ministero
dell’Università
e della Ricerca

Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA

UNIVERSITÀ
DEGLI STUDI
DELL’AQUILA