The Behaviour and Efficacy of Using Adiabatic Quantum Computing with Parameterised Quantum Circuits to Estimate the Ground State Energy of Molecules

Christopher Crampton

4th Year Project Report
Artificial Intelligence and Computer Science
School of Informatics
University of Edinburgh

2024

Abstract

Adiabatic Quantum Computing with Parameterised Quantum Circuits (AQC-PQC) is a new variational method that if found successful, boasts advantages over alternative variational methods due to the polynomial number of steps required to avoid a barren plateau at each one, compared with the exponential probability of starting in a barren plateau for larger problems. AQC-PQC can be used to estimate the ground states of molecules. This was established by testing the H_2 molecule using an adaption of the Qiskit-based Python implementation of the AQC-PQC code used in [11]. In one case, AQC-PQC performed as well as a Variational Quantum Eigensolver (VQE), using the overlap with the true ground state as the measure of success.

This paper explores the effectiveness of the AQC-PQC algorithm under varying starting conditions for the four combinations using the Jordan-Wigner and Bravyi-Kitaev mappings on the Unitary Couple-Cluster (UCCSD) and Efficient SU2 ansatze. In each combination of mapper and ansatz, the implementation of the AQC-PQC algorithm worked as intended, with a range of probabilities of success. The results of the experiments show that the choice of starting parameters is important since not all parameter vectors θ that correspond to the ground state will give an accurate final result.

The results were compared with the success of VQE on the same H_2 molecule. VQE was found to be more consistent across all experimental setups, and superior in all cases. However, this was with the acknowledgement that the experimental conditions were far from optimal due to the limited time and processing resources available. The highest probability of success with AQC-PQC was found when using the Jordan-Wigner mapping and the UCCSD ansatz combination. In this case, the performance difference between AQC-PQC and VQE was negligible: 96.8% probability of reaching the ground state with VQE versus 96.6% for AQC-PQC. However, this combination is less practical in implementation than the three alternatives. VQEs are currently the go-to method for NISQ devices. VQE was first proposed in 2014 [19] and has been subject to many refinements since then. The investigations of the more recent AQC-PQC contained in this Dissertation, represent only a first attempt at its evaluation, yet they are encouraging enough to propose that continuing the development, evaluation and dissemination of AQC-PQC implementations is a worthwhile goal.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics Committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Christopher Crampton*)

Acknowledgements

Firstly I wish to thank my supervisor Professor Petros Wallden for being so supportive and accommodating while I was getting to grips with the relevant material, and for providing invaluable insights and assistance so consistently throughout the year.

I am grateful to my family for being so encouraging throughout my academic career and for proofreading this work.

Special thanks to Dr Cate Anstoter for talking me through the chemical concepts relevant to the project and clarifying my understanding of the material, and to Ioannis Kolotouros for constructing code that was easy to work with, taking the time to talk through it and providing suggestions.

I would also like to thank the OneDrive Recycle Bin for keeping extra backups when an accidental program run overwrote every single data file I had before I committed them to Github.

Table of Contents

Chapter 1

Introduction

1.1 Motivation

Richard Feynman once said

> "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy."

Trabesinger, A. [25]

This statement referred to the inherent complexity of quantum simulation with classical methods. Namely, the overhead costs of discretely storing and manipulating complex-valued quantum states and the exponential growth when modelling quantum interactions [5]. Quantum computers address these issues by allowing operations to be carried out on a superposition of states, and by being inherently complex valued. This both encourages and necessitates using quantum computers in quantum problems, such as finding the ground state of large molecules. This is a useful task to accomplish because from the ground state of molecules, many properties can be deduced that have essential applications in fields such as drug development in medicine.

The method involves finding a global minimum amongst particle positions in space and atomic/molecular orbital occupation. In practice, this problem requires finding the local minimum electronic ground state for each of a set of nuclear positions, to determine which is the global minimum. The nuclear positions are fixed as part of the molecular Hamiltonian, and the ground state is found to be the eigenvector with the smallest corresponding eigenvalue (representing the energy). This project looks at quantum methods to find the minimum eigenvalue for a problem Hamiltonian corresponding to a set of predefined nuclear positions.

In the future, if quantum computers can reach a noiseless scale that allows for large molecular simulations to be carried out, methods such as Adiabatic Quantum Computing (AQC) [2], or Quantum Phase Estimation [8] could be used used to determine the ground states of problem Hamiltonians. AQC involves starting in an easy-to-prepare ground state with a simpler Hamiltonian, and slowly perturbing the initial Hamiltonian to the problem Hamiltonian such that the qubits remain in the ground state of each intermediate perturbed Hamiltonian until the endpoint, when the solution is reached. This requires noise mitigation to be much more effective due to the massive circuit depth required for the method and as such is infeasible for near-term devices. As a result, current methods are variational by nature, for example in [19] and [24], and use a combination of classical and quantum devices to reach the goal. This is accomplished by numerous shallower shots of quantum devices dictated by a classical optimiser.

Current methods such as VQE are subject to exponentially low probabilities of avoiding a barren plateau when initialising, but AQC-PQC is shown to only require a polynomial number of steps to stay in the ground state at each step. This suggests significantly improved scalability for future implementations, thus motivating an investigation into AQC-PQC's efficacy for ground-state energy estimation.

1.2 Problem Statement

This project aimed at determining if the hybrid algorithm AQC-PQC can be used successfully when calculating the ground state energies of molecules, in comparison with the widely adopted VQE [19], and exploring the behaviour of the AQC-PQC algorithm in finer detail.

1.3 Dissertation Structure

The relevant prerequisites are detailed in the background assessment in Chapter 2. This includes the algorithms used, the chemistry theory and some specific details about the problem being tested to provide a frame of reference to assess the findings.

Experiments were carried out using VQE as a baseline, the method and results for which are detailed in Chapter 3.

The method and results for the comparable AQC-PQC tests are found in Chapter 4. A scoping investigation (Section 4.3) is included to outline an idea that I had too close to the submission deadline to explore fully, relating to the step number and run-time, which suggested some promising preliminary results.

A Discussion is included in Chapter 5. This chapter also contains an evaluation of the methods used, their limitations and suggested areas for future work to expand on what has been done in this project. These expansions either aim to clarify and deduce the causes of the results or indicate preferred experimental setups that should give more accurate results. Section 5.6 evaluates in more detail some of the behaviour exhibited by AQC-PQC in the experiments and explores methods for determining the cause. I also propose some methods for future testing that could improve the accuracy and time complexity of the AQC-PQC algorithm itself.

The Conclusion in Chapter 6 provides a summary analysis of the work carried out and its significance to the field.

1.4 Contributions

To generate the baseline for VQE:

- I created a Chemical Problem class based on the Qiskit Nature library to facilitate the creation of electronic structure problems and used it in conjunction with the Qiskit VQE library.
- I performed statistical analysis and generated figures to display the results in a manner that is easy to interpret.

To generate and compare the results for the AQC-PQC methods:

- I used the Chemical Problem class that I created for the VQE tests and added attributes relating to the problem for analysis later.
- I encapsulated the code written by I. Kolotouros in [11] to allow for numerous runs that would save the results and important attributes of the runs to JSON files for analysis later.
- I implemented a change to the UCCSD ansatz produced with Qiskit Nature to allow for proper ground state preparation for the initial Hamiltonian. This involved the addition of RY gates to each qubit at the end as I had noticed the VQE initialisation seemed unable to reach the ground state of the initial Hamiltonian. This change is scalable for the UCCSD ansatz on any number of qubits.
- I established preference in both ansatz family and mapping with respect to accuracy and practical implementation in AQC-PQC.

To explore and analyse the failure cases of AQC-PQC in more detail:

- This Dissertation was the first use of AQC-PQC for chemistry and the first use with the problem-specific ansatz UCCSD.
- I added more attributes to my Chemical Problem class that could be read by other programs that used it.
- I extended my adaptation of I. Kolotouros' code and my Chemical Problem class to store more attributes of a given iteration to allow for more in-depth analysis.
- I fixed a problem that meant smaller parameter shifts rounded to zero on the UCCSD ansatz.
- I used the NumPy library to generate all of the eigenvalues and eigenstates for the problem Hamiltonian under each of the mappings. Thus, I verified the mappings' required properties to ensure fair tests and explore causes from observed phenomena.

To guide and ease future work on the topic:

- I developed two ideas to add to the algorithm to improve accuracy and provide a confidence proxy metric associated with the probability that the state reached was the ground state.
- I hypothesised a relationship between step number and runtime that encourages a higher step count within a certain range and ran preliminary experiments to probe this theory. The premise of the hypothesis stated that the optimiser speedup from being closer to the new ground state outweighed the added time taken from extra steps within a certain range. The results were promising.
- I made the classes that I created and modified simple to use, allowing for a vast array of easy customisability for future studies that wish to change any of the hyperparameters such as step count, iteration number, ansatz and mappings used and the molecule string to be tested.
- I created a series of scripts that can read in the saved data and visualise it appropriately for analysis.

1.5 Acronyms Used

- AQC - Adiabatic Quantum Computing: The original, continuous and fully quantum version of the algorithm being investigated.
- AQC-PQC - Adiabatic Quantum Computing with Parameterised Quantum Circuits: The quantum algorithm being explored in this paper.
- BO - Born-Oppenheimer: In reference to the Born-Oppenheimer Approximation used to simplify models of quantum interactions.
- ESU2 - Efficient SU2: The Qiskit class used to generate the Hardware-Efficient ansatz being tested.
- GS - Ground State: The lowest energy state of a system.
- NISQ - Noisy Intermediate-Scale Quantum: In reference to quantum devices, they contain up to 1000 qubits and do not have effective enough noise mitigation to achieve quantum advantage [28].
- SD - Standard Deviation: This was shortened for use on Graphs that struggled to fit on the page.
- SLSQP - Sequential Least Squares Programming: An optimiser used to find the smallest parameter shifts that satisfies the constraint in AQC-PQC and the parameters that correspond to the smallest expectation value in VQE.
- UCCSD - Unitary Coupled Cluster with Single and Double Excitations: The quantum computing implementation of an effective technique that simulates systems consisting of multiple interacting bodies.
- VQE - Variational Quantum Eigensolver: The current industry standard for NISQ devices that employ the variational principle.

Chapter 2

Background

2.1 Adiabatic Quantum Computing with Parameterised Quantum Circuits

2.1.1 Overview

Adiabatic Quantum Computing with Parameterised Quantum Circuits [11] is a new method for simulating adiabatic evolution through a sequence of discrete steps. The premise behind this approach relies on the ability to approximate to a high probability the change in the ground state for an iteratively perturbed problem Hamiltonian. Successfully estimating how the ground state will change negates the need to carry out a single, slow time evolution required in normal adiabatic quantum computing. The new ground state is found by finding the smallest parameter shift vector whose gradient is vanishing which indicates a critical point, and whose Hessian is positive-semi definite which indicates the critical point is a minimum. If the number of steps is small enough, the classical optimiser at each new perturbed Hamiltonian will correctly step towards the new ground state and not one of the other energy states. This algorithm's pseudo code is shown in Figure 2.1. Due to the high probability that the nearest local minimum at each step is the desired global minimum, I decided to use the SLSQP optimisation algorithm as it is gradient-based. This makes it better suited to converging towards the new minimum when compared with derivative-free methods.

This approach has been tested using Hardware-Efficient ansatz families on problems such as Max-Cut and Number Partitioning. This project is the first use case on chemistry problems that make use of problem-specific ansatz families and shows that the use cases are not restricted to Hardware-Efficient problems or Ising Hamiltonians as explored in [11].

The greatest advantage the AQC-PQC algorithm offers over VQE is the avoidance of barren plateaus. From initialisation in the ground state, there is a guarantee that with a polynomial number of steps (with respect to the number of qubits), each new step will be in a parameter space with gradient. This is in contrast with VQE which has an exponentially small probability of starting somewhere with a gradient that can

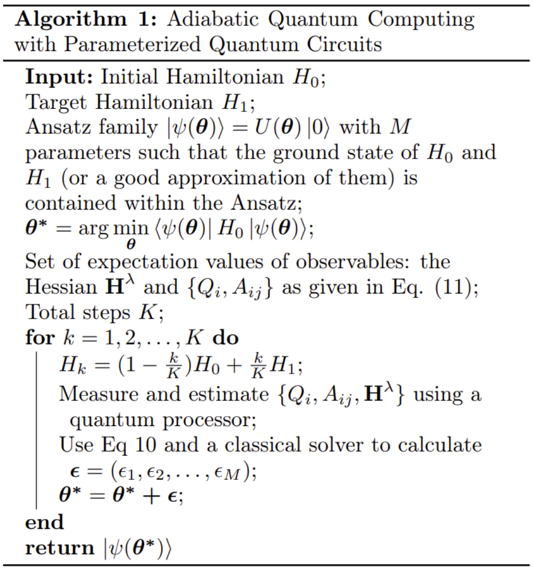

```
Algorithm 1: Adiabatic Quantum Computing with Parameterized Quantum Circuits
Input: Initial Hamiltonian $H_0$;
Target Hamiltonian $H_1$;
Ansatz family $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle$ with $M$ parameters such that the ground state of $H_0$ and $H_1$ (or a good approximation of them) is contained within the Ansatz;
$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \langle\psi(\boldsymbol{\theta})| H_0 |\psi(\boldsymbol{\theta})\rangle$;
Set of expectation values of observables: the Hessian $\mathbf{H}^\lambda$ and $\{Q_i, A_{ij}\}$ as given in Eq. (11);
Total steps $K$;
for $k = 1, 2, \dots, K$ do
    $H_k = (1 - \frac{k}{K})H_0 + \frac{k}{K}H_1$;
    Measure and estimate $\{Q_i, A_{ij}, \mathbf{H}^\lambda\}$ using a quantum processor;
    Use Eq 10 and a classical solver to calculate $\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \dots, \epsilon_M)$;
    $\boldsymbol{\theta}^* = \boldsymbol{\theta}^* + \boldsymbol{\epsilon}$;
end
return $|\psi(\boldsymbol{\theta}^*)\rangle$
```

Figure 2.1: The pseudo code algorithm of the original AQC-PQC implementation. Source: Kolotouros et al [11]

converge towards a minimum. This speedup is not seen in the small example contained in this work but is important to keep in mind for future works on larger molecules, as a significant reason to encourage adoption of the AQC-PQC method.

2.1.2 Alterations to AQC-PQC

Since the paper's initial release, I. Kolotouros, the author of the Python implementation, has updated the code to include three heuristics to improve the runtime, which were integrated into my refactored code. The only change relevant to this paper is the null-space heuristic. This involved singular value decomposition to reduce dimensionality. This allowed the optimisation algorithm to operate on a smaller set of parameters which in turn meant each iteration was performed faster. In some cases however, an empty set was returned. This could be due to matrix elements being so close to zero that the NumPy implementation equated them to be zero, or the matrix could have been rank deficient. It is also sometimes the case that after the process there are simply no significant vectors remaining. If there are no singular value indices returned, all parameters are used as a default which has the effect of ignoring the heuristic.

When running the experiments with the original algorithm, the time taken for each run with 20 steps made it infeasible to run repeated and updated experiments within the time frame for the project. The null-space heuristic was applied to improve runtime.

2.1.3 Behaviours

In this section, problematic behaviours that can be exhibited in the running of AQC-PQC are highlighted and explained.

2.1.3.1 Spectral Gap

If the spectral gap is small enough at some step k, the algorithm may erroneously shift some of the probability mass towards the higher energy state. For this to happen, the next energy state must satisfy the constraints declared by the optimiser: namely the gradient at the point must be vanishing and the Hessian at the point must be positive semi-definite (which implies a minimum and a turning point). By definition, the ground state is always guaranteed to be a turning point, but the other energy states are not. There is also a chance that small approximation errors could lead to a saddle point being mistaken as a turning point, making a higher energy state seem more likely to be the ground state. If the first energy state satisfies the criteria and is closer to the previous point in the parameter space, then it will be chosen as the next point to move to. A suggested change in the algorithm to mitigate this issue can be found in Section 5.5.

2.1.3.2 Expressibility

It is assumed when running the algorithm, that the ansatz chosen can approximate the ground state of the initial and final Hamiltonian to a suitable degree. In the case where the ground state of one of the intermediate Hamiltonians H_k cannot be reached, the entire probability mass can jump up to the next excited state. This is because when finding the smallest parameter changes that satisfy the constraints, if the optimiser is unable to find the ground state, the next suitable shift vector is likely to correspond to the next energy state up. This problem can work both ways however, meaning if it is in the first excited state of Hamiltonian $H_{k\lambda}$, and the corresponding first excited state of $H_{(k+1)\lambda}$ cannot be reached, then the optimiser may find the ground state of $H_{(k+1)\lambda}$ and shift all of the probability mass down to that. There are examples of both these phenomena in Chapter 4.

2.1.3.3 Optimiser Inaccuracy

The optimiser allows a tolerance of 1×10^{-6} when checking the constraints. This can cause tiny inaccuracies when moving to the new ground state. The inaccuracies of parameter shifts may cause some of the new ground state to be partially in one of the other states as well. This is the most likely cause of the slight inaccuracy in the Jordan-Wigner UCCSD results for AQC-PQC in Section 4.2.2.2.

2.2 Chemistry

This section summarises the background chemistry for the problem-specific domain that this dissertation investigates.

2.2.1 Energy States of Molecules

Atoms are comprised of nuclei and orbiting electrons. These electrons can be found anywhere in space with a corresponding likelihood. The locations where electrons are most likely to be found are known as orbitals. In the most simplistic view, an occupied orbital can be modelled as a discrete volume that will contain an electron.

In reality, an orbital is a wave function denoting the probability of finding an electron somewhere at a given time. The previously discretised ’orbitals’ represent the region within which there is a 95% probability of finding the electron at a given time if occupied. These orbitals are solutions to the Schrodinger equation using an atom or molecule’s Hamiltonian: $\hat{H}\psi(\mathbf{r}) = E\psi(\mathbf{r})$. The eigenvectors of the Hamiltonian are the energy states of the system, and the corresponding eigenvalues are their respective energies. The eigenvector with the smallest eigenvalue is known as the ground state, and due to being the lowest in energy, is the most stable state.
Finding the ground state of a molecule is important in fields such as medicine since from the ground state, one can infer other important properties such as drug binding mechanisms, reaction rates and molecular stability. From the ground state one can derive dipole moments, polarisabilities and whether reactions will occur. If the total energy of a molecule is calculated to be less than the sum of the individual parts, then the molecule formed is considered stable.

2.2.2 Hamiltonians

A Hamiltonian is an operator that determines the energy of a system. In this context, the word Hamiltonian could refer to three distinct but connected operators: Molecular, Fermionic or Qubit [7].

$$\hat{H}_{\text{mol}} = \sum_{i=1}^{N} \hat{h}_i(\mathbf{r}_i) + \sum_{i=1}^{N} \sum_{j=1,j\neq i}^{N} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_{I=1}^{M} \sum_{J=1,J\neq I}^{M} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} \tag{2.1}$$

Equation 2.1: The general equation for a time-independent molecular Hamiltonian whose parameters are the particles’ positions in space

$$\hat{H}_{\text{ferm}} = \sum_{pq} h_{pq} \hat{a}_p^\dagger \hat{a}_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_s \hat{a}_r \tag{2.2}$$

Equation 2.2: A simplified Hamiltonian that only considers electrons and treats nuclei as static point charges. The result of the Born-Oppenheimer Approximation applied to the molecular Hamiltonian

Molecular Hamiltonians refer to operators defined by the location and charge of the particles in a molecule, namely the electrons and nuclei. It is the summation of the electron kinetic energies and the electrostatic forces between the positive nuclei and the negative electrons (otherwise known as Coulomb interactions). To simplify the problem and make it more accessible, several approximations and simplifications are made in practice. An example of this is the Born-Oppenheimer approximation, which omits the nuclear kinetic energy terms by viewing the nuclei as static point charges, leaving electron-only terms. Due to the static assumption of the nuclei, the variables representing the position of the nuclei can be dropped and instead treated as constants. Solving a Hamiltonian involves finding the eigenvalues of a partial differential equation of dimensions proportional to the number of particles. The lower bound complexity for a classical eigensolver algorithm is n^2 in the number of variables [6]. The decrease in the quadratic terms (number of particles) through the BO approximation is replaced with numerous iterations of different nuclei positions, but this is a beneficial trade off in practice.

In this dissertation the H_2 molecule is used with two atoms containing two nuclei and two electrons, each with three spatial dimensions. This totals 12 variables for the wave function which is well within the capabilities of modern computing power. However, when scaling up to larger molecules, such as riboflavin ($C_{17}H_{20}N_4O_6$), it quickly becomes infeasible to calculate. For perspective, a single carbon atom contains six electrons and there are 17 carbon atoms in riboflavin.

2.2.3 Basis Sets and Fermionic Operators

A basis set of functions defines the way in which the partial differential equations of a Hamiltonian are handled by a computer. These basis functions represent approximations to the atomic orbitals discussed above. They vary in complexity and detail, which in turn affects the quality of the solution. For a given basis set, the wave function corresponds to the coefficients of linear combination of these orbital approximations. Basis sets do not have to refer to atomic orbitals but can be defined as plane waves or real-space approaches. Minimal basis sets are defined with a one-to-one representation of basis function to individual atomic orbital. As they are such rough approximations to wave function solutions, they are not used when performing chemical research. This dissertation does not aim to make any chemical advances and as such, I use them for their efficiency improvements. For these experiments I will only use the Gaussian-type atomic orbital minimal basis set: STO-3G.

In selecting this way to represent the wave functions of the electrons in the system, ways of spanning the subspace of solutions such that no physical laws are violated are needed. For this one can use creation and annihilation operators, respectively referring to the addition and removal of fermionic quanta. The definition of these operators ensures that in the process of finding the ground state, no new particles are added or taken from the system. Furthermore, fermionic creation and annihilation operators obey anticommutivity due to the Pauli Exclusion Principle [18]. The principle states that no two fermions can occupy the same quantum state which causes the order of operators to be important, unlike with bosonic particles. The fermionic Hamiltonian (Equation 2.2) is comprised of creation ($\hat{a}^\dagger$) and annihilation operators ($\hat{a}$).

2.2.4 Implementation of the Chemistry Side

The Qiskit Nature library contains methods to carry out all of these calculations behind the scenes, ready for use with the quantum algorithms being tested. The inputs required are the atomic locations, charges and spins, the basis used and the desired distance unit. It was important to understand what each of these were in respect to the problem field and efficiency of implementation. As previously stated in Section 1.2, the emphasis of this dissertation is an investigation into the viability of AQC-PQC in obtaining the ground state of a Hamiltonian, not producing new experimental chemical data. This meant that I chose parameters solely for improved time complexity and kept them constant throughout the experiments. I then used a NumPy eigensolver to obtain the goal ground-state against which to compare the methods used here. The library yielded an electronic structure problem which was then converted using the desired mapper to a problem Hamiltonian. The variational principle:

$$E_{\text{ground}} \leq \langle \phi | H | \phi \rangle \tag{2.3}$$

ensures that the expected value calculated is an upper bound on the ground state of the Hamiltonian, but the approximations made in the process of conversion to the qubit Hamiltonian does not have any direct guarantee that the true ground state value of the original molecular Hamiltonian is greater than, less than or equal to the final result that the method produces. I chose the bond length between the hydrogen atoms to be 0.735Å based on the work in [1]. Angstroms (Å) is a standard unit of atomic length.

2.3 Ansatz Families

An ansatz family in the context of quantum computing is a parameterised circuit that spans a portion of the Hilbert Space. The ansatz is used in this context to generate a state from a parameter vector: θ, which will then be operated upon by the Hamiltonian to gain an energy estimate. It is important that the ground state solution can be well approximated within the span of the ansatz otherwise a good solution cannot be found. For AQC-PQC, it is important that the ground states of all intermediate Hamiltonians are included as well. If not, this can cause some peculiar behaviour, as detailed in the AQC-PQC results in Chapter 4. This project investigates the effectiveness of the algorithms when using a Hardware-Efficient ansatz and a problem-specific ansatz.

2.3.1 ESU2

The hardware-specific ansatz is from the Qiskit Efficient SU2 ansatz method that generates the circuit shown in Figure 2.2 for a given number of qubits. There are some parameters that can be changed when generating an Efficient SU2 circuit, but this test case represents the most generic use case option, with no orientation towards the problem and so I kept the majority of the initialisation parameters as their defaults. To clarify, initialisation parameters in this case refer to the arguments for the method, not the θ vector.

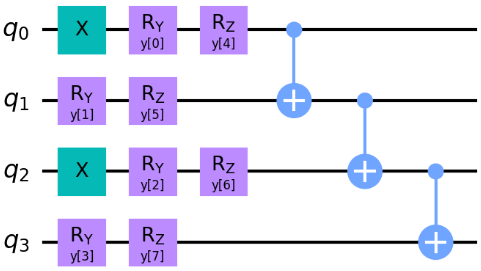

Figure 2.2: The Efficient SU2 Ansatz Circuit for a Hydrogen Molecule

2.3.2 UCCSD

The Unitary Couple Cluster Ansatz with single and double excitations (UCCSD) is widely used in quantum chemistry applications. It makes use of fermionic creation and annihilation operators, which then undergo a transformation such as the Jordan-Wigner or Bravyi-Kitaev transforms to become spin operators. The state generated by the ansatz for parameters θ is given by:

$$e^{T(\theta)-T(\theta)^\dagger}|\psi_0\rangle \tag{2.4}$$

Equation 2.4 The effect of the parameterised UCCSD ansatz on a quantum state

where T is the summation of the cluster operators:

$$T_1 = \sum_i \sum_a \theta_a^i \hat{a}_i^\dagger \hat{a}_a \tag{2.5}$$

$$T_2 = \sum_{i,j} \sum_{a,b} \theta_{ab}^{ij} \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_a \hat{a}_b \tag{2.6}$$

Equations 2.5 and 2.6 The T cluster operators used in the UCCSD ansatz

This exponentiation is then trotterised [27] to produce a quantum circuit which is handled by Qiskit. This is the first implementation of a problem-specific ansatz' use with AQC-PQC. As this use case is novel, there is currently no deterministic way to fit this ansatz in its current state to the initial Hamiltonian used in AQC-PQC. In Figure 2.3 there are three blocks which refer to two single excitations and a final double excitation. The UCCSD ansatz is only the first three blocks, the reason behind the addition of the RY gates at the end is explained in the Method Section (4.1) of the AQC-PQC Results.

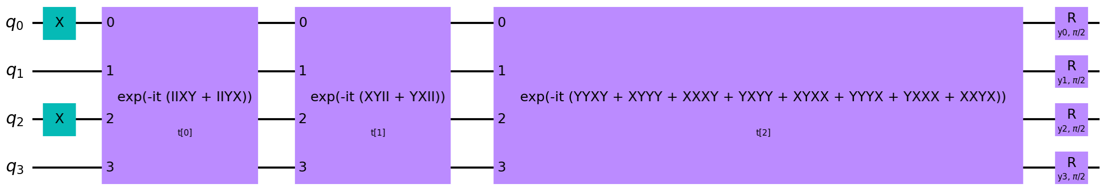

Figure 2.3: The Unitary Coupled-Cluster with Singles and Doubles Ansatz for a Hydrogen Molecule with the AQC-PQC RY modification.

2.4 Mappers

This project will use and compare two transformations: the Jordan-Wigner and Bravyi-Kitaev. Their role is to map the fermionic creation and annihilation operators of the Hamiltonian (and the UCCSD ansatz) onto Pauli gates ensuring preservation of anticommutativity properties. The main difference between Jordan-Wigner and Bravyi-Kitaev is the locality of the final solution. Jordan-Wigner requires a qubit locality of

O(n) which is difficult in practice to implement. Bravyi-Kitaev [3] on the other hand is able to perform a fermionic operation as quantum operators with $O(\log n)$ locality [23], and through use of gadgets can even be 2-local with ancilla bits [13]. Seeley at al [22] shows that Bravyi-Kitaev has superior practical efficiency for: “All quantum computations of electronic structure”. As this project uses classical simulation of quantum operations, I investigated how the final results using AQC-PQC differ with the different mappings and if there would be a trade off in practical applications between quality of results and efficiency of implementation. Figures 2.4 and 2.5 show the difference in UCCSD ansatz circuits for each of the mappers.

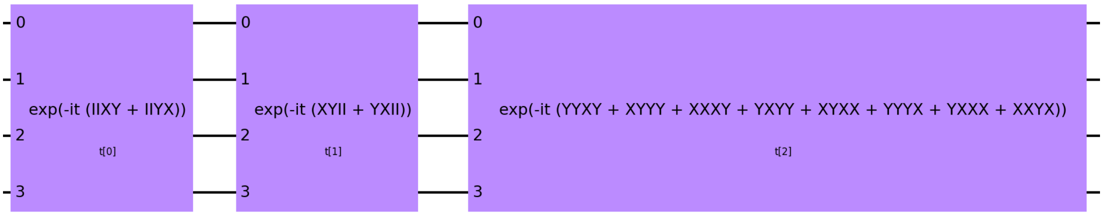

Figure 2.4: The UCCSD Ansatz for the H_2 molecule using the Jordan-Wigner mapping

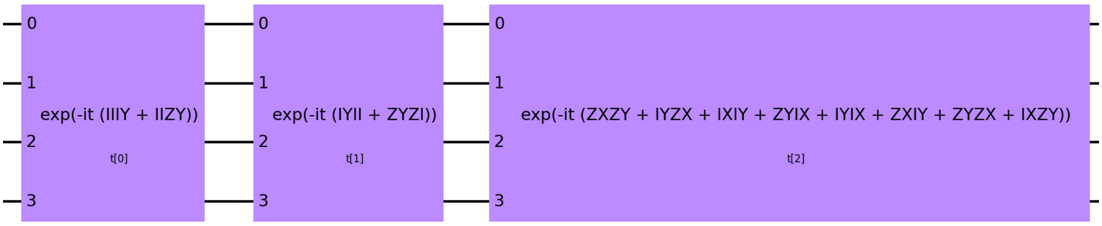

Figure 2.5: The UCCSD Ansatz for the H_2 molecule using the Bravyi-Kitaev mapping

I performed my own verification as detailed below that certain necessary properties were maintained and would not affect the quality of results. The most important property of the two mappings was that the eigenvalues of the problem Hamiltonians should be the same, although this makes no guarantee that the intermediate Hamiltonians will have corresponding eigenvalues. Through testing this was shown to be the case. Due to the small size of the Hamiltonian matrix for a Hydrogen molecule, I could use the NumPy linear algebra eigensolver. This showed identical eigenvalues but some of the corresponding eigenvectors differed. The two matrices both had a magnitude of 3.93 (2dp), but subtracting them from each other yielded a matrix of magnitude of 2.88 (2dp). The difference between these Hamiltonian representations caused differences in results as the path from the same initial Hamiltonian was different for each. A more detailed comparison of the two mappers can be found in [26].

2.5 Variational Quantum Eigensolvers

Variational Quantum Eigensolvers have been tested on numerous occasions in the context of quantum chemistry with promising results.McCaskey [12] had success in finding the ground state of alkali metal hydrides by leveraging the IBM Tokyo and Rigetti Aspen processors, and noise mitigation techniques to obtain chemical accuracy. Nam [16] made use of trapped-ion quantum hardware to estimate the ground state energy of H_2O to chemical accuracy. VQEs make use of classical optimisers on parameterised circuits and the variational principle to reach the ground state of a problem Hamiltonian. Their hybrid nature allows for a large number of calculations to be carried out but with minimal circuit depth for a given run. This has a strong mitigating effect on the damage of noise from methods that use deep circuits such as Adiabatic Quantum Computing. The method starts at a random point contained within the ansatz family, unless otherwise stated, and measures the expected value with the problem Hamiltonian at that point. The classical optimiser will use these results and perturbations of the parameters to converge on a minimum which it will then yield as its estimation of the ground state.

The biggest problem facing the use of VQE for larger problems is barren plateaus that cause convergence in the wrong place [14]. The probability of starting in a state with a vanishing gradient is exponential in the number of qubits, which implies an exponential number of expected runs before a successful one. The advantage of AQC-PQC is that a polynomial number of steps is required to ensure that the parameters start in an area with a gradient at each new step.

One of the most influential aspects of the quality of results using VQE is the classical optimiser chosen. Commonly used optimisation algorithms fall into three broad categories: Gradient-based, Derivative-Free or Reinforcement Learning [4]. Gradient-based algorithms are much better suited to converging towards minima, assuming the starting parameters are not in a barren plateau. The ground state is always guaranteed to be a turning point by definition of being a global minimum, whereas the higher energy levels might not. To maximise the comparability of testing between AQC-PQC and VQE, I elected to use the same SLSQP for optimisation with VQE as well.

VQEs are currently the go-to method for NISQ devices and as such constitute a benchmark for AQC-PQC to be compared against. This is with the consideration that VQE was first proposed in 2014 [19] compared with AQC-PQC which was first proposed in 2022 [11]. This means that later VQE implementations have had the opportunity to be improved in both efficiency and quality, with numerous experiments carried out to aid in determining the optimal conditions under which to carry out different forms of experiments. The initial VQE paper [19] had been cited 1,693 times (as of March 2024) which indicates a high level of interest over time. A proportion of those interested might adopt AQC-PQC as an alternative method to VQE for the advantages stated in Section 2.1. The initial promising results from this very early study may similarly inspire similar additions and optimisations from the academic community.

2.6 Molecule Preliminaries

The molecule being examined in this project is H_2 . It is a simple neutral molecule consisting of two nuclei and two electrons which means only four qubits are used. This makes it feasible to perform simulations on a classical computer. In this project, where the word hydrogen is used, assume the diatomic molecule, not the atom, unless explicitly stated otherwise. It is a molecule often used for testing in similar studies involving quantum computing applications for chemistry problems [1]. Its Hamiltonian is a 16 by 16 matrix (four qubits) with 10 distinct energy states, some of which are degenerate as shown in Figure 2.6 (the subscripts refer to the degree of degeneracy). The expectation value when the problem Hamiltonian is applied to the ground state is -1.857 (3dp) and has a unique eigenvector representation. The second and third energy levels are degenerate with two and three corresponding eigenvectors respectively. The full eigenspectrum to two decimal places is listed below. The eigenvectors representing the ground state when using the Bravyi-Kitaev and Jordan-Wigner mappers are identical. This has been verified using NumPy and the Chemical Problem generation module that I created for use in this project. The units used for these values are Hartrees [9]. One Hartree's metric value is equal to 4.36 $\times 10^{-18}$ Joules (2dp). This corresponds to the approximate electric potential energy of the electron in a hydrogen atom. Note that this is the hydrogen atom not the hydrogen molecule, and that the ground state for the hydrogen molecule is lower than the atomic ground state which is why hydrogen is normally found in nature in its diatomic molecular state.

$$E_S = [-1.85, -1.26_2, -1.25_3, -1.16_2, -0.88, -0.47_2, -0.35_2, -0.22, \approx 0,\ 0.21]$$

Figure 2.6: The results of the NumPy eigensolver on the problem Hamiltonian

In the case of hydrogen, due to prior knowledge of the molecule it is known that the ground state is not degenerate, since the two electrons will both occupy the single 1s orbital. This simplifies later calculations in determining the overlap with the true ground state of the Hamiltonian. However, it should be noted that degeneracy of the ground state is possible in other molecules or atoms. As an example, atomic carbon has two valence electrons and three 2p orbitals in which to place them with equal energy, and there are three distinct arrangements of this state, each with unique wave functions but the same energy. This should be accounted for in future work when determining the ground state of larger molecules as the final ground state could be a linear combination of multiple possible ground states.

It should also be noted that the ground state of a given Hamiltonian can be represented by a variety of parameter vectors, as shown in Figure 2.7 below. The different parameters will have different local energy landscapes which will result in differing optimisation successes. For AQC-PQC this has the implication that different initial parameters may cause the algorithm to follow a different path and end up with a potentially completely different result. For VQE the different landscapes may mean an iteration will get caught at a local energy minimum in the parameter space, one that would not be there for the same quantum state defined by a different set of parameters.

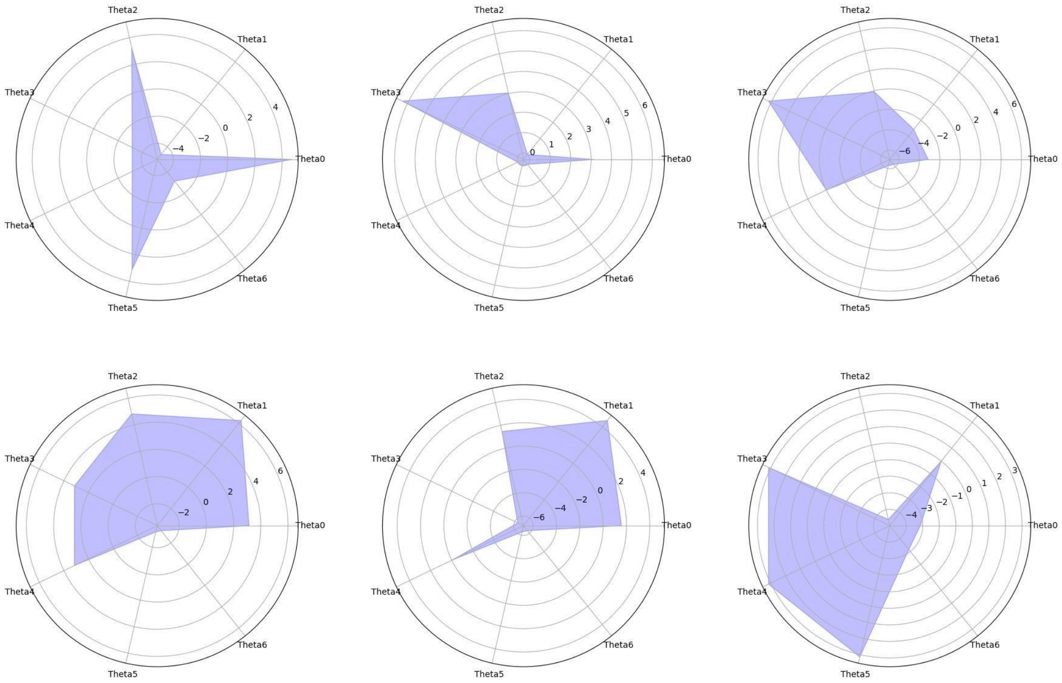

Figure 2.7: Radar chart showing different parameter values for the UCCSD ansatz that can be used to represent the ground state of the problem Hamiltonian. Each point on the star represents a parameter and is valued between -2π and 2π.

Chapter 3

VQE Analysis

In this Chapter I analyse the spread of results that I produced across the UCCSD and the ESU2 ansatz families for both the Jordan-Wigner and Bravyi-Kitaev mappings using VQE.

3.1 Method

To simulate VQE in Python I made use of the Qiskit library [20], and the same class for setting up my Chemical Problem class as used with AQC-PQC. The optimiser used was the Qiskit SLSQP algorithm. It should also be noted that despite not having the same initialisation problem as AQC-PQC for the UCCSD ansatz (see Section 4.1), the ansatz used and defined as UCCSD is the same format as shown in Figure 2.3, with the additional RY gates at the end. This allows for a fairer comparison. A thousand iterations for each experiment were chosen because it is a large enough sample size to gauge probabilities of success with reasonable confidence, but was small enough to be feasible to generate

The frequency graphs shown in Sections 3.2.1 and 3.2.2 are the summation of results truncated to five decimal places. The reason for this is that almost all the final results were unique to the accuracy provided by the implementation. When directly visualised the graphs were difficult to draw conclusions from. Truncating to five decimal places reduced the number of bins from over 800 to an average of about 10 each, although the majority of these only represented up to five counts and as such are not visible. The statistical analysis across the data points is shown in Figure 3.1. As we know the eigenvectors that correspond to each energy level from the Molecule Preliminaries in Section 2.6, the overlap of a given state to the ground state can be directly calculated. This method involves obtaining the inner product between the ansatz-generated state ($|\psi(\theta^*)\rangle$) and the ground energy state vector ($|E_0\rangle$):

$$p_0 = |\langle E_0|\psi(\theta^*)\rangle|^2 \tag{3.1}$$

This calculation was performed for every result generated, and the mean of these across the 1,000 iterations is displayed as the Ground State Probability column (GS Prob.) in Figure 3.1.

3.2 Results

Ansatz	Mapper	Mean	Median	SD	Minimum	Maximum	GS Prob.
ESU2	Bravyi-Kitaev	-1.522	-1.256	0.319	-1.837	-1.064	0.481
ESU2	Jordan-Wigner	-1.564	-1.857	0.313	-1.857	-1.064	0.531
UCCSD	Bravyi-Kitaev	-1.731	-1.837	0.229	-1.837	-1.203	0.795
UCCSD	Jordan-Wigner	-1.838	-1.837	0.069	-1.857	-1.245	0.969

Figure 3.1: Statistical Analysis of the energy results for the VQE runs with the Jordan-Wigner and Bravyi-Kitaev Mappers on the ESU2 and UCCSD Ansatz. Units: Hartrees

Across a thousand iterations, only the cases with the Jordan-Wigner Mapper were able to fully reach the ground state, although the Bravyi-Kitaev experiments came very close. This lends authenticity to the idea that the efficiency improvement from using the Bravyi-Kitaev mapper does come at the cost of weaker performance, albeit slight in this instance. This is further supported by a lower ground state probability in both ansatz families when using Bravyi-Kitaev instead of Jordan-Wigner. All approaches were able to get to a reasonable approximation of the ground state at least once ($p_0 > 0.99$).

Using the calculation shown in Equation 3.1, the best results from the imperfect Bravyi-Kitaev experiments still have a probability of 0.996 (3dp) for being measured as the ground state. Bravyi-Kitaev also yields consistently higher standard deviation, albeit not much different in the Hardware-Efficient case, but when using the problem-specific ansatz it has a much larger spread of results.

3.2.1 ESU2 Ansatz

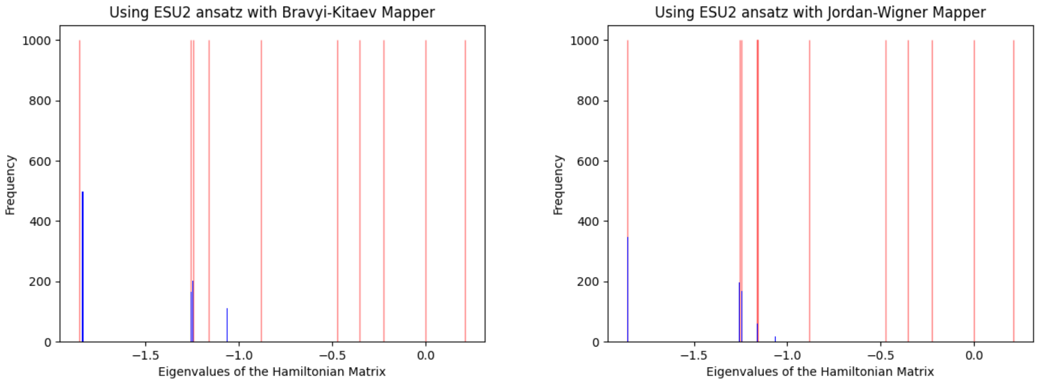

Figure 3.2: Using the ESU2 ansatz for 1000 VQE Runs

Using the Hardware-Efficient ansatz, VQE was able to achieve an acceptably close result to the ground state approximately one in every two runs as shown in Figure 3.2. Otherwise, it was getting caught in one of the subsequent energy states. The Bravyi-Kitaev runs give a larger spread and variance with a much higher probability of ending well past the first few energy states. Despite the Jordan-Wigner implementations managing to end completely in the ground state, it only achieved this around 35% of the time compared with the Bravyi-Kitaevs implementation which reached an ‘almost perfect’ result 50% of the time. The consistency with which the results line up with the energy states would suggest local minima at those locations. It should be reiterated that those energy states are degenerate and so may constitute a linear combination of eigenvectors representing the same energy level.

3.2.2 UCCSD Ansatz

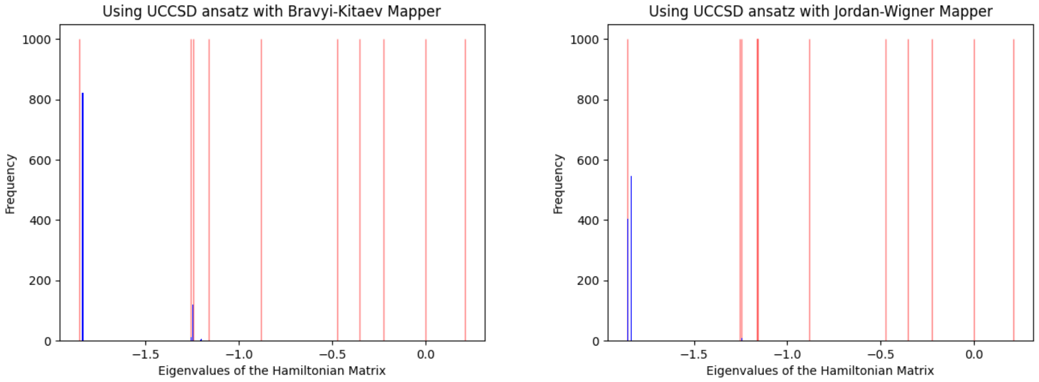

Figure 3.3: Using the UCCSD ansatz for 1000 VQE Runs

The UCCSD ansatz produced states with almost perfect overlap with the ground state 82% of the time with Bravyi-Kitaev, and 98.7% of the time with Jordan-Wigner as shown in Figure 3.3. It is noteworthy that the Jordan-Wigner runs reached the true ground state 44% of the time but more often got caught in the slightly sub optimal state at -1.837 54% of the time. Similarly to the Hardware-Efficient ansatz, the Bravyi-Kitaev mapping has a higher probability of being caught nowhere near the ground state, and in some cases not even below the first three energy states.

3.2.3 Overall

The Bravyi-Kitaev mapper consistently performed worse than the Jordan-Wigner, both in its inability to fully reach the ground state but also in the proportion of results that were far off. The random initialisation has highlighted the pitfalls of poor starting parameter choices with VQE and shows the possibility of getting caught in erroneous locations far from the ground state. However, it must also be considered that this could be due to the limitations of the optimiser. The time taken to run each of these four experiments was about 4 hours each which was significantly less than the AQC-PQC experiments. The main reason that this was so quick was that the Qiskit implementation is mature and has been worked on and improved across numerous releases [21] as discussed in Section 2.5.

Chapter 4

AQC-PQC Analysis

This chapter explores the testing of the AQC-PQC Python implementation using the null-space heuristic as described in Section 2.1.2 across the same UCCSD and ESU2 ansatz families for both the Jordan-Wigner and Bravyi-Kitaev mappings used with VQE (see Chapter 3). A scoping investigation into the speedup from varying the number of steps is described.

4.1 Method

One of the most important factors in the successful running of AQC-PQC is that the starting state is the ground state of the initial Hamiltonian (the tensor product of an X gate on each qubit). For the approach to be effective, the ground state of the initial Hamiltonian must be well prepared and constitute a good approximation to the true ground state. The Hardware-Efficient ansatz did not have obvious starting parameters for the ground state of H_0, so VQE was first used to generate starting parameters for AQC-PQC. To begin each run the program iteratively performed VQE until the difference between the resulting ground state approximation and the true ground state was below a predefined threshold: 1×10^{-6}. For UCCSD, early experiments yielded poor results due to the difficulty of finding a suitable ground state for the initial Hamiltonian with VQE. As such, I added an RY gate to each qubit at the end of the ansatz. If the initial thetas for the UCCSD part of the new ansatz are all set to 0, there will be no effect on the starting state. If each theta on the RY gates is set to $\frac{\pi}{2}$, it gives the ground state of the initial Hamiltonian. As shown in Figure 2.3, there are two X gates at the start of the ansatz, for RY gates acting on qubits with X gates acting on them at the start, setting the parameter to $\frac{-\pi}{2}$ has the same effect as setting them to $\frac{\pi}{2}$ with no initial X gates.

I was not just interested in exploring how effective AQC-PQC can be compared with VQE, but also in exploring what happens from different starting parameters or even an incorrect starting state. In the case of incorrect starts, I was interested to discover if the method could ever reconcile towards the true ground state by the time the final Hamiltonian was reached. For the following results, each mapper and ansatz pairing was run twice: once with a close approximation to the initial ground state (with initial constraint), and once with a single VQE run to approximate the initial Hamiltonian's ground state (no initial constraint). This was a deliberate decision to allow for VQE to converge at the wrong point. As mentioned in the Molecule Preliminaries section of the Background (Section 2.6), and shown in Figure 2.7, different parameter vectors can all give the same correct starting state. Kolotouros et al [11] concluded that the path taken could have a significant effect on the result, and so the no initial constraint experiments were designed to explore this.

Overall, eight distinct experiments were performed, testing the same mappings and ansatz families both with and without the initialisation constraint. Each AQC-PQC experiment run had 25 iterations of 20 steps, which took on average 9 hours each with a large variance in time taken between methods. The code used was my refactoring of Kolotouros' work [11] and my additions to facilitate data collection across multiple iterations. This was in conjunction with my Chemical Problem class that allowed easy construction of the Hamiltonians and ansatze with chemistry problems. The choice of 25 runs and 20 steps was for efficiency and is elaborated upon in Section 5.3.1. The algorithm code is based on the Qiskit library in Python and uses the SLSQP optimiser throughout. The ability to collect more detailed data across each run has allowed for a more telling visualisation of the results as discussed in Section 4.2. In this section there are segmented bar graphs that represent the initial and final states of each run so as to more easily show the moving of probability mass to or from the ground state.

Interesting examples were taken from the dataset and analysed in more detail. The proportions of the result's mass in each energy state was calculated using NumPy and Equation 3.1 for each eigenvector. In the case of degenerate states, the probability mass for eigenvectors corresponding to the same eigenvalue were summed. In the cases where the initial constraint led to the same answer repeatedly, only the starting and final states have been shown for one of the runs in each case. Equation 3.1 was also used to determine the mean probability of finding the ground state, as shown in the GS probability column of Figure 4.1.

4.2 Results

Method	**Mean**	**Median**	**SD**	**Min.**	**Max.**	**GS Probability**
BK ESU2 No Init. Cons.	-1.296	-1.250	0.235	-1.837	-0.935	0.144
BK ESU2 Init. Cons.	-1.250	-1.250	0.000	-1.250	-1.250	0.000
BK UCCSD No Init. Cons.	-1.140	-1.158	0.134	-1.387	-0.904	0.017
BK UCCSD Init. Cons.	-1.233	-1.233	0.000	-1.233	-1.233	0.000
JW ESU2 No Init. Cons.	-1.373	-1.253	0.243	-1.857	-1.133	0.214
JW ESU2 Init. Cons.	-1.406	-1.256	0.274	-1.857	-0.894	0.279
JW UCCSD No Init. Cons.	-1.837	-1.837	0.000	-1.837	-1.837	0.966
JW UCCSD Init. Cons.	-1.837	-1.837	0.000	-1.837	-1.837	0.966

Figure 4.1: Comparison of each mapper and ansatz family with and without initialisation constraints. Units: Hartrees

The results for the experiments that used identical starting parameters (UCCSD Initial Constraint) demonstrate that the algorithm is largely deterministic as expected, however the sample size and use cases are not large enough to guarantee that any randomness present in classical optimisers would not cause deviations. The standard deviation reported for the Jordan-Wigner ESU2 initially constrained run (0.274) shows the importance that the starting state can have. Both the Bravyi-Kitaev and Jordan-Wigner mapper experiments had the same threshold and displayed a variety of starting parameters and final parameters. There is noticeably worse performance with Bravyi-Kitaev as shown in the ground state probability column. This might be explained by a more irregular parameter space meaning that the ground state of an intermediate Hamiltonian cannot be reached, causing the mass to jump up to the next energy state. Another explanation might be a smaller spectral gap when using Bravyi-Kitaev that was repeatedly breached. This could be determined by looking in more detail at the state vectors and Hamiltonians at each step.

The VQE initialisations when there was no initial constraint had a large range of error, with the worst cases being further from the ground state than the worst cases in the VQE runs. This is also expected because a poor starting state from the VQE initialisation may be exacerbated by further deviation during the AQC-PQC run. However, this does not say anything about the success of AQC-PQC as an algorithm, because the initial requirement that the system starts in the initial Hamiltonian was violated. As mentioned above, the intention of these runs was merely to explore how the process behaves with imperfect starts.

4.2.1 Bravyi-Kitaev

4.2.1.1 ESU2

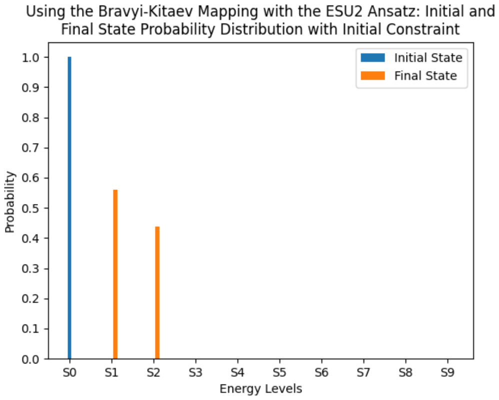

Figure 4.2: A comparison between the respective energy levels at the initial Hamiltonian and the final Hamiltonian

The experiments with the initial constraint all yielded near-identical results. Figure 4.2 shows a representative example of the state changes from start to finish for these runs. The results that start with the parameters described in the Method Section (Section 4.1) indicate that the setup works correctly. Despite starting correctly, the probability mass shifted away from the ground state and ended approximately halfway between the first and second excited states. This could mean that at some step in the process, the ground state was unreachable by the ansatz and so all of the mass shifted towards the first excited state instead. The reason for the subsequent decomposition into the second excited state could be that the energy gap between the two became small enough at one point that some of the probability mass jumped up into the higher state. This could be mitigated by increasing the number of steps, increasing ansatz expressibility or perhaps by running the experiment without the null-space heuristic as that may have limited ansatz expressibility.

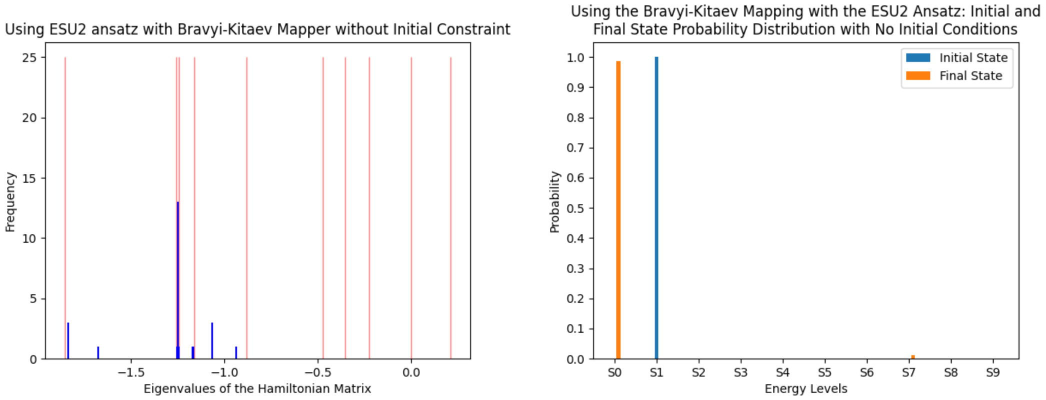

Figure 4.3: The experiment run with no initial constraints and an example of one of the successful cases

Surprisingly, there were three cases of the 25 iterations like the one shown in Figure 4.3 that were able to reach the ground state but only when starting from the first excited state. The most likely explanation is that the topology from the first excited state at the point in the parameter space was conducive to it falling into the lower energy state. Higher energy states do not necessarily have to be turning points and as such a saddle point could easily be passed by the optimiser and continue down to the minimum at the ground state. This can be evaluated by observing the Hessian at the first excited state. If the Hessian is positive definite then it is a local minima which would contradict the hypothesis. If the Hessian is indefinite then it is a saddle point which would support the explanation. There is also the chance that when approaching the spectral gap from the other side it accidentally jumps into the ground state instead of staying in the first excited state. This would be a case of the spectral gap breach working positively. This is expected to be less likely to occur if the step number increases. The final explanation could also be ansatz expressibility or null-space heuristic causing the first excited state to be unreachable and leading the optimiser back down to the first state.

4.2.1.2 UCCSD

Similarly to the Hardware-Efficient runs, all of the probability mass has slipped away from the ground state into higher energy states as shown in Figure 4.4. It seems that at some time step the ground state became unreachable causing the entire mass to shift to the next energy state, similarly to in the ESU2 experiment. Subsequently, it also

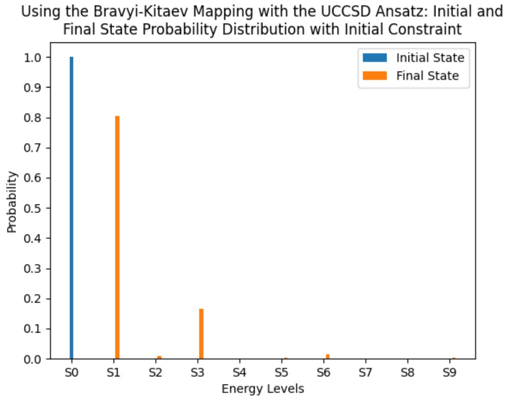

Figure 4.4: A comparison between the respective energy levels at the initial Hamiltonian and the final Hamiltonian

seems that small energy gaps caused some of that mass to slip further away to the later energy states. None of the experiments were able to yield an even partially correct result. The most successful cases occurred when the initial constraint was removed. This was unusual as both of the starting values were within the threshold for those parameters to have been found by the VQE when enforcing the constraint. This is likely a result of such a small sample size and I assume would not be the case if larger experiments were carried out. The RY additions are again shown to work successfully in producing the ground state of H_0.

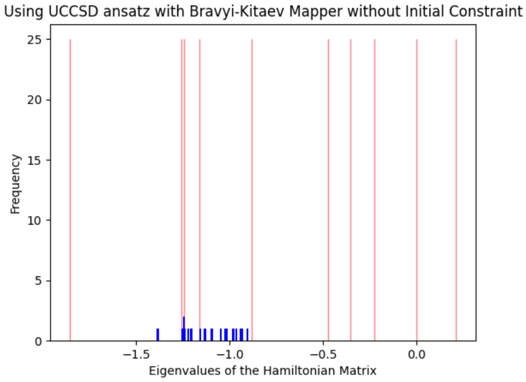

Figure 4.5: The experiment with no initial constraint

As expected, removing the constraint yielded a more diverse array of results as shown in Figure 4.5. The best case showed improvement and ended up at least 21.7% in the ground state.

4.2.2 Jordan-Wigner

4.2.2.1 ESU2

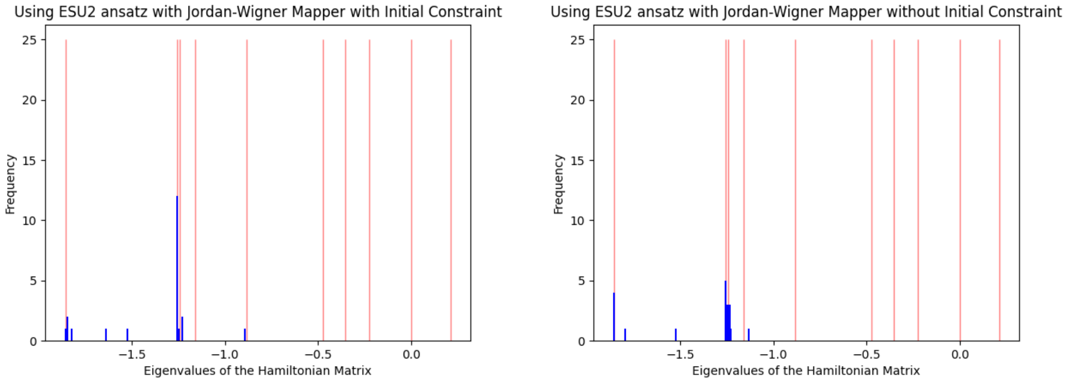

Figure 4.6: The results of the experiment with and without the initial constraint

For some reason, there was room for a large variety of results, even in the supposedly constrained experiment as shown by the left graph in Figure 4.6. This was also the only experimental setup that was able to fully reach the final ground state across those carried out. In most cases, the state seemed to be caught at the first energy level and stay there which most likely implies a small energy gap that was violated due to the small number of steps.

The experiment with no initial constraint performed worse on average, as expected. A higher proportion of the runs gave almost perfect results and the best run of the unconstrained experiment was even marginally better than the best of the constrained experiment. Conversely, the worst result was found in the constrained experiment which should not have been the case given the unconstrained experiment could start many energy levels above the ground state. However, this could be put down to a small sample size that is not representative of the underlying success probabilities.

Surprisingly, the true ground state was reached more often in the unconstrained experiment, some examples of which are shown in Figure 4.7. From inspection of some of the cases, it seems as though when starting in the first energy level or a mix of both first and ground energy levels, the algorithm was able to reach the ground state of H_1 multiple times. Obviously, the constrained version of each of the experiments were unable to explore these cases. These examples clearly highlight how important starting parameter choices are in respect to topology and proximity to final ground state parameter vectors. This means that even when starting at a higher energy level, if there is more conducive topology to getting to the ground state at one or more of the intermediate steps, that can have a bigger impact on final success than starting at the ground state and having to navigate a more irregular parameter space over the course of the algorithm. This is only expected to be the case with a smaller than ideal number of steps, since in cases with very small λ, topology would be expected to have much less of an impact as you should always start close to the ground state for the next step.

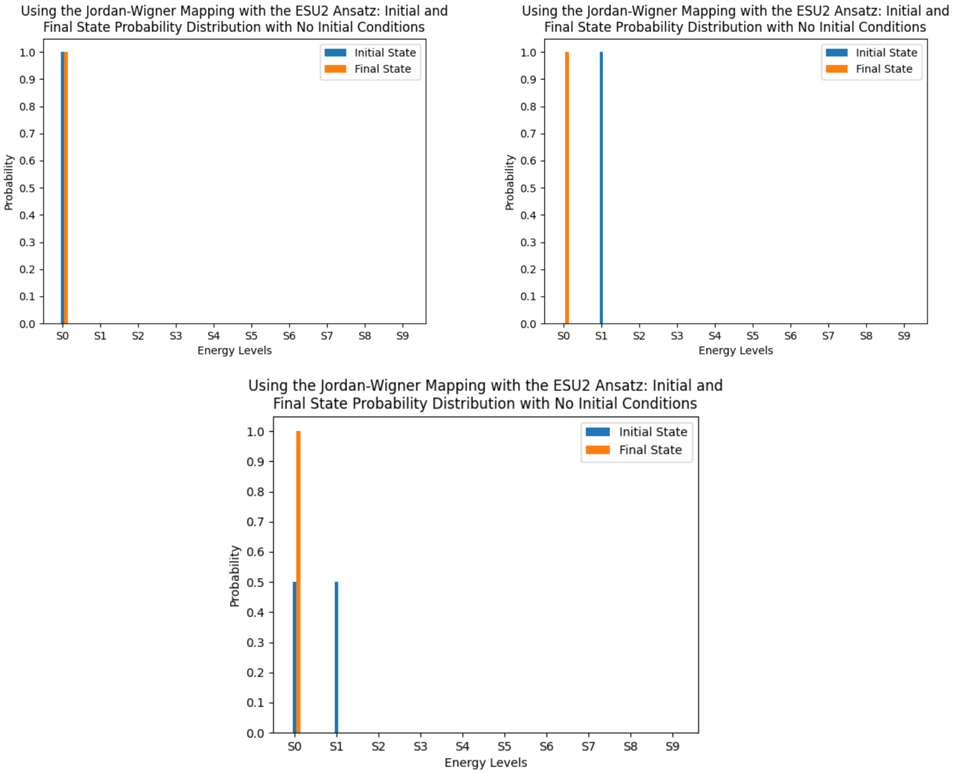
Using the Jordan-Wigner Mapping with the ESU2 Ansatz: Initial and Final State Probability Distribution with No Initial Conditions

Using the Jordan-Wigner Mapping with the ESU2 Ansatz: Initial and Final State Probability Distribution with No Initial Conditions

Using the Jordan-Wigner Mapping with the ESU2 Ansatz: Initial and Final State Probability Distribution with No Initial Conditions

Figure 4.7: Three examples of the five that achieved good results in the non-initially constrained experiment

4.2.2.2 UCCSD

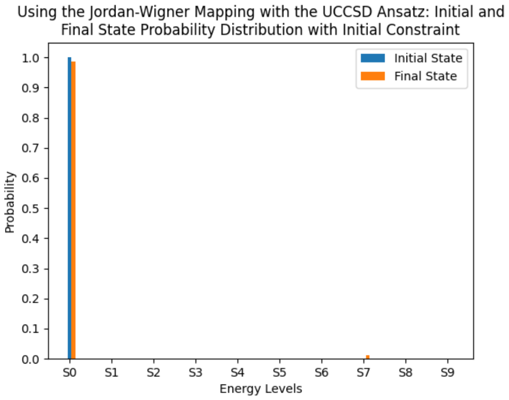

Figure 4.8: A comparison between the respective energy levels at the initial Hamiltonian and the final Hamiltonian

Figure 4.8 is a representative example of what happened on all 50 runs using the Jordan-Wigner mapping on the UCCSD ansatz, both with and without initialisation constraints. These experiments showed that AQC-PQC can work well consistently. Almost all of the probability mass remained in the ground state throughout, with a tiny bit of mass slipping all the way to the seventh energy level. It would be interesting to explore what the cause of the tiny slip was, but it is most likely the result of an imperfect optimisation step that exploits the tolerance threshold. A significantly less likely cause would be repeated slipping of probability mass down to the seventh energy level from multiple energy gap violations. Similarly to the other cases of slipping, it would be useful to plot the energy state probability masses at each point, to determine if the cause was multiple slips from a small spectral gap or if there was just one parameter shift mistake caused by an optimisation error. It is also possible that the seventh energy state may have been close to the ground state in the parameter space which might explain a direct jump.

The fact that both experiments yielded near-identical results suggests that the parameter space was regular at each step and that the spectral gap remained large enough throughout. This encouraged me to explore the parameter changes more deeply, the results are shown in Figure 4.9. As mentioned, the null space heuristic was employed for each of the experiments which limited the changing of parameters that seemed the least significant as detailed in Section 2.1.2.

Figure 4.9 shows that the RY parameters (that were only included to facilitate the starting ground state preparation) were overused. The first three parameters (0..2) represent the UCCSD original part of the ansatz, the final four parameters (3..7) are the RY gates added on. The difference made on the UCCSD parameters was negligible which might imply that the ansatz was not used to its full potential. This could explain its inability to fully finish in the ground state. The results of the VQE experiments show that the final ground state should be reachable with each ansatz under the Jordan-Wigner mapping, which would imply that the heuristic had a negative effect. If it is the case that the original UCCSD part was not being used effectively, then the impact that the four RY gates were able to have was much larger than expected. If true, this is surprising as

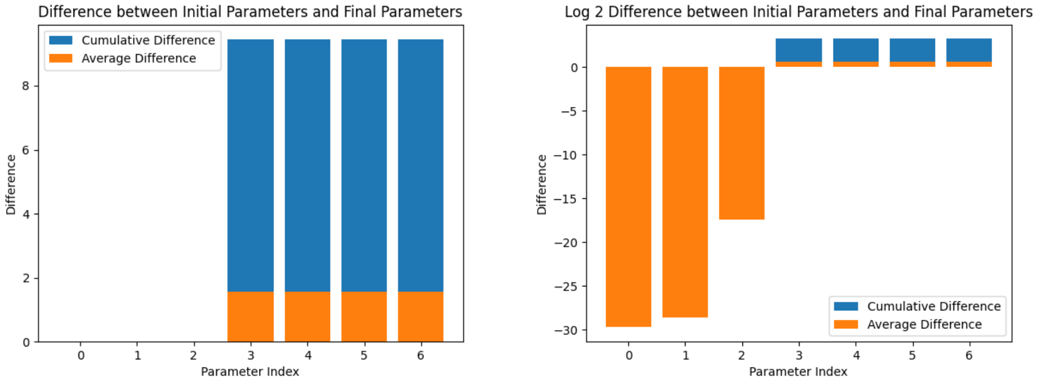

Figure 4.9: The mean and cumulative parameter changes between the start and final θ vectors

the Hardware-Efficient ansatz performed poorly across the experiments with respect to accuracy. It would be worth investigating an ansatz even simpler than the Hardware-Efficient one currently used. Alternatively, I may be significantly underestimating the impact those tiny UCCSD changes may be having. It would be worth carrying out the experiments without the null-space heuristic and performing the same test again to determine how much impact it was having.

4.3 Scoping Investigation

Whilst writing up the limitations of this work and the justifications for the number of steps chosen, I had an idea that would imply a speedup for a greater number of steps within a certain range. This investigation had the greatest potential impact out of those discussed in Section 5.6 for subsequent works. Thus I performed initial testing to probe its feasibility within the limited time I had left.

4.3.1 Theory

The theory that prompted this idea stemmed from the monotonic assumption that the magnitude of the parameter shift to the new ground state decreases as the step size λ decreases, or in layman's terms the shorter the step, the closer you already are to the ground state of the next Hamiltonian. The hypothesis is that shorter parameter shift distances would decrease the optimiser's runtime more than the extra time taken for more steps would increase runtime. This would only apply below an upper bound as there is a constant cost applied to each new step which would start to outweigh small optimiser performance increases at higher step numbers. There would also be a lower bound, as there is likely to be another energy state closer than the ground state that the optimiser would erroneously choose much faster. I estimate that the relationship between distance to the new ground state and the number of optimisation iterations is super-linear due to the complexity of the parameter topology, as it is unlikely that the optimiser would follow a straight path to the new ground state. Even if there were a linear relationship that meant the runtime stayed constant, the improved accuracy from a higher step number would be advantageous.

4.3.2 Results

I ran one iteration of each of the four scenarios and timed them with a loose initialisation constraint but without the null-space heuristic. I tested one set with 50 steps and the other with 20. The experiments were not done with a large enough sample size to form the basis of evidence that supports the hypothesis, so it should be noted that these represent only a promising indication that warrants further study. The results are shown in Figure 4.10:

Mapping	Ansatz	Result	GS Probability	50 Time	20 Time	Change(%)
BK	ESU2	-1.8370	0.966	14:03	21:34	53.5
BK	UCCSD	-1.2550	0	03:52:00	04:21:22	12.7
JW	ESU2	-1.8572	1	13:17	20:09	51.7
JW	UCCSD	-1.8180	0.935	01:18:13	01:37:21	24.5

Figure 4.10: Results of the scoping investigation. Units: Hartrees

Each of the experiments with 20 steps took longer than the experiments with 50 steps, despite having 2.5× as many steps. The improvement was most apparent with the ESU2 ansatz. This could imply that an increase in parameter shift magnitude has a greater increase on the number of optimisation iterations required for ESU2 than UCCSD.

The success of each case when run with 50 steps and no null-space heuristic, could lend support to my assumptions that some of the failure cases earlier in this chapter could be improved upon by increasing the number of steps. If the accuracy is improved upon further study, that would imply less dependency on the starting parameters. The fact that the Jordan-Wigner UCCSD case performed worse with 50 steps could be due to the lack of initialisation constraint, but that would be surprising given the initialisation constraint seemed to make no difference in the experiments detailed in Section 4.2.2.2. The single data point is not enough to draw conclusions from, but if that trend continues for a larger study, the causes should be investigated. I explore how this theory could be continued in Section 5.6.1.

Chapter 5

Discussion

5.1 Scalability and Quantum Device Implementation

There has already been a lot of research dedicated to improving the practicality of using quantum computers to estimate ground state energies. These include areas like ansatz efficiency improvement, noise mitigation techniques and compaction of the molecule's representation. An example of this is Moll et al [15] who exploited block diagonality in fermionic Hamiltonians to reduce the required qubits by a factor of at least two. This involves a classical computation carried out in advance. McClean et al [13] made use of perturbative gadgets that reduce the Bravyi-Kitaev transformed Hamiltonian into a two-local implementation for Adiabatic Quantum Computing. Fedorov et al [8] conducted a survey on VQE-implemented advancements that reduce the problem's complexity, some of which represent scalable methods that can be applied to AQC-PQC as well (as with any AQC developments). The majority of these refer to the effects of noise on the accuracy of computation which is outside the scope of this project, but bode well for future applications of the AQC-PQC method on a quantum device.

The data in Sections 3.2 and 4.2 shows that the Bravyi-Kitaev mapping gives worse results in exchange for a more efficient implementation on a quantum device. This trade-off is completely one-sided when classically simulating AQC-PQC but must be taken into consideration when discussing real-world usage and scalability. If a quantum device were to exist in the future that has full locality, the results of this project would indicate that the Jordan-Wigner mapping should be used for chemistry problems. It seems more realistic however that the Bravyi-Kitaev mapper must be used instead on NISQ devices. It was shown in this project that the AQC-PQC algorithm can be run successfully using the Bravyi-Kitaev mapper but that multiple attempts may be required to find the ground state. If, in future studies, the Bravyi-Kitaev mapping is used and found successful with a larger number of steps, there may be no trade-off element involved and the Bravyi-Kitaev mapping would be unanimously favoured for its ease of implementation.

The qubit number required for increasingly large molecules is not linear when using the Qiskit implementation. Molecular hydrogen used four qubits and three parameters (excluding the RY gates) when implementing UCCSD. Lithium Hydride with 2 nuclei and four electrons used thirteen qubits and 104 parameters when implemented in Qiskit. The scale factor is quartic in the number of spin orbitals for UCCSD implementations [10]. This kind of scaling would completely rule out any near term implementations of quantum eigensolvers on meaningfully large molecules. Kandala et al [10] was able to use mathematical frameworks and physical principles to derive a six qubit implementation of Beryllium Hydride (three nuclei and six electrons) using ”a compact encoding of fermionic Hamiltonians”. This method used Hardware-Efficient ansatz in its approximations. The scaling factor for the Hardware-Efficient ansatz is likely to be worse than quartic for large molecules if sufficient expressibility is to be reached, unless similar methods of compacting the problem can be applied on a larger scale.

The other main constraint on the effectiveness of AQC-PQC in practice is the time taken to perform a run. There would be a speedup when implementing the algorithm on a quantum device as the classical simulation by comparison is slow, and would be impossible for large molecules. It is important to investigate the proportion of the runtime dedicated to the classical and quantum parts respectively, to better assess the impact that the speedup would have. From there, estimates of the expected running time with a quantum device could be made, which would better assess the algorithm’s feasibility in practice. It is reasonable to assume that the speedup would be significant. From frequent observations of the methods that the processor was running while using Google Colab, it is clear that the majority of the processing occurs in the optimisation section.

Finally, it is important to mention that this project was simulated in a noiseless scenario. This is a crucial factor to take into account when evaluating the scalability of the method. The scope of this project only included the evaluation of the method on a noiseless device. Further study would need to be carried out to investigate the effects of noise on the algorithm.

5.2 Ansatz Expressiveness and Mappers

The data shows that the success of the algorithm does not just depend on the starting state vector for the initial Hamiltonian, but can depend on the starting parameters that give that state. This is because the path taken is important [11]. In some experiments, it seems the energy levels of the state vectors in that part of the parameter space were unreachable. This caused the whole probability mass to shift up or down accordingly. This is directly correlated to the expressiveness of the ansatz, however, this could have been exacerbated by the use of the null-space heuristic. Thus, the expressiveness of the ansatz with respect to all of the intermediate Hamiltonians H_λ, not just H_0 and H_1, is important to consider when choosing an ansatz in future work.

5.3 Limitations

5.3.1 Time and Processing Power

For this study, I chose 25 runs as a trade-off between time and maximising sample size. It was also useful to investigate how the algorithm would behave with a fewer than ideal step number, as large-scale implementations of this algorithm may be unable to implement with the full polynomial step number required to guarantee the minimum energy gap would not be violated. Each experiment of 25 runs took an average of 9 hours to complete on my laptop, which was a substantial operating constraint.

I explored using Google Colab but there was not much difference in runtime and it occasionally would time out. I also looked into using the GPU cluster offered at the university but was informed that the parallel processing benefit from this is not noticeable for such a small number of qubits. Due to the exploratory nature of this project, I found that it was often necessary to repeat the experiments with some code tweak or alternate data collection method to better understand what was happening. This was one of the most constraining factors on the number of iterations I used. I estimate I have performed the data generation for AQC-PQC in its entirety at least five times. New iterations were normally due to me being curious to explore a new facet of the algorithm, such as the parameters at the start and the end, the state vectors produced or the statistical analysis performed across all of the final results. This learning process relating to which factors were important to take note of brought me closer and closer to the deadline. This has limited my ability to generate a large data set of a suitable sample size with the latest version, but increased my confidence in the method by eliminating several areas of concern.

Another cause of the time limitation was due to the steep learning curve from my not having any prior experience in quantum computing or the chemistry field to which this project relates. This impeded my ability to work effectively on the project in the first semester. Afterwards, my productivity and working efficiency increased as the second semester went on, but this limited the time that I had to perform the most recent experimental setups.

Given the opportunity to continue this work beyond the scope of the dissertation, I would perform the experiments again on a larger scale and keep track of each of the intermediate Hamiltonians, state vectors, and Hessians to better understand what is happening at each step.

5.4 Evaluation of Methods Used

The study was a useful indication of what might happen when using a less than ideal experimental scenario. The reason I describe it as less than ideal is that many changes were required to be made to accommodate the time limitation. These changes were expected to have had detrimental effects on accuracy. Below I list some improvements to the methods described in Sections 4.1 and 3.1 that could be made to improve accuracy.

The first improvement would be not using the null-space heuristic, which indicated a detrimental effect on the UCCSD ansatz in this study. I believe that an improved version of the study should be carried out without the heuristic.

Another improvement would be collecting many more than 25 results for each AQC-PQC experiment. If there had been more time I would have generated at least 100 iterations for each. Ideally the sample size would match that of the VQE method's. The iteration number limits the confidence of conclusions one can draw about a more generalised probability of success for the algorithm. However, it was sufficient for exploring some of the behaviours of the algorithm and demonstrating that it can work in this problem-specific context. I therefore evaluate that the project was satisfactory for establishing the viability of the method in the problem space, but that sample size for the AQC-PQC experiments was too small to derive confident conclusions beyond indications, of the average performance compared with VQE.

An improvement that could be made that was not limited by time would be maximising the expressibility of the ansatze to allow for more energy states to be reached. This would mitigate the failure cases involving the entire probability mass shifting down that are supposed to be due to poor ansatz expressibility for the ground state of some intermediate Hamiltonian H_λ. This would likely come at the cost of increased parameters which would increase the time complexity of the experiment. This may not prove necessary if improvements are found by increasing the step count or removing the null-space heuristic.

Finally, the ground state preparation method for the Bravyi-Kitaev experiments was poor. The results display that the method was able to work better without the constraint ground state preparation method, which should not have been the case. This could be down to irregular variations across a small sample size which I am inclined to believe as some of the successful runs should have been found via the constrained initialisation. A study with a larger sample size should be able to determine this more accurately.

5.5 Proposed Algorithm Amendments

5.5.1 Extended Search Space with the Optimiser

If the spectral gap is small at a point there is a chance that the parameter shift that the optimiser finds is for the first energy level instead of the ground state. To mitigate this, a heuristic could be applied that will span further than the first satisfying point, to find the next one (within a certain range limit) and compare the Hessians to evaluate which is a more pronounced turning point. Note that this assumes non-degeneracy. If some or many of the values of the Hessian are close to zero, this might imply that the point chosen is less likely to be the ground state, and more likely a saddle point at the first energy level. This algorithm addition would increase the time complexity significantly as the optimisation is already the bottleneck. It may however be able to provide more certainty that the new energy state found at each point is the ground state. This is by no means a guarantee as there is always the chance that the first excited state is also a turning point. Alternatively, when the two values are found, the energy could be estimated with the Hamiltonian at that step and be used to determine which of the two is the ground state. As mentioned before, there would need to be a limit on the optimisation's range otherwise it would waste needless time for steps where the minimum energy gap is large.

5.5.2 Probability of Success Proxy

This metric would be the result of a supervised learning algorithm. It would use data such as the magnitudes and directions of the Hessians and gradients, and the magnitude of the parameter shift at each step to create a feature template. This feature template would then be trained using the state shifting as a loss function, with no loss implying remaining in the ground state. The supervised examples would consist of algorithm runs that are small enough to classically calculate the eigenvectors and values at each perturbed Hamiltonian to generate the loss function. Once this metric was trained for use as an estimate of the likelihood that the state has moved from the ground state, the same model could be tested for unseen examples. The premise of this would be hoping that chemistry problems have distinctive signatures in their parameter topologies that would follow patterns when coming from the initial Hamiltonian.

5.6 Areas for Future Work

5.6.1 Continuing with the Scoping Investigation

The first thing that should be done is to repeat the scoping experiment from Section 4.3 in a similar way to the main AQC-PQC experiments but with a larger sample size and a variety of step numbers. For a more direct indication that the hypothesis is true, the average number of optimisation iterations for each run should be stored, alongside the chosen parameter shift magnitude at each step. If the average parameter shift magnitude is lower for experiments with higher step numbers, that would support the monotonic assumption made. The number of optimisation iterations should be plotted against the parameter shift vector magnitude for each step across all experiments to determine if there is a correlation, and if so what the nature of the relationship is (linear, polynomial or other). If my hypothesis is found to be true, it would improve the scalability potential for the algorithm because at the moment one of the biggest drawbacks is that in large cases with small spectral gaps, a large number of steps would need to be used to avoid bypassing it, but this theory would indicate that the problem might not be as costly as currently assumed.

5.6.2 Exploring the Hessians

To explore the topology of each point, I would export the state vector, Hamiltonian and Hessian at each time step of a run. The changes to the code would not be long or difficult, but they would require the rerunning of all the experiments which would be time-consuming. As mentioned in the Results Section 4.2 for AQC-PQC, Hessians and gradients would be useful for analysing the topology of the parameter space to ascertain whether some errors in the process are the result of topological issues, or from passing through the spectral gap. It is clear from the data that the path taken is important as some initial parameters that give the initial ground state can reach the final ground state, whereas some cannot, despite resulting in the same state vector.

5.6.3 Alternative Data Visualisation and Success Probability Metric

There are many ways that I would have liked to illustrate some of the theories I had regarding what is happening at certain points, specifically concerning the topology. These could be visualised in future studies by plotting the values of the Hessians at various points to explore if they behave as expected. It would also be interesting to plot the state vector paths at each step to see how the probability mass shifts, particularly on erroneous runs.

5.6.4 Investigation into the Classical vs Quantum Runtime

To better understand the scalability of the algorithm, studies should be carried out to determine the proportion of the total runtime dedicated to the classical parts and quantum parts. If my hypothesis is correct and the classical part is cheap, then the time gained when performing the calculations with a quantum device should be significant. This would significantly increase the feasibility of this algorithm's use in practice.

5.6.5 Step Increase

As discussed, this method was carried out with an insufficient number of steps to guarantee the minimum energy gap would not be crossed. The preliminary results in Section 4.3 suggest that a step count increase would be both faster and provide more accurate results, but more testing is required to investigate the extent and scope of the runtime hypothesis. It can be said with some confidence that increasing the step number sufficiently would improve the accuracy of the results.

5.6.6 Not using Null-Space

Preliminary testing without using the null space heuristic was taking an infeasible amount of time, but it likely would have given better results as indicated by the investigation in Section 4.3. Many of the experiments resulted in an inability to reach the ground state at some point. It was verified by looking at the parameter changes with UCCSD from start to finish that the UCCSD-specific parameters did not change much. This could be due to the null-space heuristic limiting the optimisation search space to exclude them. If that is the case then the ansatz expressibility may have been so significantly penalised in favour of time improvement, that it was unable to reach the ground state at one of the intermediate steps. If this is the case then we can expect some of the failure cases, that had the entire probability mass change energy level, to disappear on a rerun without the null-space heuristic.

5.6.7 Other Molecules

AQC-PQC should be further tested upon other molecules within classically simulatable range to see if the trends found in this study generalise to chemistry problems, or if they were Hamiltonian specific. These molecules could be inspired by other works that were able to efficiently compact larger molecules such as Beryllium Hydride in [10].

Chapter 6

Conclusions

6.1 Research Question

This project was successfully completed in line with the Problem Statement in Section 1.2. The AQC-PQC algorithm was shown to be viable and the behaviour was explored in finer detail. In addition, a scoping investigation was carried out (see Section 4.3) that indicated the number of steps could be increased both to improve the quality of the results and speed up the processing. A set of changes were suggested that could be pursued to improve the accuracy of results for future studies.

6.2 Main Findings

Analysis of the results of the experiments has determined that the method can work in each of the cases tried other than Bravyi-Kitaev with the UCCSD ansatz. This does not imply that using the Bravyi-Kitaev mapping and the UCCSD ansatz cannot work, but that it does not appear to work with this experimental setup. This could suggest it is the least effective pairing of those tested. This might be rectified by trying the changes suggested in Section 5.4. These results can still be improved upon with further analysis and research and there is no indication that AQC-PQC cannot outperform VQE with superior experimental setups that lack the same time restraint. There are also proposed amendments and suggestions for how to perform the algorithm with a smaller step number that would mitigate energy gap breaches.

We have seen that in the cases of both VQE and AQC-PQC, the Jordan-Wigner mapping and the UCCSD ansatz yielded higher probabilities of success than either of their alternatives. The exception to this was the Bravyi-Kitaev AQC-PQC experiments which performed better with the ESU2 ansatz than the UCCSD ansatz. It is suspected that a single spectral gap breach caused consistent failure in the UCCSD implementation, so this would likely not remain the case in further studies of more suitable hyperparameters. The practical implications of using each of the mappers and ansatze were explored which yielded the inverse preference. Thus, which setup would be used in practice on a NISQ device depends on its architecture and noise tolerance. It seems prudent in the early stages that the trade-off between practicality and result quality should be explored more thoroughly.

Regarding the behaviour of the algorithm, it was noted that the path taken, dictated by the starting point, has a sizeable impact on the outcome. As many parameter vectors can correspond to the same state, it can be difficult to know which starting parameter vector will be able to reach the correct solution. However, the difference may be less apparent, or disappear completely with more appropriate starting hyperparameters such as more steps or not using the null-space heuristic.

This project has established AQC-PQCs viability for use in chemistry problems and explored its scalability potential for future uses. The behaviour of the algorithm under limited conditions has been investigated as a projection for restrained cases in the future. The code produced for this paper and ideas that came from it can be used to facilitate future studies and springboard further research into the use of AQC-PQC for chemistry problems.

6.3 Reproduction of Work

All the code used in the project has been made public and open source. There is a github repository that can be found at https://github.com/cramptced/AQC-PQC, which contains all of the code at the different stages used to produce the results referenced in this dissertation. It should be noted that due to randomness in variables such as parameter initialisation and optimisation results, not all experiments performed will match up exactly with what I have recorded. This is especially likely to be the case given the small sample size used in this study. Within this repository there is also a requirements.txt file that can be used by pip to replicate the dependencies through the command:

```
pip install -r requirements.txt
```

Bibliography

[1] Ali Abu-Nada. Quantum computing simulation of the hydrogen molecular ground-state energies with limited resources. *Open Physics*, 19(1):628–633, 2021.

[2] R. Babbush, P. Love, and A. Aspuru-Guzik. Adiabatic quantum simulation of quantum chemistry. *Sci Rep*, 4:6603, 2014. Received 21 May 2014; Accepted 22 September 2014; Published 13 October 2014.

[3] S. B. Bravyi and A. Y. Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210–226, 2002.

[4] 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. *Nature Reviews Physics*, 3(9):625–644, August 2021.

[5] Hans Hon Sang Chan, Richard Meister, Tyson Jones, David P. Tew, and Simon C. Benjamin. Grid-based methods for chemistry simulations on a quantum computer. *Science Advances*, 9(9):eabo7484, 2023.

[6] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 3rd edition, 2009.

[7] R. de Wolf. Quantum computing: Lecture notes. *arXiv:1907.09415*, 2023.

[8] D.A. Fedorov, B. Peng, and N. et al. Govind. Vqe method: a short survey and recent developments. *Mater Theory*, 6(2), 2022.

[9] D. R. Hartree. The wave mechanics of an atom with a non-coulomb central field. part i. theory and methods. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(1):89–110, 1928.

[10] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, and et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549:242–246, 2017.

[11] I. Kolotouros, I. Petrongonas, M. Prokop, and P. Wallden. Adiabatic quantum computing with parameterised quantum circuits. 2023.

[12] A.J. McCaskey, Z.P. Parks, and J. et al. Jakowski. Quantum chemistry as a benchmark for near-term quantum computers. *npj Quantum Inf*, 5:99, 2019.

[13] Jarrod R. McClean, Ryan Babbush, Peter J. Love, and Alán Aspuru-Guzik. Exploiting locality in quantum computation for quantum chemistry. *J. Phys. Chem. Lett.*, 5:4368–4380, 2014.

[14] J.R. McClean, S. Boixo, and V.N. et al. Smelyanskiy. Barren plateaus in quantum neural network training landscapes. *Nat Commun*, 9:4812, 2018.

[15] N. Moll, A. Fuhrer, P. Staar, and I. Tavernelli. Optimizing qubit resources for quantum chemistry simulations in second quantization on a quantum computer. *J Phys A*, 49:295301, 2016.

[16] Y. Nam, JS. Chen, and N.C. et al. Pisenti. Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Inf*, 6:33, 2020.

[17] NIST. 2018 codata value: Bohr radius. `https://physics.nist.gov/cgi-bin/cuu/Value?bohrrada0`, 2019. The NIST Reference on Constants, Units, and Uncertainty. Accessed on 20 May 2019.

[18] W. Pauli. "Über den zusammenhang des abschlusses der elektronengruppen im atom mit der komplexstruktur der spektren". *Zeitschrift für Physik*, 31(1):765–783, 1925.

[19] 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. *Nature Communications*, 5(1):4213, 2014.

[20] Qiskit, 2024. The API documentation for the Qiskit SDK as of 3/2024.

[21] Qiskit, 2024. The webpage for qiskit nature's release history as of 3/2024.

[22] J. T. Seeley, M. J. Richard, and P. J. Love. The bravyi-kitaev transformation for quantum computation of electronic structure. *J Chem Phys*, 137:224109, 2012.

[23] Mark Steudtner and Stephanie Wehner. Fermion-to-qubit mappings with varying resource requirements for quantum simulation. *New Journal of Physics*, 20(6):063010, jun 2018.

[24] Michael Struwe. *Variational Methods: Applications to Nonlinear Partial Differential Equations and Hamiltonian Systems*. Springer Berlin, Heidelberg, 1 edition, 1990. Published: 17 April 2013.

[25] Andreas Trabesinger. Quantum simulation. *Nature Physics*, 8:263, 2012.

[26] Andrew Tranter, Peter Love, Florian Mintert, and Peter Coveney. A comparison of the bravyi-kitaev and jordan-wigner transformations for the quantum simulation of quantum chemistry. *Journal of Chemical Theory and Computation*, 14, 09 2018.

[27] H. F. Trotter. On the product of semi-groups of operators. *Proc. Amer. Math. Soc.*, 10:545–551, 1959.

[28] Wikipedia, 2024. The wikipedia page defining NISQ.

Appendix A

Appendix

A.1 Dictionary

Ansatz A family of functions which differ by the parameters they are given. An eloquent example is choosing an appropriate degree polynomial for curve fitting where the coefficient of each term represents a parameter that can be optimised. As in the example above, if the original Ansatz chosen is poor, you will not get an answer that is correct or even necessarily in the right area.

Anti-commutation Commutativity is the property that the order of operators does not matter and will yield the same result: AB = BA. Anti-commutation is the property that says they must equal the negative of each-other but have the same magnitude: AB = -BA. This is necessary in ensuring the Pauli exclusion principle is adhered to in fermionic creation and annihilation operators.

Aufbau Principle In the ground state, electrons will fill up the lowest energy shell available.

Bohr Radius A physical constant named after Niels Bohr that represents the most probable distance between the nucleus and the electron in a hydrogen atom.[17] Coherent Quantum errors – Gate miscalibration under/over rotation

Control Precision“Dynamic range of field values which a device must be able to resolve in order to embed the intended eigen spectrum to a desired accuracy” [2]

Cost Function A function to evaluate the quality of a result given. This is used in the optimisation stage of variational algorithms through gradient descent and other methods to determine how to improve the subsequent iteration.

Eigen spectrum The set of eigenvalues of a linear operator Eigenvalue – The scalar value by which the associated eigenvector is scaled in magnitude upon application of the linear operator

Eigenvector The vector that changes only in magnitude when a given linear operator is applied to it.

Fermi Mode The mode of a system of fermionic particles. In reference to quantum chemistry the number of modes is the number of electrons in orbitals. (See Mode)

Fermionic Molecular Hamiltonian Operator that determines the energy of a system for a given input state.

Gadgets In the context of Hamiltonians, gadgets are tools which allow the locality of a target Hamiltonian to be reduced. For example, using the log(n)-locality of the Bravyi-Kitaev Transform we can apply gadgets to produce a two-local Hamiltonian.

Hartree Atomic units that are designed for convenience in computational settings. One Hartree is approximately the electric potential energy of the hydrogen atom in its ground state and, by the virial theorem, approximately twice its ionization energy; the relationships are not exact because of the finite mass of the nucleus of the hydrogen atom and relativistic corrections. [9]

Incoherent Quantum errors Environmental factors causing bitflips and depolarisation

Locality The scope of how many qubits in a quantum system can interact with each other. Local interactions are when qubits can only interact immediately adjacently with eachother.

Mode In reference to oscillation, the mode of a system is a pattern of motion in which all components oscillate at the same frequency, phase relationship and in a sin waves motion.

Normal OperatorA matrix that is commutative with it's conjugate transpose

ObservablesHermitian matrices that represent a decomposed part of a Hamiltonian

Pauli Exclusion Principle Two or more identical particles cannot be in the same quantum state at the same time in the same quantum system [18].

Perturbation Theory A method of finding approximate solutions to problems by using the exact solution of related and simpler problems. In adiabatic quantum computing it refers to incrementally adding another Hamiltonian which represents a discretised time slice towards the final problem Hamiltonian.