CHALMERS
UNIVERSITY OF TECHNOLOGY

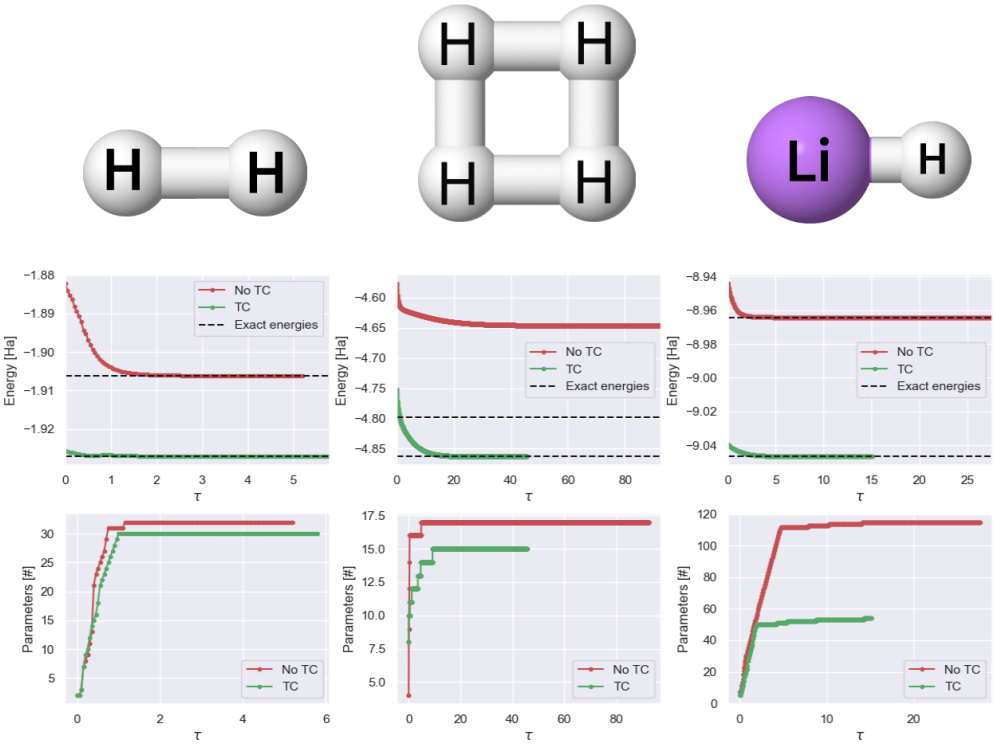

Combining the Transcorrelated Method with Adaptive Ansatzes for Near-term Quantum Computation

Transcorrelated Adaptive Variational Quantum Imaginary Time Evolution

Master's thesis in Theoretical Chemistry

Erika Magnusson

DEPARTMENT OF CHEMISTRY AND CHEMICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2023
www.chalmers.se

Master's thesis 2023

Combining the Transcorrelated Method with Adaptive Ansatzes for Near-term Quantum Computation

Transcorrelated Adaptive Variational Quantum Imaginary Time Evolution

Erika Magnusson

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Chemistry and Chemical Engineering
Division of Physical Chemistry
Rahmlab group
Chalmers University of Technology
Gothenburg, Sweden 2023

Combining the Transcorrelated Method with Adaptive Ansatzes for Near-term Quantum Computation
Transcorrelated Adaptive Variational Quantum Imaginary Time Evolution
Erika Magnusson

Supervisor: Werner Dobrautz, Department of Chemistry and Chemical Engineering
Examiner: Martin Rahm, Department of Chemistry and Chemical Engineering

Master's Thesis 2023
Department of Chemistry and Chemical Engineering
Division of Physical Chemistry
Rahmlab group
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: A selection of representative results for TC-AVQITE along with modelled systems. For more information, see chapter 4.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2023

Combining the Transcorrelated Method with Adaptive Ansatzes for Near-term Quantum Computation
Transcorrelated Adaptive Variational Quantum Imaginary Time Evolution
Erika Magnusson
Department of Chemistry and Chemical Engineering
Chalmers University of Technology

Abstract

The fundamental problem of quantum chemistry is that the cost of solving the electronic Schrödinger equation scales exponentially with system size. One way to potentially circumvent this scaling is using quantum computers, whose properties might enable an exponential computational speedup. However, near-term quantum computers struggle to compete with conventional quantum chemistry methodology. Near-term quantum processors are both small regarding the number of available qubits, and the number of operations that can be run on them is limited by systematic noise. Despite this, various methods attempt to find quantum advantage in the current noisy intermediate-scale quantum (NISQ) regime. Variational quantum imaginary time evolution iteratively moves toward the energy ground state by evolving the quantum state in imaginary time. By splitting the evolution into steps and iterating, the necessary circuit depth is limited.
This thesis evaluates the combination of variational quantum imaginary time evolution with two approaches aimed at decreasing quantum costs: the transcorrelated (TC) method of Boys and Handy and the adaptive ansatzes of Grimsley *et al.*
The resulting method, transcorrelated adaptive variational quantum imaginary time evolution (TC-AVQITE), is evaluated through simulations of near-term quantum devices. The results for small systems (H_2, quadratic H_4, and lithium hydride) indicate that the combination works well, especially for increasing system sizes. This development takes us a step closer to chemically relevant calculations on quantum computers.

Keywords: Quantum computer, Quantum chemistry, NISQ, QITE, VarQITE, ADAPT, ADAPT-VarQITE, Transcorrelated method.

Acknowledgements

I want to thank Gomes *et al.* for allowing me to clone the code used in your paper! How lucky I am to have a supervisor with contacts; writing AVQITE from scratch would have taken quite a lot of time. Speaking of supervisor, I want to thank Werner Dobrautz for being a great help during all steps of the process! As a natural continuation of this, I want to thank Martin Rahm and the rest of the group for company and feedback. Finally, I also want to thank my opponent Oskar Olander.

Erika Magnusson, Gothenburg, June 2023

Contents

1
Introduction

In quantum chemistry, the aim is to calculate atomic and molecular properties as faithfully as possible. From these fundamental calculations, one can obtain the ground state energy for the quantum system, from which chemically relevant properties can be derived. The usual approach involves simulating the elementary particles in the molecules by making relevant approximations, and then using methods such as density functional theory or coupled cluster approaches. Here, a new method is gradually appearing – simulating chemical systems on a quantum computer. This method's promise is great, as the fundamental properties of quantum computers allow us to bypass that the wave function is exponentially scaling with particle number and system size. However, current quantum computers have many issues, such as systematic noise from imperfect quantum gates or a limited number of quantum bits (qubits).

Despite the limits of current quantum computers, several methods are developed for quantum computation. Furthermore, there are also methods developed that cut down on quantum costs. Two of these are the transcorrelated method of Boys and Handy [1], as well as adaptive ansatzes developed by Grimsley *et al.* [2]. In this thesis, a combination of these two methods for cutting quantum computation costs is evaluated. Specifically, we want to investigate whether this combination outperforms simulations with only one of the methods and if the methods work well together.

Quantum chemistry is a fundamental field whose developments could have long-reaching impacts. As quantum computation could theoretically let us calculate fundamental properties of molecules, it could in the future come to be used for developing new medicines or other useful chemical compounds. This could then have wide-reaching consequences for society at large. As it is now, a quantum revolution is not yet in sight, but by continued research it may one day be reached.

2

Theory

The theory chapter will start with the fundamental problem – the electronic structure problem. After that, quantum computing is introduced and motivated, and the specific methods used in the thesis are detailed.

2.1 Background

2.1.1 The Electronic Structure Problem

Fundamentally, computational quantum chemistry is all about the problem of calculating the states of electrons in atoms and molecules. This problem is known as the electronic structure problem. Determining the features of the electronic energy surface, i.e. the collective behaviour of the electrons, is important for several properties such as reactivity and reaction rates, which motivates why the problem should be studied [3]. The problem can (in most cases) essentially be reduced to solving the stationary non-relativistic Schrödinger equation for various molecules, i.e.,

$$\hat{H}_{\mathrm{m}}(\vec{r})\Psi(\vec{r},t) = E\Psi(\vec{r},t), \tag{2.1}$$

where the molecular Hamiltonian in atomic units is given as

$$\hat{H}_{\mathrm{m}} = -\underbrace{\sum_i \frac{\nabla_i^2}{2}}_{\text{KE of } e} - \underbrace{\sum_I \frac{\nabla_I^2}{2M_I}}_{\text{KE of } n} - \underbrace{\sum_{i,I} \frac{Z_I}{|\vec{r}_i - \vec{R}_I|}}_{e\text{-}n \text{ attraction}} + \underbrace{\frac{1}{2}\sum_{i\neq j} \frac{1}{|\vec{r}_i - \vec{r}_j|}}_{e\text{-}e \text{ repulsion}} + \underbrace{\frac{1}{2}\sum_{I\neq J} \frac{Z_I Z_J}{|\vec{R}_i - \vec{R}_j|}}_{n\text{-}n \text{ repulsion}}. \tag{2.2}$$

Here, the uppercase indices go over the nucleons (concisely written as n in the equation), while the lowercase indices go over the electrons (e). The significance of each term is described under them – kinetic energies (KE), attraction, and repulsion. The charge and mass numbers Z and M are given in units of electron charges and masses.

Under the Born-Oppenheimer approximation, we can decouple the electronic and nuclear Hamiltonians by assuming the nuclei to be stationary. The electronic Hamiltonian then becomes

$$\hat{H} = -\underbrace{\sum_i \frac{\nabla_i^2}{2}}_{\text{KE of } e} - \underbrace{\sum_{i,I} \frac{Z_I}{|\vec{r}_i - \vec{R}_I|}}_{e\text{-}n \text{ attraction}} + \underbrace{\frac{1}{2}\sum_{i\neq j} \frac{1}{|\vec{r}_i - \vec{r}_j|}}_{e\text{-}e \text{ repulsion}}. \tag{2.3}$$

It is possible to completely reformulate this Hamiltonian into the so-called second quantisation (or, more descriptively, the occupation number representation) by projecting onto a basis set that spans the relevant Fock space. This basis set can be chosen freely if it spans the relevant space. There are several options, each with advantages and drawbacks (more about basis sets in Section 2.2.2). After choosing a basis set, one can construct creation and annihilation operators $a_p^\dagger$, a_p, which add and remove electrons in the spin orbital p respectively. These are constructed in accordance with the anti-commutative property of fermions,

$$\{a_p, a_q\} = \{a_p^\dagger, a_q^\dagger\} = 0, \quad \{a_p, a_q^\dagger\} = \delta_q^p. \tag{2.4}$$

In the second quantisation, the electronic Hamiltonian thus looks like

$$\hat{H} = \underbrace{\sum_{pq} h_p^q a_p^\dagger a_q}_{\text{one-body terms}} + \underbrace{\frac{1}{2} \sum_{pqrs} V_{pq}^{rs} a_p^\dagger a_q^\dagger a_r a_s}_{\text{two-body terms}}, \tag{2.5}$$

with

$$h_p^q = \int \mathrm{d}\vec{x} \phi_p^*(\vec{x}) \left(-\frac{\nabla^2}{2} - \sum_I \frac{Z_I}{|\vec{r} - \vec{R}_I|} \right) \phi_q(\vec{x}) \tag{2.6}$$

and

$$V_{pq}^{rs} = \int \mathrm{d}\vec{x}_1 \mathrm{d}\vec{x}_2 \frac{\phi_p^*(\vec{x}_1)\phi_q^*(\vec{x}_2)\phi_r(\vec{x}_1)\phi_s(\vec{x}_2)}{|\vec{r}_1 - \vec{r}_2|}, \tag{2.7}$$

where the basis functions are $\phi(\vec{x})$. To clarify, the indices $p, q...$ here go through all spin orbitals.
The second quantisation will be assumed throughout the rest of the thesis. The approach is a good match with quantum computing, as it often gives simpler Hamiltonians and thus requires fewer qubits in general [4].

2.1.2 Why Quantum Computing?

Various methods exist for solving this problem, ranging from mean-field theories like Hartree-Fock to expensive *ab initio* methods. One new and promising method is using quantum computers, as their properties are particularly well suited for quantum simulation [3], which in the end, is what we want to do in computational chemistry. That they match so well is why the electronic structure problem is often thought to be one of the first useful applications of quantum computers [5].
As a field, quantum chemistry can benefit immensely from powerful quantum computers, which would let us solve problems which are intractable on normal computers. In particular, problems that explicitly require the wave function due to the simulated properties, problems with a high degree of entanglement, and problems involving highly correlated systems, could benefit immensely [5, 6]. Furthermore, other fields, like materials science, can also benefit from the methods developed for quantum chemistry [3]. For example, a highly correlated system in physics would be high-temperature superconductors, which are studied in the context of the Hubbard model. The potential of quantum simulation is clear to see – and one day, it may become more accurate to theory than measurements of actual experiments. This would be revolutionary and enable designing useful molecules from scratch [5].

2.1.3 Caveats

There are, of course, a few caveats to consider. Firstly, for the truly groundbreaking quantum simulation, it is estimated that orders of magnitude more qubits are needed than what are available today, plus error correction and lower error rates [5]. Thus, in the near term, quantum computers cannot quite deliver what is prophesied, but could still outperform classical options for highly correlated systems [6]. Simulations have shown that certain approaches can reach competitiveness with classical methods with only a few qubits [7].
Secondly, quantum computing with today's quantum hardware is difficult, and most computations done today correspond to simulations done decades ago on classical hardware. The calculations done today are not quite yet pushing the limits of science as it were, but are more proof of concept.
Thirdly and finally, not everything is suited for quantum computers. Quantum advantage – when a quantum computer outperforms a classical computer – is not universal, but problem specific. An arbitrary simulation will therefore not necessarily be sped up when done on a quantum computer, and one must instead identify which problems are worth considering [6, 3]. Generic ground state determination lies in complexity class QMA, which is a class of problems whose relative efficiency of exponential speedup on a quantum computer is not theoretically guaranteed [8, 9]. Fortunately, recent studies suggest that the problem as addressed in quantum chemistry can be made BQP-complete, which thus in turn suggests quantum advantage [10, 11].

2.2 Quantum Computing

As mentioned, quantum computing is an emerging alternative to classical computation strategies. A large part of their benefit lies in the inherent advantage that quantum computers can store the entire wave function, which is exponentially scaling for classical computers. This is possible because N qubits can store 2^N bits of information, circumventing the exponential scaling issue. There are two ways of using quantum hardware to simulate real systems – analogue and digital [5]. Analogue quantum simulation would be done by constructing hardware mimicking a system of interest, while digital computation maps a problem to a more generic quantum computer. This thesis is strictly about digital quantum computation.

2.2.1 Encoding

To do a (digital – from now on omitted) quantum computation, one must somehow transcribe the problem to the quantum computer, which is not a trivial task. The largest complication is that all known quantum computers have distinguishable qubits, but electrons, which we wish to simulate, are indistinguishable in nature [6]. Several mappings exist, but the simplest one is the Jordan-Wigner (JW) mapping, where a qubit directly represents the occupation number of a spin-orbital. As an example, in the electronic structure problem, the Hartree-Fock-state for a two-level

system with two electrons would be

$$|\Psi_{\mathrm{HF}}\rangle \doteq |0011\rangle\,, \tag{2.8}$$

where the registers from right to left are spin down ground state, spin up ground state, spin down excited state, and spin up excited state, respectively. This is one possible implementation of the qubit order for the JW mapping – one can as easily order them in some other way. This has no real impact on the simulated systems though, only on the algorithmic implementation. Exciting the spin-down electron in this implementation of JW mapping would give

$$a^\dagger_{1\downarrow} a_{0\downarrow} |0011\rangle = |0110\rangle\,. \tag{2.9}$$

For an illustration of these two states, see Figure 2.1.

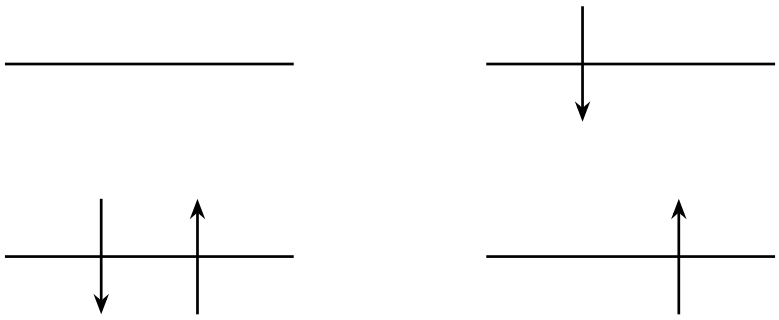

Figure 2.1: An illustration of the two states represented in the JW mapping in equations (2.8) (left) and (2.9) (right).

The advantage of the JW mapping is its simplicity. However, other, less straightforward mappings are more efficient, such as parity mapping [12]. Amongst other features, the parity mapping makes use of inherent symmetries of the problem, in this case that the number of up- and down-spin electrons are conserved, which allows us to reduce the necessary number of qubits to encode the system by two compared to the JW mapping [12]. This is a particularly large advantage in near-term computation, as the number of qubits is limited.

2.2.2 Basis sets

The mapping is not the only thing that needs to be considered; one must also consider which basis set is used in equation (2.5). As mentioned, these basis sets have various properties that make them attractive in certain scenarios but inadvisable in others. One relevant class of basis sets is the minimal ones, particularly the STO-nG family of basis sets. This stands for Slater Type Orbital, n Gaussians. Of these, STO-6G is both readily available and as accurate as possible while still being minimal. Due to this, it was used for all systems except the smallest, the hydrogen dimer.
A minimal basis set is the smallest possible set of basis functions to describe a system. Another relevant class of basis sets are the Pople basis sets, which are split-valence basis sets that allow for greater flexibility in the Hartree-Fock optimisation. A Pople basis set was used for the hydrogen dimer – specifically, 6-31G.

2.2.3 Near-term approaches

Once the electronic structure problem is mapped to the quantum computer, there are several algorithms for solving it. However, due to the size of near-term quantum computers, most methods are still inaccessible to us [6]. For the near-term devices, the focus is on methods that could work well on the noisy and small-scale hardware available – these devices are known as Noisy Intermediate Scale Quantum (NISQ) computers. These methods try to be NISQ-friendly by offloading parts of the computation to classical computers and only using quantum resources when necessary. As a result, circuit depth and coherence time requirements decrease.

2.2.4 Variational quantum eigensolver

The most straightforward of these methods is the Variational Quantum Eigensolver (VQE). The variational quantum eigensolver works by preparing a trial state, an ansatz, dependent on a set of parameters $\vec{\theta}$ on the quantum hardware. This ansatz is then used to calculate the expected energy, which according to the variational theorem will always be higher or equal to the ground state energy [13]. The parameters are then passed to a classical computer, where an optimisation procedure follows. Afterwards, a new set of parameters are obtained, which will be used to prepare a new trial state. This procedure is repeated until a minimum is reached, which is hopefully the ground state energy. Figure 2.2 depicts the algorithm.

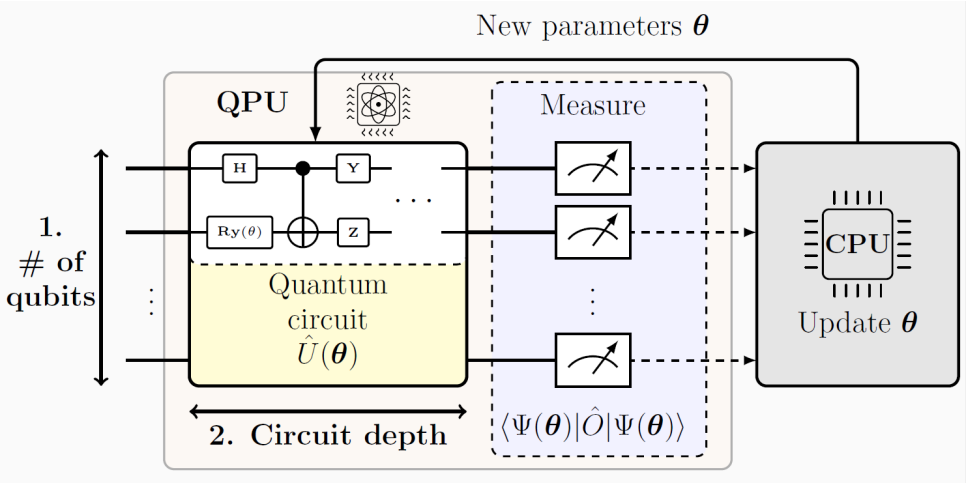

Figure 2.2: An overview of the VQE algorithm. Taken with permission from Werner Dobrautz.

While incredibly simple in idea and comparatively easy to execute, some issues must be considered. Firstly, the aforementioned ansatz needs a non-vanishing overlap with the actual ground state wave function, or the converged energy will not be the ground state energy [6]. One can intuitively see that the quality of the ansatz has a large impact, so one would want as physical of an ansatz as possible from this perspective. However, the state must also be comparatively easy to prepare on quantum hardware, or it will become too expensive for NISQ computers. Trying to

balance this leads to the two classes of physically motivated and hardware-efficient ansatzes, where one property is prioritised over the other.
A long-time friend of chemists is coupled cluster theory, which in the quantum computer is realised as the physically motivated unitary coupled cluster (UCC), often truncated at double excitations (UCCSD) [14]. An example of a hardware-efficient ansatz is the RY ansatz [15], which tries to be as minimal as possible to implement, but might not conserve physical properties such as the number of electrons during the optimisation. On their own, hardware-efficient ansatzes often lead to problems in the optimisation step of VQE, which makes them a comparatively poor match [16].
Finally, the VQE algorithm unsurprisingly fails if the Dirac-Frenkel variational principle (i.e., the "normal" variational principle) does not hold, for example if the Hamiltonian were to become non-Hermitian [17].

2.3 Quantum Imaginary Time Evolution

Quantum Imaginary Time Evolution (QITE) is a quantum computer implementation of the concept of imaginary time evolution [18]. This is a nonphysical mathematical trick used in various fields such as statistical mechanics, cosmology, and quantum mechanics [19]. Imaginary time evolution is a way of determining the ground state energy by replacing time in normal time evolution (i.e., dynamics simulation) with imaginary time: $t \to i\tau$. Given an initial state at starting time $|\psi(0)\rangle$, the normalised imaginary time evolution to imaginary time τ is

$$|\psi(\tau)\rangle = \frac{e^{-\hat{H}\tau}\,|\psi(0)\rangle}{\sqrt{\langle\psi(0)|\,e^{-2\hat{H}\tau}\,|\psi(0)\rangle}}. \tag{2.10}$$

As $\tau \to \infty$, the state is guaranteed to go to the ground state for the Hamiltonian $\hat{H}$ given that the initial state overlaps with the ground state [19]. An equivalent formulation to this is to solve the Wick-rotated Schrödinger equation instead [19]:

$$\frac{\partial}{\partial\tau}\,|\psi(\tau)\rangle = -(\hat{H} - E_\tau)\,|\psi(\tau)\rangle\,, \quad E_\tau = \langle\psi(\tau)|\,\hat{H}\,|\psi(\tau)\rangle\,. \tag{2.11}$$

While the guarantee of convergence sounds tempting, there are implementation issues. It is sadly not straightforward to decompose this evolution into a sequence of unitary gates on a quantum computer, and is thus hard to implement as a circuit [19]. Furthermore, these circuits would likely be very deep and thus not NISQ-friendly. Compared to VQE, this is thus not something immediately useful.

2.3.1 Variational QITE

However, there is a way to reformulate QITE in a variational manner [19]. The resulting algorithm, VarQITE, is a hybrid quantum-classical algorithm similar to VQE. In doing this, we sadly lose the strict guarantee of convergence due to approximation. Figure 2.3 depicts the algorithm.

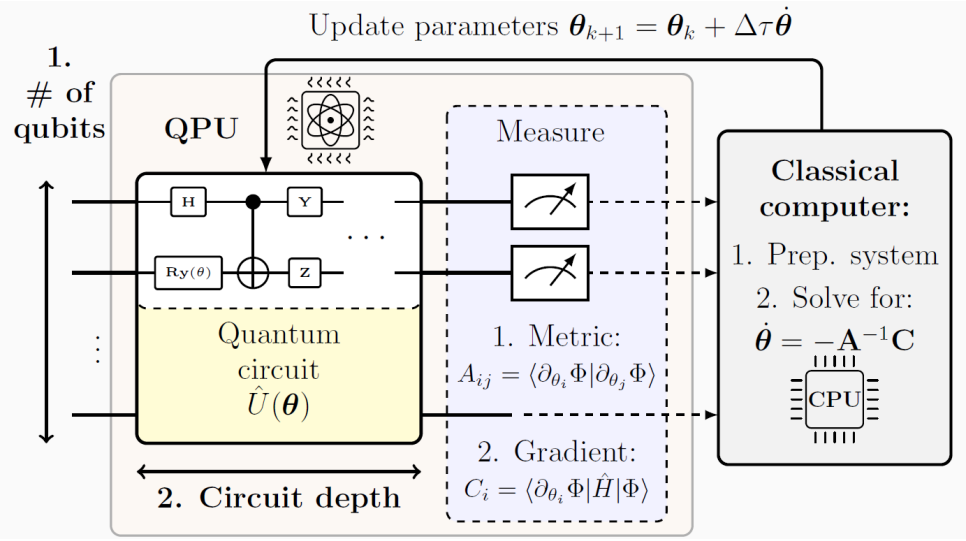

Figure 2.3: An overview of the VarQITE algorithm. Taken with permission from Werner Dobrautz.

To get VarQITE, one approximates the starting state $|\psi\rangle$ as an ansatz, depending on a set of parameters $\vec{\theta}$,

$$|\psi(\tau)\rangle \approx \left|\phi(\vec{\theta}(\tau))\right\rangle \doteq |\phi(\tau)\rangle\,. \tag{2.12}$$

The imaginary time evolution is then approximated using McLachlan's variational principle [20],

$$\delta\left|\left(\frac{\partial}{\partial\tau} + \hat{H} - E_\tau\right)|\psi(\tau)\rangle\right| = 0, \quad ||\psi\rangle| = \sqrt{\langle\psi|\psi\rangle}. \tag{2.13}$$

By expanding in the parameter space and simplifying, one obtains the following differential equation:

$$\sum_j A_{ij}\dot{\theta}_j = C_i, \tag{2.14}$$

where

$$A_{ij} = \mathrm{Re}\left(\frac{\partial\,\langle\phi(\tau)|}{\partial\theta_i}\frac{\partial\,|\phi(\tau)\rangle}{\partial\theta_j}\right), \quad C_i = \mathrm{Re}\left(-\frac{\partial\,\langle\phi(\tau)|}{\partial\theta_i}\hat{H}\,|\phi(\tau)\rangle\right). \tag{2.15}$$

A and C depend on the imaginary time τ. They can be interpreted as the metric $A(\tau)$, related to the quantum Fisher information matrix [21], and the gradient $C(\tau)$, respectively. It is possible to implement VarQITE in such a way that you do not need to measure the energy at every iteration, as long as you track these two objects (as in Figure 2.3) [7].
With these, one can approximate the imaginary time evolution iteratively with, for example, the Euler method [19]. Solving this requires inverting the metric matrix A, which is often ill-conditioned. This can be avoided with Tikhonov regularisation [19].
The reason why McLachlan's variational principle is used rather than perhaps the more familiar Dirac-Frenkel variational principle is that it is the most appropriate

for QITE [22]. This is, amongst other reasons, because it will always give real solutions given real parameters. Given complex parameters the principles are equivalent, but limiting ourselves to real wave functions is thus possible with McLachlan's formulation [22].
An important measure is the McLachlan distance L^2, defined as

$$L^2 = \sum_{i,j} A_{ij}\dot{\theta}_i\dot{\theta}_j - 2\sum_i C_i\dot{\theta}_i + 2\,\mathrm{Var}(\hat{H}). \tag{2.16}$$

The significance of the McLachlan distance is that it shows how accurate the simulation is – a smaller McLachlan distance means a smaller difference between variational and normal QITE, and thus the "optimal" path to the ground state. Having a high McLachlan distance can lead to "walking off the path" to the ground state, and thus ending up in a higher energy state.
VarQITE has several advantages compared to VQE. Firstly, it avoids the optimisation problems that plague VQE when the parameter landscape is mostly flat, and generally performs better than VQE in the presence of noise [19]. Furthermore, it is possible to apply to non-Hermitian problems, and implementation of it can be optimised for this [17]. Overall, it is thus a competitive alternative to VQE, and a NISQ-friendly way of handling non-Hermitian problems, if such should arise.

2.4 Adaptive Ansatzes

Adaptive ansatzes are a class of ansatzes that are iteratively built for the problem to be simulated. The first adaptive ansatz was the Fermionic-ADAPT-VQE by Grimsley *et al.* [2], and several versions have followed from it, optimised for various purposes. The origin of adaptive ansatzes lies in the realisation that the choice of ansatz has a large impact on the quality of a variational quantum simulation – "a simulation is only as good as the ansatz" [2].

2.4.1 Fermionic-ADAPT-VQE

The original adaptive ansatz, ADAPT-VQE (short for Adaptive Derivative-Assembled Psuedo-Trotter), constructs an ansatz from a predefined operator pool consisting of UCCSD excitations, which are chosen depending on their contribution to the gradient to get a maximally compact ansatz.
The algorithm to develop an ansatz in ADAPT-VQE begins with a few preparatory steps:

1. Compute the one-body and two-body integrals in equations (2.6) and (2.7) on classical hardware.
2. Define the operator pool, for example, all UCCSD excitations.
3. Initialise the qubits at some reference state, for example, the HF state.

This is then followed by repeating the following steps until a specified tolerance is reached:

1. Prepare a trial state on the quantum hardware with the current ansatz.
2. Measure the commutator of each operator in the operator pool to get the gradient.

3. If the norm of the gradient is smaller than the tolerance, exit. Otherwise, find the operator in the pool with the largest contribution to the gradient and add it to the ansatz.
4. Reoptimise all parameters and restart.

The reason for measuring the gradient is to find the best operator to add to the ansatz, as the operator with the largest gradient will have the largest impact on energy minimisation. The ADAPT-VQE successfully decreases the circuit depth compared to including the entire operator pool, but the cost paid is more measurements. This is good for NISQ hardware, where circuit depth is one of the limiting factors. Interestingly enough, the ADAPT-VQE ansatz performs better than including all operators in the pool regarding the accuracy, while also minimising parameter count. [2]

2.4.2 Qubit-ADAPT-VQE

An important realisation was made by Tang *et al.* [23]: A parameter-efficient pool is not necessarily gate-efficient. Thus, Fermionic-ADAPT-VQE does not minimise the ansatz depth. Furthermore, it is unclear how to construct the operator pool or whether the constructed ansatz will be good enough for convergence. These are the main factors driving the development of a hardware-efficient ADAPT algorithm called Qubit-ADAPT-VQE.

The main difference lies in the choice of operators in the operator pool – Qubit-ADAPT-VQE has replaced the fermionic operators in its predecessor with a minimised set of Pauli strings, series of repeating Pauli gates. Thus, the amount of gates is decreased. The pool is chosen to be as small as possible while still guaranteeing completeness, i.e., that it can adequately describe the ground state. This minimum pool size is shown to scale linearly with the number of qubits.

The cost for minimising circuit depth is that the number of parameters is no longer strictly minimised, which naturally leads to larger parameter counts than fermionic ADAPT. [23]

2.4.3 ADAPT-QITE – AVQITE

AVQITE is what you get if you apply the adaptive ansatz formalism to variational quantum imaginary time evolution, which was done by Gomes *et al.* [24]. However, as VQE and VarQITE are quite different, there are many differences in implementation specifics.

The initial steps are similar to a normal VarQITE calculation but with an arbitrary first ansatz. During the evolution, the number of operators is allowed to be expanded. This is done by appending operators after each imaginary time step if the conditions to grow the ansatz are met. In ADAPT-VQE, it is simple to decide when to append an operator and which operator to pick, but the same evaluation method does not work in AVQITE. Instead, the ansatz is grown by adding operators which minimise the McLachlan distance (Eq. (2.16)), thus keeping the evolution as close to normal QITE as possible. As an extension of this, operators are added when the McLachlan distance becomes too large.

AVQITE can immediately be realised with a hardware-efficient operator pool like Qubit-ADAPT-VQE by, for example, choosing all the Pauli strings in the UCCSD ansatz as the operator pool. This pool was chosen by Gomes *et al.*, whose results showed over one order of magnitude shallower circuits than the UCCSD ansatz. [24]

2.5 Transcorrelated Method

The transcorrelated (TC) method was introduced by Boys and Handy [1], and is an explicitly correlated method based on factorising the electronic wave function in Jastrow form [25],

$$|\psi\rangle = e^J |\phi\rangle . \tag{2.17}$$

To be a proper Jastrow ansatz, J has to satisfy

$$J = \sum_{i<j} u(\vec{r_i}, \vec{r_j}), \tag{2.18}$$

where u is a symmetric correlation function over electron pairs. Solving the Schrödinger equation with this appearance of $|\psi\rangle$ gives us

$$\hat{H} |\psi\rangle = E |\psi\rangle \Rightarrow \hat{H} e^J |\phi\rangle = E e^J |\phi\rangle \Rightarrow \underbrace{e^{-J} \hat{H} e^J}_{\bar{H}} |\phi\rangle = E |\phi\rangle . \tag{2.19}$$

We have rephrased the equation in terms of $|\phi\rangle$, and if J is chosen smartly we can simplify the problem. Worth noting is that this is not an approximation but an exact similarity transform [26].
In quantum chemistry, explicitly correlated methods are first and foremost applied to deal with Kato's cusp condition [27], when two electrons approach each other and give rise to a dip in the wave function. This sharp, non-differentiable feature of the wave function is hard to simulate with smooth basis functions properly, and thus, many basis functions must be included [7]. The R12 formalism or its improved version F12 is normally used in the field to try and combat this behaviour [28]. With the transcorrelated method, one can lift out the non-differentiable behaviour of $|\psi\rangle$ with a smart Jastrow ansatz J. We have thus effectively transferred the description from the wave function to the Hamiltonian. This would have immense benefits, as one would need significantly fewer quantum resources to simulate the problem – the basis set can be smaller, which means fewer qubits are necessary as every orbital needs a qubit to be represented [28, 17].
However, it is not a free simplification. The price we need to pay is that the Hamiltonian becomes more complex. We get non-Hermiticity and three-body terms – compare to equation (2.5) [28]:

$$\bar{H} = \underbrace{\sum_{pq} h_p^q a_p^\dagger a_q}_{\text{one-body terms}} + \underbrace{\frac{1}{2} \sum_{pqrs} (V_{pq}^{rs} - K_{pq}^{rs}) a_p^\dagger a_q^\dagger a_r a_s}_{\text{two-body terms}} - \underbrace{\frac{1}{6} \sum_{pqrstu} L_{pqr}^{stu} a_p^\dagger a_q^\dagger a_r^\dagger a_s a_t a_u}_{\text{three-body terms}} . \tag{2.20}$$

Here K and L arise from the similarity transformation and are dependent on the choice of u.

As anticipated, we have non-Hermiticity. This invalidates VQE, and we must use QITE for any quantum computation with the transcorrelated method [17]. Previous work has shown that the TC method combines well with projective methods such as QITE. For example, Dobrautz *et al.* [26] found that the approach leads to compact right eigenvectors, which has a great positive effect on efficacy. (Due to the non-Hermiticity, the left and right eigenvectors are now different: $\bar{H}\,|\phi\rangle = E_R\,|\phi\rangle$, $\langle\phi|\,\bar{H} = E_L\,\langle\phi|$.)
As the transcorrelated method introduces extra computational cost by adding three-body terms to the Hamiltonian, it is not obvious that applying the method will be beneficial. To test the method, it has been applied to systems such as the Hubbard model [17] and molecular systems [7]. The result for the Hubbard model is that the method allows highly accurate results to be reached with much shallower ansatz circuits. Furthermore, in all tested examples, Solokov *et al.* [17] found that the TC converges faster than the non-TC test cases. This shows the power of the TC approaches and verifies that the three body terms are not an insurmountable issue. Furthermore, if the implementation of VarQITE allows it, one can to some degree ignore the impact of three body terms at every single iteration until convergence is close. This is possible as convergence can be tracked by measuring the norm of the gradient and metric, as mentioned earlier.
In the molecular systems case, Dobrautz *et al.* [7] found that dissociation energy of lithium hydride could be reached within *experimental* accuracy with less than ten qubits, using a UCCSD ansatz. In the non-TC case, at least ten times more qubits would be needed for this accuracy [7].
Finally, the transcorrelated method has been shown to scale well for larger systems. As smaller basis sets are required for good results, the TC method is thought to require orders of magnitude fewer measurements than non-TC methods. This highlights the strength of the TC approach on NISQ computers, as it may be used to perform actual reliable quantum chemistry calculations on current and near-term quantum hardware [29].

3
Methods and Implementation

The primary goal of the thesis is to evaluate the performance of TC-AVQITE. This will be done by comparing simulations of TC-AVQITE with corresponding AVQITE simulations. As the goal is to show the potential of the algorithm, all simulations will be done noise-free and assuming perfect gates. We want to reach as high accuracy as possible compared to the exact energy, i.e., the lowest eigenenergy obtained from exact diagonalisation. Additionally, we want to append as few operators to the ansatz as possible to minimise circuit depth or, equivalently, minimise the parameter count. This Section will cover various relevant implementation details, review the iteration method, and describe the modelled systems.

3.1 Implementation details

Combining the transcorrelated approach with adaptive ansatzes followed naturally from simply sequentially performing the two. As the TC method produces a Hamiltonian, it can then be used directly in the AVQITE code. The fermionic TC Hamiltonian is non-Hermitian, so as a consequence the Hamiltonian after mapping to qubits will be complex valued. This complication had to be taken into account. The AVQITE code used in this thesis was a modified version of Gomes *et al.*'s code, as used in [24, 30]. The code for performing TC-AVQITE was written in Python and used the two quantum computation specific packages Qiskit [31, 32] and Qutip [33] for convenience.

3.1.1 Encoding

Regarding encoding, only two mappings were considered: The Jordan-Wigner mapping and the parity mapping. Due to its simplicity, the JW mapping was used during the first runs of the code to verify functionality, but eventually, the parity mapping was more or less exclusively used for qubit efficiency reasons. The mappings were compared to each other to look for unexpected differences in Appendix A.

3.1.2 Basis set

The minimal basis set STO-6G was used for all simulated systems except the hydrogen molecule, which was studied using the 631-G basis. This limitation was chosen due to the smaller qubit requirements needed for simulating the systems, as well as the excellent combination with the transcorrelated method, which excels in

smaller basis sets. However, as a result, the simulated systems have little chemical significance. In the field of quantum chemistry, "chemical accuracy" is often used as a measurement of simulation quality. It is the accuracy required to make realistic predictions, and is commonly defined as an error of 1 kcal/mol ($\sim 10^{-4}$ Ha) to the solution in the complete basis set (CBS) limit/experiment [34]. We follow the convention of Ref.[[34]] and use the term "computational accuracy" here when the difference of a quantum calculation and the exact solution in a given (finite) basis set do not exceed 10^{-4} Ha.
These terms are often mixed up in quantum computing literature which can be somewhat misleading – especially as these energies can differ significantly [35].

3.1.3 Operator pools

For every modelled system, an operator pool was generated, dependent on the number of qubits and encoding. The generation was done in two different ways, leading to two types of pools. The first type consisted of all valid Pauli strings on the space spanned by the number of qubits, minus those strings with an odd number of Y operators. These strings can safely be excluded as they have to behave anti-symmetrically.
These pools were usually fairly large, which led to high flexibility in choosing an operator to append in order to decrease the McLachlan distance. However, the cost of a large pool comes in having to iterate over the entire pool when appending an operator, as outlined in section 2.4. This lead to this type of pool scaling poorly with system size. These pools will be referred to as non-UCCSD pools.
The second type of pool was instead created from Qiskit's implementation of the UCCSD ansatz, where all the Pauli strings from a generated UCCSD ansatz were chosen to form the pool. This is the same method outlined in Gomes *et al.*'s paper [24], and the size of these pools was orders of magnitude smaller than the previous type. This meant that iterating over these pools when a new operator was to be appended took much shorter time. However, it can generally lead to higher parameter counts and, thus, a deeper circuit than the other pool type. The reason is that there are fewer operators to choose from to minimise the McLachlan distance, leading to more inefficient choices. These pools will be referred to as UCCSD pools.

3.1.4 Hamiltonian generation

To verify that AVQITE performs better in conjunction with the TC method, we have to be able to supply both a normal and a TC Hamiltonian for each studied system, basis, and mapping. Thus, code had to be written to generate Hamiltonians for every run. The normal Hamiltonians were generated by a code based on Qiskit's problem drivers by letting the Python package PySCF [36] self-consistently calculate the Hamiltonian.
The transcorrelated Hamiltonians were generated by Dobrautz following the workflow outlined in Refs. [7, 37].
The parameters of the Jastrow factor, Eq. (2.18), were optimised with the variational Monte Carlo (VMC) method [38, 39, 37] using the `CASINO` package [40, 41]. Then

the `TCHint` library [42] was used to calculate the 2- and 3-body integrals required to construct the TC molecular Hamiltonian in second quantisation, Eq. (2.20).

3.1.5 Parameters

The code has several parameters which dictate the behavior of the simulation. The most relevant ones are listed below.

- The Mclachlan distance cutoff r_{cut}: if the (square root of the strictly positive) McLachlan distance L is greater than r_{cut}, a new operator is to be appended to the ansatz. Adjusting this number will dramatically modify the reached accuracy and number of parameters in a codependent fashion. Due to this, comparing runs at different values of r_{cut} can be likened to comparing different "levels of theory". The used value of r_{cut} in all simulations was fixed to 10^{-5}, as it was found to give sufficient accuracy.
- The maximum operator addition per iteration cycle $i_{\max}$: The maximum amount of times the code was allowed to check whether to add an operator to the ansatz after reaching the operator addition conditions. A higher value means that the ansatz can very quickly be optimised but at the risk of adding extra operators, which may turn out to be unnecessary for good convergence. The value used in most of the simulations was 8, but different values were studied in Section 4.2.
- The minimum numbers of iteration cycles per operator addition `delay`: like an inverse to $i_{\max}$, how often the ansatz was allowed to add an operator even though the conditions to do so were met. A higher delay means that the imaginary time evolution had to keep evolving with its current operators instead of building the ansatz. The reason for the delay parameter is to let the current ansatz propagate in imaginary time without expanding the ansatz. If this is not done, possibly unnecessary operators could be added. For most simulations, it was set to 0. Together with $i_{\max}$, these two parameters are referred to as the operator addition parameters, and were studied in Section 4.2.
- The imaginary time step $\Delta\tau$: how far we evolve in imaginary time each iteration. This was fixed to 0.05 for all simulated systems.

Ideally, we would want TC-AVQITE to be robust – to perform well regardless of parameter settings. Because of this, comparing results with different parameter settings is prudent – especially the operator addition parameters, as their significance is purely algorithmic.

3.2 Iteration

The initial state was chosen to be the Hartree-Fock state due to it being easy to construct and it having relatively high accuracy.
After constructing a proper operator pool and Hamiltonian, the iteration procedure follows the one outlined in Section 2.4.3.

1. Calculate the energy with $E = \langle\psi| \hat{H} |\psi\rangle$.
2. Check if energy has converged, i.e., if it has not changed by more than 10^{-8}. If yes, finish.

3. If not, calculate the matrix A_{ij} and parameter vector C_i following equation (2.15).
4. Evaluate the McLachlan distance by equation (2.16). If the expansion criterion is met, expand the trial state. This can be repeated $i_{\max}$ times, but if the trial state was expanded less than `delay` iteration steps ago, none will be added.
5. Update the parameters by equation (2.14). Repeat.

The trial state was expanded when the McLachlan distance reached a certain upper threshold – when the (square root of) the McLachlan distance is larger than r_{cut}. The appended operator was chosen from the pool as the one which minimises the McLachlan distance, thus the best fit. An additional criterion was also tried: only if the energy difference between two steps had converged to a tolerance of 10^{-3} Ha was an operator to be appended. See Section 4.2.2 for details.
The reason why the energy is calculated at every time step even though it is not strictly necessary (as detailed in Section 2.3.1, Figure 2.3) is to have simple, conventional convergence checks and to have plenty of data for plotting. This can potentially be effectivised, though.

3.3 Modelled systems

To evaluate TC-AVQITE, three molecules were used as test cases. These were the hydrogen molecule, H_2, quadratic hydrogen, H_4, and lithium hydride, LiH. The molecules are depicted in Figure 3.1.

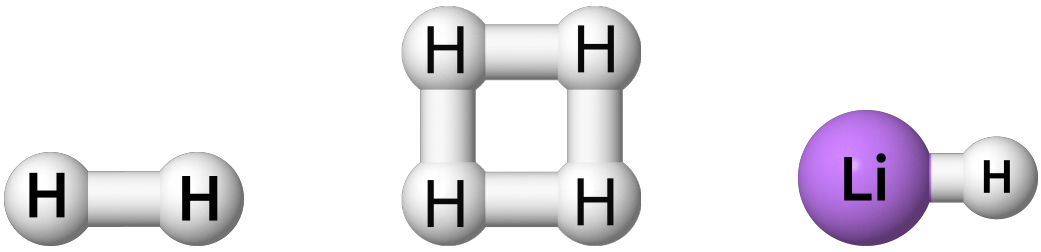

(a) Hydrogen molecule. **(b)** Quadratic H_4. **(c)** Lithium hydride.

Figure 3.1: The molecules treated.

For these molecules, different bond lengths to study were chosen from their potential energy plots to capture different bond characteristics, like bonded, broken bond, and half-broken bond regimes. These distances were 0.7 Å, 1.2 Å, and 3.0 Å for H_2, see Figure 3.2, 1.0 Å, 1.23 Å, and 2.0 Å for H_4, see Figure 3.3, and finally 1.6 Å and 3.0 Å for LiH, see Figure 3.4. For H_4 and LiH, the half-broken bond regime is not captured. For H_4, the regime of "too short" bond is captured instead, and for LiH only two bond lengths were studied. This was due to LiH being considerably slower to simulate due to the exponential scaling of classically simulating qubits.
The hydrogen molecule is an easy system to study as its size makes it the simplest molecule to simulate. Due to this reason, it was chosen for the initial evaluation of TC-AVQITE. However, it has very characteristic properties that do not necessarily hold for larger systems. To illustrate, in a minimal basis with parity mapping and

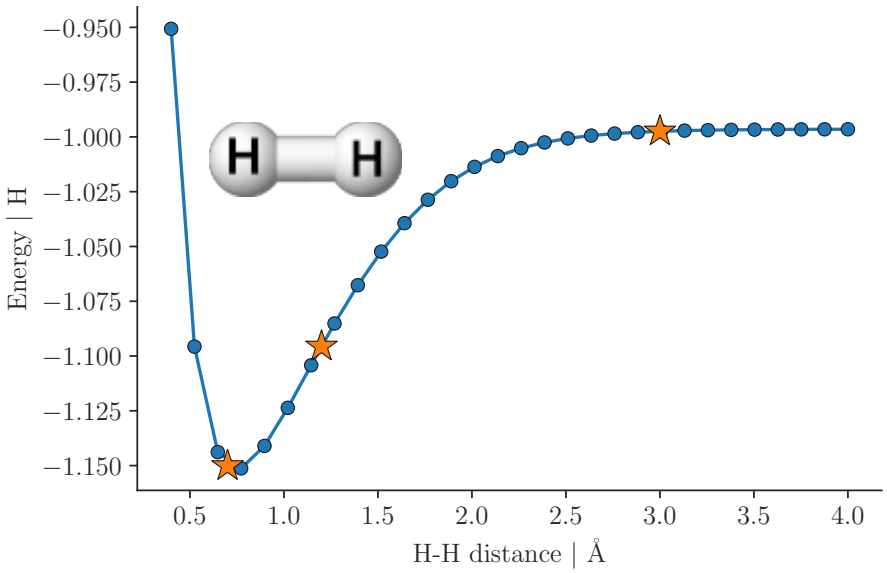

Figure 3.2: Potential energy plot for H_2, with the chosen bond lengths to study marked with a star. These are 0.7 Å, 1.2 Å, and 3.0 Å.

taking point group symmetries into account, one needs as little as one qubit to adequately simulate it [43]. It is thus clear that larger systems need to be investigated as well.

Quadratic H_4 is a very artificial system from a chemist's point of view but is a very interesting system nonetheless. It consists of four hydrogen atoms arranged in a square, and is a highly correlated system. This makes it difficult to study normally, thus a system where the TC method may excel. It was thus natural to use it for evaluating TC-AVQITE.

Lithium hydride was chosen simply due to being one of the simplest heterogeneous molecules to simulate. It is a larger (compared to H_2) molecule that is not highly correlated, which is why it is a useful comparison in combination with H_4.

For every system, the method's performance was evaluated by comparing it with a non-transcorrelated AVQITE simulation. Furthermore, for the transcorrelated method, all systems had evolution by the left and right eigenvector compared. For quadratic hydrogen, TC-AVQITE was also compared to TC-VarQITE for all bond lengths – i.e., removing the adaptive parts of the algorithm.

The algorithm's robustness was probed by adjusting the operator addition parameters $i_{\max}$ and `delay` for all H_2 and H_4 systems. Due to computational cost, LiH was not studied here. The effect of various other parameters was also studied. This included evaluating the operator pools, changing the operator addition criterion, comparing the Jordan-Wigner and parity mappings, and changing the McLachlan distance cutoff. These effects were studied for both the TC and non-TC case for quadratic H_4 with bond distance 1.0 Å to make comparisons, and can be seen in the following Section and in Appendix A.

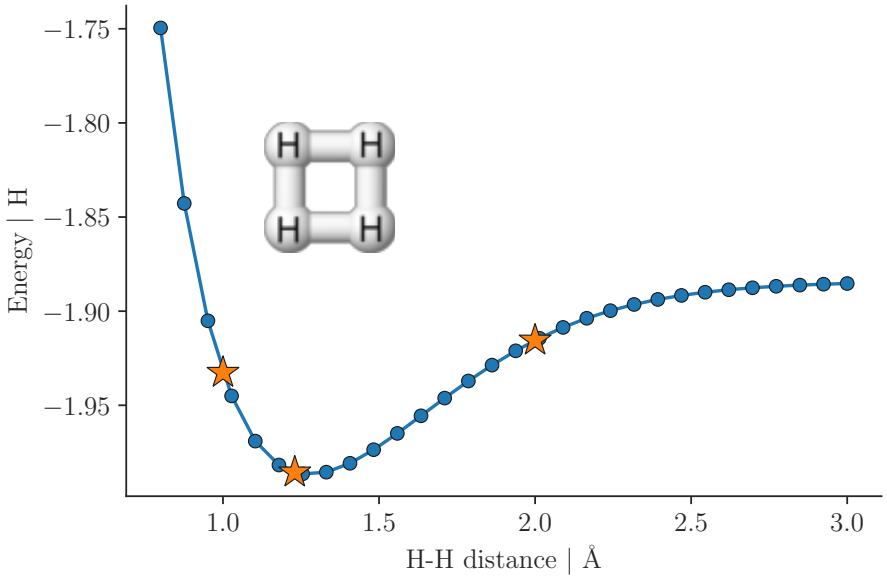

Figure 3.3: Potential energy plot for H_4, with the chosen bond lengths to study marked with a star. These are 1.0 Å, 1.23 Å, and 2.0 Å.

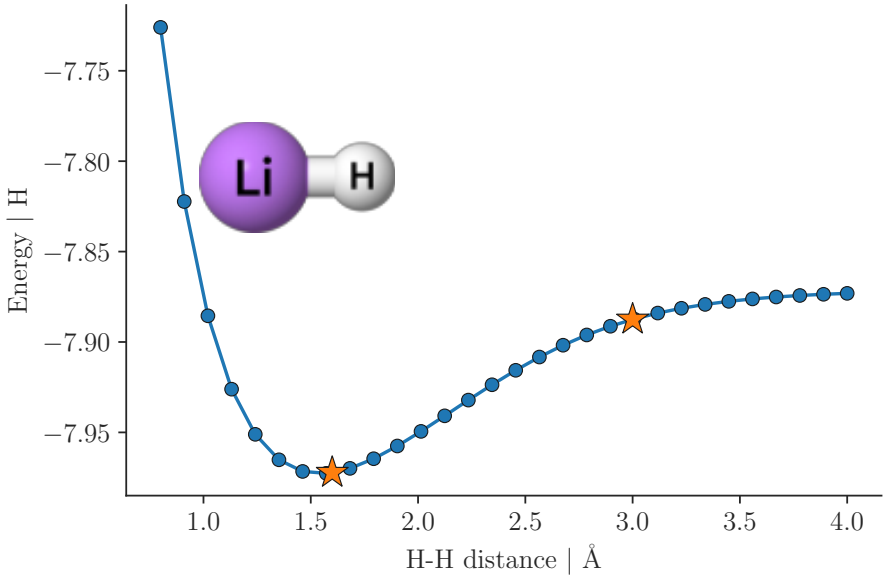

Figure 3.4: Potential energy plot for LiH, with the chosen bond lengths to study marked with a star. These are 1.6 Å and 3.0 Å.

4

Results and Discussion

4.1 Transcorrelated advantage

For every system, TC-AVQITE was compared to a corresponding AVQITE simulation. Furthermore, a comparison was made between the left and right eigenstate evolutions for TC-AVQITE. The left eigenstate evolution usually requires more parameters than the right eigenstate evolution, and characteristically overshoots the target energy to then approach it from below. This likely happens because the transcorrelated method makes the right eigenstate more compact, but therefore makes the left eigenstate less compact than the eigenstate before the transformation.

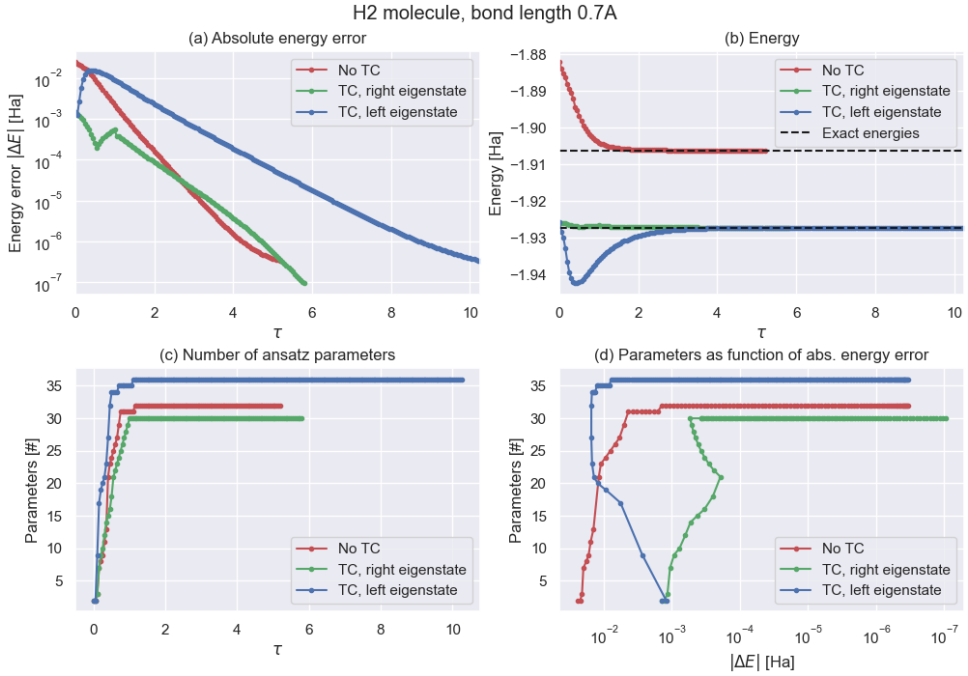

Figure 4.1: Comparing TC-AVQITE and AVQITE for H_2, with bond length 0.7 Å. Subplot (a) shows the absolute difference between the energy and the exact energy as a function of imaginary time τ. Subplot (b) shows the computed energies at imaginary time τ, with the exact energies as dotted lines. Subplot (c) shows the number of parameters as a function of imaginary time τ, what we want to minimise. Subplot (d) is the number of parameters and absolute energy error plotted against each other.

In Figure 4.1, we see a clear advantage for TC-AVQITE in green compared to the normal AVQITE in red. A higher accuracy is reached (as seen in the subplot a) but with fewer parameters (subplot c, or alternatively both properties can be seen at once in subplot d by noting that the TC-AVQITE always is closer to the bottom right corner). This is a positive result, as it means that TC-AVQITE outperforms AVQITE with respect to the number of ansatz parameters and reached accuracy.
Additionally, in subplot b, one of the main strengths of the transcorrelated method is visible – that the ground state energy will be lower for the transcorrelated Hamiltonian. This means that the "exact energy" (which, as mentioned, is only exact with respect to basis set) is closer to the experimental value. The similarity transform causes this; when parts of the complexity of the wave function are moved to the Hamiltonian, a smaller basis set is needed for the same quality of results. Consequently, the same basis set size will have an exact energy closer to the chemical exact energy for the transcorrelated Hamiltonian than for the non-transcorrelated Hamiltonian.

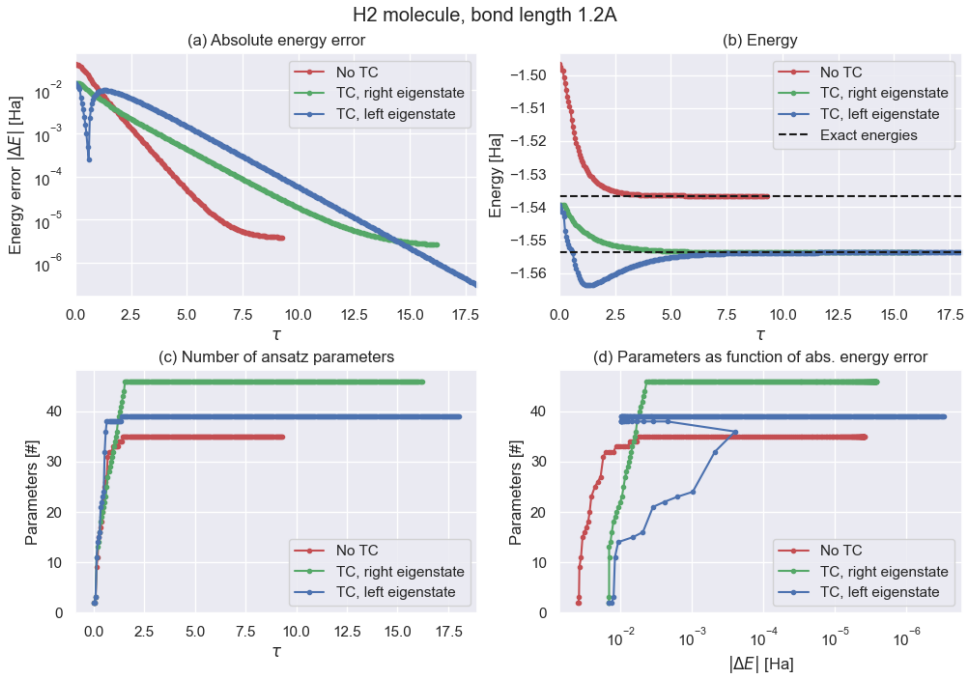

Figure 4.2: Comparing TC-AVQITE and AVQITE for H_2, with bond length 1.2 Å. The sudden, sharp spikes in accuracy are caused by the evolution crossing the exact energy, which can happen due to the non-Hermitian Hamiltonian.

For the half-broken bond regime for H_2, bond length 1.2 Å (Figure 4.2), the picture looks different. Here TC-AVQITE under the right eigenstate reaches only a comparable accuracy to AVQITE, but with *more* parameters, and slower. It is tempting to think that it performs worse than AVQITE for this system, but due to the exact energy being lower, it is still more accurate with respect to experiment. Curiously, the left eigenstate evolution of TC-AVQITE outperforms the right eigenstate evolution, and even reaches a higher accuracy than AVQITE. While this initially is surprising, it is likely caused by the increased compactness of the right eigenstate

being a disadvantage in this half-broken bond regime, and conversely, the less local left eigenstate is advantaged. The reason why TC-AVQITE is slower to converge than AVQITE is the three-body terms in the Hamiltonian (Eq. (2.20)).

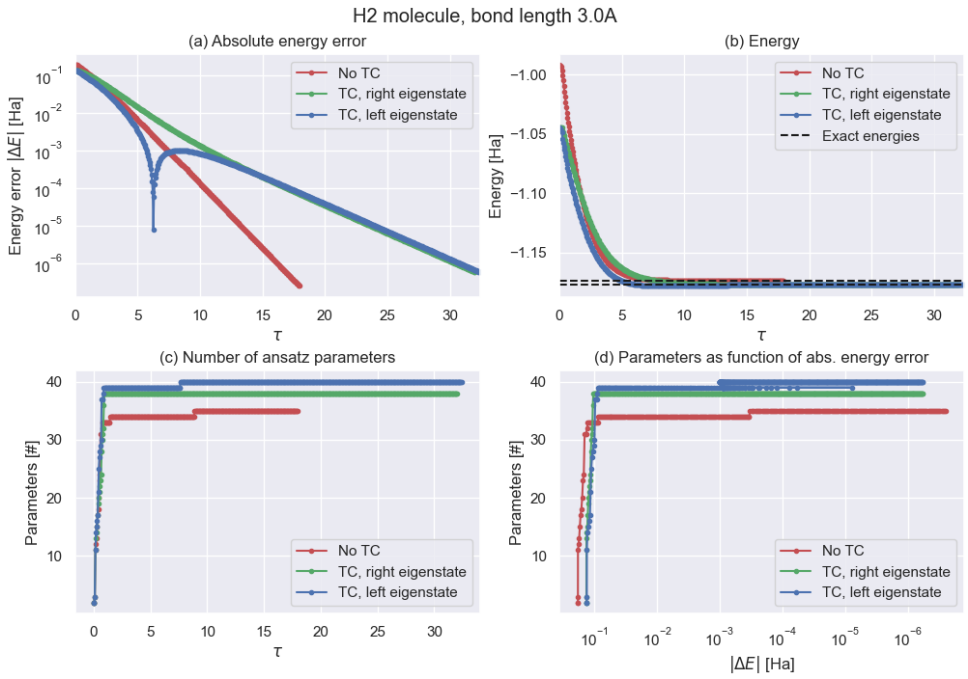

Figure 4.3: Comparing TC-AVQITE and AVQITE for H_2, with bond length 3.0 Å.

In Figure 4.3, we see AVQITE outperforming TC-AVQITE for both eigenstate evolutions, in the sense of imaginary time τ, reached accuracy, and in the number of parameters. The only metric in which TC-AVQITE brings a benefit, in this case, is the total energy, and it is only a slight improvement. In this case, transcorrelation has very little advantage, and as it constitutes extra costs due to the three-body terms of the Hamiltonian, it may not be worth the trouble. Even though the convergence is better for AVQITE, note that the final accuracies of TC-AVQITE are well below 10^{-4} Ha – what in quantum computing context is often referred to as "chemical accuracy" (although, as mentioned in Section 3.1.2, this is not a proper use of terminology due to comparing to the exact energy of the basis set, not experiment). One suggested cause for the transcorrelated method converging worse in this case, is that due to its overall lower energy, it may be harder to properly reach the exact energy than for the non-transcorrelated case.

To summarise, for the hydrogen molecule, there seem to be cases where there is less to gain with transcorrelation for AVQITE. However, as mentioned, this system is very small and symmetric, and one can simulate it to a reasonable accuracy with only one qubit [43]. Therefore, more difficult systems will also need to be treated to draw any conclusions.

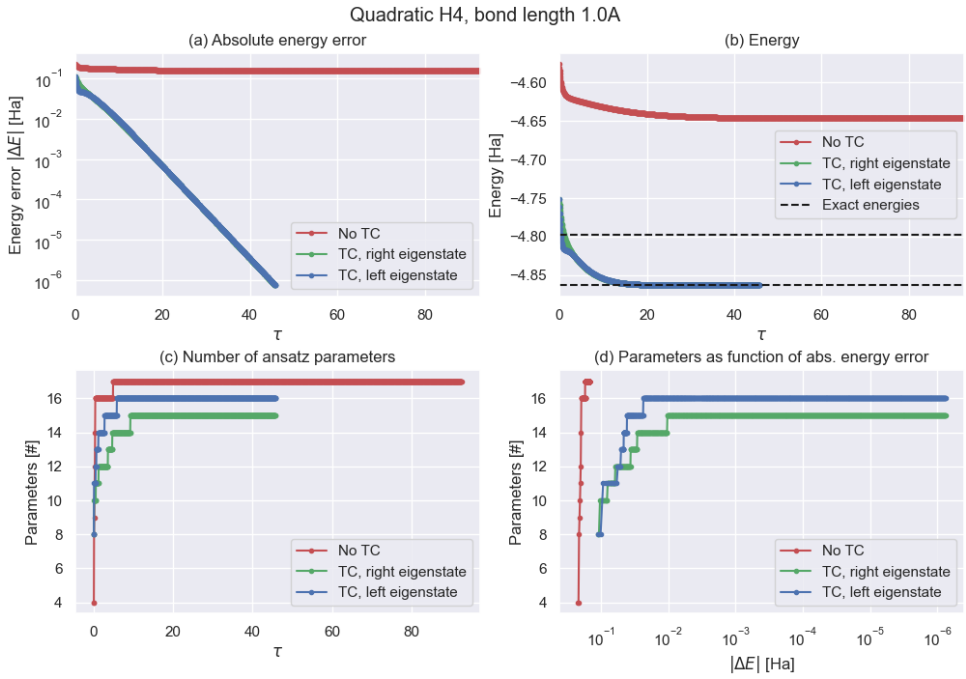

Figure 4.4: Comparing TC-AVQITE and AVQITE for H_4, with bond length 1.0 Å.

The next system studied was quadratic H_4, starting with the compressed regime. In Fig. 4.4, we see significant results – AVQITE converges with a large energy error of 0.1 Ha, while TC-AVQITE, for both eigenstates, approaches the exact energy (specifically, they converge at a difference of 10^{-6} Ha). This shows off the strength of the transcorrelated method; by only transforming the Hamiltonian, much better results are observed. Furthermore, in the number of parameters, TC-AVQITE also outperforms AVQITE, which in combination with the exact energy difference between the non-transcorrelated and transcorrelated Hamiltonian leads to triple the benefit – convergence to the correct result, lower parameter count and thus circuit depth, and higher accuracy to experiment.
In Figure 4.5, the bond length is increased to bonding regime. Like in Fig. 4.4, both TC-AVQITE runs converge well, but here AVQITE manages to converge well too. This may be due to this configuration being less highly correlated than the 1.0 Å case. In the first subfigure, one can see the impact of the three-body terms in the different slopes of AVQITE and TC-AVQITE, and similar to the previous configuration, TC-AVQITE is more parameter efficient while still having a significantly lower total energy.
In Fig. 4.6, we see that the same conclusions appear to hold in the disassociated regime – TC-AVQITE always requires fewer parameters than AVQITE for quadratic H_4. This is not the same result as what was observed for the hydrogen molecule! If this is primarily due to the larger and more complicated system, or if it is due to the system being highly correlated is unclear; a larger, less correlated system will have to be studied as well. Furthermore, are the good convergences to the exact basis set energies in Figs. 4.4, 4.5, and 4.6 for TC-AVQITE caused by the transcorrelation itself, or is there a positive synergy between the transcorrelated method and the adaptive ansatzes? This will be studied further in Section 4.3.

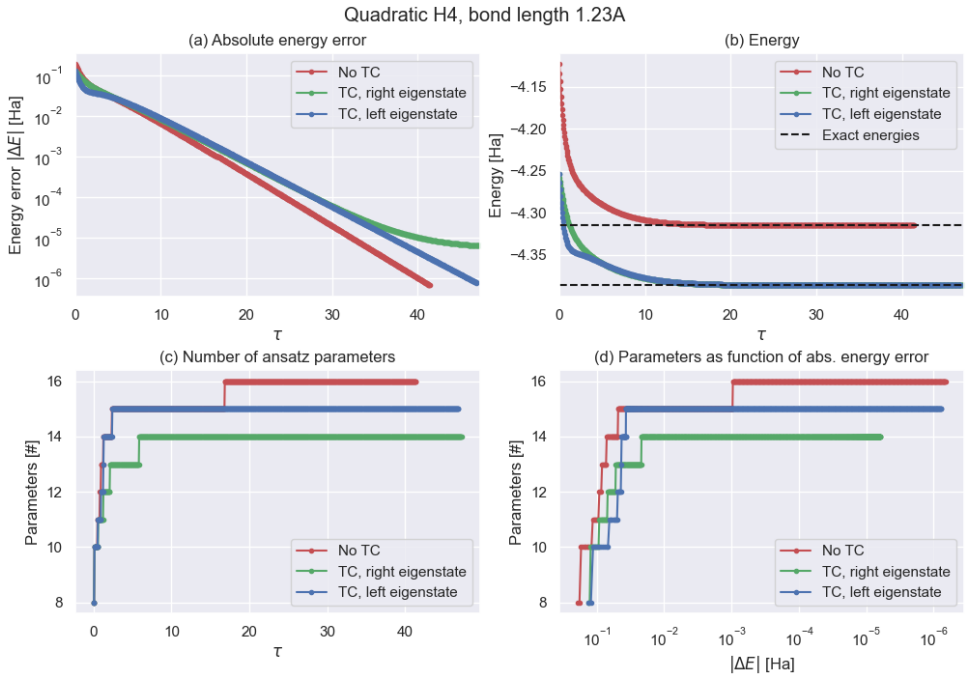

Figure 4.5: Comparing TC-AVQITE and AVQITE for H_4, with bond length 1.23 Å.

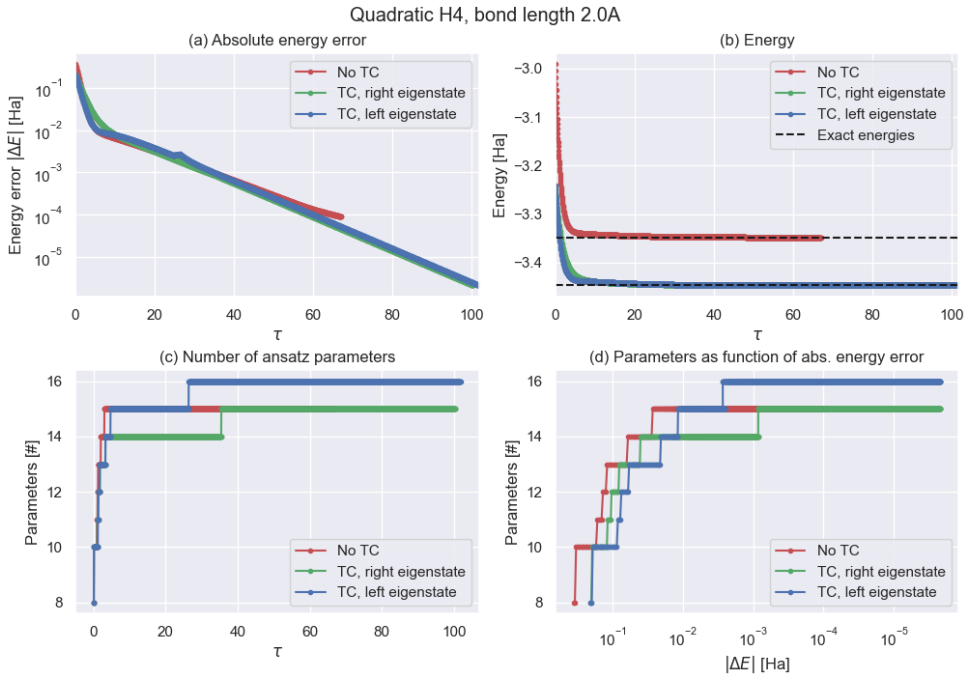

Figure 4.6: Comparing TC-AVQITE and AVQITE for H_4, with bond length 2.0 Å.

The final two systems studied were lithium hydride in the bonding and disassociated regime. As we can see in Fig. 4.7, all methods converge with energy errors well within 10^{-4} Ha, but critically TC-AVQITE needs less than half of AVQITE's parameter count. As lithium hydride is not very strongly correlated, this seems to imply that TC-AVQITE's parameter count advantage is system size dependent. Similar results

are achieved in the disassociated regime shown in Fig. 4.8. With these results, it is clear that TC-AVQITE has an advantage over AVQITE. To continue, the algorithm's robustness under parameter changes must be probed.

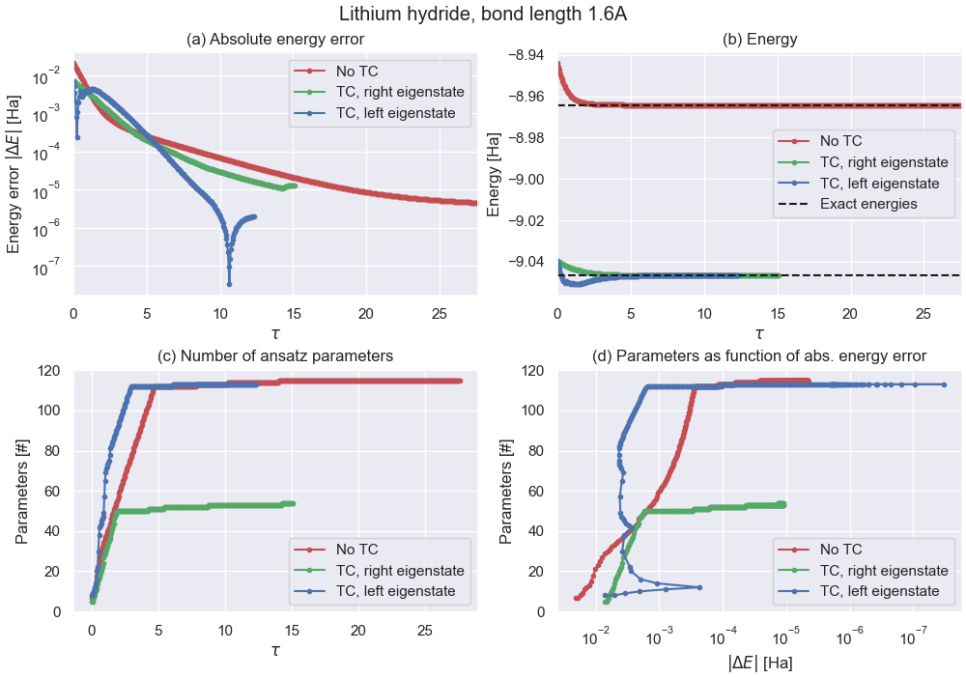

Figure 4.7: Comparing TC-AVQITE and AVQITE for LiH, with bond length 1.6 Å.

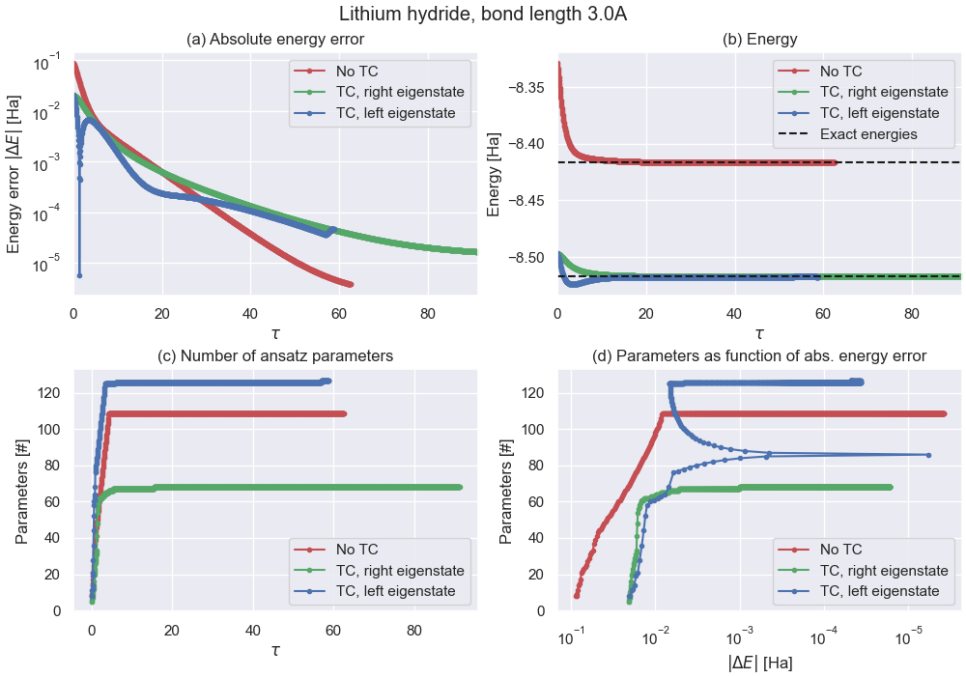

Figure 4.8: Comparing TC-AVQITE and AVQITE for LiH, with bond length 3.0 Å.

4.2 Robustness of the algorithm

4.2.1 Operator addition parameters

To investigate the algorithm's sensitivity to parameter changes, simulations were run with various parameter configurations. Primarily, the parameters studied were the operator addition parameters, i.e., the maximum amount of operators per iteration, $i_{\max}$ and the numbers of iterations between operator addition, `delay`. This was studied for two example systems: H_2 at bond length 0.7 Å and H_4 at bond length 1.0 Å. For both systems, the effect was evaluated for both the UCCSD as well as the non-UCCSD pool detailed in Section 3.1.3.
In Figs.4.9 and 4.10, we see that while changing these parameters has some effect, it is not changing the relation between TC-AVQITE and AVQITE significantly. The final energy errors for all data points are more or less the same, with one outlier for the case with the highest delay for H_2. In this outlier, the delay between adding operators was large enough that the simulation converged without ever adding a new operator, which naturally resulted in a high error and low parameter count.

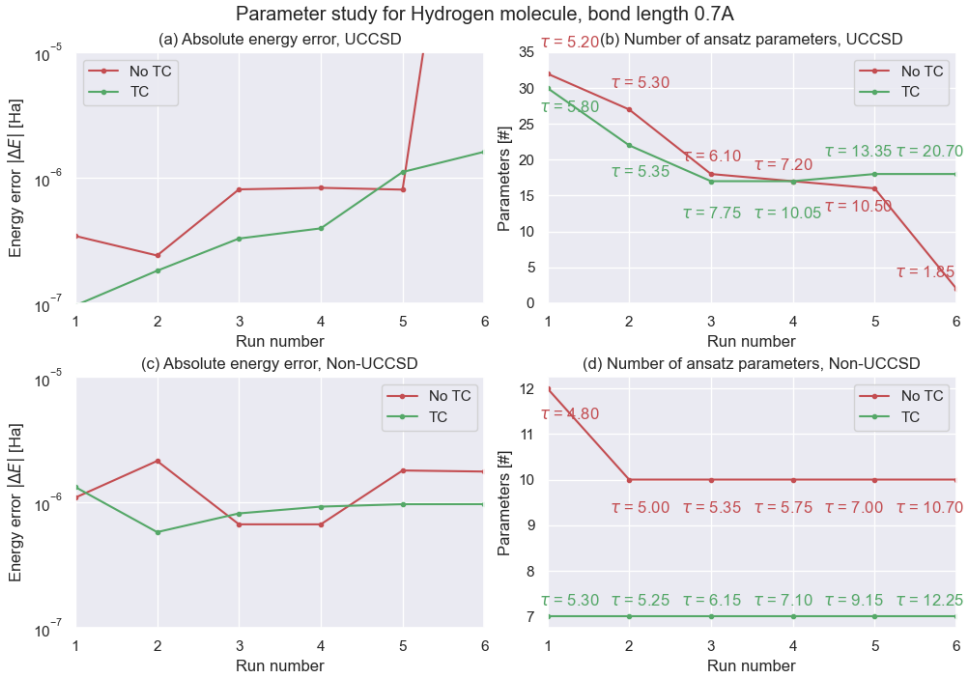

Figure 4.9: A study of the effect of the operator addition parameters for the hydrogen molecule. Run 1 had no delay and a maximum of 8 operators added per iteration. Run 2 had no delay, and a maximum of 1 operator. Run 3 had a delay of 3, run 4 had a delay of 5, run 5 had a delay of 10, and run 6 had a delay of 20. The τ at every point is the imaginary time necessary for convergence.

There is some variation in the final parameter counts, and the trend showed that longer delays seemed to decrease parameter counts at the cost of longer simulations in imaginary time. However, it did not seem to favour either TC-AVQITE or AVQITE, which is why it could be safely ignored. To conclude, the effect of modi-

fying the operator addition parameters did influence the total parameter count but appeared to do so equally between the compared methods. As this was at the cost of lengthier calculations, as well as risking premature convergence like AVQITE run 6 in figure 4.9, it was not investigated further.

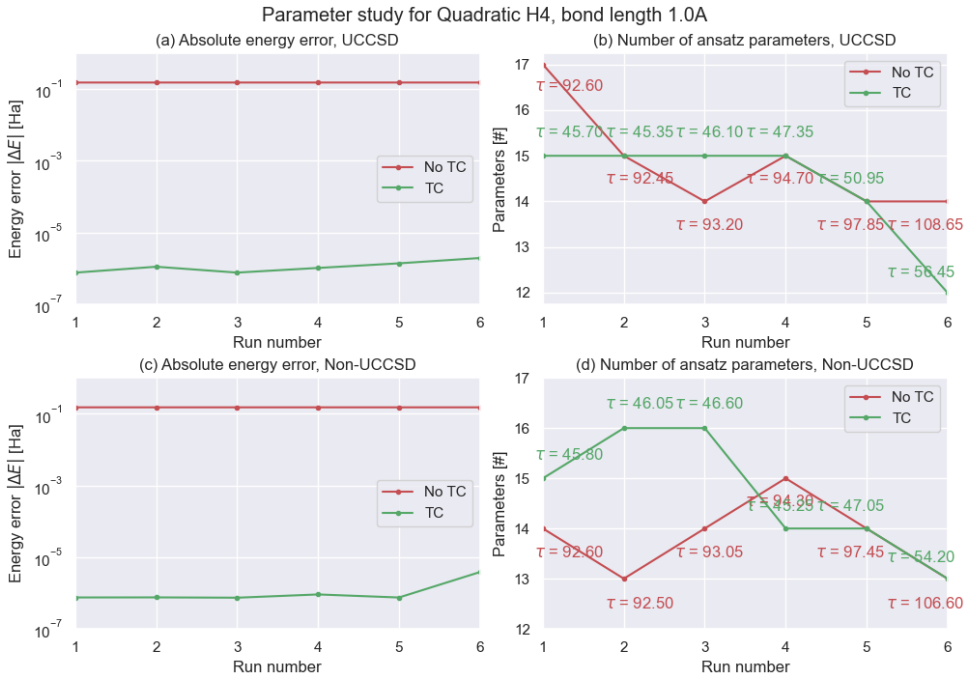

Figure 4.10: A study of the effect of the operator addition parameters for quadratic H_4. Run 1 had no delay and a maximum of 8 operators added per iteration. Run 2 had no delay, and a maximum of 1 operator. Run 3 had a delay of 5, run 4 had a delay of 10, run 5 had a delay of 20, and run 6 had a delay of 50. The τ at every point is the imaginary time necessary for convergence.

4.2.2 Operator addition criterion

The ordinary method of deciding when an operator should be added was by comparing the McLachlan distance with some cutoff value r_{cut}, as detailed in Section 3.1.5. An attempt was made to optimise this operator addition criterion by adding an extra requirement. This extra requirement was that the energy difference between two iterations should be smaller than 10^{-3} Ha before any operator should be appended. This different operator addition criterion (labeled "Convergence") was then compared with the ordinary one (labeled "McLachlan"), as seen in Figures 4.11, 4.12, and 4.13.

As is evident from Figs 4.11, 4.12, and 4.13, the new method performs about the same as the ordinary method for H_4, but tends towards a slightly lower parameter count at the cost of some accuracy. As the benefits were slight, it was not studied further.

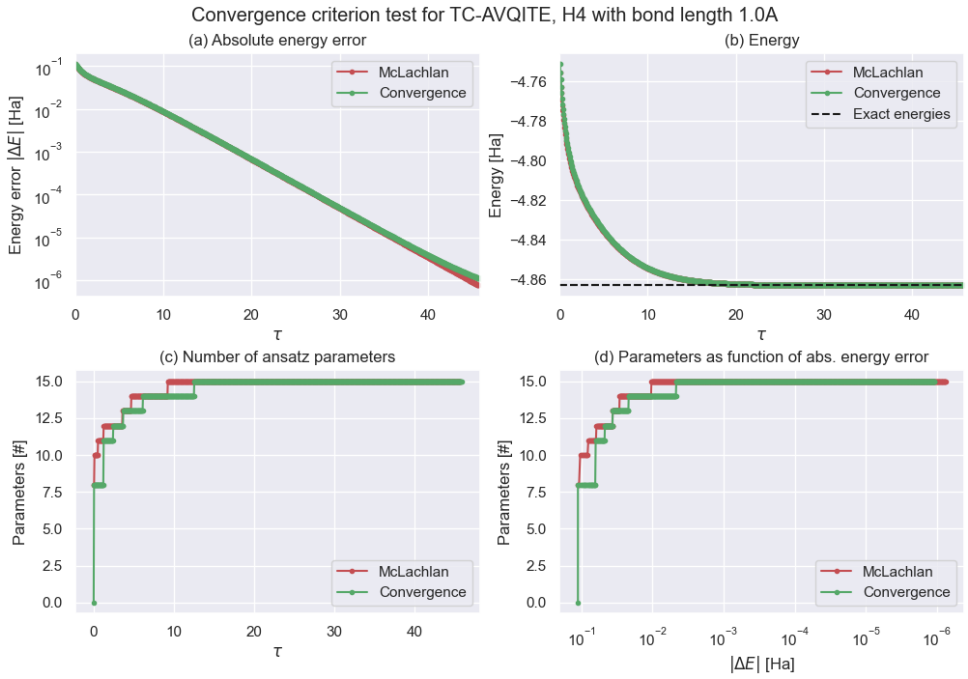

Figure 4.11: Comparing the two different methods of ansatz expansion for H_4 with bond length 1.0 Å. The ordinary method of only considering the McLachlan distance has strictly better convergence for this case.

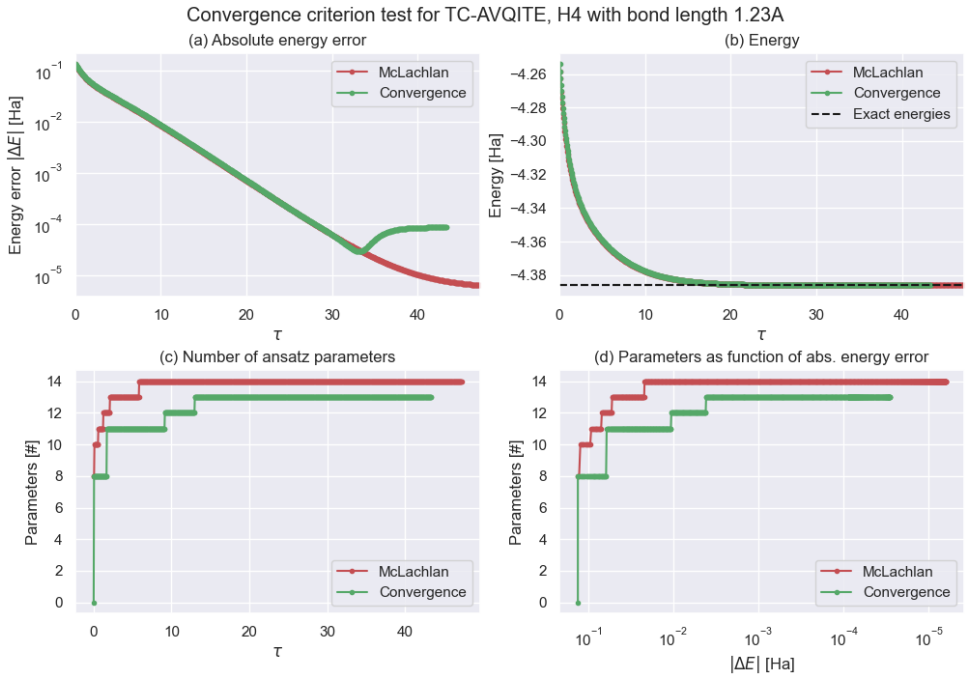

Figure 4.12: Comparing the two different methods of ansatz expansion for H_4 with bond length 1.23 Å. Here, the new method has one less parameter at the cost of accuracy.

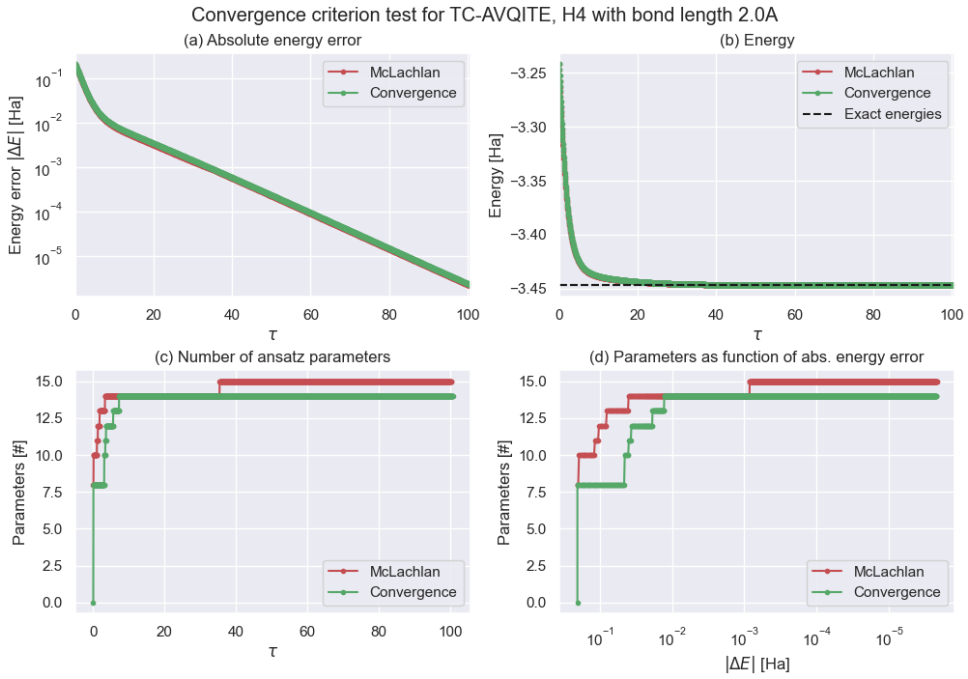

Figure 4.13: Comparing the two different methods of ansatz expansion for H_4 with bond length 2.0 Å. For this system the new method has one less parameter at no cost.

4.3 AVQITE vs VarQITE

In Figs. 4.4, 4.5, and 4.6, TC-AVQITE was observed to outcompete AVQITE. The question was raised whether this is due to only transcorrelation, or a synergy between the transcorrelated method and the adaptive ansatzes. To study this, regular TC-VarQITE (Note the absence of A – adaptive) was simulated for H_4, by disabling the adaptive algorithm and instead adding all parameters at once in the start of the simulation. The results are shown in Figs. 4.14, 4.15, and 4.16.
From the figures, we can see that TC-AVQITE outperforms TC-VarQITE in terms of final energy error. Naturally, the circuit depth of the adaptive algorithm is lower, as the circuit for VarQITE is fixed as the maximum length AVQITE can reach. With this result and the results in Section 4.1, it becomes clear that there is some positive effect of using both adaptive ansatzes and the transcorrelated method together, as transccorrelation alone was not enough to replicate the results of TC-AVQITE.

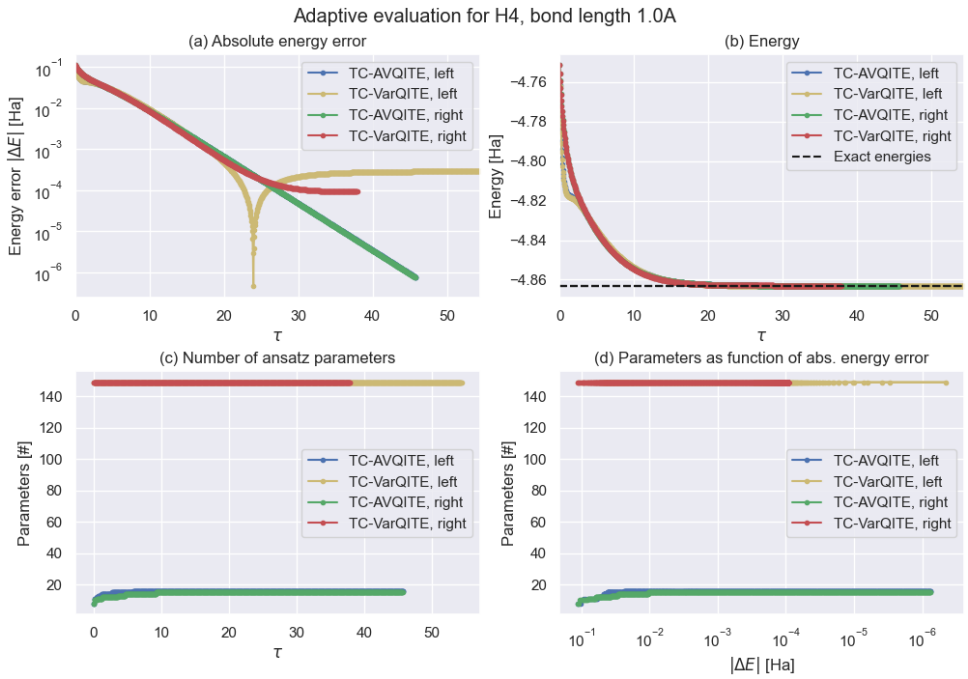

Figure 4.14: Evaluating the adaptive part of the algorithm for H_4 with bond length 1.0 Å by comparing to a run where all operators are immediately appended.

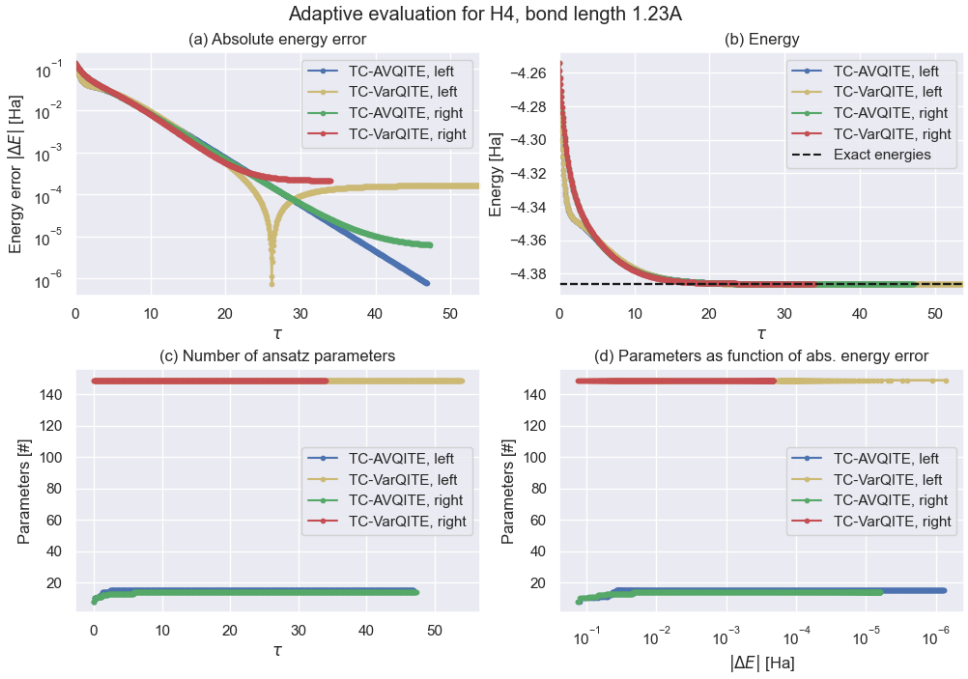

Figure 4.15: Evaluating the adaptive part of the algorithm for H_4 with bond length 1.23 Å by comparing to a run where all operators are immediately appended.

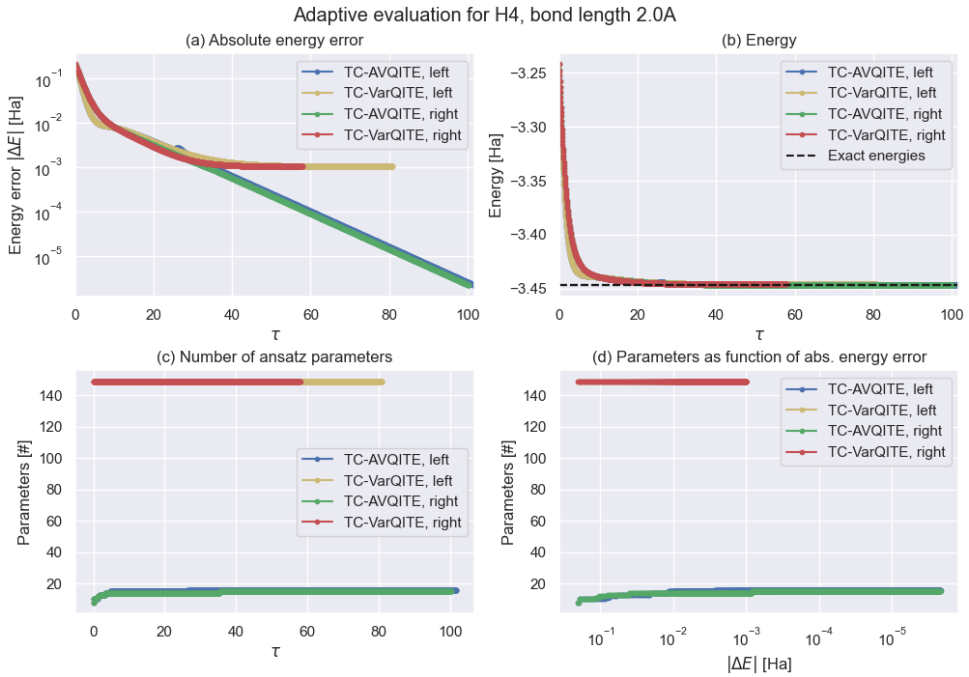

Figure 4.16: Evaluating the adaptive part of the algorithm for H_4 with bond length 2.0 Å by comparing to a run where all operators are immediately appended.

5
Conclusion

With the results in Section 4.1, we have now shown that TC-AVQITE is a good method for calculating ground state energies for small molecular systems, and that the advantage appears to scale with system size. Furthermore, in Section 4.2 we showed that the method is fairly insensitive to choice of parameters such as how operators are added. In Section 4.3, we have also shown that the transcorrelated method and adaptive ansatzes appear to benefit from each other as well. With this, the main questions of the thesis has been answered.

The version of TC-AVQITE developed in this thesis is fairly simple, and while several parameters have been studied, it is not necessarily optimised. Dobrautz *et al.* raise in their paper that one could possibly ignore the three-body terms in the Hamiltonian by tracking the metric A_{ij} and gradient C_i, which would help the TC-AVQITE method further [7]. Furthermore, it might be possible to parallelise the code for faster simulations.

There is naturally room for further evaluation of TC-AVQITE. A large restriction of the thesis has been that the effects were studied in either minimal bases or on very small, symmetric systems. To generalise findings, larger systems would be good to study, and in larger basis sets. As the benefits of TC-AVQITE over AVQITE appear to scale, larger systems should see an even bigger benefit.

Finally, no noise has been taken into consideration. This could be amended by adding simulated noise. This would be interesting to probe as noise is one of the great limiters of NISQ computation, as well as that variational QITE is said to perform well under noisy conditions [19]. After this, running the method on actual quantum hardware seems a logical next step for method maturity.

Bibliography

[1] Samuel Francis Boys and Nicholas Charles Handy. "A condition to remove the indeterminacy in interelectronic correlation functions". In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 309.1497 (1969-03), pp. 209–220. DOI: 10.1098/rspa.1969.0038. URL: https://doi.org/10.1098/rspa.1969.0038.

[2] Harper R. Grimsley et al. "An adaptive variational algorithm for exact molecular simulations on a quantum computer". In: *Nature Communications* 10.1 (2019-07). DOI: 10.1038/s41467-019-10988-2. URL: https://doi.org/10.1038/s41467-019-10988-2.

[3] Bela Bauer et al. "Quantum Algorithms for Quantum Chemistry and Quantum Materials Science". In: *Chemical Reviews* 120.22 (2020-10), pp. 12685–12717. DOI: 10.1021/acs.chemrev.9b00829. URL: https://doi.org/10.1021/acs.chemrev.9b00829.

[4] Alan Aspuru-Guzik et al. "Simulated Quantum Computation of Molecular Energies". In: *Science* 309.5741 (2005-09), pp. 1704–1707. DOI: 10.1126/science.1113479. URL: https://doi.org/10.1126/science.1113479.

[5] Sam McArdle et al. "Quantum computational chemistry". In: *Rev. Mod. Phys.* 92 (1 2020-03), p. 015003. DOI: 10.1103/RevModPhys.92.015003. URL: https://link.aps.org/doi/10.1103/RevModPhys.92.015003.

[6] Yudong Cao et al. "Quantum Chemistry in the Age of Quantum Computing". In: *Chemical Reviews* 119.19 (2019-08), pp. 10856–10915. DOI: 10.1021/acs.chemrev.8b00803. URL: https://doi.org/10.1021/acs.chemrev.8b00803.

[7] Werner Dobrautz et al. *Ab Initio Transcorrelated Method enabling accurate Quantum Chemistry on near-term Quantum Hardware*. 2023. DOI: 10.48550/ARXIV.2303.02007. URL: https://arxiv.org/abs/2303.02007.

[8] A Yu (alexei Yu) Kitaev. *Classical and Quantum Computation*. en. Graduate Studies in Mathematics. Providence, RI: American Mathematical Society, 2002-08.

[9] Seunghoon Lee et al. "Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry". In: *Nature Communications* 14.1 (2023-04). DOI: 10.1038/s41467-023-37587-6. URL: https://doi.org/10.1038/s41467-023-37587-6.

[10] Sevag Gharibian et al. *Improved Hardness Results for the Guided Local Hamiltonian Problem*. 2022. arXiv: 2207.10250 [quant-ph].

[11] Sevag Gharibian and François Le Gall. "Dequantizing the Quantum Singular Value Transformation: Hardness and Applications to Quantum Chemistry and the Quantum PCP Conjecture". In: STOC 2022. Rome, Italy: Association for

Computing Machinery, 2022, pp. 19–32. ISBN: 9781450392648. DOI: 10.1145/3519935.3519991. URL: https://doi.org/10.1145/3519935.3519991.

[12] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. "The Bravyi-Kitaev transformation for quantum computation of electronic structure". In: *The Journal of Chemical Physics* 137.22 (2012-12), p. 224109. DOI: 10.1063/1.4768229. URL: https://doi.org/10.1063/1.4768229.

[13] David J Griffiths and Darrell F Schroeter. *Introduction to Quantum Mechanics*. 3rd ed. Cambridge, England: Cambridge University Press, 2018-08.

[14] Jonathan Romero et al. "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz". In: *Quantum Science and Technology* 4.1 (2018-10), p. 014008. DOI: 10.1088/2058-9565/aad3e4. URL: https://doi.org/10.1088/2058-9565/aad3e4.

[15] Abhinav Kandala et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets". In: *Nature* 549.7671 (2017-09), pp. 242–246. DOI: 10.1038/nature23879. URL: https://doi.org/10.1038/nature23879.

[16] Jarrod R. McClean et al. "Barren plateaus in quantum neural network training landscapes". In: *Nature Communications* 9.1 (2018-11). DOI: 10.1038/s41467-018-07090-4. URL: https://doi.org/10.1038/s41467-018-07090-4.

[17] Igor O. Sokolov et al. *Orders of magnitude reduction in the computational overhead for quantum many-body problems on quantum computers via an exact transcorrelated method*. 2022. DOI: 10.48550/ARXIV.2201.03049. URL: https://arxiv.org/abs/2201.03049.

[18] Mario Motta et al. "Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution". In: *Nature Physics* 16.2 (2019-11), pp. 205–210. DOI: 10.1038/s41567-019-0704-4. URL: https://doi.org/10.1038/s41567-019-0704-4.

[19] Sam McArdle et al. "Variational ansatz-based quantum simulation of imaginary time evolution". In: *npj Quantum Information* 5.1 (2019-09). DOI: 10.1038/s41534-019-0187-2. URL: https://doi.org/10.1038/s41534-019-0187-2.

[20] A.D. McLachlan. "A variational solution of the time-dependent Schrodinger equation". In: *Molecular Physics* 8.1 (1964-01), pp. 39–44. DOI: 10.1080/00268976400100041. URL: https://doi.org/10.1080/00268976400100041.

[21] James Stokes et al. "Quantum Natural Gradient". In: *Quantum* 4 (2020-05), p. 269. ISSN: 2521-327X. DOI: 10.22331/q-2020-05-25-269. URL: https://doi.org/10.22331/q-2020-05-25-269.

[22] Xiao Yuan et al. "Theory of variational quantum simulation". In: *Quantum* 3 (2019-10), p. 191. DOI: 10.22331/q-2019-10-07-191. URL: https://doi.org/10.22331/q-2019-10-07-191.

[23] Ho Lun Tang et al. "Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor". In: *PRX Quantum* 2 (2 2021-04), p. 020310. DOI: 10.1103/PRXQuantum.2.020310. URL: https://link.aps.org/doi/10.1103/PRXQuantum.2.020310.

[24] Niladri Gomes et al. “Adaptive Variational Quantum Imaginary Time Evolution Approach for Ground State Preparation”. In: *Advanced Quantum Technologies* 4.12 (2021-10), p. 2100114. DOI: 10.1002/qute.202100114. URL: https://doi.org/10.1002/qute.202100114.

[25] Robert Jastrow. “Many-Body Problem with Strong Forces”. In: *Physical Review* 98.5 (1955-06), pp. 1479–1484. DOI: 10.1103/physrev.98.1479. URL: https://doi.org/10.1103/physrev.98.1479.

[26] Werner Dobrautz, Hongjun Luo, and Ali Alavi. “Compact numerical solutions to the two-dimensional repulsive Hubbard model obtained via nonunitary similarity transformations”. In: *Phys. Rev. B* 99 (7 2019-02), p. 075119. DOI: 10.1103/PhysRevB.99.075119. URL: https://link.aps.org/doi/10.1103/PhysRevB.99.075119.

[27] Russell T Pack and W. Byers Brown. “Cusp Conditions for Molecular Wavefunctions”. In: *The Journal of Chemical Physics* 45.2 (1966-07), pp. 556–559. DOI: 10.1063/1.1727605. URL: https://doi.org/10.1063/1.1727605.

[28] Aron J. Cohen et al. “Similarity transformation of the electronic Schrödinger equation via Jastrow factorization”. In: *The Journal of Chemical Physics* 151.6 (2019-08), p. 061101. DOI: 10.1063/1.5116024. URL: https://doi.org/10.1063/1.5116024.

[29] *IBM Quantum Computing | Roadmap — ibm.com*. https://www.ibm.com/quantum/roadmap. [Accessed 09-Mar-2023].

[30] Yongxin Yao. *avqite*. https://gitlab.com/gqce/avqite. 2022.

[31] Qiskit contributors. *Qiskit: An Open-source Framework for Quantum Computing*. 2023. DOI: 10.5281/zenodo.2573505.

[32] The Qiskit Nature developers and contributors. “Qiskit Nature 0.6.0”. In: (2023-04). Qiskit Nature has some code that is included under other licensing. These files have been removed from the zip repository provided here and are only available via Github. See https://github.com/Qiskit/qiskit-nature for more details. DOI: 10.5281/zenodo.7828768.

[33] F. Nori J. R. Johansson P. D. Nation. *QuTiP 2: A Python framework for the dynamics of open quantum systems*. 2013. DOI: 10.1016/j.cpc.2012.11.019.

[34] Phalgun Lolur et al. “Reference-State Error Mitigation: A Strategy for High Accuracy Quantum Computation of Chemistry”. In: *Journal of Chemical Theory and Computation* 19.3 (2023-01), pp. 783–789. DOI: 10.1021/acs.jctc.2c00807. URL: https://doi.org/10.1021/acs.jctc.2c00807.

[35] NIST Computational Chemistry Comparison and Benchmark Database. *NIST Standard Reference Database Number 101*. Ed. by Russell D. Johnson III. 2022-05. URL: http://cccbdb.nist.gov/.

[36] Qiming Sun et al. “Recent developments in the PySCF program package”. In: *The Journal of Chemical Physics* 153.2 (2020-07), p. 024109. DOI: 10.1063/5.0006074. URL: https://doi.org/10.1063/5.0006074.

[37] J. Philip Haupt et al. “Optimizing Jastrow factors for the transcorrelated method”. In: *The Journal of Chemical Physics* 158.22 (2023-06). DOI: 10.1063/5.0147877. URL: https://doi.org/10.1063/5.0147877.

[38] D. Ceperley, G. V. Chester, and M. H. Kalos. “Monte Carlo simulation of a many-Fermion study”. In: *Phys. Rev. B* 16.7 (1977-10), pp. 3081–3099. DOI:

10.1103/physrevb.16.3081. URL: https://doi.org/10.1103/physrevb.16.3081.

[39] W. M. C. Foulkes et al. "Quantum Monte Carlo simulations of solids". In: *Rev. Mod. Phys.* 73 (1 2001-01), pp. 33–83. DOI: 10.1103/RevModPhys.73.33. URL: https://link.aps.org/doi/10.1103/RevModPhys.73.33.

[40] R. J. Needs et al. "Variational and diffusion quantum Monte Carlo calculations with the CASINO code". In: *J. Chem. Phys.* 152.15 (2020-04), p. 154106. DOI: 10.1063/1.5144288. URL: https://doi.org/10.1063/1.5144288.

[41] P. López Ríos et al. "Inhomogeneous backflow transformations in quantum Monte Carlo calculations". In: *Phys. Rev. E* 74 (6 2006-12), p. 066701. DOI: 10.1103/PhysRevE.74.066701. URL: https://link.aps.org/doi/10.1103/PhysRevE.74.066701.

[42] Pablo López Ríoz et al. "TCHINT Library". to be published. 2023.

[43] Kanav Setia et al. "Reducing Qubit Requirements for Quantum Simulations Using Molecular Point Group Symmetries". In: *Journal of Chemical Theory and Computation* 16.10 (2020-08), pp. 6091–6097. DOI: 10.1021/acs.jctc.0c00113. URL: https://doi.org/10.1021/acs.jctc.0c00113.

A

Additional results

A.1 Operator pool evaluation

To justify discarding the non-UCCSD pools, the performance of the UCCSD and non-UCCSD pools were compared for H_4 at bond length 1.23 Å. This system was picked due to both TC-AVQITE and AVQITE converging, which means both results are interesting to compare. The comparison is shown in figure A.1.

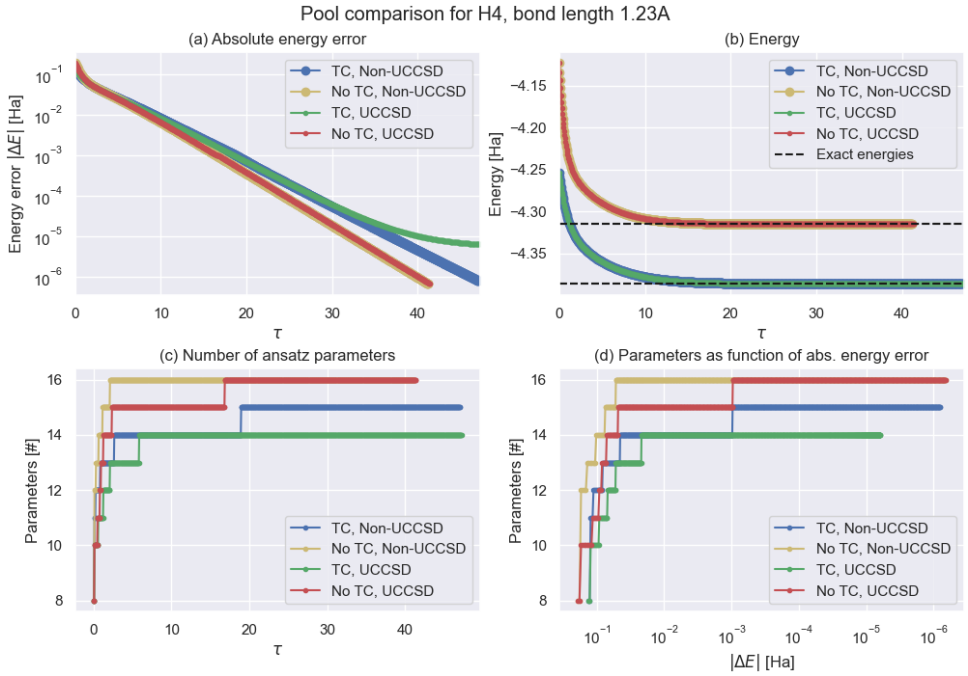

Figure A.1: The performance of the different pool types in both the normal and transcorrelated Hamiltonian case for quadratic H_4 with bond length 1.23 Å.

In subfigure a, one can see that there is only a small difference in performance between the pools. When it comes to number of ansatz parameters, the TC-AVQITE run appears to benefit slightly while AVQITE has the same performance. The cost for this is paid in earlier convergence. All in all, the difference is very slight though, and as mentioned in section 3.1.3 the non-UCCSD pools are computationally more expensive. Due to no clear advantage, we thus only used the UCCSD pools for easier computation.

A.2 Mapping evaluation

The parity mapping and JW mapping were also compared for H_4 with bond length 1.0 Å. The results are shown in figure A.2. As is evident in subfigure c, the transcorrelated method decrease the number of necessary operators to be added in both cases. This decrease is larger for the JW mapping, but as the JW mapping requires two additional qubits this relative circuit depth gain was completely negated by higher quantum resource cost.

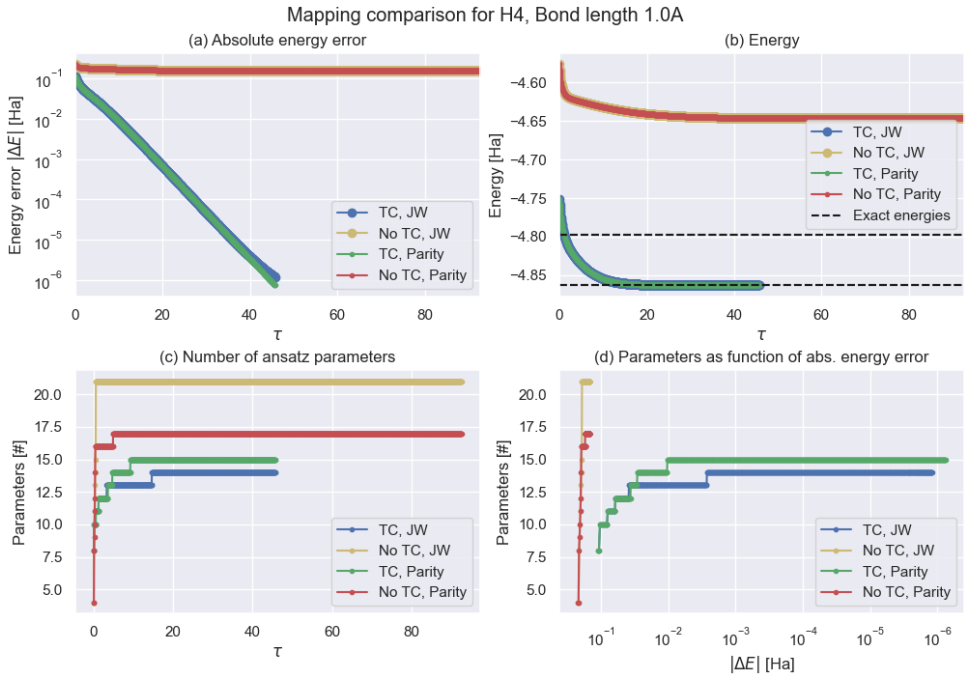

Figure A.2: The performance of the different mappings in both the normal and the transcorrelated Hamiltonian case for quadratic H_4 with bond length 1.0 Å.

A.3 Condition numbers of A-matrix during AVQITE and VarQITE comparison

As we can see in figures A.3, A.4, and A.5, the condition numbers are overall quite large, but significantly larger in the VarQITE cases. This can to some extent explain the poor performance of the VarQITE cases. Additionally, the condition numbers of the left eigenstate evolution were always larger than the more compact right eigenstate evolution. The large condition numbers motivate the use of Tikhonov regularisation, if nothing else.

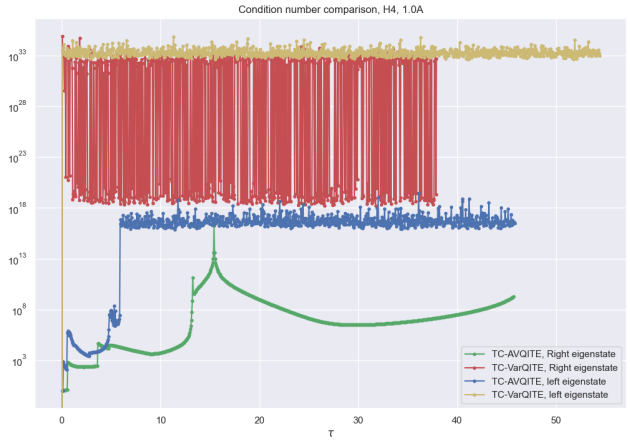

Figure A.3: The condition numbers of the runs in figure 4.14.

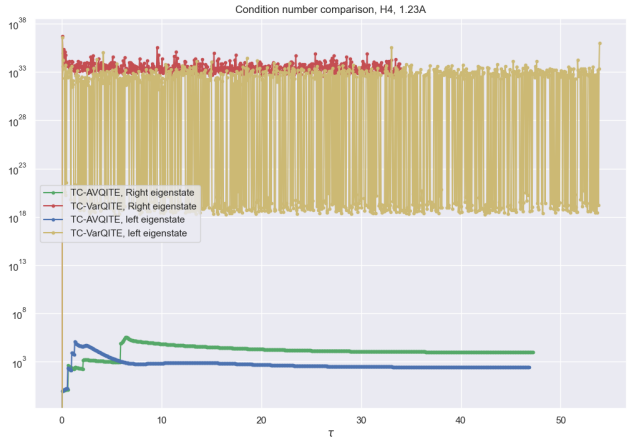

Figure A.4: The condition numbers of the runs in figure 4.15.

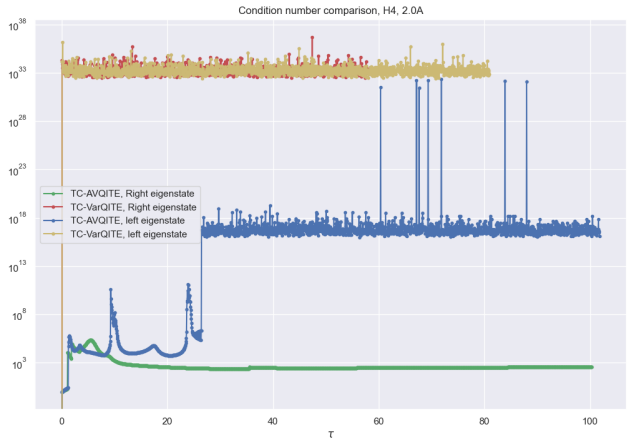

Figure A.5: The condition numbers of the runs in figure 4.16.

DEPARTMENT OF CHEMISTRY AND CHEMICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY